PLAUSIBLE FUTURE, WRONG CHOICE: STEERING WORLD-ACTION MODELS TOWARD INTENDED FUTURE

Anonymous authors
Paper under double-blind review



Figure 1: **From plausible futures to reliable stage progression. Top (left):** Overlapping task trajectories can induce Stage Confusion: WAMs may follow plausible yet incorrect execution routes. **Top (right):** StageBench systematically evaluates this challenge through controlled overlapping multi-stage trajectories across eight task families and more than 2.2 million demonstration frames. **Bottom:** Visualization of TaskWAM performing reliable long-horizon manipulation across diverse robot embodiments and stage progressions. Visually similar stages are highlighted with matching color borders to indicate potential stage confusion.

ABSTRACT

World-Action Models (WAMs) have emerged as a promising paradigm for embodied generalist policies by jointly modeling future world evolution and robot actions across diverse manipulation tasks. However, representing diverse futures is only the first step: a capable WAM must also generate futures that match the requirements of the current task and execution progress. In long-horizon manipulation where similar observation-action segments occur across different tasks and stages, existing approaches may generate locally plausible yet misdirected rollouts to drift toward wrong task stage, *i.e.*, a phenomenon we term *Stage Confusion*. To address this problem, we propose **TaskWAM**, a closed-loop, stage-aware framework that represents fine-grained task requirements with *Mode Tokens* and tracks execution progress with *State Tokens*, forming a structured task state to steer world–action generation toward intended future. We further introduce **StageBench**, a controlled multi-stage manipulation benchmark with parameterized stage variations and automated evaluation. Extensive experiments on both simulation benchmarks and real-world tasks demonstrate that TaskWAM effectively reduces confusion and achieves competitive performance with state-of-the-art methods. Moreover, TaskWAM provides strong plug-and-play transfer across heterogeneous WAM architectures and robot embodiments, including single-arm robots, dual-arm grippers, dexterous hands, and humanoid robots.

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.

1 INTRODUCTION

Recently, World-Action Models (WAMs) (Bi et al., 2026; Ye et al., 2026; Li et al., 2026; Yuan et al., 2026; Chen et al., 2026a) have shown promise for embodied generalist policies by jointly modeling future world evolution and robot actions, with the goal of supporting diverse manipulation tasks in everyday environments(Zhu et al., 2025; Li et al., 2025; 2026). However, modeling diverse futures is only the first step. Once models are capable of representing multiple possible futures, the next challenge lies in generating future that is consistent with the task requirements and execution context. Existing WAMs typically infer stage progression implicitly from task instructions, observations, and execution histories (Lin et al., 2026; Zhao et al., 2026). This becomes particularly challenging in long-horizon household manipulation, where similar observation-action interactions can correspond to different next stages depending on fine-grained task requirements, *e.g.*, ordering, timing, repetition, or embodiment assignment, as illustrated in Figure 1 (*Top left*).

Consequently, WAMs may generate futures that are visually plausible and physically feasible, yet lead execution to the wrong stage. We refer to this failure as *stage confusion in world-action modeling* (hereafter, *stage confusion*). A few recent WAMs Yang et al. (2026); Cai et al. (2026) have explored longer-term history modeling, but stage confusion is not merely a memory-length problem, leaving two key challenges: (1) similar interaction histories may still require different next stages, making instruction-specific guidance essential for resolving stage ambiguity; and (2) simply extending history introduces substantial redundant context, which can obscure critical execution cues for stage tracking while increasing the architectural burden of WAMs.

To address these challenges, we propose **TaskWAM**, a closed-loop, stage-aware framework that explicitly grounds world-action generation in both task requirements and execution progress for long-horizon manipulation. Specifically, we introduce *Structured Task-State Grounding* with two complementary representations: *Mode Tokens* retrieve mode representations from a learnable mode codebook and ground them with text context to encode fine-grained execution requirements, while *State Tokens* track execution progress through causal temporal modeling of closed-loop action-observation transitions. We further propose *Task-State Conditioned Future Steering*, which steers future prediction by inject the structured task state into world and action generation through attention conditioning and residual flow correction. TaskWAM thus bridges diverse future modeling and intended future generation, offering two main advantages: (i) future generation grounding in closed-loop robot-environment interaction; (ii) plug-and-play steering mechanism applicable across heterogeneous WAM backbones.

Additionally, existing benchmarks (Mees et al., 2022; Liu et al., 2023; Chen et al., 2026b; Nasiriany et al., 2026) has mainly evaluated tasks in relatively independent settings, lacking sufficient data and evaluation protocols to assess this capability. To support systematic evaluation, we further introduce **StageBench**, a controlled multi-stage manipulation benchmark built on LIBERO Liu et al. (2023) and RoboTwin Chen et al. (2026b). StageBench constructs overlapping observation-action trajectories across tasks and execution stages by varying stage compositions and orders, covering visual grounding, sequential reasoning, embodiment coordination, *etc.* Extensive experiments in simulation and diverse real-world robotic tasks demonstrate that TaskWAM significantly reduces stage confusion and improves multi-stage manipulation over state-of-the-art methods. We further validate its generalization on RoboTwin 2.0 benchmark Chen et al. (2026b) and its plug-and-play applicability across heterogeneous WAM architectures. Moreover, our method can be applied across diverse robot embodiments, including single-arm robots, dual-arm grippers, dexterous hands, and humanoid platforms. Our contributions are summarized as follows:

- We identify *stage confusion in world-action modeling*, where WAMs generate locally plausible yet stage-inconsistent futures under overlapping observation-action trajectories.
- We propose TaskWAM, which explicitly grounds world-action generation in task requirements and execution progress through Structured Task-State Grounding and task-state-conditioned future steering for stage-consistent future generation.
- We introduce StageBench, a controlled multi-stage manipulation benchmark that systematically constructs overlapping trajectories through parameterized task-stage variations and provides automated stage-aware evaluation.

- Extensive experiments across simulation and real-world tasks demonstrate that TaskWAM significantly reduces stage confusion and improves multi-stage manipulation across heterogeneous WAM backbones and diverse robot embodiments.

2 Related Work

World-Action Models. Generalist robot learning aims to develop unified policies that can handle diverse manipulation tasks within a single model. Vision-Language-Action models (Black et al., 2025a; NVIDIA et al., 2025) have advanced this direction by scaling language-conditioned robot policies. World-Action Models (WAMs) (Bi et al., 2026; Zhu et al., 2025; Ye et al., 2026; Li et al., 2026; Yuan et al., 2026; Cai et al., 2026; Chen et al., 2026a) further extend this paradigm by jointly modeling future observations and robot actions. Recent studies explores inference-time imagination, world-action information exchange, and more efficient generation schedules (Ye et al., 2026; Li et al., 2026; Yuan et al., 2026), such as DreamZero (Ye et al., 2026) and LingBot-VA (Li et al., 2026). However, existing WAMs still typically learn world-action futures from unified conditioning contexts (Lin et al., 2026; Zhao et al., 2026), leaving stage progression to be inferred implicitly from instructions, observations, and histories. Consequently, a predicted future can remain locally plausible while leading the robot to the wrong stage, *i.e.*, *plausible future, yet wrong choice*. In contrast, TaskWAM explicitly grounds future generation in task requirements and execution progress, targeting stage confusion in long-horizon manipulation.

Historical Context Modeling. Temporal context has been widely used to support long-horizon decision making in language modeling (Dai et al., 2019), navigation (Fang et al., 2019), video generation (Henschel et al., 2025), and robot manipulation (Guhur et al., 2023). Existing robot policies incorporate history through temporal abstraction (Belkhale et al., 2024), visual memory (Fang et al., 2025), retrieval (Shah et al., 2026), or explicit state representations (Chung et al., 2026). However, task instructions and longer histories is insufficient to guarantee that task requirements remain effective throughout execution; policies may still rely on easier visual-action correlations (Lin et al., 2026; Zhao et al., 2026). Moreover, longer execution contexts often contain redundant information, which can dilute critical execution cues while increasing design complexity and computational cost (Sridhar et al., 2026). Rather than extending raw history, TaskWAM extracts progress-relevant evidence from causal action-observation transitions for explicit execution tracking.

Long-Horizon Multi-Stage Manipulation. Long-horizon household tasks are common in daily life and remain a key challenge for reliable robot deployment (Li et al., 2023a). A common characteristic of these tasks is that similar observation-action interactions often recur within the same scene across different tasks and execution stages (Nasiriany et al., 2026). Reliable execution therefore requires robots to distinguish such overlapping trajectories and complete the required stages in demanding order (Mees et al., 2022; Li et al., 2023a). However, prior work has largely learned and evaluated tasks in relatively isolated settings (O'Neill et al., 2024), leaving this capability underexplored. In this work, we target this setting with a stage-aware approach and introduce a controlled multi-stage benchmark for systematic evaluation.

3 TaskWAM Framework

3.1 Problem Formulation

Stage Confusion. For task i, we denote a rollout by $\mathcal{T}_i = \{\ell_i, \xi_i\}$, where $\xi_i = \{(o_t^i, a_t^i)\}_{t=1}^{T_i}$ is segmented into a stage sequence $\{s_i^1, \dots, s_i^{n_i}\}$. At the transition from stage k, let $\mathcal{H}_i^k$ denote the observed execution history and s_i^{k+1} the subsequent stage. Two stage-transition points (i, k) and (j, k') may have similar local histories, $\mathcal{H}_i^k \approx \mathcal{H}_j^{k'}$, but require different subsequent stages, $s_i^{k+1} \neq s_j^{k'+1}$. WAMs may generate a locally plausible future $\hat{s}_i^{k+1}$ yet advances execution to the wrong stage, which we refer to as *Stage Confusion*. It typically occurs in two cases: ❶ *Cross-task confusion*, where the model continued toward a wrong stage from another task, *i.e.*, $\hat{s}_i^{k+1} = s_j^{k'}$, $j \neq i$; ❷ *Within-task confusion*, where the model mistakes similar stages within the same task for one another, leading to stage repetition or skipping, *i.e.*, $\hat{s}_i^{k+1} = s_i^{k'}$, $k' \neq k+1$.

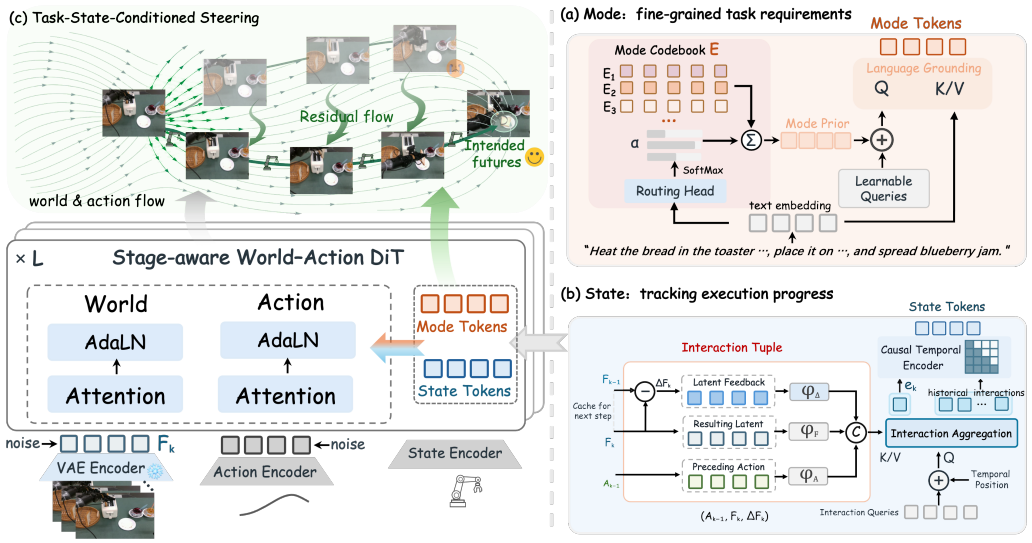

Figure 2: **Overall framework.** TaskWAM augments a World-Action DiT through explicit task-state grounding and conditioning. **(a) Mode Tokens** retrieve mode representations from a learnable codebook and ground them with text context to encode fine-grained task requirements. **(b) State Tokens** aggregate action-observation interactions into a causal interaction history that tracks execution progress. **(c) Task-state-conditioned steering** conditions the world and action streams through attention and modulation, while residual flow guidance steers generation toward intended futures.

Task-State Grounding for Future Generation. Stage progression depends on two complementary task factors: task requirements specify how execution should proceed, while closed-loop interaction history provides evidence of the progress achieved so far. We represent them as an explicit task state $\mathbf{T}_t = [\mathbf{M}; \mathbf{S}_t]$, where $\mathbf{M}$ encodes fine-grained task requirements grounded in instruction ℓ, and $\mathbf{S}_t$ tracks execution progress from history $\mathcal{H}_t$. We condition future world evolution $\mathbf{W}_t^+$ and robot actions $\mathbf{A}_t^+$ on this structured state:

$$(\widehat{\mathbf{W}}_t^+, \widehat{\mathbf{A}}_t^+) \sim p_\theta\big(\mathbf{W}_t^+, \mathbf{A}_t^+ \mid \ell, o_t, \mathcal{H}_t, \mathbf{T}_t\big). \tag{1}$$

Our goal is to explicitly model stage progression in world-action generation rather than leaving it to be implicitly inferred from an unstructured conditioning context.

