

BRIDGE-WA: Predicting Where and How the World Changes for Robotic Action

Yongjie Bai^{1,2} Hanting Wang¹ Mingtong Dai^{2,3} Qijun Zhong¹ Yang Liu^{1,✉} Liang Lin^{1,2,4}

¹Sun Yat-sen University ²Pengcheng Laboratory

³Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences ⁴X-Era AI Lab

✉ Corresponding Author

Abstract: General-purpose vision-language-action models benefit from large vision-language priors, but effective manipulation also requires anticipating action-relevant scene changes. Existing world-action models often rely on large generative world models or dense future rollouts, which are expensive and spend capacity on visual details weakly coupled to control. We present BRIDGE-WA, a lightweight world-action framework that distills a frozen future-change teacher into three compact priors: future tokens for intended outcomes, change maps for intervention support, and motion-flow maps for local transition direction. A WORLDBRIDGE conditions the action transformer on these priors through multi-source attention memories and spatial-temporal biases, while the teacher model is removed at inference. Across VLABench, RoboTwin2.0, LIBERO-Plus and real-robot evaluations, BRIDGE-WA improves task success, progress, and robustness, with particularly clear gains under out-of-distribution visual shifts. By focusing action generation on where and how the scene will change, BRIDGE-WA suppresses nuisance appearance factors such as background, lighting, and distractors, leading to better generalization without deployment-time dense future-image generation. Code and visualizations are available at: [🌐 BRIDGE-WA](#).

Keywords: Robot manipulation, World models, Vision-Language-Action models

1 Introduction

General-purpose manipulation requires more than recognizing objects and parsing instructions. A robot must also reason about intervention: after it pushes, grasps, pours, or inserts an object, which regions should change, how should they move, and which visual details should remain irrelevant? Recent embodied foundation models and vision-language-action (VLA) policies have made strong progress by scaling robot data and reusing web-scale visual-language priors [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11]. Action chunking, diffusion policies, and flow-based policies further improve temporal consistency and continuous control [12, 13, 14, 15]. Yet most direct VLA policies still learn a reactive mapping from current observation and language to actions, leaving future consequences weakly represented and making policies vulnerable to shortcut correlations in backgrounds, colors, and camera-specific appearance.

World models offer a natural way to close this gap: predict how the environment will evolve, then use that prediction to guide action. Latent dynamics, imagination-based agents, visual foresight, and robotic video prediction have shown that anticipated futures can support planning and manipulation [16, 17, 18, 19, 20, 21, 22, 23, 24, 25]. Recent world-action models extend this idea by jointly modeling future states and actions [26, 27, 28, 29]. The challenge is practical: running or decoding a large generative world model can be expensive, and dense image/video prediction allocates capacity to many pixels that do not affect control. For manipulation, the useful future is often smaller than a photorealistic rollout: it is a structured *change field* that tells the policy where the intended interaction takes effect and how the affected regions move.

We propose BRIDGE-WA, a lightweight world-action framework that conditions action generation on future world priors. Before policy distillation, we pre-train a future-prediction model on real-robot manipulation data to predict future observation structure from current observations, language instructions, and robot state. This separate pre-training stage gives the model manipulation-specific priors about object motion, contact-induced change, and task-conditioned outcomes. We then freeze the model and use its predictions to build three cached supervision targets for downstream policy training: future tokens T_f for global outcome information, a change map M_c for intervention-relevant spatial support, and a motion-flow map M_o for expected movement direction and magnitude. A lightweight predictor learns these priors from the current context, and WORLDBRIDGE uses them to condition the action transformer through multi-source memory and change/flow-guided attention. At test time, the predictive model and cached targets are removed; only the predictor and world-prior-conditioned policy remain.

This design occupies a middle ground between reactive VLAs and full world-action models. Compared with direct VLA policies, it provides explicit future-aware structure. Compared with WAMs that require large future-model components at deployment, it keeps the online path close to a direct action transformer. By emphasizing change and motion rather than raw appearance, the policy is encouraged to ignore visually salient but causally irrelevant scene factors, a property aligned with prior work on robot visual representations and robustness under distribution shift [30, 31, 32, 33, 34].

Our contributions are threefold:

- We propose BRIDGE-WA, a lightweight world-action framework that first pre-trains a 5B future-change world teacher to acquire manipulation-specific predictive priors, and then distills its future-world knowledge into compact priors during policy training, removing the need for expensive generative world models at inference.
- We design the WORLDBRIDGE block to condition an action transformer on these priors via multi-source attention memories and spatial-temporal biases, using a coarse-to-fine schedule that focuses the policy on action-relevant causal dynamics rather than nuisance appearance variations.
- We validate BRIDGE-WA on VLABench, RoboTwin2.0, LIBERO-Plus, and real-robot evaluations, showing strong gains in long-horizon manipulation, cross-category generalization, and out-of-distribution robustness.

2 Related Work

Vision-language-action models. Large-scale vision-language-action (VLA) models have shown that web-scale semantic priors and robot demonstrations can be combined for generalist robot control. Language-grounded systems such as SayCan, Inner Monologue, and Code-as-Policies use foundation models for planning or skill composition, while PaLM-E and Gato broaden the embodied-modeling interface beyond a single robot task [1, 2, 3, 4, 5]. RT-1 and RT-2 transfer visual-language and large-scale robot data into robotic actions, Open X-Embodiment and RT-X scale robot data across embodiments, and OpenVLA, Octo, X-VLA, BC-Z, VIMA, and RoboCat demonstrate increasingly general robot policies or prompt-conditioned manipulation [6, 7, 8, 9, 10, 35, 36, 37, 38]. In parallel, action-chunking and generative action models, including ACT, Implicit BC, Behavior Transformers, Diffusion Policy, and $\pi_0/\pi_{0.5}$, improve temporal consistency or multi-modal action prediction by predicting structured action distributions rather than a single low-level command [14, 12, 13, 15, 11, 39]. These policies are efficient and increasingly capable, but their core training signal is still an observation-to-action mapping. As noted by recent WAM surveys, such reactive policies do not explicitly model how the physical world evolves under intervention [27]. This gap motivates preserving the fast action-transformer interface of VLA policies while adding explicit representations of future outcomes, scene change, and motion.

World models for robotics. World models provide predictive representations for planning, policy learning, evaluation, and synthetic data generation in robotics [26]. Visual foresight, video prediction, and future-generation methods show that anticipated observations can guide manipulation

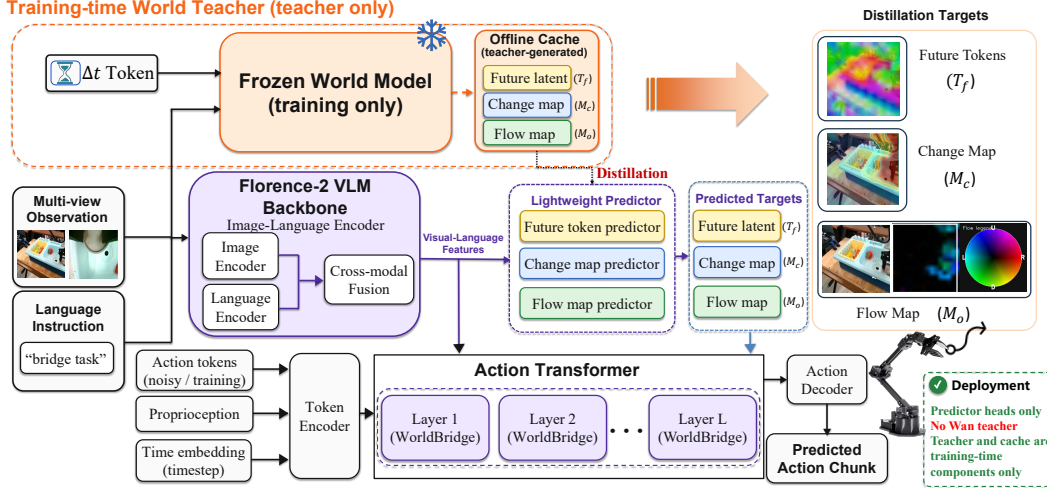


Figure 1: BRIDGE-WA represents the action-relevant future with three compact world-change representations: future tokens for the intended outcome, change maps for where the scene should change, and motion-flow maps for how the change should move. A lightweight predictor estimates these representations from the current robot context, and WORLDBRIDGE conditions the action transformer on them for world-aware action generation. The training-time future model and offline cache are removed at deployment.

and improve generalization [21, 22, 23, 24, 25, 40, 41]. Recent world-action models further couple future-state modeling with action generation through latent future reasoning or object-addressable states [27, 28, 29]. However, full image/video rollouts or large latent world modules add training and inference overhead and may spend capacity on visual details weakly coupled to control. BRIDGE-WA follows the same motivation but changes the interface: the deployed policy does not generate future images or run a large world model, but uses compact priors that summarize outcome, change, and motion.

Change-centric grounding. A complementary line of work shows that manipulation benefits from explicit spatial and actionable structure. Transporter Networks, CLIPort, and Where2Act model spatial displacement, separate semantic “what” from spatial “where”, or predict likely interaction regions [42, 43, 44], while PerAct, RVT, Act3D, and VoxPoser use 3D structure, view aggregation, and value maps for spatially precise control [45, 46, 47, 48]. Visual representation and interpretability studies further suggest that policies should retain task-relevant structure while suppressing spurious appearance correlations under distribution shift [30, 31, 32, 49]. BRIDGE-WA builds on this insight at the world-prior level: instead of only predicting where to act, it predicts where the scene should change and how that change should move, providing an action-relevant bottleneck that preserves useful foresight while filtering nuisance appearance.

3 BRIDGE-WA Model

Overall framework. BRIDGE-WA separates training-time world-prior acquisition from deployable action generation, as illustrated in Figure 1. A frozen future-change teacher produces three cached supervision targets: future tokens for the intended outcome, change maps for intervention support, and motion-flow maps for local movement. A lightweight predictor learns to recover these priors from the current robot context, and WORLDBRIDGE injects the predicted priors into the action transformer as memory tokens and attention guidance. At inference, the teacher and offline cache are removed; the policy only runs the predictor and the world-prior-conditioned action transformer.

3.1 Problem Formulation and Design Principle

We consider language-conditioned manipulation from multi-view RGB observations, proprioception, and language:

$$x_t = (\{I_t^v\}_{v=1}^V, q_t, \ell), \quad \mathbf{a}_{t:t+H-1} = (a_t, \dots, a_{t+H-1}). \quad (1)$$

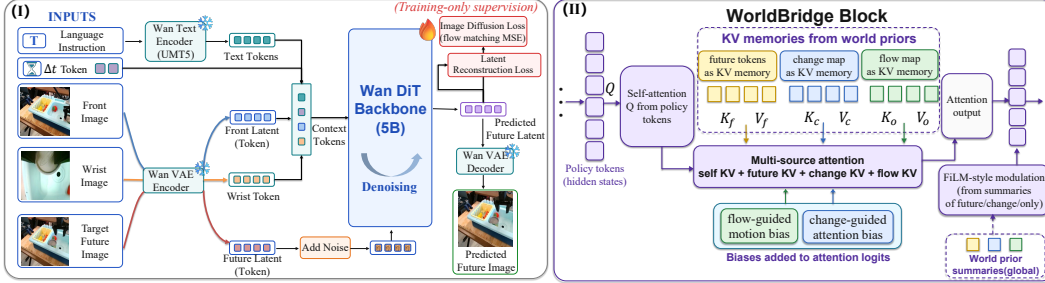


Figure 2: (I) Robot-state-conditioned teacher world model based on the Wan2.2-5B generative backbone. (II) WORLDBRIDGE conditioning block. Predicted future, change, and flow representations are projected into memory tokens and attention biases that condition the action transformer without invoking the training-time teacher world model.

A direct VLA learns $\pi_\theta(\mathbf{a}_{t:t+H-1} \mid x_t)$. BRIDGE-WA instead predicts a compact world prior

$$\hat{Z}_t = g_\phi(x_t), \quad Z_t = (T_{f,t}, M_{c,t}, M_{o,t}), \quad \pi_{\theta,\phi}(\mathbf{a}_{t:t+H-1} \mid x_t) = \pi_\theta(\mathbf{a}_{t:t+H-1} \mid x_t, \hat{Z}_t). \quad (2)$$

Here T_f is a future-outcome token, $M_c \in [0, 1]^{V \times h \times w}$ is a spatial change map, and $M_o \in \mathbb{R}^{V \times h \times w \times 2}$ is an image-coordinate motion-flow map. The design principle is to replace dense future-image prediction with an action-sufficient world-change summary: outcome, intervention support, and local motion. This preserves predictive information useful for control while filtering background, lighting, and other nuisance appearance factors.

3.2 World-Prior Distillation and Prediction

During training, a frozen world teacher \mathcal{W}_ψ defines the predictive supervision space for BRIDGE-WA. \mathcal{W}_ψ is not an off-the-shelf video generator: we construct, as illustrated in the Fig. 2(I), a robot-state-conditioned world model based on the Wan2.2-5B generative backbone, so that its prediction target is aligned with manipulation-induced scene transitions. Before any Bridge-WA distillation, \mathcal{W}_ψ is pre-trained on real-robot manipulation trajectories from BridgeData V2 [50], learning to map current images, language, and robot state to future observation structure; this stage is independent of the action policy parameters. After pre-training, \mathcal{W}_ψ is frozen and used only for downstream cache construction.

Given current observations, language, and robot state, \mathcal{W}_ψ provides three supervision targets: T_f^* , a compact summary of the predicted future; M_c^* , a soft intervention-support map measuring where the future differs from the current scene; and M_o^* , a motion field describing how changed regions move toward the future state. These targets can be extracted from \mathcal{W}_ψ latents, decoded futures, or induced correspondences, and are stored as an offline cache. The deployable component is a lightweight predictor g_ϕ that infers $\hat{Z}_t = (\hat{T}_{f,t}, \hat{M}_{c,t}, \hat{M}_{o,t})$ from the current images, proprioception, and instruction. It is trained with a structured distillation objective: latent regression and cosine alignment for future tokens, spatial agreement for change maps, and vector-field agreement for motion flow. Thus policy training does not back-propagate through \mathcal{W}_ψ or repeatedly decode future rollouts, and deployment only requires recovering action-useful world summaries rather than photorealistic future images. Appendix give the cache construction and loss details.

3.3 World-Conditioned Action Learning and Deployment

The predicted priors condition the action transformer through WORLDBRIDGE layers, as illustrated in Fig. 2(II). WORLDBRIDGE converts future tokens, change maps, and motion-flow maps into policy-readable conditioning memories: future tokens provide global outcome context, change tokens identify the spatial support of intervention, and flow tokens encode directional motion. Queries remain policy-centered, so the action transformer decides how strongly to read from each prior source. Change and flow maps also provide attention guidance toward likely-changing and motion-consistent regions, without requiring future-pixel reconstruction.

The conditioning scheme is layered rather than uniform. For each modality $m \in \{f, c, o\}$, we define an active layer range by a start layer s_m and a layer budget n_m :

$$\mathcal{L}_m = \{l \mid s_m \leq l < \min(L, s_m + n_m)\}, \quad \mathcal{A}_l = \{m \in \{f, c, o\} \mid l \in \mathcal{L}_m\}, \quad (3)$$

where L is the transformer depth and \mathcal{A}_l is the set of active priors at layer l . This schedule follows a coarse-to-fine order: future conditioning is used earlier as a task-level outcome prior, change conditioning grounds the intended interaction in scene regions, and flow conditioning is routed closer to action decoding where local displacement is most useful.

At an active layer, inactive priors are omitted from both the projected-conditioning pathway and the modulation path. The WORLDBRIDGE attention update is

$$\text{Attn}_{\text{wb}}^l = \text{softmax} \left(\frac{Q^l[K_s^l; \{K_m^l\}_{m \in \mathcal{A}_l}]^\top}{\sqrt{d}} + \sum_{m \in \mathcal{A}_l \cap \{c, o\}} B_m^l \right) [V_s^l; \{V_m^l\}_{m \in \mathcal{A}_l}], \quad (4)$$

where K_s^l, V_s^l are the standard self-attention key/value, K_m^l, V_m^l are projected world-prior memories, and B_c^l, B_o^l are change- and flow-derived attention biases. This progressive interface keeps the policy’s self-attention stream intact while presenting outcome, change, and motion information at compatible abstraction levels.

Training combines action imitation with world-prior distillation. The action term can use a flow, diffusion, deterministic, or autoregressive action head; BRIDGE-WA only requires the action transformer to condition on the predicted world-prior summary. At inference time, the teacher model and offline cache are removed, leaving one lightweight world-prior prediction followed by world-prior-conditioned action generation.

4 Experiments

Simulation setup. We evaluate BRIDGE-WA on three simulated benchmarks with complementary evaluation targets. All policies use RGB observations, robot state, and language instructions, and output continuous action chunks. VLABench [51] is used as the primary long-horizon manipulation benchmark: we adopt its multi-view setting with front, secondary, and wrist observations, train and evaluate on five tracks, and report success rate (SR), intention score (IS), and progress score (PS). RoboTwin 2.0 [52] evaluates bimanual manipulation under domain-randomized scenes; we follow its official Easy/Hard evaluation protocol. Fig. 4(I), (II), and (III) present the success rate of each task and the average success rate on the 15-task subset of RoboTwin2.0 under the easy and hard settings. LIBERO-Plus [53], built on LIBERO [54], is used as a zero-shot robustness benchmark: the base language-conditioned manipulation tasks are kept fixed while the evaluation distribution is independently perturbed along camera viewpoint, robot initial state, language, lighting, background, image noise, and object layout. All methods are evaluated without training on LIBERO-Plus data, and Tab. 3 reports success rates for each perturbation dimension and the average.

Real-robot setup. We evaluate on the DoBot Nova2 platform with an Intel RealSense D455 third-person camera and a D405 wrist camera. Demonstrations are stored in LeRobot format with synchronized third-person/wrist RGB, cartesian robot states, and absolute cartesian actions, with RGB standardized to 640×480 . The suite contains five tasks: `PickGrape`, `StackBowls`, `PushButtons`, `CollectFruits`, and `PutItemInDrawer`, each with 50 training demonstrations. Easy evaluation uses 10 near in-distribution trials per task; Hard evaluation uses 15 trials per task across tablecloth, distractor-object, and lighting shifts, with five trials per factor, as illustrated in Fig. 3(I). All methods use matched initial states, language instructions, and success criteria, and Tab. 4 and Fig. 4(V) report per-task and average success rates for each track.