3.2 Overview

TaskWAM grounds world–action generation in task requirements and execution progress to reduce stage confusion. As illustrated in Figure 2, it consists of two core components: (1) *Structured Task-State Grounding* constructs Mode Tokens through codebook retrieval and text grounding, and State Tokens through interaction aggregation and causal temporal encoding; and (2) *Task-State Conditioned Future Steering* incorporates the resulting task state into backbone conditioning and velocity predictions, steering generation toward the intended future.

3.3 Structured Task-State Grounding

Following Sec. 3.1, we construct Mode Tokens $\mathbf{M} \in \mathbb{R}^{N_m \times d}$ and State Tokens $\mathbf{S}_t \in \mathbb{R}^{N_s \times d}$ to form the structured task state $\mathcal{T}_t = [\mathbf{M}; \mathbf{S}_t]$.

Mode Tokens: instruction-grounded requirements. Tasks can share the same manipulation steps while differing in fine-grained requirements, *e.g.*, ordering, timing, repetition, and embodiment assignment (Liu et al., 2025). To keep these distinctions available throughout execution, we retrieve mode representations from a learnable codebook and ground them in the full instruction context.

Given instruction ℓ, the pretrained T5 encoder (Raffel et al., 2020) produces text embedding $\mathbf{C} = \mathcal{E}_{\text{text}}(\ell) \in \mathbb{R}^{L\times D}$. A routing head then predicts a soft assignment over K_{prior} learnable modes:

$$\mathbf{c}_{\text{pool}} = \frac{\sum_{i=1}^{L} \text{mask}_i \mathbf{c}_i}{\sum_{i=1}^{L} \text{mask}_i}, \qquad \boldsymbol{\alpha} = \text{softmax}(W_2\, \text{SiLU}(W_1 \mathbf{c}_{\text{pool}} + b_1) + b_2), \tag{2}$$

where mask_i identifies valid tokens and $\boldsymbol{\alpha} \in \mathbb{R}^{K_{\text{prior}}}$.

The mode codebook $\mathbf{E} \in \mathbb{R}^{K_{\text{prior}}\times N_m \times D}$ stores N_m learnable slots per mode. We retrieve the corresponding mode slots through soft routing $\mathbf{m}_l^{\text{raw}} = \sum_{k=1}^{K_{\text{prior}}} \alpha_k \mathbf{E}_{k,l}$ for $l = 1, \ldots, N_m$. The retrieved slots provide a coarse execution-requirement prior, serving as a prior for constructing Mode Tokens. We then use these slots as queries over the original language-token sequence to incorporate instruction-specific requirements, such as the precise ordering and timing:

$$\mathbf{m}_l = \mathcal{Q}_{\text{ground}}\left(\mathbf{m}_l^{\text{raw}}, \mathbf{C}\right), \qquad \mathbf{M} = [\mathbf{m}_1, \ldots, \mathbf{m}_{N_m}] \in \mathbb{R}^{N_m \times d}. \tag{3}$$

where $\mathcal{Q}_{\text{ground}}$ performs language cross-attention and projects its output to dimension d. The resulting Mode Tokens combine the retrieved prior with detailed task requirements and are retained across stages, providing persistent task guidance throughout world-action generation.

State Tokens: causal temporal execution tracking. Mode Tokens specify the task requirements but do not indicate which stages have been completed. We therefore construct State Tokens from actions and their observed outcomes in closed-loop robot-environment interaction, rather than encoding raw visual history alone. A causal execution-state encoder first aggregates each interaction into a compact token and then models the sequence over time.

Specifically, for each history slot k, we pair the preceding action chunk $\mathbf{A}_{k-1}$ with the visual latent $\mathbf{F}_k$ of the resulting observation, encoded by the Wan Video VAE (Team Wan et al., 2025), and its change $\Delta\mathbf{F}_k = \mathbf{F}_k - \mathbf{F}_{k-1}$. After modality-specific projection and positional encoding, learnable queries $\mathbf{q}_k$ aggregate these signals through a Q-Former (Li et al., 2023b):

$$\mathbf{e}_k = \mathcal{Q}_{\text{s}}\left(\mathbf{q}_k,\, [\phi_v(\mathbf{F}_k); \phi_\Delta(\Delta\mathbf{F}_k); \phi_a(\mathbf{A}_{k-1})]\right) \in \mathbb{R}^d. \tag{4}$$

where ϕ_v, ϕ_Δ, and ϕ_a are modality-specific projections. A causal temporal encoder then processes the most recent K interaction tokens in execution order:

$$\mathbf{S}_t = \mathcal{E}_{\text{causal}}\left(\mathbf{e}_{t-K+1:t}\right) \in \mathbb{R}^{K\times d}. \tag{5}$$

with $N_s = K$. The positional encodings identify interaction order, while the causal attention mask restricts each slot to current and preceding interactions, allowing State Tokens to summarize how interactions unfolded over time.

During deployment at step t, the robot applies the predicted actions $\mathbf{A}_t$ and encodes the next observation $\mathbf{F}_{t+1}$, the new transition $(\mathbf{F}_{t+1}, \Delta\mathbf{F}_{t+1}, \mathbf{A}_t)$ is then appended, and $\mathbf{S}_{t+1}$ is updated. This supplies subsequent world–action generation with progress information from robot–environment feedback. Detailed settings of both token constructions are provided in Appendix A.1.2.

3.4 Task-State Conditioned Future Steering

Given the structured task state $\mathbf{T}_t$, TaskWAM steers world–action generation toward the intended future through backbone conditioning and residual velocity correction.

Task-State Conditioning. We append $\mathbf{T}_t$ to the key-value context of both streams. World and action tokens can therefore attend to the same task requirements and execution progress throughout generation. To keep the active execution mode effective across backbone layers, we further pool the Mode Tokens into $\bar{\mathbf{m}} \in \mathbb{R}^d$ and apply layer-wise AdaLN modulation (Peebles & Xie, 2023):

$$[\boldsymbol{\gamma}_l^m, \boldsymbol{\beta}_l^m] = g_l(\bar{\mathbf{m}}), \qquad \widetilde{\mathbf{H}}_l = (1 + \boldsymbol{\gamma}_l^m) \odot \text{LN}(\mathbf{H}_l) + \boldsymbol{\beta}_l^m, \tag{6}$$

where g_l denotes a layer-specific modulation network. This carries instruction-specific requirements across generation layers while attention incorporates the updated execution state. The design preserves the backbone's native world-action attention relations, supporting integration across heterogeneous WAM architectures (Sec. 5.2).

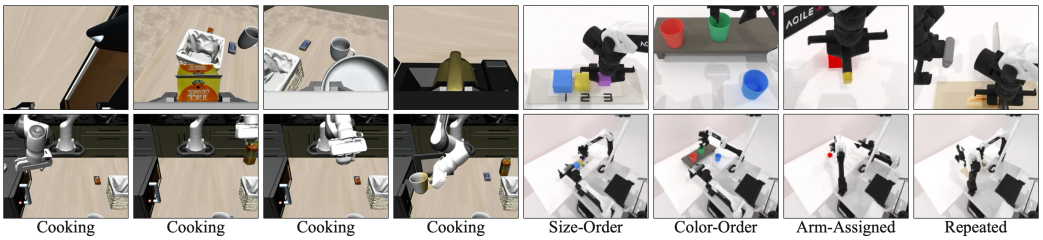

Figure 3: The visualization of task families in StageBench (see Appendix B.1 for more details).

World-Action Flow Steering. Beyond backbone conditioning, we apply task-state-conditioned residual corrections (Black et al., 2025b) to steer world and action generation toward the intended future. For flow-based backbones, let $\mathbf{H}^b_{t,\tau}$ and $\mathbf{v}^b_{t,\tau}$ denote the hidden states and velocity predictions at rollout step t and flow time τ, where $b \in \{w, a\}$ denotes the world or action stream. Each lightweight cross-attention head $\mathcal{R}_b$ uses the corresponding hidden states to query $\mathbf{T}_t$ and predict a velocity residual:

$$\Delta\mathbf{v}^b_{t,\tau} = \mathcal{R}_b(\mathbf{H}^b_{t,\tau}; \mathbf{T}_t), \qquad \hat{\mathbf{v}}^b_{t,\tau} = \mathbf{v}^b_{t,\tau} + \gamma_b \Delta\mathbf{v}^b_{t,\tau}, \tag{7}$$

where γ_b is a learnable bounded scale. Both flows are corrected with the same task state, steering the intended stage progression rather than merely producing mutually consistent world and action predictions. For Cosmos Policy, the residual instead corrects denoised latent predictions under its native EDM formulation.

Training Objective. We train TaskWAM end-to-end using each backbone's vanilla generative objective, *i.e.*, flow matching for flow-based WAMs and EDM denoising for Cosmos Policy. The overall training objective is $\mathcal{L} = \mathcal{L}_{\text{world}} + \mathcal{L}_{\text{action}}$. More implementation details are provided in Appendix A.2.

4 Stage-Confusion Benchmark

To support the systematic evaluation of stage confusion, we construct *StageBench* on the single-arm LIBERO and bimanual RoboTwin 2.0 platforms. StageBench includes diverse kitchen and household scenarios with meaningful multi-stage manipulation routines, comprising 10 task families with visual grounding, sequential reasoning, temporal reasoning, and embodiment coordination, *etc.*, as illustrated in Fig. 3. Additional details are provided in Appendix B.1.

Task design and data collection. Each task family is instantiated into multiple *task variants* built from the same scene, objects, and manipulation stages, with matched stage counts and comparable difficulty. Across variants, we construct different stage sequences by changing their order, timing, arm assignment, or repetition to target *cross-task confusion*. Within each variant, repeated or similar interaction stages appear at different execution positions to target *within-task confusion*. We collect demonstration trajectories through 3D mouse teleoperation and automated data generation, resulting in over 2.20 million frames in total, with approximately 12 hours of manual teleoperation.

Evaluation. We build an automatic evaluator that reconstructs stage-wise execution from simulator states. During each rollout, it records task-relevant object states, contacts, arm actions, and their timestamps, and matches these events against the stage definitions of the corresponding task variant. This allows us to summarize which stages are entered and completed throughout execution, from which we derive task completion, execution route, and stage-level progress. For each task, we manually inspect 100 evaluation results to confirm the reliability of the automatic evaluator.

5 Experiment

5.1 Experimental Setup

Simulation Benchmarks. We evaluate our method in two robotic simulation benchmarks: (1) In **StageBench**, a controlled multi-stage benchmark with 8 task families spanning single-arm and bimanual manipulation, as detailed in Sec. 4. (2) In **RoboTwin 2.0** Chen et al. (2026b), we follow

Table 1: **Performance comparison on diverse long-horizon manipulation tasks of StageBench.** Our method outperforms all baselines in Success Rate (*Suc.*, %↑). Best results are in **bold**.

Method	Kitchen Manipulation				Bimanual Manipulation				Avg. *Suc.*
	Cooking & Cheese	Cooking & Juice	Cooking & Cleanup	Cooking & Breakfast	Size-Order Cubes	Color-Order Cups	Arm-Assigned Hammering	Repeated Bread Cutting	
GR00T N1.7	43	33	9	28	60	28	55	43	37.4
$\pi_{0.5}$	52	54	43	40	**69**	46	75	**57**	54.5
Motus	27	28	27	12	15	10	47	20	23.3
Cosmos Policy	35	53	35	55	18	7	45	13	32.6
w/ TaskWAM	57	61	45	63	29	22	56	31	45.5
Δ	**+22**	**+8**	**+10**	**+8**	**+11**	**+15**	**+11**	**+18**	**+12.9**
FastWAM	59	63	51	48	11	16	51	28	40.9
w/ TaskWAM	**74**	**78**	**72**	**73**	65	**54**	**83**	56	**69.4**
Δ	**+15**	**+15**	**+21**	**+25**	**+54**	**+38**	**+32**	**+28**	**+28.5**

Table 2: **Performance comparison on RoboTwin 2.0 Benchmark (Clean vs. Randomized, 50 tasks).** We show the representative tasks and full comparison results is in Appendix B.2.