4.1 Main Results

Overall comparison. We evaluate BRIDGE-WA across VLABench, RoboTwin 2.0, LIBERO-Plus, and a Dobot real-robot suite, covering long-horizon manipulation, bimanual domain randomization, controlled zero-shot robustness, and real visual shifts. Across these settings, the predictive

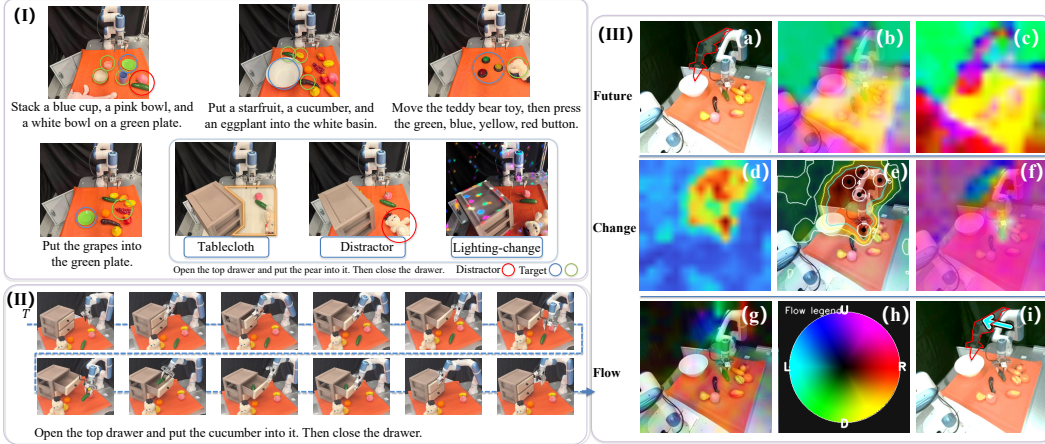


Figure 3: Real-robot visualization analysis. (I) Five manipulation tasks and OOD evaluation scenes, with target and distractor regions marked. (II) Execution sequence for “Open the top drawer and put the cucumber into it. Then close the drawer.” (III) Teacher-target visualizations from the Dobot Nova2 robot scene, show that BRIDGE-WA distills the future into outcome, change, and motion priors that support world-aware action generation. In (e), the numbered circles denote high-response peaks in the change map, corresponding to localized regions most likely to undergo task-induced change. In (h), the color wheel is the flow legend, where hue encodes motion direction and the L/R/U/D labels indicate left, right, up, and down displacements.

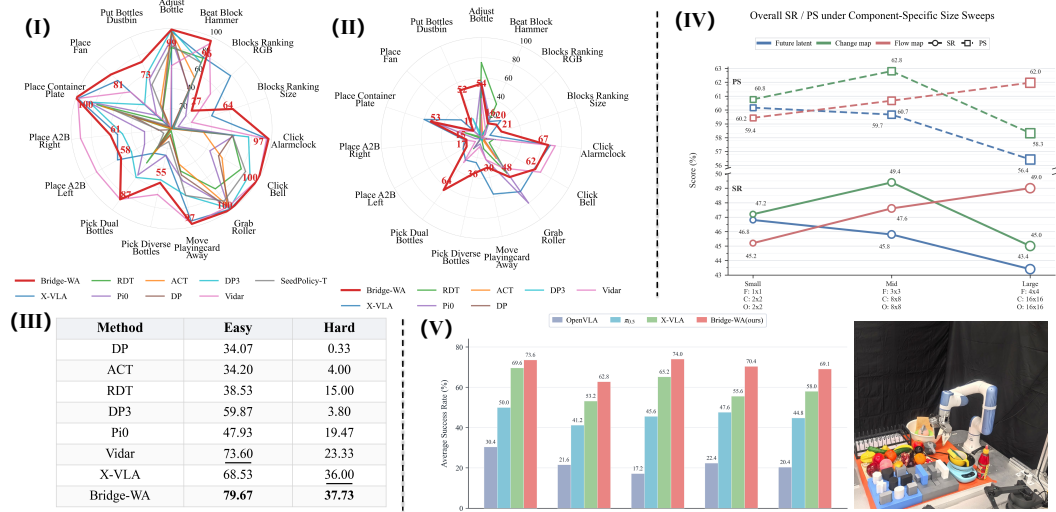


Figure 4: Experimental results across simulation, ablation, and real-robot evaluations. (I)-(III) RoboTwin 2.0 results on the 15-task subset, including per-task radar plots under the Easy and Hard tracks and average success rates for each model across both tracks. (IV) VLABench-5 ablation results, reporting the average success rate (SR) and progress score (PS) for each world-prior and conditioning configuration. Here, 1×1 , etc. denote the spatial resolution of future, change, or flow priors before flattening into conditioning tokens or spatial biases. (V) Dobot real-robot evaluation results, reporting the average success rate under each evaluation track.

world-prior interface improves average performance over baselines. More **real-robot platforms (Franka)** results are provided in the **Appendix**.

Simulation results. On VLABench, BRIDGE-WA achieves the best average SR (52.8%) in Tab. 1, improving over the strongest SR baseline by 9.7 points, and also obtains the best average IS/PS (71.2%/64.0%) in Tab. 2. The gains are strongest on in-distribution, common-sense, and semantic-instruction tracks, indicating that predicted change and motion help the policy make progress toward the intended outcome. On RoboTwin 2.0, Fig. 4(I)–(III) shows that BRIDGE-WA reaches the best

Method	In-Dist.	Cross-Cat.	Common Sense	Semantic Instr.	Unseen Texture	Avg. SR
$\pi_{0.5}$ [39]	40.6	22.6	18.0	16.1	25.6	24.6
π_0 [11]	47.0	21.2	29.1	17.3	32.2	29.4
π_0 -Fast [55] (relative chunk)	29.1	18.1	21.1	19.9	23.6	22.4
X-VLA [35]	40.0	11.0	33.0	39.0	24.0	29.4
π_0 -Fast [55] (delta chunk)	<u>51.2</u>	34.2	<u>41.7</u>	<u>44.2</u>	<u>44.4</u>	<u>43.1</u>
BRIDGE-WA (Ours)	78.0	<u>23.0</u>	51.1	67.0	45.0	52.8

Table 1: VLABench SR results. The SR block reports five public categories: in-distribution, cross-category, common-sense, semantic-instruction, and unseen-texture; Avg. SR averages these five categories.

Method	In-Dist.		Cross-Cat.		Common Sense		Semantic Instr.		Unseen Texture		Avg.	
	IS	PS	IS	PS	IS	PS	IS	PS	IS	PS	IS	PS
π_0^\diamond [11]	67.8	62.7	44.0	33.6	54.9	43.0	<u>58.0</u>	38.7	50.6	42.5	55.0	44.1
$\pi_{0.5}^\diamond$ [39]	75.0	60.8	<u>49.6</u>	<u>35.3</u>	57.5	41.6	57.1	30.3	62.0	47.4	60.2	43.1
ACoT-VLA $^\diamond$ [56]	79.8	66.1	54.1	38.9	52.3	37.8	56.8	39.6	<u>74.6</u>	<u>54.6</u>	63.5	47.4
X-VLA [35]	88.0	<u>66.5</u>	42.0	24.0	<u>66.0</u>	<u>51.7</u>	82.0	<u>62.5</u>	73.0	51.2	<u>70.2</u>	<u>51.2</u>
BRIDGE-WA (Ours)	<u>85.0</u>	85.8	39.0	28.8	74.2	64.4	82.0	80.3	76.0	60.3	71.2	64.0

Table 2: VLABench IS/PS comparison. IS is intention score and PS is progress score, reported over the five public VLABench categories. Rows marked with \diamond are reproduced from the 60K-step frozen-backbone setting in ACoT-VLA [56]; X-VLA and BRIDGE-WA use the same evaluation categories and metrics.

average score on the 15-task subset, with 79.7% Easy and 37.7% Hard success; compared with X-VLA, the mean across Easy/Hard improves from 52.3% to 58.7%. Together, these results show that the world-prior representation improves both long-horizon manipulation and domain-randomized bimanual control.

On VLABench, BRIDGE-WA achieves the best average SR (52.8%) in Tab. 1, improving over the strongest SR baseline by 9.7 points, and also obtains the best average IS/PS (71.2%/64.0%) in Tab. 2. The gains are strongest on in-distribution, common-sense, and semantic-instruction tracks, indicating that predicted change and motion help the policy make progress toward the intended outcome. On RoboTwin 2.0, Fig. 4(I)–(III) shows that BRIDGE-WA reaches the best average score on the 15-task subset, with 79.7% Easy and 37.7% Hard success; compared with X-VLA, the mean across Easy/Hard improves from 52.3% to 58.7%. On LIBERO-Plus, Tab. 3 shows that BRIDGE-WA achieves the highest average zero-shot success rate (72.1%). The clearest advantages appear under robot-state and lighting perturbations, where BRIDGE-WA reaches 92.8% and 96.7%, and it also achieves the best background robustness at 94.3%. This pattern suggests that world-prior conditioning helps separate task-caused scene changes from nuisance appearance variation. Performance under camera-view perturbation remains lower than the strongest camera-view result, indicating that viewpoint robustness still depends on stronger geometric or view-invariant representations. Together, these results show that the world-prior representation improves long-horizon manipulation, domain-randomized bimanual control, and controlled zero-shot robustness.