Task	Motus		$\pi_{0.5}$		LingBot-VA		LaWAM		AHA-WAM		Fast-WAM		**w/ TaskWAM**	
	Clean	Rand.	Clean	Rand.	Clean	Rand.	Clean	Rand.	Clean	Rand.	Clean	Rand.	Clean (Δ)	Rand.(Δ)
Open Laptop	95	91	90	96	92	94	**100**	**100**	98	97	98	**100**	**100**	**100**
Place Bread Basket	91	94	77	64	97	95	92	85	89	94	91	93	**98**	**97**
Place Bread Skillet	86	83	85	66	**95**	90	90	83	91	87	90	93	94	**98**
Place Object Scale	88	85	86	80	**96**	95	95	88	**96**	87	90	97	**96**	**99**
...								...						
Put Object Cabinet	88	71	80	79	85	87	90	82	91	89	**94**	89	90	**90**
Rotate QRcode	89	73	89	87	**96**	**91**	94	89	90	90	93	89	87	**91**
Beat Block Hammer	95	88	96	93	96	**98**	90	93	**100**	94	99	97	**100**	**98**
Average (%)	88.66	87.02	82.74	76.76	92.90	91.50	92.52	89.48	**93.40**	92.20	91.88	91.78	92.70 (+0.82)	**93.06** (+1.28)

the same experimental settings of (Bi et al., 2026; Yuan et al., 2026; Li et al., 2026), training all models on 2,500 demonstrations collected in clean scenes (50 per task) and 25,000 demonstrations collected in randomized scenes (500 per task).

Baselines. We compare our approach against the state-of-the-art (SOTA) baselines, including WAMs: Motus (Bi et al., 2026), Cosmos Policy (Kim et al., 2026), DreamZero (Ye et al., 2026), Fast-WAM (Yuan et al., 2026), LingBot-VA (Li et al., 2026), LaWAM (Chen et al., 2026a), and AHA-WAM (Cai et al., 2026). We further include classical VLA baselines, including $\pi_{0.5}$ (Black et al., 2025a) and GR00T N1.7 (NVIDIA et al., 2025). We evaluate both final task success and stage-aware progression, including stage-wise completion and Stage Progress Rate (*SPR*), to measure reliable multi-stage execution beyond overall completion.

Implementation Details. We build our framework upon Fast-WAM (Yuan et al., 2026) and Cosmos Policy (Kim et al., 2026), preserving their native inference structures while injecting task-state conditioning. For StageBench, all methods are trained and evaluated under the same data and evaluation settings. For RoboTwin 2.0, we follow its standard training and evaluation protocols. Experiments are trained on 16× NVIDIA H100 80GB GPUs, and all simulation results are evaluated over 100 trials. More details are provided in the supplementary material.

5.2 Main Results

Results on StageBench. Table 1 shows that TaskWAM consistently improves both WAM backbones across all eight task families. On FastWAM, it raises average success from 40.9% to 69.4% (+28.5%), surpassing the strongest standalone baseline, $\pi_{0.5}$, at 54.5%. The gains range from 15–25% on kitchen tasks and 28–54% on bimanual tasks, with the largest improvement on Size-Order Cubes (11% to 65%). TaskWAM also improves Cosmos Policy from 32.6% to 45.5% (+12.9%), demonstrating that task-state grounding provides consistent benefits across backbones and different stage-confusion scenarios.

Generalization to Robotwin 2.0. Table 2 compares different methods on 50 RoboTwin 2.0 tasks under clean and randomized settings. TaskWAM consistently improves the Fast-WAM backbone and achieves highly competitive performance against recent SOTA methods, while obtaining the best result under the randomized setting. These results demonstrate the effectiveness of TaskWAM across diverse manipulation tasks and environmental variations.

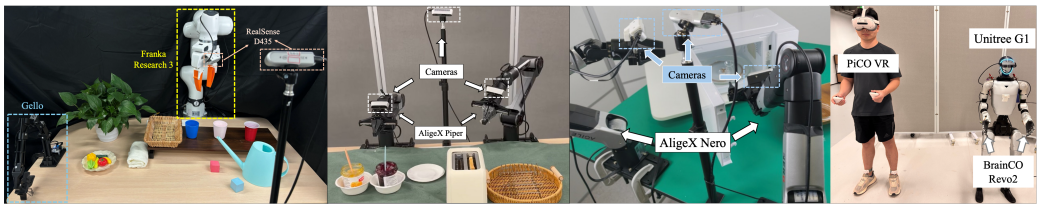

Figure 5: **Real-world robot platform.** (a) Franka Research 3 with third-person and wrist-view RealSense D435 cameras; (b-c) dual-arm AgileX Piper (6-DoF) and Nero (7-DoF) equipped with two wrist cameras and one head-mounted camera, respectively; (c) Unitree G1 humanoid robot equipped with BrainCo Revo2 dexterous hands and operated via PICO VR for whole-body teleoperation.

5.3 Ablation and Analysis

Component Ablation. Table 3 validates the contribution of each TaskWAM component under the same StageBench setting. Starting from the 40.9% baseline, progressively adding Mode Tokens raises *Suc.* to 56.4%, State Tokens further improve it to 61.6%, and residual flow steering yields the full-model result of 69.4%. These consistent gains confirm that each component contributes to intended future generation and reliable stage progression.

Table 3: **Component Ablation.**

Method	Mode	State	Residual Flow	*Suc.*↑
Baseline				40.9
+ Mode Tokens	✓			56.4
+ State Tokens	✓	✓		61.6
Full TaskWAM	✓	✓	✓	**69.4**

Effectiveness of Task-State Grounding. Beyond final success rates, we provide complementary benchmark-level and transition-level analyses. More details are provided in Appendix B.1. ❶ **TaskWAM reduces stage confusion.** *SPR* measures the fraction of all eligible transitions that complete the correct next stage without filtering failed rollouts; higher is better. TaskWAM raises *SPR* from 68.4% to 87.1%, demonstrating more reliable completion of the intended next stage. ❷ **TaskWAM sustains progression across stage transitions.** In the Cooking & Cleanup case study (Figure 4), TaskWAM achieves higher completion across all transitions, with its advantage widening at later stages, indicating less accumulation of stage confusion over the execution horizon.

Table 4: **Stage Progression Rate** (*SPR*, %↑) on StageBench.

Method	*Suc.*↑	*SPR*↑
FastWAM Baseline	40.9	68.4
w/ TaskWAM	69.4	87.1

Moreover, explicit task-state grounding also improves training efficiency: TaskWAM converges faster and achieves lower action loss than FastWAM under the same training setting. The convergence curves and full analysis are provided in Appendix C.3 (Figure 13).

Finally, we verify that temporal ordering and execution history are essential for reliable progress tracking. Additional analysis of State Token design is provided in Appendix C.2 (Table 10).

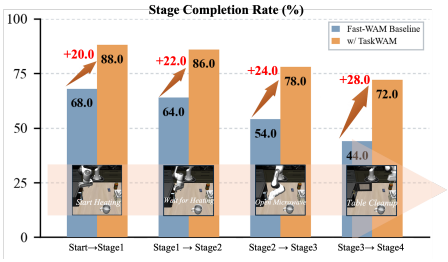

Figure 4: **Stage-wise completion** over long-horizon execution. TaskWAM maintains higher completion across all stage transitions.

5.4 Real-World Experiments

Robot Platform. As shown in Figure 5, we evaluate our method on gripper-arm, dexterous-hand, and humanoid platforms. For each task, we collect 100-150 teleoperated demonstrations using Gello Wu et al. (2024) for gripper arms, TypeTele Lin et al. (2025) for dexterous hands, and SONIC Luo et al. (2026) for the humanoid platform. Visual observations are captured by RealSense D435 cameras, providing wrist-mounted and third-person views.

Multi-Stage Real-World Tasks. We evaluate our method on four multi-stage tasks drawn from three representative real-world scenarios (bottom in Figure1): ❶ **Sequential Cup Nesting**; Breakfast Preparation, including ❷ **Bread Toasting & Spreading** and ❸ **Bread Cutting**; and ❹ **Laboratory Fluid Manipulation**. These settings naturally reuse objects and local interactions across tasks and execution stages, providing real-world evaluation of stage confusion under varying task

Table 5: **Results of real-world experiments.** We report the stage-wise success rate for each task, S_k denotes cumulative success through stage k. Best results are **bolded**.

Task	Setting	Cosmos Policy	FastWAM	TaskWAM	Δ
		$S_1 \to S_2 \to S_3$	$S_1 \to S_2 \to S_3$	$S_1 \to S_2 \to S_3$	
Sequential Cup Nesting	Single-Arm Franka	45.0 → 40.0 → 30.0	55.0 → 45.0 → 35.0	**70.0 → 65.0 → 55.0**	+20.0
Bread Cutting	Dexterous Hands	95.0 → 80.0 → 20.0	**100.0** → 80.0 → 35.0	**100.0 → 95.0 → 70.0**	+35.0
Lab. Fluid Manipulation	Bi-Gripper Nero	65.0 → 40.0 → 40.0	80.0 → 55.0 → 50.0	**90.0 → 70.0 → 70.0**	+20.0
Bread Toasting & Spreading	Bi-Gripper Piper	$S_1 \to S_2 \to S_3 \to S_4 \to S_5$	$S_1 \to S_2 \to S_3 \to S_4 \to S_5$	$S_1 \to S_2 \to S_3 \to S_4 \to S_5$	
		90.0 → 20.0 → 20.0 → 0.0 → 0.0	**95.0** → 25.0 → 25.0 → 10.0 → 5.0	**95.0 → 45.0 → 45.0 → 30.0 → 20.0**	+15.0

requirements and execution progress. Each task contains 2–3 variants, with each variant evaluated over 10 trials. Full details are provided in Appendix B.3.

Baseline. We compare our method with two baseline WAMs including Cosmos PolicyKim et al. (2026) and Fast-WAM Yuan et al. (2026). We report stage-wise completion rates, where S_k denotes cumulative success through stage k, with qualitative stage definitions shown in Figure 1 (bottom).

Results Analysis. Table 5 evaluates whether TaskWAM can preserve reliable stage progression in real-world environments. Across four tasks and diverse embodiments, TaskWAM consistently achieves higher final-stage success than FastWAM, improving by 15–35%. The gains become larger in longer tasks with more stage transitions, where stage confusion accumulates over execution. For instance, in the challenging nearly 2-minute Bread Toasting & Spreading task, both methods achieve similar early-stage completion such as toasting. However, after the long waiting stage, FastWAM often confuses the visually similar retrieval state with earlier stages, reducing final-stage success to 5.0%, while TaskWAM maintains 20.0% completion. A representative failure case involving visually similar stages after long waiting periods is further analyzed in Appendix D.

5.5 Results on Heterogeneous Embodiments

To assess cross-embodiment adaptability, we deploy TaskWAM on four real-world platforms with the same framework while adapting only robot-specific observations and action spaces, as demonstrated in Figure 1 (bottom). As shown by Table 5, TaskWAM consistently improves final-stage success by 15–35% over the baseline and exhibits strong adaptability to heterogeneous embodiments, including single-arm robots, 6-DoF and 7-DoF dual-arm platforms, and humanoid with bi-dexterous hands. Additional details are provided in Appendix B.3.

5.6 Performance-Efficiency Analysis

Table 6 reports task performance and end-to-end inference latency under identical hardware setting. With 30 warm-up and 100 timed runs on a single RTX 5090 GPU, TaskWAM improves *Suc.* from $40.9\% \to 69.4\%$ with only 19 ms additional latency over FastWAM, demonstrating an effective performance-efficiency trade-off validated by real-world deployment.

Table 6: **Performance-efficiency trade-off.** TaskWAM achieves significant gains in multi-stage reliability with minimal inference overhead.

Method	*Suc.*↑	*SPR.*↑	Latency (s)↓
FastWAM	40.9	68.4	0.190
TaskWAM	**69.4**	**87.1**	0.209

6 Conclusion

Reliable stage progression is a critical capability for long-horizon robotic manipulation, especially when different tasks and execution stages share similar observation-action interactions. In this work, we identify *stage confusion* in World-Action Models, where locally plausible yet misdirected rollouts to drift toward wrong task stage, and propose TaskWAM to address this problem. Our key insight is to ground future generation in explicit task requirements and realized execution progress through mode codebook retrieval, causal execution tracking, and backbone-agnostic future steering. To support systematic evaluation, we further introduce StageBench with overlapping multi-stage trajectories and stage-aware evaluation. Extensive experiments in simulation and real-world settings validate TaskWAM across heterogeneous WAM architectures and diverse robot embodiments. In summary, our results highlight explicit stage-aware future generation as an important step toward reliable generalist robot behavior in long-horizon tasks.

AI USE STATEMENT

In this work, we used generative AI tools for checking the consistency of technical writing, and identifying possible typographical or formatting errors in the manuscript. We have not used generative AI tools for developing the research methodology, and tasks involving datasets, experiments, implementation, qualitative data analysis, and the interpretation of experimental results are not applicable to this work. Additionally, we used generative AI tools for improving the grammar and readability of the manuscript. We have reviewed all AI-assisted work. We reviewed all AI-assisted revisions and manually verified that they preserved the intended technical meaning and did not alter the methodology, experimental results, or conclusions. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

We provide the implementation of our method and detailed experimental configurations in the supplementary material. We include model architectures, training settings, hyperparameters, benchmark construction details, and evaluation protocols to facilitate faithful reproduction of our results.

REFERENCES

Suneel Belkhale, Tianli Ding, Ted Xiao, Pierre Sermanet, Quan Vuong, Jonathan Tompson, Yevgen Chebotar, Debidatta Dwibedi, and Dorsa Sadigh. RT-H: Action Hierarchies using Language. In *Proceedings of Robotics: Science and Systems*, Delft, Netherlands, July 2024. doi: 10.15607/RSS.2024.XX.049.

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. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 35101–35113, June 2026.

Kevin Black, Noah Brown, James Darpinian, Karan Dhabalia, Danny Driess, Adnan Esmail, Michael Robert 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. In Joseph Lim, Shuran Song, and Hae-Won Park (eds.), *Proceedings of The 9th Conference on Robot Learning*, volume 305 of *Proceedings of Machine Learning Research*, pp. 17–40. PMLR, 27–30 Sep 2025a.

Kevin Black, Manuel Galliker, and Sergey Levine. Real-time execution of action chunking flow policies. In *Advances in Neural Information Processing Systems*, volume 38, Main Conference, pp. 33383–33407. Curran Associates, Inc., 2025b. doi: 10.52202/085713-1122.

Jisong Cai, Long Ling, Shiwei Chu, Zhongshan Liu, Jiayue Kang, Zhixuan Liang, Wenjie Xu, Yinan Mao, Weinan Zhang, Xiaokang Yang, Ru Ying, Ran Zheng, and Yao Mu. Aha-wam:asynchronous horizon-adaptive world-action modeling with observation-guided context routing, 2026. URL https://arxiv.org/abs/2606.09811.

Jialei Chen, Kai Wang, Kang Chen, Shuaihang Chen, Feng Gao, Wenhao Tang, Zhiyuan Li, Weilin Liu, Zhuyu Yao, Boxun Li, Yuanbo Xu, and Chao Yu. Lawam: Latent world action models for efficient dynamics-aware robot policies, 2026a. URL https://arxiv.org/abs/2606.15768.

Tianxing Chen, Zanxin Chen, Baijun Chen, Zijian Cai, Yibin Liu, Zixuan Li, Qiwei Liang, Xianliang Lin, Yiheng Ge, Zhenyu Gu, Weiliang Deng, Yubin Guo, Tian Nian, Xuanbing Xie, Qiangyu Chen, Kailun Su, Tianling Xu, Guodong Liu, Mengkang Hu, Huan ang Gao, Kaixuan Wang, Zhixuan Liang, Yusen Qin, Xiaokang Yang, Ping Luo, and Yao Mu. Robotwin 2.0: A scalable data generator and benchmark with strong domain randomization for robust bimanual robotic manipulation. In *Forty-third International Conference on Machine Learning*, 2026b.

Nhat Chung, Taisei Hanyu, Toan Nguyen, Huy Le, Frederick Bumgarner, Duy Minh Ho Nguyen, Khoa Vo, Kashu Yamazaki, Chase Rainwater, Tung Kieu, Anh Nguyen, and Ngan Le. Rethinking progression of memory state in robotic manipulation: An object-centric perspective. *Proceedings of the AAAI Conference on Artificial Intelligence*, 40(5):3407–3415, 2026. doi: 10.1609/aaai.v40i5.37337.

Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc Le, and Ruslan Salakhutdinov. Transformer-XL: Attentive language models beyond a fixed-length context. In Anna Korhonen, David Traum, and Lluís Màrquez (eds.), *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 2978–2988, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1285.

Haoquan Fang, Markus Grotz, Wilbert Pumacay, Yi Ru Wang, Dieter Fox, Ranjay Krishna, and Jiafei Duan. SAM2Act: Integrating visual foundation model with a memory architecture for robotic manipulation. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu (eds.), *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 15925–15942. PMLR, 13–19 Jul 2025.

Kuan Fang, Alexander Toshev, Li Fei-Fei, and Silvio Savarese. Scene memory transformer for embodied agents in long-horizon tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

Pierre-Louis Guhur, Shizhe Chen, Ricardo Garcia Pinel, Makarand Tapaswi, Ivan Laptev, and Cordelia Schmid. Instruction-driven history-aware policies for robotic manipulations. In Karen Liu, Dana Kulic, and Jeff Ichnowski (eds.), *Proceedings of The 6th Conference on Robot Learning*, volume 205 of *Proceedings of Machine Learning Research*, pp. 175–187. PMLR, 14–18 Dec 2023.

Roberto Henschel, Levon Khachatryan, Hayk Poghosyan, Daniil Hayrapetyan, Vahram Tadevosyan, Zhangyang Wang, Shant Navasardyan, and Humphrey Shi. StreamingT2V: Consistent, dynamic, and extendable long video generation from text. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2568–2577, June 2025.

Moo Jin Kim, Yihuai Gao, Tsung-Yi Lin, Yen-Chen Lin, Yunhao Ge, Grace Lam, Percy Liang, Shuran Song, Ming-Yu Liu, Chelsea Finn, and Jinwei Gu. Cosmos policy: Fine-tuning video models for visuomotor control and planning. In *International Conference on Learning Representations*, volume 2026, pp. 71531–71552, 2026.

Chengshu Li, Ruohan Zhang, Josiah Wong, Cem Gokmen, Sanjana Srivastava, Roberto Martín-Martín, Chen Wang, Gabrael Levine, Michael Lingelbach, Jiankai Sun, Mona Anvari, Minjune Hwang, Manasi Sharma, Arman Aydin, Dhruva Bansal, Samuel Hunter, Kyu-Young Kim, Alan Lou, Caleb R Matthews, Ivan Villa-Renteria, Jerry Huayang Tang, Claire Tang, Fei Xia, Silvio Savarese, Hyowon Gweon, Karen Liu, Jiajun Wu, and Li Fei-Fei. Behavior-1k: A benchmark for embodied ai with 1,000 everyday activities and realistic simulation. In Karen Liu, Dana Kulic, and Jeff Ichnowski (eds.), *Proceedings of The 6th Conference on Robot Learning*, volume 205 of *Proceedings of Machine Learning Research*, pp. 80–93. PMLR, 14–18 Dec 2023a.

Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 19730–19742. PMLR, 23–29 Jul 2023b.

Lin Li, Qihang Zhang, Yiming Luo, Shuai Yang, Ruilin Wang, Luyao Zhang, Mingrui Yu, Zelin Gao, Nan Xue, Boyu Zhou, Xing Zhu, Mingyu Ding, Yujun Shen, and Yinghao Xu. Causal World Modeling for Robot Control. In *Proceedings of Robotics: Science and Systems*, Sydney, Australia, July 2026. doi: 10.15607/RSS.2026.XXII.016.

Shuang Li, Yihuai Gao, Dorsa Sadigh, and Shuran Song. Unified Video Action Model. In *Proceedings of Robotics: Science and Systems*, LosAngeles, CA, USA, June 2025. doi: 10.15607/RSS.2025.XXI.074.

Yihan Lin, Jiawei He, Shifeng Bao, Chen Zhao, Yang Li, Xiaobo Wang, Yan Wang, Cheng Chi, and Jing Zhang. JEPA-WAM: Learning vision-language-action policies with joint-embedding world modeling, 2026.

Yuhao Lin, Yi-Lin Wei, Haoran Liao, Mu Lin, Chengyi Xing, Hao Li, Dandan Zhang, Mark Cutkosky, and Wei-Shi Zheng. Typetele: Releasing dexterity in teleoperation by dexterous manipulation types. In Joseph Lim, Shuran Song, and Hae-Won Park (eds.), *Proceedings of The 9th Conference on Robot Learning*, volume 305 of *Proceedings of Machine Learning Research*, pp. 4975–4993. PMLR, 27–30 Sep 2025.

Bo Liu, Yifeng Zhu, Chongkai Gao, Yihao Feng, Qiang Liu, Yuke Zhu, and Peter Stone. Libero: Benchmarking knowledge transfer for lifelong robot learning. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 44776–44791. Curran Associates, Inc., 2023. doi: 10.52202/075280-1939.

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. In Y. Yue, A. Garg, N. Peng, F. Sha, and R. Yu (eds.), *International Conference on Learning Representations*, volume 2025, pp. 29982–30009, 2025.

Zhengyi Luo, Ye Yuan, Tingwu Wang, Chenran Li, Fernando Castañeda, Sirui Chen, Zi-Ang Cao, Jiefeng Li, David Minor, Qingwei Ben, Jinhyung Park, David Sami, Zi Wang, Xingye Da, Runyu Ding, Cyrus Hogg, Lina Song, Edy Lim, Eugene Jeong, Tairan He, Haoru Xue, Wenli Xiao, Simon Yuen, Jan Kautz, Yan Chang, Umar Iqbal, Linxi Jim Fan, and Yuke Zhu. SONIC: supersizing motion tracking for natural humanoid whole-body control. *Sci. Robotics*, 11(117):eaed4592, 2026. doi: 10.1126/SCIROBOTICS.AED4592.

Oier Mees, Lukas Hermann, Erick Rosete-Beas, and Wolfram Burgard. Calvin: A benchmark for language-conditioned policy learning for long-horizon robot manipulation tasks. *IEEE Robotics and Automation Letters*, 7(3):7327–7334, 2022. doi: 10.1109/LRA.2022.3180108.

Soroush Nasiriany, Sep Nasiriany, Abhiram Maddukuri, and Yuke Zhu. Robocasa365: A large-scale simulation framework for training and benchmarking generalist robots. In *International Conference on Learning Representations*, volume 2026, pp. 98643–98667, 2026.

NVIDIA, :, Johan Bjorck, Fernando Castañeda, Nikita Cherniadev, Xingye Da, Runyu Ding, Linxi "Jim" Fan, Yu Fang, Dieter Fox, Fengyuan Hu, Spencer Huang, Joel Jang, Zhenyu Jiang, Jan Kautz, Kaushil Kundalia, Lawrence Lao, Zhiqi Li, Zongyu Lin, Kevin Lin, Guilin Liu, Edith Llontop, Loic Magne, Ajay Mandlekar, Avnish Narayan, Soroush Nasiriany, Scott Reed, You Liang Tan, Guanzhi Wang, Zu Wang, Jing Wang, Qi Wang, Jiannan Xiang, Yuqi Xie, Yinzhen Xu, Zhenjia Xu, Seonghyeon Ye, Zhiding Yu, Ao Zhang, Hao Zhang, Yizhou Zhao, Ruijie Zheng, and Yuke Zhu. Gr00t n1: An open foundation model for generalist humanoid robots, 2025. URL https://arxiv.org/abs/2503.14734.

Abby O'Neill, Abdul Rehman, Abhiram Maddukuri, Abhishek Gupta, Abhishek Padalkar, Abraham Lee, Acorn Pooley, Agrim Gupta, Ajay Mandlekar, Ajinkya Jain, Albert Tung, Alex Bewley, Alexander Herzog, Alex Irpan, Alexander Khazatsky, Anant Rai, Anchit Gupta, Andrew E. Wang, Anikait Singh, Animesh Garg, Aniruddha Kembhavi, Annie Xie, Anthony Brohan, Antonin Raffin, Archit Sharma, Arefeh Yavary, Arhan Jain, Ashwin Balakrishna, Ayzaan Wahid, Ben Burgess-Limerick, Beomjoon Kim, Bernhard Schölkopf, Blake Wulfe, Brian Ichter, Cewu Lu, Charles Xu, Charlotte Le, Chelsea Finn, Chen Wang, Chenfeng Xu, Cheng Chi, Chenguang Huang, Christine Chan, Christopher Agia, Chuer Pan, Chuyuan Fu, Coline Devin, Danfei Xu, Daniel Morton, Danny Driess, Daphne Chen, Deepak Pathak, Dhruv Shah, Dieter Büchler, Dinesh Jayaraman, Dmitry Kalashnikov, Dorsa Sadigh, Edward Johns, Ethan Paul Foster, Fangchen Liu, Federico Ceola, Fei Xia, Feiyu Zhao, Freek Stulp, Gaoyue Zhou, Gaurav S. Sukhatme, Gautam Salhotra, Ge Yan, Gilbert Feng, Giulio Schiavi, Glen Berseth, Gregory Kahn, Guanzhi Wang, Hao Su, Haoshu Fang, Haochen Shi, Henghui Bao, Heni Ben Amor, Henrik I. Christensen, Hiroki Furuta, Homer Walke, Hongjie Fang, Huy Ha, Igor Mordatch, Ilija Radosavovic, Isabel Leal, Jacky Liang, Jad Abou-Chakra, Jaehyung Kim, Jaimyn Drake, Jan Peters, Jan Schneider, Jasmine Hsu, Jeannette Bohg, Jeffrey T. Bingham, Jeffrey Wu, Jensen Gao, Jiaheng Hu, Jiajun Wu, Jialin Wu, Jiankai Sun, Jianlan Luo, Jiayuan Gu, Jie Tan, Jihoon Oh, Jimmy Wu, Jingpei Lu, Jingyun Yang, Jitendra Malik, João Silvério, Joey Hejna, Jonathan Booher, Jonathan