Real-robot results. On the Dobot Nova2 platform, BRIDGE-WA achieves the best average Easy and Hard success rates, 73.6% and 69.1%, as shown in Tab. 4. Compared with X-VLA, this improves the Easy and Hard averages by 4.0 and 11.1 points, with larger gains under OOD distractors, lighting, and tablecloth shifts. These results indicate that change- and motion-conditioned action decoding improves robustness to nuisance appearance.

4.2 Ablation Studies

We ablate the three world-prior sources and the conditioning design on the VLABench-5 subset. Each configuration is trained with the same five tasks and the same optimization budget, and Tab. 5 reports the evaluation track with 100 episodes per task. We further perform a single-factor size sweep for the three world-prior interfaces: future latent tokens, change maps, and motion-flow maps, in Fig. 4(IV).

Component analysis. The ablation shows that the three world priors are not interchangeable. Change-only improves the no-world baseline from 46.3% to 48.4%, and Future+Flow reaches

Method	Camera	Robot	Language	Light	Background	Noise	Layout	Avg.
OpenVLA [9]	0.8	3.5	23.0	8.1	34.8	15.2	28.5	15.6
WorldVLA [57]	0.1	27.9	41.6	43.7	17.1	10.9	38.0	25.0
NORA [58]	2.2	37.0	65.1	45.7	58.6	12.8	62.1	39.0
UniVLA [59]	1.8	46.2	69.6	69.0	81.0	21.2	31.9	42.9
FastWAM [60]	16.4	44.5	68.9	78.2	53.7	37.7	60.7	51.5
π_0 [11]	13.8	6.0	58.8	85.0	81.4	79.0	68.9	53.6
π_0 -Fast [55]	65.1	21.6	61.0	73.2	73.2	74.4	68.8	61.6
RIPT-VLA [61]	55.2	31.2	77.6	88.4	91.6	73.5	74.2	68.4
OpenVLA-OFT [62]	56.4	31.9	79.5	88.7	93.3	75.8	74.2	69.6
BRIDGE-WA (Ours)	25.0	92.8	77.2	96.7	94.3	64.8	69.3	72.1

Table 3: Zero-shot evaluation on LIBERO-Plus. All methods are evaluated without training on LIBERO-Plus data. Results are success rates across perturbation dimensions.

Method	PickGrape					StackBowls					PushButtons				
	Easy	Dist.	Light	T-Cloth	Hard	Easy	Dist.	Light	T-Cloth	Hard	Easy	Dist.	Light	T-Cloth	Hard
OpenVLA [9]	3/10	1/5	2/5	1/5	4/15	1.2/10	0.9/5	0.3/5	0.6/5	1.8/15	8/10	<u>3.5/5</u>	<u>2/5</u>	<u>4/5</u>	<u>9.5/15</u>
$\pi_{0.5}$ [11]	<u>8/10</u>	<u>3/5</u>	<u>4/5</u>	<u>5/5</u>	<u>12/15</u>	4/10	1.5/5	2.9/5	2.1/5	6.5/15	3/10	2.5/5	0.5/5	1.5/5	4.5/15
X-VLA [35]	<u>8/10</u>	<u>3/5</u>	<u>4/5</u>	<u>3/5</u>	10/15	<u>3.4/10</u>	1.2/5	2.3/5	1.6/5	5.1/15	8/10	<u>3.5/5</u>	<u>2/5</u>	<u>4/5</u>	<u>9.5/15</u>
BRIDGE-WA (Ours)	9/10	4/5	5/5	5/5	14/15	3/10	0.9/5	2/5	2.1/5	5/15	9/10	4.5/5	4.5/5	<u>3.5/5</u>	12.5/15

Method	CollectFruits					PutItemInDrawer					Avg. SR				
	Easy	Dist.	Light	T-Cloth	Hard	Easy	Dist.	Light	T-Cloth	Hard	Easy	Dist.	Light	T-Cloth	Hard
OpenVLA [9]	2/10	0/5	0/5	0/5	0/15	1/10	0/5	0/5	0/5	0/15	30.4%	21.6%	17.2%	22.4%	20.4%
$\pi_{0.5}$ [11]	7/10	3.3/5	4/5	<u>3.3/5</u>	10.6/15	<u>3/10</u>	<u>0/5</u>	0/5	0/5	0/15	50.0%	41.2%	45.6%	47.6%	44.8%
X-VLA [35]	<u>8.4/10</u>	<u>3.6/5</u>	<u>5/5</u>	<u>3.3/5</u>	11.9/15	7/10	<u>2/5</u>	<u>3/5</u>	<u>2/5</u>	7/15	69.6%	53.2%	65.2%	55.6%	58.0%
BRIDGE-WA (Ours)	8.8/10	4.3/5	<u>5/5</u>	<u>4/5</u>	13.3/15	7/10	<u>2/5</u>	<u>2/5</u>	<u>3/5</u>	7/15	73.6%	62.8%	74.0%	70.4%	69.1%

Table 4: Dobot real-robot evaluation results. Easy uses 10 tests per task; Hard contains three tracks with 5 tests each: distractors, lighting generalization, and tablecloth generalization. Each PushButtons test is scored as 0, 0.5, or 1; each CollectFruits and StackBowls test is scored as 0, 0.3, 0.6, or 1; all other tasks are scored as 0 or 1. Hard reports the total score over the three hard tracks.

48.6%, while Flow-only drops to 45.2%. This suggests that motion cues are useful only when grounded by outcome or spatial-change information, and that simply adding more priors can introduce competition inside a fixed-capacity action transformer.

The conditioning design is equally important. Attention-only conditioning falls below the no-world baseline, no-gate conditioning improves to 49.3%, and layered conditioning performs best at 50.4% (+4.1 points). In the layered design, future tokens are introduced in earlier transformer layers as global outcome context, change maps are injected in middle layers to localize interaction regions, and flow maps are emphasized closer to action decoding to guide local displacement. The largest gains appear on Tracks 2 and 3, indicating that this coarse-to-fine routing is especially useful for cross-category transfer and common sense grounding.

World-prior Size Sweep Analysis The size sweep in Fig. 4(IV) shows that each prior needs a different spatial resolution budget. Future latents work best as compact global summaries: 1×1 reaches 46.8% SR and 60.2% PS, while 4×4 drops to 43.4% SR and 56.4% PS. Change maps peak at the intermediate 8×8 scale (49.4% SR / 62.8% PS), balancing localization and stability. Flow maps benefit from finer resolution, improving from 45.2% SR at 2×2 to 49.0% SR at 16×16 . Overall, BRIDGE-WA benefits from asymmetric capacity allocation: compact future tokens, medium-resolution change maps, and higher-resolution flow maps.

4.3 Visualization Analysis

The learned world-prior interface is visualized on real-robot rollouts in Fig. 3. Panels (I) and (II) show the task context and execution sequence, while Panel (III) decomposes the prediction into three views: future tokens capture the intended outcome, change maps localize the object and contact regions that should change, and flow maps indicate the direction of local motion. The alignment among these views shows that BRIDGE-WA provides the policy with a compact, action-relevant description of scene evolution, helping it focus on controllable changes rather than background, lighting, or distractor appearance.

Configuration	World token			Success rate (%)				
	F	C	O	Track 1	Track 2	Track 3	Track 4	Avg.
<i>Token-combination matrix</i>								
WorldBridge w/o world tokens				51.5	27.3	50.0	53.5	46.3
Future only	✓			57.7	30.0	50.0	48.5	47.2
Change only		✓		60.8	27.3	53.8	48.5	48.4
Flow only			✓	55.4	26.4	50.8	45.4	45.2
Future + Change	✓	✓		56.2	23.6	51.5	43.8	44.6
Future + Flow	✓		✓	60.8	26.4	56.2	47.7	48.6
Change + Flow		✓	✓	53.1	31.8	53.1	48.1	47.1
Future + Change + Flow	✓	✓	✓	56.9	29.1	50.0	47.7	46.6
<i>Other conditioning designs</i>								
Attention-only conditioning	✓	✓	✓	53.8	24.5	50.8	46.9	44.8
No-gate conditioning	✓	✓	✓	56.2	29.1	56.9	51.9	49.3
Layered conditioning	✓	✓	✓	<u>59.2</u>	<u>30.9</u>	57.7	50.8	50.4

Table 5: VLABench-5 ablation results. F, C, and O denote future tokens, change maps, and motion-flow maps. Tracks 1–4 refer to in-distribution, cross-category, common-sense, and semantic-instruction, respectively.

5 Conclusion and Limitations

This work shows that useful world modeling for robot control need not take the form of deployment-time future video generation. By distilling a frozen future-change teacher into compact future, change, and motion-flow priors, BRIDGE-WA gives an action policy access to the parts of the predicted future that matter for control: what outcome to reach, where the scene should change, and how that change should move. The results across VLABench, RoboTwin2.0, and Dobot real-robot evaluations suggest that such structured world priors improve task success and progress, with particularly clear gains under OOD visual shifts where policies must ignore nuisance appearance and focus on task-caused change. BRIDGE-WA is still limited by the quality and coverage of the training-time teacher, the short-horizon image-space nature of the priors, and the extra cost of building the offline cache. Future work should extend the framework to geometry-aware priors, longer-horizon abstractions, and broader cross-embodiment evaluation.

References

- [1] M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, C. Fu, K. Gopalakrishnan, K. Hausman, et al. Do as i can, not as i say: Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022.
- [2] W. Huang, F. Xia, T. Xiao, H. Chan, J. Liang, P. Florence, A. Zeng, J. Tompson, I. Mordatch, Y. Chebotar, et al. Inner monologue: Embodied reasoning through planning with language models. In *Conference on Robot Learning*, pages 1769–1782. PMLR, 2023.
- [3] J. Liang, W. Huang, F. Xia, P. Xu, K. Hausman, B. Ichter, P. Florence, and A. Zeng. Code as policies: Language model programs for embodied control. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023.
- [4] D. Driess, F. Xia, M. S. Sajjadi, C. Lynch, A. Chowdhery, B. Ichter, A. Wahid, J. Tompson, Q. Vuong, T. Yu, et al. Palm-e: an embodied multimodal language model. In *Proceedings of the 40th International Conference on Machine Learning*, pages 8469–8488, 2023.
- [5] S. Reed, K. Zolna, E. Parisotto, S. G. Colmenarejo, A. Novikov, G. Barth-maroon, M. Giménez, Y. Sulsky, J. Kay, J. T. Springenberg, et al. A generalist agent. *Transactions on Machine Learning Research*.
- [6] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu, et al. Rt-1: Robotics transformer for real-world control at scale. *Robotics: Science and Systems XIX*, 2023.
- [7] B. Zitkovich, T. Yu, S. Xu, P. Xu, T. Xiao, F. Xia, J. Wu, P. Wohlhart, S. Welker, A. Wahid, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. In *Conference on Robot Learning*, pages 2165–2183. PMLR, 2023.
- [8] A. O’Neill, A. Rehman, A. Maddukuri, A. Gupta, A. Padalkar, A. Lee, A. Pooley, A. Gupta, A. Mandlekar, A. Jain, et al. Open x-embodiment: Robotic learning datasets and rt-x models: Open x-embodiment collaboration 0. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6892–6903. IEEE, 2024.
- [9] M. J. Kim, K. Pertsch, S. Karamcheti, T. Xiao, A. Balakrishna, S. Nair, R. Rafailov, E. P. Foster, P. R. Sanketi, Q. Vuong, et al. Openvla: An open-source vision-language-action model. In *8th Annual Conference on Robot Learning*.
- [10] O. M. Team, D. Ghosh, H. Walke, K. Pertsch, K. Black, O. Mees, S. Dasari, J. Hejna, T. Kreiman, C. Xu, et al. Octo: An open-source generalist robot policy. *arXiv preprint arXiv:2405.12213*, 2024.
- [11] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, L. Groom, K. Hausman, B. Ichter, et al. Pi0: A vision-language-action flow model for general robot control. *arXiv preprint arXiv:2410.24164*, 2024.
- [12] P. Florence, C. Lynch, A. Zeng, O. A. Ramirez, A. Wahid, L. Downs, A. Wong, J. Lee, I. Mordatch, and J. Tompson. Implicit behavioral cloning. In *Conference on robot learning*, pages 158–168. PMLR, 2022.
- [13] N. M. Shafiullah, Z. Cui, A. A. Altanzaya, and L. Pinto. Behavior transformers: Cloning k modes with one stone. *Advances in neural information processing systems*, 35:22955–22968, 2022.
- [14] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn. Learning fine-grained bimanual manipulation with low-cost hardware. *arXiv preprint arXiv:2304.13705*, 2023.

- [15] C. Chi, Z. Xu, S. Feng, E. Cousineau, Y. Du, B. Burchfiel, R. Tedrake, and S. Song. Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research*, 44(10-11):1684–1704, 2025.
- [16] D. Ha and J. Schmidhuber. World models. *arXiv preprint arXiv:1803.10122*, 2018.
- [17] D. Hafner, T. Lillicrap, I. Fischer, R. Villegas, D. Ha, H. Lee, and J. Davidson. Learning latent dynamics for planning from pixels. In *International conference on machine learning*, pages 2555–2565. PMLR, 2019.
- [18] D. Hafner, T. Lillicrap, J. Ba, and M. Norouzi. Dream to control: Learning behaviors by latent imagination. In *International Conference on Learning Representations*, 2020.
- [19] D. Hafner, J. Pasukonis, J. Ba, and T. Lillicrap. Mastering diverse domains through world models. *arXiv preprint arXiv:2301.04104*, 2023.
- [20] J. Schrittwieser, I. Antonoglou, T. Hubert, K. Simonyan, L. Sifre, S. Schmitt, A. Guez, E. Lockhart, D. Hassabis, T. Graepel, et al. Mastering atari, go, chess and shogi by planning with a learned model. *Nature*, 588(7839):604–609, 2020.
- [21] C. Finn and S. Levine. Deep visual foresight for planning robot motion. In *2017 IEEE international conference on robotics and automation (ICRA)*, pages 2786–2793. IEEE, 2017.
- [22] F. Ebert, C. Finn, S. Dasari, A. Xie, A. Lee, and S. Levine. Visual foresight: Model-based deep reinforcement learning for vision-based robotic control. *arXiv preprint arXiv:1812.00568*, 2018.
- [23] M. Babaeizadeh, M. T. Saffar, S. Nair, S. Levine, C. Finn, and D. Erhan. Fitvid: Overfitting in pixel-level video prediction. *arXiv preprint arXiv:2106.13195*, 2021.
- [24] Y. Du, S. Yang, B. Dai, H. Dai, O. Nachum, J. Tenenbaum, D. Schuurmans, and P. Abbeel. Learning universal policies via text-guided video generation. *Advances in neural information processing systems*, 36:9156–9172, 2023.
- [25] S. Zhou, Y. Du, J. Chen, Y. Li, D.-Y. Yeung, and C. Gan. Robodreamer: Learning compositional world models for robot imagination. In *International Conference on Machine Learning*, pages 61885–61896. PMLR, 2024.
- [26] B. Hou, G. Li, J. Jia, T. An, X. Guo, S. Leng, H. Geng, Y. Ze, T. Harada, P. Torr, et al. World model for robot learning: A comprehensive survey. *arXiv preprint arXiv:2605.00080*, 2026.
- [27] S. Wang, J. Shi, Z. Fu, X. He, F. Liu, C. Yang, Y. Zhou, Z. Fei, J. Gong, J. Fu, et al. World action models: The next frontier in embodied ai. *arXiv preprint arXiv:2605.12090*, 2026.
- [28] H. Luo, W. Zhang, Y. Feng, S. Zheng, H. Xu, C. Xu, Z. Xi, Y. Fu, and Z. Lu. Being-h0. 7: A latent world-action model from egocentric videos. *arXiv preprint arXiv:2605.00078*, 2026.
- [29] Y. Liu, P. Sun, S. Li, Y. Xie, L. Zhang, X. Chao, S. Dong, F. Chen, X.-P. Zhang, and W. Ding. Oa-wam: Object-addressable world action model for robust robot manipulation. *arXiv preprint arXiv:2605.06481*, 2026.
- [30] S. Nair, A. Rajeswaran, V. Kumar, C. Finn, and A. Gupta. R3m: A universal visual representation for robot manipulation. In *Conference on Robot Learning*, pages 892–909. PMLR, 2023.
- [31] Y. J. Ma, S. Sodhani, D. Jayaraman, O. Bastani, V. Kumar, and A. Zhang. Vip: Towards universal visual reward and representation via value-implicit pre-training. In *The Eleventh International Conference on Learning Representations*, 2023.