Tompson, Jonathan Yang, Jordi Salvador, Joseph J. Lim, Junhyek Han, Kaiyuan Wang, Kanishka Rao, Karl Pertsch, Karol Hausman, Keegan Go, Keerthana Gopalakrishnan, Ken Goldberg, Kendra Byrne, Kenneth Oslund, Kento Kawaharazuka, Kevin Black, Kevin Lin, Kevin Zhang, Kiana Ehsani, Kiran Lekkala, Kirsty Ellis, Krishan Rana, Krishnan Srinivasan, Kuan Fang, Kunal Pratap Singh, Kuo-Hao Zeng, Kyle Hatch, Kyle Hsu, Laurent Itti, Lawrence Yunliang Chen, Lerrel Pinto, Li Fei-Fei, Liam Tan, Linxi Jim Fan, Lionel Ott, Lisa Lee, Luca Weihs, Magnum Chen, Marion Lepert, Marius Memmel, Masayoshi Tomizuka, Masha Itkina, Mateo Guaman Castro, Max Spero, Maximilian Du, Michael Ahn, Michael C. Yip, Mingtong Zhang, Mingyu Ding, Minho Heo, Mohan Kumar Srirama, Mohit Sharma, Moo Jin Kim, Naoaki Kanazawa, Nicklas Hansen, Nicolas Heess, Nikhil J. Joshi, Niko Sünderhauf, Ning Liu, Norman Di Palo, Nur Muhammad (Mahi) Shafiullah, Oier Mees, Oliver Kroemer, Osbert Bastani, Pannag R. Sanketi, Patrick Tree Miller, Patrick Yin, Paul Wohlhart, Peng Xu, Peter David Fagan, Peter Mitrano, Pierre Sermanet, Pieter Abbeel, Priya Sundaresan, Qiuyu Chen, Quan Vuong, Rafael Rafailov, Ran Tian, Ria Doshi, Roberto Martín-Martín, Rohan Baijal, Rosario Scalise, Rose Hendrix, Roy Lin, Runjia Qian, Ruohan Zhang, Russell Mendonca, Rutav Shah, Ryan Hoque, Ryan Julian, Samuel Bustamante-Gomez, Sean Kirmani, Sergey Levine, Shan Lin, Sherry Moore, Shikhar Bahl, Shivin Dass, Shubham D. Sonawani, Shuran Song, Sichun Xu, Siddhant Haldar, Siddharth Karamcheti, Simeon Adebola, Simon Guist, Soroush Nasiriany, Stefan Schaal, Stefan Welker, Stephen Tian, Subramanian Ramamoorthy, Sudeep Dasari, Suneel Belkhale, Sungjae Park, Suraj Nair, Suvir Mirchandani, Takayuki Osa, Tanmay Gupta, Tatsuya Harada, Tatsuya Matsushima, Ted Xiao, Thomas Kollar, Tianhe Yu, Tianli Ding, Todor Davchev, Tony Z. Zhao, Travis Armstrong, Trevor Darrell, Trinity Chung, Vidhi Jain, Vincent Vanhoucke, Wei Zhan, Wenxuan Zhou, Wolfram Burgard, Xi Chen, Xiaolong Wang, Xinghao Zhu, Xinyang Geng, Xiyuan Liu, Liangwei Xu, Xuanlin Li, Yao Lu, Yecheng Jason Ma, Yejin Kim, Yevgen Chebotar, Yifan Zhou, Yifeng Zhu, Yilin Wu, Ying Xu, Yixuan Wang, Yonatan Bisk, Yoonyoung Cho, Youngwoon Lee, Yuchen Cui, Yue Cao, Yueh-Hua Wu, Yujin Tang, Yuke Zhu, Yunchu Zhang, Yunfan Jiang, Yunshuang Li, Yunzhu Li, Yusuke Iwasawa, Yutaka Matsuo, Zehan Ma, Zhuo Xu, Zichen Jeff Cui, Zichen Zhang, and Zipeng Lin. Open x-embodiment: Robotic learning datasets and RT-X models : Open x-embodiment collaboration. In *IEEE International Conference on Robotics and Automation, ICRA 2024, Yokohama, Japan, May 13-17, 2024*, pp. 6892–6903. IEEE, 2024. doi: 10.1109/ICRA57147.2024.10611477.

William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 4195–4205, October 2023.

Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67, 2020.

Rutav Shah, Yisu Li, Femi Bello, Yuke Zhu, and Roberto Martín-Martín. Memory Retrieval in Visuomotor Policies for Long-Horizon Robot Control. In *Proceedings of Robotics: Science and Systems*, Sydney, Australia, July 2026. doi: 10.15607/RSS.2026.XXII.010.

Ajay Sridhar, Jennifer Pan, Satvik Sharma, and Chelsea Finn. Scaling up memory for robotic control via experience retrieval. In *International Conference on Learning Representations*, volume 2026, pp. 97142–97166, 2026.

Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, Jianyuan Zeng, Jiayu Wang, Jingfeng Zhang, Jingren Zhou, Jinkai Wang, Jixuan Chen, Kai Zhu, Kang Zhao, Keyu Yan, Lianghua Huang, Mengyang Feng, Ningyi Zhang, Pandeng Li, Pingyu Wu, Ruihang Chu, Ruili Feng, Shiwei Zhang, Siyang Sun, Tao Fang, Tianxing Wang, Tianyi Gui, Tingyu Weng, Tong Shen, Wei Lin, Wei Wang, Wei Wang, Wenmeng Zhou, Wente Wang, Wenting Shen, Wenyuan Yu, Xianzhong Shi, Xiaoming Huang, Xin Xu, Yan Kou, Yangyu Lv, Yifei Li, Yijing Liu, Yiming Wang, Yingya Zhang, Yitong Huang, Yong Li, You Wu, Yu Liu, Yulin Pan, Yun Zheng, Yuntao Hong, Yupeng Shi, Yutong Feng, Zeyinzi Jiang, Zhen Han, Zhi-Fan Wu, and Ziyu Liu. Wan: Open and advanced large-scale video generative models, 2025.

Philipp Wu, Yide Shentu, Zhongke Yi, Xingyu Lin, and Pieter Abbeel. GELLO: A general, low-cost, and intuitive teleoperation framework for robot manipulators. In *IEEE/RSJ International*

Conference on Intelligent Robots and Systems, IROS 2024, Abu Dhabi, United Arab Emirates, October 14-18, 2024, pp. 12156–12163. IEEE, 2024. doi: 10.1109/IROS58592.2024.10801581.

Sizhe Yang, Juncheng Mu, Tianming Wei, Chenhao Lu, Xiaofan Li, Linning Xu, Zhengrong Xue, Zhecheng Yuan, Dahua Lin, Jiangmiao Pang, and Huazhe Xu. Memorywam: Efficient world action modeling with persistent memory, 2026.

Seonghyeon Ye, Yunhao Ge, Kaiyuan Zheng, Shenyuan Gao, Sihyun Yu, George Kurian, Suneel Indupuru, You Liang Tan, Chuning Zhu, Jiannan Xiang, Ayaan Malik, Kyungmin Lee, William Liang, Nadun Ranawaka, Jiasheng Gu, Yinzhen Xu, Guanzhi Wang, Fengyuan Hu, Avnish Narayan, Johan Bjorck, Jing Wang, Gwanghyun Kim, Dantong Niu, Ruijie Zheng, Yuqi Xie, Jimmy Wu, Qi Wang, Ryan Julian, Danfei Xu, Yilun Du, Yevgen Chebotar, Scott Reed, Jan Kautz, Yuke Zhu, Linxi "Jim" Fan, and Joel Jang. World action models are zero-shot policies, 2026. URL https://arxiv.org/abs/2602.15922.

Tianyuan Yuan, Zibin Dong, Yicheng Liu, and Hang Zhao. Fast-wam: Do world action models need test-time future imagination?, 2026.

Ruiteng Zhao, Zhengshen Zhang, Yue Su, Wenshuo Wang, Jiahui Li, Zhiyuan Yang, Francis E. H. Tay, Marcelo H. Ang Jr., and Haiyue Zhu. SG-WAM: Self-guided world modeling in geometry-aware policy space, 2026.

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. In *Proceedings of Robotics: Science and Systems*, LosAngeles, CA, USA, June 2025. doi: 10.15607/RSS.2025.XXI.015.

APPENDIX

This appendix provides additional details on the implementation, experimental settings, quantitative analyses, qualitative results, and discussion to complement the main paper.

- **Section A** describes the implementation details of TaskWAM, including model architectures, task-state token construction, training procedures, and inference settings.
- **Section B** provides detailed experimental settings, including StageBench construction, RoboTwin 2.0 configurations, and real-world evaluation protocols.
- **Section C** presents supplementary quantitative analyses, including detailed benchmark results, additional ablations, and further analysis of task-state grounding.
- **Section D** provides supplementary qualitative results, including additional multi-stage task executions, stage confusion cases, and real-world manipulation demonstrations.
- **Section E** discusses additional insights into the motivation behind TaskWAM.

A IMPLEMENTATION DETAILS

A.1 MODEL DETAILS

A.1.1 MODE TOKENS.

Reliable stage progression requires interpreting execution evidence under the fine-grained requirements specified by the instruction. When tasks share local interactions, even an accurate history may support different continuations. For example, after a microwave has started, one instruction may require another preparation step during heating, whereas another requires waiting until heating finishes. The observed history can be identical, but the appropriate next stage differs because of the timing requirement. Execution evidence therefore needs to be considered together with the instruction constraints that determine how the task should proceed.

We construct a persistent, instruction-grounded representation of these constraints, which we refer to as *Mode Tokens*. Here, an execution mode describes how manipulation should proceed under fine-grained requirements, including ordering, timing, repetition, and arm assignment. Mode Tokens represent these requirements, while State Tokens provide evidence of the progress actually made. Their complementary roles allow generation to use both the intended execution requirements and the observed execution context.

The construction comprises two steps. Instruction-dependent routing retrieves query slots from a learnable codebook, providing a coarse execution-requirement prior. Full-text grounding then refines these queries using the specific requirements expressed in the current instruction.

Instruction-Dependent Routing. Given a task instruction ℓ, the text encoder produces $\mathbf{C} = \mathcal{E}_{\text{text}}(\ell) \in \mathbb{R}^{L\times D}$, where L is the language-sequence length and D is the text-feature dimension. We obtain a global instruction feature through masked average pooling:

$$\mathbf{c}_{\text{pool}} = \frac{\sum_{i=1}^{L} \text{mask}_i \mathbf{c}_i}{\sum_{i=1}^{L} \text{mask}_i}, \tag{8}$$

where mask_i indicates a valid language token. A two-layer MLP predicts routing logits and normalized weights:

$$\begin{aligned} \boldsymbol{\ell}_{\text{prior}} &= W_2\, \text{SiLU}\left(W_1 \mathbf{c}_{\text{pool}} + b_1\right) + b_2, \\ \boldsymbol{\alpha} &= \text{softmax}\left(\boldsymbol{\ell}_{\text{prior}}\right) \in \mathbb{R}^{K_{\text{prior}}}, \end{aligned} \tag{9}$$

where K_{prior} is the number of codebook entries. The weights determine an instruction-dependent combination of the query slots used for subsequent requirement extraction.

Learnable Requirement Codebook. We maintain a codebook

$$\mathbf{E} \in \mathbb{R}^{K_{\text{prior}}\times N_m \times D}, \tag{10}$$

with N_m learnable query slots per entry. The routing weights combine corresponding slots across entries:

$$\mathbf{m}_l^{\text{raw}} = \sum_{k=1}^{K_{\text{prior}}} \alpha_k \mathbf{E}_{k,l}, \qquad l = 1, \ldots, N_m. \tag{11}$$

The codebook parameters are shared across instructions, while their combination adapts to the current language context. The retrieved slots serve as a coarse execution-requirement prior by initializing the queries that read the instruction.

This construction does not require a predefined one-to-one correspondence between codebook entries and tasks. Individual entries are not manually assigned to specific execution requirements; their representations are learned jointly through the world and action prediction objectives. The retrieved mixture is subsequently refined against the full instruction, so the final Mode Tokens depend on more than the routing weights alone.

Fine-Grained Instruction Grounding. Routing supplies an instruction-dependent query initialization, but its pooled input does not preserve the full token-level structure of the language sequence. We therefore use the retrieved slots as queries over the complete text context through a multi-layer Q-Former:

$$\begin{aligned} \mathbf{M}^{\text{raw}} &= [\mathbf{m}_1^{\text{raw}}, \ldots, \mathbf{m}_{N_m}^{\text{raw}}], \\ \widetilde{\mathbf{M}} &= \mathcal{Q}_{\text{ground}}\left(\mathbf{M}^{\text{raw}}, \mathbf{C}\right). \end{aligned} \tag{12}$$

The language tokens provide the key-value context, giving the queries access to the words and contextual relationships that specify ordering, timing, repetition, and arm assignment. Grounding thus adapts the retrieved prior to the detailed requirements of the current instruction.

We project the grounded queries to the backbone conditioning dimension:

$$\mathbf{M} = \mathcal{P}_m\left(\widetilde{\mathbf{M}}\right) \in \mathbb{R}^{N_m \times d}, \tag{13}$$

where $\mathcal{P}_m$ denotes the output projection. In the implementation described here, each codebook entry contains $N_m = 8$ query slots, and the module outputs eight Mode Tokens with dimension $d = 4096$. The number of query slots N_m is distinct from the number of codebook entries K_{prior}.