- [32] S. Karamcheti, S. Nair, A. S. Chen, T. Kollar, C. Finn, D. Sadigh, and P. Liang. Language-driven representation learning for robotics. *Robotics: Science and Systems*, 2023.
- [33] D. Hendrycks and T. Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *International Conference on Learning Representations*, 2019.
- [34] D. Hendrycks, S. Basart, N. Mu, S. Kadavath, F. Wang, E. Dorundo, R. Desai, T. Zhu, S. Parajuli, M. Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 8320–8329. IEEE, 2021.
- [35] J. Zheng, J. Li, Z. Wang, D. Liu, X. Kang, Y. Feng, Y. Zheng, J. Zou, Y. Chen, J. Zeng, et al. X-vla: Soft-prompted transformer as scalable cross-embodiment vision-language-action model. *International Conference on Learning Representation*, 2026.
- [36] E. Jang, A. Irpan, M. Khansari, D. Kappler, F. Ebert, C. Lynch, S. Levine, and C. Finn. Bc-z: Zero-shot task generalization with robotic imitation learning. In *5th Annual Conference on Robot Learning*.
- [37] Y. Jiang, A. Gupta, Z. Zhang, G. Wang, Y. Dou, Y. Chen, L. Fei-Fei, A. Anandkumar, Y. Zhu, and L. Fan. Vima: Robot manipulation with multimodal prompts. In *International Conference on Machine Learning*, pages 14975–15022. PMLR, 2023.
- [38] K. Bousmalis, G. Vezzani, D. Rao, C. M. Devin, A. X. Lee, M. B. Villalonga, T. Davchev, Y. Zhou, A. Gupta, A. Raju, et al. Robocat: A self-improving generalist agent for robotic manipulation. *Transactions on Machine Learning Research*.
- [39] K. Black, N. Brown, J. Darpinian, K. Dhabalia, D. Driess, A. Esmail, M. R. Equi, C. Finn, N. Fusai, M. Y. Galliker, et al. Pi0.5: a vision-language-action model with open-world generalization. In *9th Annual Conference on Robot Learning*.
- [40] H. Wu, Y. Jing, C. Cheang, G. Chen, J. Xu, X. Li, M. Liu, H. Li, and T. Kong. Unleashing large-scale video generative pre-training for visual robot manipulation. In *International Conference on Learning Representations*, volume 2024, pages 10641–10662, 2024.
- [41] K. Zhou, Y. Chen, F. Zhan, H. Hua, G. Chen, X. Chang, A. Qu, Y. Du, Z. Liu, P. P. Liang, and M. Wang. Gem-4d: Geometry-enhanced video world models for robot manipulation. *arXiv preprint arXiv:2605.22882*, 2026.
- [42] A. Zeng, P. Florence, J. Tompson, S. Welker, J. Chien, M. Attarian, T. Armstrong, I. Krasin, D. Duong, V. Sindhwani, et al. Transporter networks: Rearranging the visual world for robotic manipulation. In *Conference on Robot Learning*, pages 726–747. PMLR, 2021.
- [43] M. Shridhar, L. Manuelli, and D. Fox. Cliport: What and where pathways for robotic manipulation. In *Conference on robot learning*, pages 894–906. PMLR, 2022.
- [44] K. Mo, L. J. Guibas, M. Mukadam, A. Gupta, and S. Tulsiani. Where2act: From pixels to actions for articulated 3d objects. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6813–6823, 2021.
- [45] M. Shridhar, L. Manuelli, and D. Fox. Perceiver-actor: A multi-task transformer for robotic manipulation. In *Conference on Robot Learning*, pages 785–799. PMLR, 2023.
- [46] A. Goyal, J. Xu, Y. Guo, V. Blukis, Y.-W. Chao, and D. Fox. Rvt: Robotic view transformer for 3d object manipulation. In *Conference on Robot Learning*, pages 694–710. PMLR, 2023.
- [47] T. Gervet, Z. Xian, N. Gkanatsios, and K. Fragkiadaki. Act3d: 3d feature field transformers for multi-task robotic manipulation. In J. Tan, M. Toussaint, and K. Darvish, editors, *Proceedings of The 7th Conference on Robot Learning*, volume 229 of *Proceedings of Machine Learning Research*, pages 3949–3965. PMLR, 06–09 Nov 2023. URL <https://proceedings.mlr.press/v229/gervet23a.html>.

- [48] W. Huang, C. Wang, R. Zhang, Y. Li, J. Wu, and L. Fei-Fei. Voxposer: Composable 3d value maps for robotic manipulation with language models. In *Conference on Robot Learning*, pages 540–562. PMLR, 2023.
- [49] H. Zhang, M. Xu, A. Dhafer, S. Yue, H. Dong, and Z. D. Hao. Embodied interpretability: Linking causal understanding to generalization in vision-language-action models. *arXiv preprint arXiv:2605.00321*, 2026.
- [50] H. R. Walke, K. Black, T. Z. Zhao, Q. Vuong, C. Zheng, P. Hansen-Estruch, A. W. He, V. Myers, M. J. Kim, M. Du, et al. Bridgedata v2: A dataset for robot learning at scale. In *7th Annual Conference on Robot Learning*.
- [51] S. Zhang, Z. Xu, P. Liu, X. Yu, Y. Li, Q. Gao, Z. Fei, Z. Yin, Z. Wu, Y.-G. Jiang, et al. Vlabench: A large-scale benchmark for language-conditioned robotics manipulation with long-horizon reasoning tasks, 2024a. URL <https://arxiv.org/abs/2412.18194>.
- [52] T. Chen, Z. Chen, B. Chen, Z. Cai, Y. Liu, Z. Li, Q. Liang, X. Lin, Y. Ge, Z. Gu, et al. Robotwin 2.0: A scalable data generator and benchmark with strong domain randomization for robust bimanual robotic manipulation. *arXiv preprint arXiv:2506.18088*, 2025.
- [53] S. Fei, S. Wang, J. Shi, Z. Dai, J. Cai, P. Qian, L. Ji, X. He, S. Zhang, Z. Fei, et al. Libero-plus: In-depth robustness analysis of vision-language-action models. *arXiv preprint arXiv:2510.13626*, 2025.
- [54] B. Liu, Y. Zhu, C. Gao, Y. Feng, Q. Liu, Y. Zhu, and P. Stone. Libero: Benchmarking knowledge transfer for lifelong robot learning. *Advances in Neural Information Processing Systems*, 36:44776–44791, 2023.
- [55] K. Pertsch, K. Stachowicz, B. Ichter, D. Driess, S. Nair, Q. Vuong, O. Mees, C. Finn, and S. Levine. Fast: Efficient action tokenization for vision-language-action models. *arXiv preprint arXiv:2501.09747*, 2025.
- [56] L. Zhong, Y. Liu, Y. Wei, Z. Xiong, M. Yao, S. Liu, and G. Ren. Acot-vla: Action chain-of-thought for vision-language-action models. *The IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2026.
- [57] J. Cen, C. Yu, H. Yuan, Y. Jiang, S. Huang, J. Guo, X. Li, Y. Song, H. Luo, F. Wang, et al. Worldvla: Towards autoregressive action world model. *arXiv preprint arXiv:2506.21539*, 2025.
- [58] C.-Y. Hung, Q. Sun, P. Hong, A. Zadeh, C. Li, U. Tan, N. Majumder, S. Poria, et al. Nora: A small open-sourced generalist vision language action model for embodied tasks. *arXiv preprint arXiv:2504.19854*, 2025.
- [59] Q. Bu, Y. Yang, J. Cai, S. Gao, G. Ren, M. Yao, P. Luo, and H. Li. Univla: Learning to act anywhere with task-centric latent actions. *arXiv preprint arXiv:2505.06111*, 2025.
- [60] T. Yuan, Z. Dong, Y. Liu, and H. Zhao. Fast-wam: Do world action models need test-time future imagination? *arXiv preprint arXiv:2603.16666*, 2026.
- [61] S. Tan, K. Dou, Y. Zhao, and P. Kraehenbuehl. Interactive post-training for vision-language-action models. In *Workshop on Foundation Models Meet Embodied Agents at CVPR 2025*.
- [62] M. J. Kim, C. Finn, and P. Liang. Fine-tuning vision-language-action models: Optimizing speed and success. *arXiv preprint arXiv:2502.19645*, 2025.
- [63] Y. Ze, G. Zhang, K. Zhang, C. Hu, M. Wang, and H. Xu. 3d diffusion policy: Generalizable visuomotor policy learning via simple 3d representations. *Robotics: Science and Systems*, 2024.

- [64] S. Liu, L. Wu, B. Li, H. Tan, H. Chen, Z. Wang, K. Xu, H. Su, and J. Zhu. Rdt-1b: a diffusion foundation model for bimanual manipulation. In *International Conference on Learning Representations*, volume 2025, pages 29982–30009, 2025.
- [65] Y. Feng, H. Tan, X. Mao, C. Xiang, G. Liu, S. Huang, H. Su, and J. Zhu. Vidar: Embodied video diffusion model for generalist manipulation. *arXiv preprint arXiv:2507.12898*, 2025.
- [66] Y. Gui, Y. Zhou, S. Cheng, X. Yuan, H. Fan, P. Cheng, and S. Liu. Seedpolicy: Horizon scaling via self-evolving diffusion policy for robot manipulation. *arXiv preprint arXiv:2603.05117*, 2026.
- [67] J. Sun, W. Zhang, Z. Qi, S. Ren, Z. Liu, H. Zhu, G. Sun, X. Jin, and Z. Chen. Vla-jepa: Enhancing vision-language-action model with latent world model. *arXiv preprint arXiv:2602.10098*, 2026.

Appendix

A Real-Robot Experimental Setup, Task Definitions, and Evaluation Protocol




Figure 5: Dobot and Franka real-world setup.

Overview. We evaluate the real-robot task execution capability and OOD generalization capability of BRIDGE-WA on two single-arm platforms: a Dobot Nova2 robot and a Franka FR3 robot. Both platforms use LeRobot-style data with synchronized third-person RGB, wrist RGB, Cartesian robot state, Cartesian end-effector action, timestamps, and task instructions. The action/state vector is seven-dimensional, $[x, y, z, \text{roll}, \text{pitch}, \text{yaw}, g]$, where $g \in [0, 1]$ denotes the gripper command. Demonstrations are collected at 30 Hz. The Dobot setup uses an Intel RealSense D455 as the third-person camera and an Intel RealSense D405 as the wrist camera, with 480×480 videos for both views. The Franka setup uses an ORBBEC Femto Bolt as the third-person camera at 720×720 resolution and an Intel RealSense D405 as the wrist camera at 480×480 resolution. At training time, images are resized to the policy resolution and the same action-chunk interface is used across embodiments. The robot is reset to a fixed home pose before each rollout. Object poses are randomly sampled within task-specific regions, while preserving reachability and avoiding initial collisions. The dataset contains 50 demonstrations per task.