Joint Learning for World-Action Generation. We jointly train the routing head, codebook, grounding module, and WAM using the backbone's world and action prediction objectives, without a separate codebook pretraining stage. As described in Sec. 3.4, Mode Tokens participate in task-state attention, layer-wise modulation, and residual prediction. These pathways allow the instruction-grounded representations to influence both world and action generation.

Prediction losses backpropagate through the conditioning pathways to the Mode-Token construction modules. The learned queries are therefore optimized according to their contribution to downstream world-action prediction, connecting instruction grounding with the generative execution objective.

Persistent Conditioning During Execution. Mode Tokens are retained while the instruction remains unchanged, keeping the task-specific execution requirements accessible across stage transitions. State Tokens are updated from actual action–observation interactions, providing evolving evidence of execution progress. Both generation streams access these complementary representations to condition subsequent world-action predictions.

At inference time, the routing weights are predicted from the task instruction. Constructing Mode Tokens does not require a user-provided task identifier or codebook index.

A.1.2 State Tokens.

Mode Tokens specify how the task should proceed, but do not indicate which requirements have already been fulfilled. When similar observations recur across stages, execution progress must be inferred from the preceding interactions. An executed action records what the robot attempted, while the subsequent observation and its change provide evidence of the resulting outcome. We therefore construct State Tokens from temporally ordered action–observation transitions, providing execution evidence for subsequent world–action generation.

Interaction Representation. We use a fixed history window of $K = 8$ interaction slots. Each slot k pairs the preceding executed action chunk $\mathbf{A}_{k-1}$ with the resulting visual latent $\mathbf{F}_k$ and its

change $\Delta\mathbf{F}_k = \mathbf{F}_k - \mathbf{F}_{k-1}$. These inputs provide complementary evidence: $\mathbf{A}_{k-1}$ describes the executed behavior, $\mathbf{F}_k$ describes the resulting scene, and $\Delta\mathbf{F}_k$ captures the observed change between interactions. Their joint representation allows the encoder to use both the executed actions and their visual feedback when tracking progress.

We independently project the three modalities into a 256-dimensional feature space. The same learnable slot-position embedding is added to all modality features belonging to slot k, associating them with the same interaction.

Within-Slot Aggregation. A 2-layer, 8-head Q-Former aggregates the multimodal features within each interaction slot. Its query is formed as

$$\mathbf{q}_k = \mathbf{q}_{\text{base}} + \mathbf{p}^q_k, \tag{14}$$

where $\mathbf{q}_{\text{base}}$ is a shared learnable query and $\mathbf{p}^q_k$ is a learnable embedding for the slot position. The Q-Former attends to the projected visual, visual-change, and action features. Its output is mapped to a 1024-dimensional interaction representation $\mathbf{e}_k$, summarizing the executed behavior and the observed scene information associated with that interaction.

Causal Temporal Encoding. The eight interaction representations are arranged in execution order and processed by a 2-layer, 8-head causal Transformer. A lower-triangular attention mask restricts each slot to its own and preceding interactions. This preserves the chronology needed to distinguish similar scene states reached through different execution histories.

The encoder outputs

$$\mathbf{S}^{\text{enc}}_t \in \mathbb{R}^{B\times 8\times 1024}, \tag{15}$$

where B denotes the batch size. We retain all eight temporally ordered tokens, allowing subsequent generation to access representations associated with different positions in the recent interaction history. Each later token can incorporate a longer prefix of the available window. A backbone-specific projection maps these representations to the conditioning dimension d:

$$\mathbf{S}_t = \mathcal{P}_s\left(\mathbf{S}^{\text{enc}}_t\right) \in \mathbb{R}^{B\times 8\times d}. \tag{16}$$

History Initialization. At the beginning of an episode, fewer than eight interaction slots may be available. We pad missing slots with the earliest available observation and set their action inputs to zero. For the first valid interaction slot, the latent-difference input is also set to zero because no preceding observation is available. This keeps the input length fixed without carrying history across episode boundaries.

Closed-Loop State Updates. After executing action chunk $\mathbf{A}_t$, the robot receives the next observation and encodes its visual latent $\mathbf{F}_{t+1}$. The new interaction $(\mathbf{F}_{t+1}, \Delta\mathbf{F}_{t+1}, \mathbf{A}_t)$ is appended to the history window, the oldest slot is removed when necessary, and the State Tokens are recomputed. Subsequent generation therefore uses execution evidence updated from the robot's actual interactions, alongside the persistent requirements represented by Mode Tokens.

A.1.3 Task-State Conditioning and Generation Steering.

The grounded task state $\mathbf{T}_t = [\mathbf{M}; \mathbf{S}_t]$ provides both task requirements and execution evidence for world-action generation. TaskWAM preserves the native backbone architecture and injects $\mathbf{T}_t$ into both world and action streams through two complementary mechanisms.

First, we append the task state to the key-value context of both streams, allowing future world and action tokens to directly access the current execution requirements and progress. In addition, pooled Mode Tokens provide layer-wise AdaLN modulation:

$$\widetilde{\mathbf{H}}_l = (1 + \boldsymbol{\gamma}^m_l) \odot \mathrm{LN}(\mathbf{H}_l) + \boldsymbol{\beta}^m_l. \tag{17}$$

Second, we introduce task-state-conditioned residual correction to directly adjust the predicted world-action evolution. For each stream $b \in \{w, a\}$, the residual is generated by querying the task state with the corresponding hidden states:

$$\Delta\mathbf{v}^b_{t,\tau} = R_b(\mathbf{H}^b_{t,\tau}; \mathbf{T}_t), \qquad \widehat{\mathbf{v}}^b_{t,\tau} = \mathbf{v}^b_{t,\tau} + \gamma_b \Delta\mathbf{v}^b_{t,\tau}, \tag{18}$$

where γ_b is a learnable bounded scale. For Cosmos Policy, the residual operates on denoised latent predictions following its native EDM formulation.

Table 7: Training configuration on kitchen manipulation tasks of StageBench.

Hyperparameter	Setting
Backbone	Wan2.2-TI2V-5B + FastWAM Action DiT
Training GPUs	16× NVIDIA H100 80GB
Per-GPU batch size	16
Gradient accumulation steps	1
Effective batch size	128
Optimizer updates	58,720
Optimizer	AdamW
Learning rate	1×10^{-4}
Learning-rate schedule	Cosine decay
Warm-up ratio	5%
Weight decay	0.01
Gradient clipping norm	1.0
Precision	BF16
Random seed	42
Image views	Two 224×224 views
Concatenated image resolution	224×448
Video training window	33 frames
Action-to-video frequency ratio	4
Action / proprioception dimensions	7 / 8
History length	8 interaction slots
History observation-anchor interval	16 action steps
Executed-action chunk length	16 steps

A.2 Training and Inference Details

Training. TaskWAM is implemented on Fast-WAM in PyTorch, retaining the backbone's native inference structure. Following the open-source training setup of Fast-WAMYuan et al. (2026), Table 7 summarizes the training configuration for long-horizon kitchen manipulation tasks constructed in LIBERO for StageBench. The simulation experiments are trained on 16 NVIDIA H100 80GB GPUs with an effective batch size of 128 for 10 epochs. Flow-based WAMs use their flow-matching generative objectives, whereas Cosmos Policy uses EDM denoising. The world and action generative terms are denoted $\mathcal{L}_{\text{world}} + \mathcal{L}_{\text{action}}$. The routing head, codebook, grounding module, and WAM are trained jointly in one stage.

Closed-loop inference. For a fixed instruction, Mode Tokens are computed once and retained throughout the rollout. The policy uses the instruction, current visual observation, and causal visual–action history to condition world and action generation. After an action chunk is executed and the next observation is encoded, the new visual–action transition is appended to the history and the Execution-State Tokens are refreshed for the next rollout step. Future-world guidance informs action steering only when the original backbone supports it.

B Experimental Details

B.1 StageBench

Task families and modes. The eight StageBench task families detailed here comprise four single-arm LIBERO families and four bimanual RoboTwin 2.0 families. Each task family includes multiple instruction-defined *task modes*, which differ in subtask order, arm assignment, execution timing, or repetition count. Each mode consists of up to five execution stages. Table 8 summarizes the task families and their corresponding modes.

Table 8: Task families and instruction-defined modes in StageBench.

ID	Task family	Task modes
Kitchen Manipulation		
K1	Cooking & Cheese	Interleaved; Sequential
K2	Cooking & Juice	Interleaved; Sequential
K3	Cooking & Cleanup	Interleaved; Sequential
K4	Cooking & Breakfast	Interleaved; Sequential
Bimanual Manipulation		
B1	Size-Order Cubes	Small-to-Large; Large-to-Small
B2	Color-Order Cups	RGB; BGR; RBG; GRB
B3	Arm-Assigned Hammering	Left Arm; Right Arm
B4	Repeated Bread Cutting	Bread-first; Knife-first; instruction-specified cutting counts

StageBench Task Variants. The eight task families in Table 8 instantiate multiple task variants by modifying fine-grained instruction requirements. All variants preserve the same scene, objects, and local manipulation primitives while varying the execution constraints that determine the correct continuation. Figure 6 visualizes representative prompts from each family. The stage-aware evaluator distinguishes physical task completion from instruction-defined execution routes based on task-specific events.

Environment Extensions. To support realistic multi-stage manipulation and stage-aware evaluation, we extend the LIBERO simulation environments with additional interactive household assets. For kitchen scenarios, we introduce functional microwave assets with stateful behaviors, including door-triggered heating activation and timer-based heating completion. These extensions enable meaningful state transitions and temporal dependencies beyond simple rearrangements of manipulation primitives. All environment extensions and task trajectories are manually validated over multiple rollouts to ensure reliable state transitions and consistent event detection for automatic evaluation.

Microwave-based Cooking Tasks. These task families share the same interaction structure: placing the mug into the microwave, starting heating, and performing a collection operation. The variant variable is the temporal relation between collection and microwave reopening, as shown in Figure 7. The prompt template is: "Put the mug in the microwave and start heating it; {collection timing}, then come back and reopen the microwave." For *Interleaved*, {collection timing} is instantiated as "while it is heating, place {target} in the basket"; for *Sequential*, it is instantiated as "wait for it to finish heating, reopen the microwave, and only then place {target} in the basket". The target variable is instantiated as cream cheese, orange juice, cream cheese and orange juice, or bowl and butter for the four cooking families.

Size-Order Cubes. The task-variant variable is the required size order. The *Small-to-Large* prompt is "Arrange the blocks in a row at the center of the table, starting with the smallest block and ending with the largest." The *Large-to-Small* prompt is "Arrange the blocks in a row at the center of the table, starting with the largest block and ending with the smallest." Both variants arrange the same blocks; evaluation checks their placement order against the instruction.

Color-Order Cups. The task-variant variable is the prescribed color sequence. For *RGB*, the prompt is "Arrange the colored cups on the horizontal stand in this order: red, green, blue." For *RBG*, it is "Arrange the colored cups on the horizontal stand in this order: red, blue, green." For *BGR*, it is "Arrange the colored cups on the horizontal stand in this order: blue, green, red." For *GRB*, it is "Arrange the colored cups on the horizontal stand in this order: green, red, blue." The variants share the initial scene and cup-placement operations. Evaluation compares the observed color-specific placement sequence with the requested prompt, as illustrated in Figure 8.

Arm-Assigned Hammering. The variant variable is the instructed manipulation arm. All variants share the same target object and manipulation procedure. The prompt template is: "Pick up the hammer with the {arm} arm and strike the block." where {arm} is instantiated as the left or right arm.

Repeated Bread Cutting. The task requirement is the specified repetition count. The prompt is: "Use both robot arms to make exactly three cuts in the bread." The evaluator checks whether the robot performs the required number of cuts before termination.

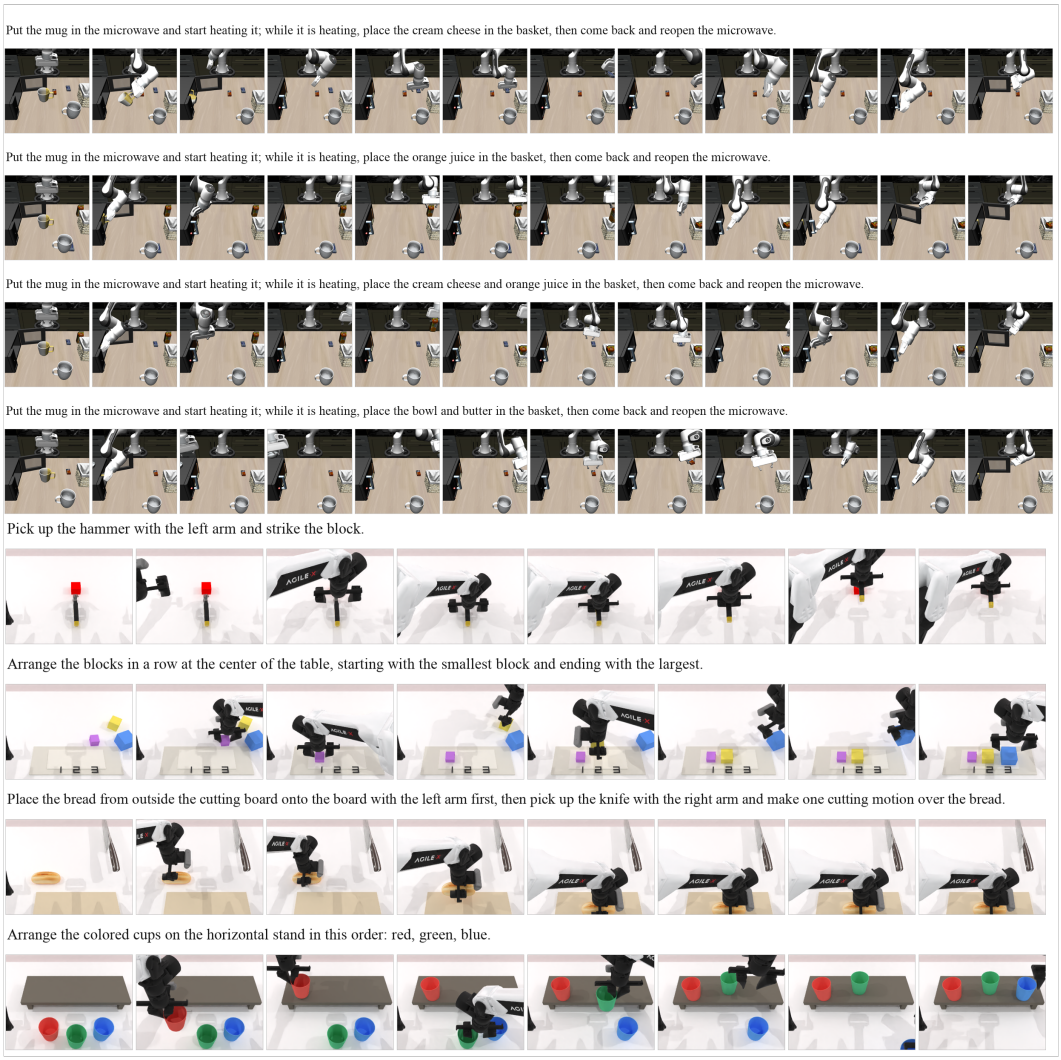

Figure 6: **StageBench benchmark visualization.** Representative task variants and their multi-stage execution processes are illustrated, highlighting diverse stage structures and instruction- dependent task requirements.

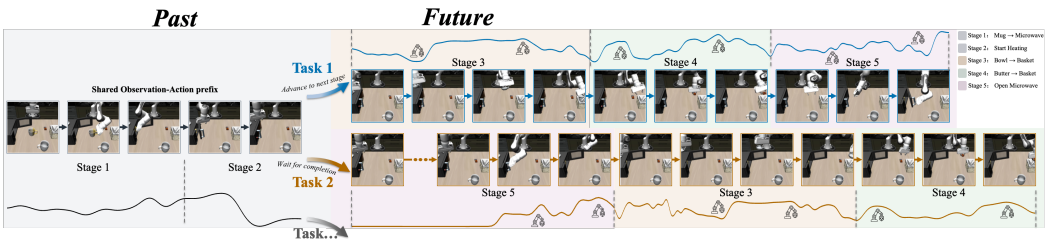

Figure 7: **Execution orders in a LIBERO household task.** Both modes share the mug-placement and heating prefix. The *Interleaved* route collects the bowl and butter during heating and then reopens the microwave. The *Sequential* route waits for heating to finish and reopens the microwave before collecting the objects.

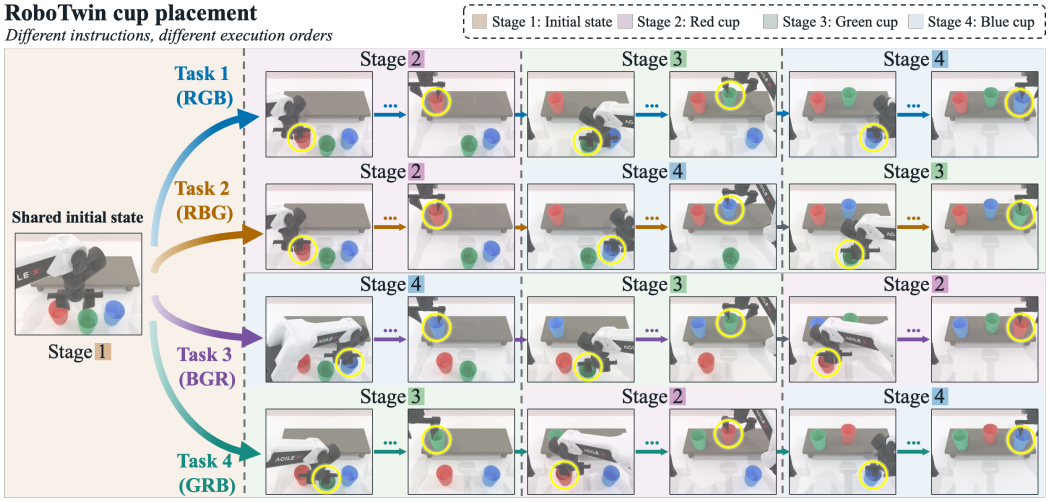

Figure 8: **Instruction-dependent cup placement in RoboTwin.** From a shared initial state, the robot places the red (R), green (G), and blue (B) cups in four prescribed orders: RGB, RBG, BGR, and GRB. Stage labels identify cup-specific operations, whose execution order varies with the instruction.

Data collection. For the LIBERO tasks, we collect demonstrations through 3D human teleoperation. For the four RoboTwin task families, we use the official data generation pipeline. We collect 50 demonstration trajectories for each task mode. The resulting dataset contains 2.20 million frames across both platforms, with approximately 12 hours spent on human teleoperation.

Event-based stage annotation. To support stage-level evaluation, we record task-relevant simulator states at every environment step and convert them into timestamped task events. The events are defined separately for each task family according to its execution requirements. For example, the LIBERO heating tasks record heating start, heating completion, post-heating microwave reopening, target-object manipulation, and whether each target object has entered the basket. A stage may depend on multiple events rather than a single simulator condition.

For RoboTwin, the recorded events depend on the task family, such as object placement order for R1 and R2, hammer manipulation by the instructed arm for R3, and completed cutting motions for R4. These event traces provide a common interface for tracking execution progress while retaining task-specific stage definitions.

Automatic evaluation. We separately evaluate physical task completion and instruction-level execution consistency. Physical success is determined using the native success criteria of the corresponding simulation platform. Meanwhile, deterministic finite-state machines analyze recorded event sequences to reconstruct execution routes and track stage completion.

For the LIBERO heating tasks, let t_{collect} denote the time step at which the collection branch is executed and t_{open} the time step at which the microwave is reopened after heating. An execution is classified as *Interleaved* when $t_{\text{collect}} < t_{\text{open}}$, and as *Sequential* when $t_{\text{open}} < t_{\text{collect}}$. Heating-start and heating-completion events are additionally retained to verify the corresponding physical stages and route consistency.

For RoboTwin, the evaluator applies task-specific rules to verify instruction-defined constraints, including object ordering, arm roles, and repetition counts.

Stage Progression Rate. To complement final task success, we report Stage Progression Rate (SPR), which measures whether a policy completes the instructed next stage after completing the current stage. For task family f and mode m, let $\mathcal{E}_{f,m}$ denote the set of eligible transition opportunities across evaluated rollouts. An opportunity becomes eligible once the preceding required stages have been completed in order. Let $\mathcal{C}_{f,m} \subseteq \mathcal{E}_{f,m}$ contain opportunities for which the first recognizable subsequent stage entry matches the instructed next stage and that stage is completed before rollout termination. We define

$$\text{SPR}_{f,m} = \frac{|\mathcal{C}_{f,m}|}{|\mathcal{E}_{f,m}|} \times 100\%. \tag{19}$$

The denominator retains all eligible opportunities, including incorrect stage entries, correct entries without completion, and cases with no recognizable subsequent entry. Task-level SPR pools these counts across modes. Higher SPR indicates more reliable progression along the instructed execution route.

Evaluation protocol. For each LIBERO task family, the two task modes are evaluated using the same fixed set of 100 unique initial states. Each task–mode pair is therefore evaluated on 100 initial states, and all compared models use the same initial states and the same task-specific maximum number of environment steps. For every rollout, we save the detected events and their time steps, physical task success, identified execution mode, and stage-level progress.

B.2 RoboTwin 2.0 Setup

Benchmark. We evaluate on all 50 bimanual manipulation tasks in RoboTwin 2.0 (Chen et al., 2026b) under the Clean and Randomized settings. The latter introduces variations in scene appearance and configuration to assess robustness to environmental changes. We report per-task success rates and their average across all 50 tasks separately for each setting. In the accompanying task-wise comparison, Δ denotes the absolute percentage-point improvement over Fast-WAM under the corresponding setting. The reported TaskWAM scores are aggregated from the best result for each task within that setting, rather than from a single evaluation run.

The Clean training set contains 2,500 demonstrations (50 per task), and the Randomized training set contains 25,000 demonstrations (500 per task). Simulation experiments are conducted over 100 trials.

Configuration. For this benchmark, TaskWAM augments Fast-WAM (Yuan et al., 2026) with hierarchical task-state modeling and task-future flow shaping. The policy receives the task instruction, current visual observation, and causal visual–action history. Task-Mode Tokens are retained while the instruction remains unchanged, whereas Execution-State Tokens are refreshed after each executed action chunk. Both token sets condition world and action generation and the corresponding velocity-field corrections, with predicted future-world features additionally informing action correction.

As shown in Table 9, TaskWAM gains are measured against Fast-WAM Clean and Randomized averages, respectively. TaskWAM Clean2 and Rand. (Hard Best) are verified same-setting per-task-best aggregates, not a single 5000-episode run.

B.3 Real-World Setup

Platforms and data collection. The real-world platforms include a single-arm Franka Research 3, dual-arm AgileX Piper (6-DoF) and AgileX Nero (7-DoF), and a Unitree G1 humanoid equipped with BrainCo Revo2 dexterous hands. Franka is equipped with third-person and wrist-mounted RealSense D435 cameras, while the dual-arm platforms use wrist cameras (and a head-mounted camera

for Nero) for visual observations. The humanoid platform is operated via PICO VR for whole-body teleoperation. Teleoperation interfaces include Gello Wu et al. (2024) for gripper-based robots, TypeTele Lin et al. (2025) for dexterous hands, and SONIC Luo et al. (2026) for the humanoid platform. We collect 100–150 teleoperated demonstrations per task.

Multi-Stage Real-World Tasks. We evaluate four multi-stage tasks spanning three representative real-world scenarios. These tasks evaluate cross-task and within-task stage confusion under shared objects, repeated interactions, and instruction-dependent continuations. ① **Sequential Cup Nesting** requires nesting red, pink, and blue cups in the instructed order, with three order variants. The breakfast-preparation scenario includes two tasks: ② **Bread Toasting & Spreading**, which requires toasting bread, transferring it to a plate, and applying the specified blueberry or yellow-peach jam; and ③ **Bread Cutting**, which requires making exactly three cuts and then stopping. ④ **Laboratory Fluid Manipulation** requires transferring a reagent into the receiving container specified by the instruction. These settings naturally reuse objects and local interactions across tasks and execution stages, providing real-world evaluation of both cross-task and within-task stage confusion. Additionally, we evaluate Bread Toasting & Spreading under a modified tabletop background to examine robustness against scene variations.

(1) Sequential Cup Nesting. The task-variant variable is the prescribed color sequence. All variants share the same scene and manipulation procedure. The prompt template is: "Assemble the colored cups by nesting them one by one on the tabletop platform in the following color order: {color order}, and place the assembled cup stack into the basket." where {color order} is instantiated as *pink, red, blue* (PRB), *red, blue, pink* (RBP), and *red, pink, blue* (RPB).

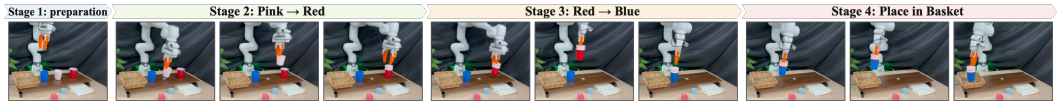

Figure 9: **Visualization of Sequential Cup Nesting.**

(2) Bread Toasting & Spreading. The task-variant variable is the specified spreading ingredient. All variants share the same preparation procedure. The prompt template is: "Heat the bread in the toaster, take it out when toasted, place it on a plate, and spread {ingredient} jam." where {ingredient} is instantiated as *blueberry* and *yellow peach*.



Figure 10: **Visualization of Bread Toasting & Spreading.**

(3) Laboratory Fluid Manipulation. The task-variant variable is the receiving container. Both variants share the same manipulation sequence. The prompt template is: "Use the right arm to remove the chemical liquid container from the microwave and pour its contents into the {container} for collection, then use the left arm to close the microwave door." where {container} is instantiated as *Petri dish* and *gel plate*.