A.1 Real-Robot Task Definitions

Pick Grape. On the table, there is a bunch of grapes, other fruits, and a plate. The robot needs to pick up and place the grapes. Language instruction: Put the grapes into the [COLOR] plate.

Collect Fruits. There are various fruits on the table. The robot needs to collect the specified fruits into the specific target. Language instruction: Put a [FRUIT], a [FRUIT], and an [FRUIT] into the [COLOR] [TARGET].

Push Buttons/Press Buttons. There are multiple buttons of different colors on the table. The robot needs to press the buttons in a specific sequence. Language instruction: Press the [COLOR] button, then the [COLOR] button, followed by the [COLOR] button, and finally the [COLOR] button.

Stack Bowls. There are a cup, some bowls and a plate on the table. The robot needs to stack a cup, the bowls and plate in a specific order. Language instruction: Stack a [COLOR] cup, a [COLOR] bowl and a [COLOR] bowl on a [COLOR] plate.

Put Item In Drawer. There is a cabinet on the table. The robot needs to open the specified drawer, place the target object inside, and then close the drawer. Language instruction: Open the [Top/Middle/Bottom] drawer and put the [OBJECT] into it.

Fold Towel. There is a towel on the table. The robot needs to pick up one corner of the towel and then fold it. Language instruction: Fold the towel.

Open Lip. There is a box on the table. The robot needs to pick up the lid and open it. Language instruction: Open the plastic lid of the biscuit case.

Hard-track perturbations. For Hard-track evaluation, we introduce controlled visual and scene perturbations while keeping the task goal and success criterion unchanged. In the distractor setting, additional irrelevant objects are placed on the tabletop or in the surrounding scene, such as extra fruit

items, plush bears, and other non-target objects. In the tablecloth setting, we change the tablecloth color and texture to alter the workspace appearance. In the lighting setting, we turn on dynamic colored lights to create non-stationary illumination changes and color shifts. These perturbations are designed to test whether the policy can focus on the task-relevant object, contact region, and intended scene change rather than relying on background appearance.

A.2 Real-Robot Evaluation Protocol

Evaluation contains one Easy track and three Hard tracks. Easy uses 10 trials per task with near-indistribution object poses and lighting. Each Hard track uses 5 trials per task and isolates one visual shift: distractor objects, lighting generalization, or tablecloth generalization. All methods are evaluated from matched initial states with the same language instruction sequence and the same manual reset procedure. For binary tasks, a trial receives 1 for success and 0 otherwise. `PushButtons` receives 0, 0.5, or 1 according to the number of correctly pressed targets. `CollectFruits` and `StackBowls` receive 0, 0.3, 0.6, or 1 according to task-specific completion stages. The Hard score is the sum over the three hard tracks, normalized as an average success rate.

B More Design Details of BRIDGE-WA

At time t , the policy observes multi-view RGB images, proprioception, and a language instruction, and predicts an H -step action chunk:

$$x_t = (\{I_t^v\}_{v=1}^V, q_t, \ell), \quad \mathbf{a}_{t:t+H-1} = (a_t, \dots, a_{t+H-1}). \quad (5)$$

BRIDGE-WA augments a direct action policy with a compact world-prior variable,

$$\hat{Z}_t = g_\phi(x_t), \quad Z_t = (T_{f,t}, M_{c,t}, M_{o,t}), \quad \pi_{\theta,\phi}(\mathbf{a}_{t:t+H-1} | x_t) = \pi_\theta(\mathbf{a}_{t:t+H-1} | x_t, \hat{Z}_t), \quad (6)$$

where $T_f \in \mathbb{R}^{N_f \times d}$ stores future-outcome tokens, $M_c \in [0, 1]^{V \times h \times w}$ is a spatial change map, and $M_o \in \mathbb{R}^{V \times h \times w \times 2}$ is an image-coordinate motion-flow map.

B.1 World Teacher Model

The teacher \mathcal{W}_ψ is a robot-state-conditioned future-change model built from a Wan2.2-TI2V-5B generative backbone. It is used only during training and cache construction. Unlike a generic text-image-to-video model, the teacher receives robot observations, proprioception, and language. Its input is the current multi-view robot context $x_t = (\{I_t^v\}, q_t, \ell)$ and its target is the future observation at a fixed temporal offset. The implementation uses 256×256 visual inputs, sparse text-image-to-video conditioning, and a temporal offset controlled by `future_index`. For the VLABench teacher run used as the reference configuration, `future_index=30`.

For each training sample, the teacher produces a future representation at offset Δ , including a latent representation $F_{t+\Delta}$ and its decoded future view $\tilde{I}_{t+\Delta}$ used for flow construction:

$$(\tilde{F}_{t+\Delta}, \tilde{I}_{t+\Delta}) = \mathcal{W}_\psi(x_t, \Delta), \quad (7)$$

where Δ is the teacher future offset. The teacher output is not consumed as a dense future video by the final policy. Instead, we extract three compact supervision signals:

$$\begin{aligned} T_f^* &= P_f(\tilde{F}_{t+\Delta}), \\ M_c^* &= \text{Norm}\left(P_c\left(|\tilde{F}_{t+\Delta} - F_t|\right)\right), \\ M_o^* &= \text{Resize}\left(\text{Flow}(I_t, \tilde{I}_{t+\Delta})\right). \end{aligned} \quad (8)$$

Here F_t is the current visual latent, P_f pools the teacher latent into future tokens, P_c pools latent differences into the change-map grid, and $\text{Flow}(\cdot, \cdot)$ estimates image-space motion between the current and teacher-predicted future views. This factorization is the central compression step: the policy receives outcome, intervention support, and local motion rather than all pixels of a generated future.

B.2 Offline Cache Design

The cache stores teacher-derived priors at sample granularity. For each training sample, we save the sample key, the teacher future offset, future tokens, pooled change maps, pooled flow maps, and metadata needed to align the cache entry with the policy input. The reference cache uses 4 GPUs, batch size 128, `diffusion_steps=10` for teacher future decoding, and `num_actions=30` for the action horizon. The cache contains half-precision tensors where possible to reduce storage and data-loading cost. During policy training, the cache is read; no teacher forward pass or future decoding is required.

$$\mathcal{C}[k] = \{k, \Delta, T_f^*, M_c^*, M_o^*, \text{meta}(x_t)\}, \quad (9)$$

where k is the dataset sample key and $\text{meta}(x_t)$ stores the camera order, time index, task id, and other alignment metadata. This makes world-prior supervision deterministic across policy-training.

B.3 World-Prior Predictor

The predictor g_ϕ maps the current policy context to $\hat{Z}_t = (\hat{T}_{f,t}, \hat{M}_{c,t}, \hat{M}_{o,t})$. It shares the policy’s encoded visual-language context and adds lightweight heads for the three world-prior targets. Each branch is supervised by both regression and cosine-alignment terms. The default weights are 0.1 for future regression, 0.1 for future cosine alignment, 0.1 for change regression, 0.1 for change cosine alignment, 0.1 for flow regression, and 0.1 for flow cosine alignment.

The cosine distance used by the predictor losses is

$$D_{\cos}(u, v) = 1 - \frac{\langle u, v \rangle}{\|u\|_2 \|v\|_2 + \epsilon}. \quad (10)$$

For cached targets from Eq. 9, the world-prior distillation loss is

$$\begin{aligned} \mathcal{L}_{\text{prior}} = & \lambda_f^r \|\hat{T}_{f,t} - T_{f,t}^*\|_2^2 + \lambda_f^c D_{\cos}(\hat{T}_{f,t}, T_{f,t}^*) \\ & + \lambda_c^r \|\hat{M}_{c,t} - M_{c,t}^*\|_1 + \lambda_c^c D_{\cos}(\hat{M}_{c,t}, M_{c,t}^*) \\ & + \lambda_o^r \|\hat{M}_{o,t} - M_{o,t}^*\|_1 + \lambda_o^c D_{\cos}(\hat{M}_{o,t}, M_{o,t}^*). \end{aligned} \quad (11)$$

The final policy objective combines action imitation with this prior loss,

$$\mathcal{L}_{\text{train}} = \mathcal{L}_{\text{act}}(\pi_{\theta, \phi}(\mathbf{a}_{t:t+H-1} | x_t), \mathbf{a}_{t:t+H-1}^*) + \mathcal{L}_{\text{prior}}. \quad (12)$$

The predictor is trained jointly with the action policy. At inference, the teacher and cache are removed, and g_ϕ becomes the only source of world priors. This makes deployment close to a direct VLA model: the extra computation is one compact prior-prediction pass and several additional attention projections inside WORLDBRIDGE.