Figure 11: **Visualization of Laboratory Fluid Manipulation.**

(4) Bread Cutting. The task requirement is the repetition count. The prompt is: "Use both robot arms to make exactly three cuts in the bread."

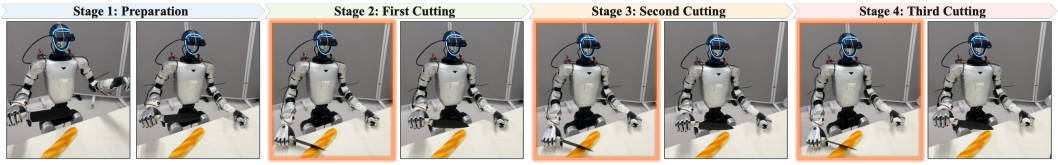

Figure 12: **Visualization of Bread Cutting.**

Evaluation. Cumulative stage completion and overall task success are defined as successful trials divided by total attempts. A trial counts as completing stage k only after it has completed the required sequence through that stage. The evaluation pools counts across the variants within each task family.

C Supplementary Quantitative Analysis

C.1 Detailed Results

The complete RoboTwin 2.0 task-wise comparison under the Clean and Randomized settings is reported in Table 9.

C.2 Additional Ablations and Analysis

Effectiveness of temporally structured State Tokens. Table 10 examines whether State Tokens benefit from structured execution history beyond the latest state. Shuffling token order preserves their content but destroys chronology, reducing success from 69.4% to 50.3%; repeating the latest token removes state evolution and yields 58.1%. These degradations confirm that both correct temporal order and evolving execution history are essential for reliable progress tracking.

Table 10: **Analysis of State Tokens.**

Setting	*Suc.*↑
TaskWAM	**69.4**
w/ Shuffled Temporal Order	50.3
w/ Repeated Latest State	58.1

C.3 Training Convergence Analysis

As shown in Figure 13, explicitly grounding world-action generation in structured task states enables faster convergence and a lower action loss than FastWAM. This result suggests that TaskWAM eases optimization by reducing the burden of implicitly recovering task requirements and execution progress from an undifferentiated context.

D Supplementary Qualitative Analysis

D.1 Real-World Stage Confusion Case Study

We further analyze a representative stage confusion case from the real-world Bread Toasting & Spreading task. As shown in Figure 14, the task contains a long waiting stage during bread toasting, followed by a retrieval stage with a visual configuration similar to the initial loading stage.

During the waiting period, the observation history contains many redundant frames with limited progress information. When the robot reaches the retrieval stage, visually similar states can therefore correspond to different execution stages depending on the preceding interactions. Without explicit execution-progress modeling, baseline WAMs may confuse the completed toasting stage with earlier preparation stages, leading to incorrect behaviors such as repeating the heating procedure or failing to advance to subsequent stages.

TaskWAM alleviates this ambiguity by combining instruction-grounded task requirements from Mode Tokens with causal execution evidence from State Tokens, enabling more reliable stage progression under long-horizon interactions (Table 5).

Table 9: **Complete RoboTwin 2.0 task-wise comparison.** Fast-WAM is positioned immediately before TaskWAM as the direct baseline. TaskWAM uses the primary Clean2 and Hard Best aggregates reported in the main paper.

Task	Motus		$\pi_{0.5}$		LingBot-VA		LaWAM		AHA-WAM		Fast-WAM		**w/ TaskWAM (ours)**	
	Clean	Rand.	Clean	Rand.	Clean	Rand.	Clean	Rand.	Clean	Rand.	Clean	Rand.	Clean / Δ	Rand. / Δ
Adjust Bottle	89	93	**100**	99	90	94	**100**	**100**	**100**	**100**	**100**	**100**	**100**	**100**
Beat Block Hammer	95	88	96	93	96	**98**	90	93	**100**	94	99	97	**100**	**98**
Blocks Ranking RGB	99	97	92	85	99	98	97	**100**	**100**	98	**100**	**100**	**100**	99
Blocks Ranking Size	75	63	49	26	94	96	93	89	**98**	**98**	94	**98**	93	93
Click Alarmclock	**100**	**100**	98	89	99	**100**	**100**	**100**	**100**	**100**	**100**	**100**	**100**	**100**
Click Bell	**100**	**100**	99	66	**100**	**100**	**100**	**100**	**100**	**100**	**100**	**100**	**100**	**100**
Dump Bin Bigbin	95	91	92	**97**	89	96	97	95	**99**	92	97	96	97	96
Grab Roller	**100**	**100**	**100**	**100**	**100**	**100**	**100**	**100**	**100**	**100**	**100**	**100**	**100**	**100**
Handover Block	86	73	66	57	**99**	78	96	**87**	96	85	95	81	82	86
Handover Mic	78	63	98	97	94	96	93	98	99	99	99	**100**	**100**	99
Hanging Mug	38	38	18	17	40	28	51	43	**85**	**78**	58	62	55	61
Lift Pot	96	99	96	85	**100**	99	**100**	99	**100**	**100**	**100**	**100**	**100**	**100**
Move Can Pot	34	74	51	55	94	**97**	**98**	93	78	83	90	88	95	96
Move Pillbottle Pad	93	96	84	61	99	**99**	97	90	**100**	**99**	**100**	**99**	98	**99**
Move Playingcard Away	**100**	96	96	84	**100**	99	**100**	**100**	**100**	**100**	**100**	**100**	**100**	**100**
Move Stapler Pad	83	85	56	42	91	79	**94**	**87**	88	77	77	64	82	75
Open Laptop	95	91	90	96	92	94	**100**	**100**	98	97	98	**100**	**100**	**100**
Open Microwave	**95**	**91**	34	77	82	86	41	43	40	46	62	45	79	77
Pick Diverse Bottles	90	**91**	81	71	89	82	**91**	88	89	85	80	85	81	83
Pick Dual Bottles	96	90	93	63	**100**	**99**	**100**	95	98	**99**	**100**	96	86	90
Place A2B Left	88	79	87	82	97	93	98	91	**100**	92	95	93	97	**97**
Place A2B Right	91	87	87	84	**97**	95	89	94	96	92	93	**99**	96	95
Place Bread Basket	91	94	77	64	97	95	92	85	89	94	91	93	**98**	**97**
Place Bread Skillet	86	83	85	66	**95**	90	90	83	91	87	90	93	94	**98**
Place Burger Fries	98	98	94	87	97	95	93	96	**99**	98	96	**99**	**99**	97
Place Can Basket	81	76	62	62	81	84	**92**	65	81	**94**	71	69	72	80
Place Cans Plasticbox	98	94	94	84	**100**	99	**100**	95	**100**	**100**	99	96	99	97
Place Container Plate	98	99	99	95	99	97	**100**	**100**	**100**	96	96	**100**	99	99
Place Dual Shoes	93	87	75	75	94	89	**98**	94	91	**99**	94	88	93	86
Place Empty Cup	99	98	**100**	99	**100**	**100**	99	**100**	**100**	**100**	**100**	**100**	**100**	**100**
Place Fan	91	87	87	85	99	93	92	93	**100**	87	96	**96**	98	94
Place Mouse Pad	66	68	60	39	**93**	**96**	91	84	**93**	82	83	89	**93**	**96**
Place Object Basket	81	87	80	76	91	88	**92**	**90**	81	80	89	88	**92**	89
Place Object Scale	88	85	86	80	**96**	95	95	88	**96**	87	90	97	**96**	**99**
Place Object Stand	98	**97**	91	85	**99**	96	92	93	94	93	90	94	96	93
Place Phone Stand	87	86	81	81	97	97	93	94	**98**	96	97	**99**	90	98
Place Shoe	99	97	92	93	98	98	**100**	**100**	91	96	96	99	88	98
Press Stapler	93	98	87	83	85	82	**98**	97	**98**	**99**	90	97	95	93
Put Bottles Dustbin	81	79	84	79	87	91	94	92	83	**96**	**95**	90	88	90
Put Object Cabinet	88	71	80	79	85	87	90	82	91	89	**94**	89	90	**90**
Rotate QRcode	89	73	89	87	**96**	**91**	94	89	90	90	93	89	87	**91**
Scan Object	67	66	72	65	**96**	91	**96**	90	94	90	89	**92**	95	89
Shake Bottle	**100**	97	99	97	**100**	97	**100**	**100**	**100**	**100**	**100**	**100**	**100**	**100**
Shake Bottle Horizontally	**100**	98	99	99	**100**	99	**100**	**100**	**100**	**100**	**100**	**100**	**100**	**100**
Stack Blocks Three	91	95	91	76	**99**	**98**	90	75	**99**	**98**	95	97	95	**98**
Stack Blocks Two	**100**	98	97	**100**	**100**	98	**100**	97	**100**	99	**100**	**100**	**100**	**100**
Stack Bowls Three	79	**87**	77	71	86	83	**90**	80	79	82	80	81	87	79
Stack Bowls Two	98	98	95	96	94	98	**100**	99	**100**	**100**	92	98	97	96
Stamp Seal	93	92	79	55	96	**97**	89	88	**98**	90	90	94	80	88
Turn Switch	**84**	**78**	62	54	44	45	47	56	70	74	61	59	73	74
Average (%)	88.66	87.02	82.74	76.76	92.90	91.50	92.52	89.48	**93.40**	92.20	91.88	91.78	92.70 (+0.82)	**93.06** (+1.28)

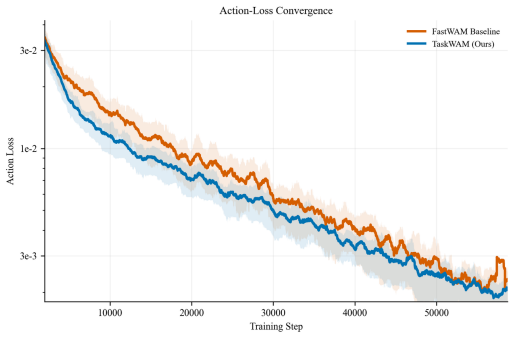

Figure 13: **Training convergence.** TaskWAM converges faster and reaches a lower action loss than FastWAM under the same training setup.

D.2 Additional Qualitative Analysis

Figures 6, 7, and 8 visualize the instruction-defined task variants and their corresponding multi-stage execution structures in StageBench. These examples illustrate how StageBench creates challenging

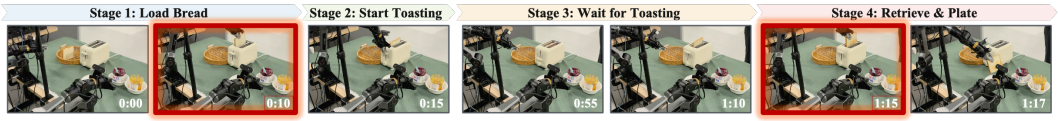

Figure 14: **Real-world stage confusion in Bread Toasting & Spreading.** Stage 1 (loading bread) and Stage 4 (retrieving toasted bread) exhibit similar visual configurations, while Stage 3 introduces a long waiting period with redundant observations. Baseline WAMs may confuse the current progress with earlier stages, resulting in repeated or incorrect execution.

yet meaningful ambiguities by preserving shared scenes, objects, and local interactions while varying task requirements and execution progress.

Figure 6 presents representative task trajectories from the benchmark task families, covering diverse single-arm and bimanual manipulation scenarios. Figure 7 shows LIBERO variants that share early-stage operations, such as mug placement and microwave heating, but differ in later-stage requirements, including whether the robot should continue with collection or reopen the microwave after heating. These variants create cross-task ambiguity where similar observations require different continuations depending on the instruction.

Figure 8 presents four cup-placement variants from the same initial configuration, where only the requested color order changes. Although the manipulation primitives remain identical, the robot must maintain the instruction-defined ordering constraint throughout the execution. These examples correspond to the event-order and stage-completion checks described in Sec. B.1, and demonstrate how StageBench evaluates reliable progression under overlapping observation-action trajectories.

E Discussion

Implicit Stage Modeling in Vanilla WAMs. For a flow-based WAM, the world and action velocity fields are conditioned on the instruction ℓ, current observation o_t, and optional history $\mathcal{H}_t$, and trained to match the target velocities of demonstrated futures:

$$\mathcal{L}_{\text{wam}} = \mathbb{E}\Big[\sum_{b\in\{w,a\}} \big\|\mathbf{v}_\theta^b(\mathbf{x}_{t,\tau}^b, \tau \mid \ell, o_t, \mathcal{H}_t) - \mathbf{u}_{t,\tau}^b\big\|_2^2\Big], \tag{20}$$

where $b \in \{w, a\}$ denotes the world or action stream, τ is the flow timestep, $\mathbf{x}_{t,\tau}^b$ is the corresponding flow input, and $\mathbf{u}_{t,\tau}^b$ is the target velocity.

While demonstrated futures directly supervise world-action prediction, stage progression is implicitly inferred from mixed conditioning signals, including instructions, observations, and execution histories. Under overlapping trajectories, similar contexts may correspond to different task requirements or execution progress, causing WAMs to generate locally plausible continuations that advance to incorrect stages. This motivates explicit task-state modeling for reliable stage progression.