B.4 Conditioning and Layer Routing

WORLDBRIDGE conditions the action transformer on world priors through two routes. First, future, change, and flow priors are projected into memory key/value tokens that can be read by policy queries. Second, change and flow maps are converted into additive attention biases so that self-attention is softly guided toward regions expected to change or move. For each prior source $m \in \{f, c, o\}$, we define a layer range

$$\mathcal{L}_m = \{l \mid s_m \leq l < \min(L, s_m + n_m)\}, \quad \mathcal{A}_l = \{m \in \{f, c, o\} \mid l \in \mathcal{L}_m\}, \quad (13)$$

where L is the action-transformer depth, s_m is the start layer, n_m is the layer budget, and \mathcal{A}_l is the set of active priors at layer l . At an active layer, each prior is projected into conditioning memories:

$$K_m^l = P_{K,m}^l(\hat{Z}_{m,t}), \quad V_m^l = P_{V,m}^l(\hat{Z}_{m,t}), \quad B_m^l = P_{B,m}^l(\hat{M}_{m,t}), \quad m \in \{c, o\}. \quad (14)$$

Here $\hat{Z}_{f,t} = \hat{T}_{f,t}$, $\hat{Z}_{c,t} = \hat{M}_{c,t}$, and $\hat{Z}_{o,t} = \hat{M}_{o,t}$. B_c^l and B_o^l are additive attention biases derived from change and flow maps.

The resulting world-prior-conditioned attention update is

$$\text{Attn}_{\text{wb}}^l = \text{softmax} \left(\frac{Q^l[K_s^l; \{K_m^l\}_{m \in \mathcal{A}_l}]^\top}{\sqrt{d}} + \sum_{m \in \mathcal{A}_l \cap \{c, o\}} B_m^l \right) [V_s^l; \{V_m^l\}_{m \in \mathcal{A}_l}], \quad (15)$$

where $\text{Attn}_{\text{wb}}^l$ denotes the WORLD BRIDGE-conditioned attention update at layer l , Q^l is the policy query, K_s^l, V_s^l are the standard self-attention key/value tensors, and K_m^l, V_m^l are world-prior memory tensors. Inactive priors are omitted from both memory projection and attention-bias computation.

The conditioning is layered rather than uniform. Future tokens are used earlier as a global outcome prior. Change tokens are routed in middle-to-late layers to ground the intended interaction region. Flow tokens are routed closest to action decoding, where local motion direction is most useful. The default configuration uses modulation scale 0.1, change-bias scale 1.0, flow-bias scale 1.0, and blends teacher guidance with predicted guidance using a 0.5 ratio during training. Inactive priors are omitted from both memory projection and attention-bias computation, preventing unused priors from increasing context length or adding noisy modulation.

C Experimental Details

C.1 Training Recipe

Training has two stages. First, we prepare the future-change teacher with a two-step teacher-training schedule. The teacher is initialized from Wan2.2-TI2V-5B and first pre-trained on BridgeData V2 robot trajectories to learn manipulation-induced future observation structure from current images, language, and robot state. We then post-train, or fine-tune, the same teacher on the target benchmark or real-robot dataset before cache construction, so that its future predictions match the embodiment, camera layout, and task distribution used by the downstream policy. The target-dataset post-training stage resumes from the BridgeData-pretrained checkpoint and updates only the 5B DiT backbone; the Wan VAE, text encoder, and other non-DiT components are kept frozen. Between the two training stages, the frozen teacher is applied to each downstream training set to build the world-prior cache. The process uses the downstream dataset metadata, task camera order, future-view selection, action mode, action horizon, and teacher future offset. Second, the policy is trained with imitation learning plus world-prior distillation. The action horizon is 30, matching `num_actions=30` in cache generation.

C.2 Baselines

OpenVLA [9] is an open-source vision-language-action model that maps RGB observations and language instructions to robot actions using large pretrained vision-language priors. π_0 [11] is a vision-language-action flow model for general robot control, while $\pi_{0.5}$ [39] extends the same family toward broader open-world generalization. π_{fast} [55] uses efficient action tokenization, and we report both relative-chunk and delta-chunk variants on VLABench because the action representation materially changes long-horizon manipulation performance. **X-VLA** [35] is a scalable cross-embodiment VLA based on soft-prompted transformers. **ACoT-VLA** [56] adds action chain-of-thought style supervision to VLA training. **ACT** [14], **Diffusion Policy (DP)** [15], and **DP3** [63] are non-VLA imitation-learning controllers that respectively model temporally coherent action chunks, denoising-diffusion action sequences, and diffusion policies conditioned on compact 3D representations. **RDT** [64], **Vidar** [65], and **SeedPolicy** [66] provide strong recent references for large diffusion transformers, video-diffusion-based generalist manipulation, and horizon-scaled self-evolving diffusion policies.

C.3 Additional Quantitative Results

We provide detailed per-task evaluations for both VLABench and RoboTwin 2.0. For VLABench, we report the performance of BRIDGE-WA in Tab. 7 across the in-distribution setting and four generalization settings, including cross-category transfer, common-sense instructions, semantic instruction variation, and unseen textures; success rate (SR), intention score (IS), and progress score (PS) are included to characterize both final task completion and intermediate execution quality. For RoboTwin 2.0, we report the 15-task per-task results of BRIDGE-WA in Tab. 6; the policy is trained only on the Easy split, and the Hard results are obtained by directly evaluating the same Easy-trained checkpoint on the Hard split without any Hard-split training.

Task	Easy	Hard
Adjust Bottle	99.00	54.00
Beat Block Hammer	96.00	16.00
Blocks Ranking RGB	27.00	20.00
Blocks Ranking Size	64.00	21.00
Click Alarmclock	97.00	67.00
Click Bell	100.00	62.00
Grab Roller	100.00	48.00
Move Playingcard Away	97.00	30.00
Pick Diverse Bottles	55.00	36.00
Pick Dual Bottles	87.00	64.00
Place A2B Left	58.00	17.00
Place A2B Right	61.00	15.00
Place Container Plate	100.00	53.00
Place Fan	81.00	11.00
Put Bottles Dustbin	73.00	52.00
Average	79.67	37.73

Table 6: Per-task RoboTwin 2.0 success rates of BRIDGE-WA on the 15-task evaluation subset. Easy and Hard denote evaluation splits; the model is trained only on Easy demonstrations and is not trained on Hard demonstrations.

Task	In-Dist.			Cross-Cat.			Common Sense			Semantic Instr.			Unseen Texture		
	SR	IS	PS	SR	IS	PS	SR	IS	PS	SR	IS	PS	SR	IS	PS
select_painting	40.0	100.0	40.0	40.0	80.0	40.0	40.0	100.0	40.0	50.0	100.0	50.0	40.0	100.0	40.0
select_book	60.0	100.0	80.0	0.0	0.0	0.0	70.0	100.0	85.0	70.0	100.0	85.0	30.0	100.0	50.0
select_drink	90.0	100.0	95.0	10.0	40.0	25.0	80.0	100.0	90.0	80.0	100.0	90.0	60.0	100.0	80.0
select_chemistry_tube	70.0	30.0	85.0	0.0	0.0	0.0	50.0	60.0	75.0	80.0	30.0	90.0	10.0	10.0	40.0
select_poker	100.0	100.0	100.0	20.0	40.0	23.3	20.0	60.0	23.3	100.0	100.0	100.0	50.0	80.0	53.3
select_mahjong	100.0	100.0	100.0	0.0	20.0	0.0	11.1	22.2	11.1	100.0	100.0	100.0	40.0	70.0	45.0
select_toy	100.0	20.0	100.0	70.0	0.0	75.0	70.0	0.0	85.0	50.0	0.0	75.0	70.0	20.0	85.0
select_fruit	90.0	100.0	95.0	60.0	80.0	70.0	60.0	100.0	80.0	60.0	100.0	80.0	50.0	90.0	65.0
add_condiment	90.0	100.0	93.3	0.0	30.0	10.0	80.0	100.0	90.0	60.0	90.0	73.3	70.0	100.0	80.0
insert_flower	40.0	100.0	70.0	30.0	100.0	45.0	30.0	100.0	65.0	20.0	100.0	60.0	30.0	90.0	65.0

Table 7: Per-task VLABench results of BRIDGE-WA. SR, IS, and PS denote success rate, intention score, and progress score.

For the Franka real-robot platform, we evaluate the Easy track. The per-task success rates for BRIDGE-WA, $\pi_{0.5}$ [39], and VLA-JEPA [67] are reported in Tab. 8.

Method	Collect Fruits	Press Buttons	Open Lip	Fold Towel	Put Item In Drawer	Average
$\pi_{0.5}$ [39]	50%	80%	50%	30%	25%	47.0%
VLA-JEPA [67]	75%	90%	30%	25%	25%	49.0%
BRIDGE-WA (Ours)	75%	100%	60%	80%	50%	73.0%

Table 8: Per-task success rates on the Franka real-robot evaluation.

D Additional Visualizations

We provide additional qualitative visualizations to contextualize the benchmark tasks, the learned world-prior representations, and the real-robot execution behavior. The benchmark-scene overview in Fig. 6 visualizes the ten VLABench tasks and the 15 RoboTwin 2.0 tasks used in our experiments. The world-prior visualization in Fig. 7 illustrates how BRIDGE-WA decomposes predicted scene evolution into future, change, and motion priors on both simulation benchmarks. Finally, the Dobot rollout sequence in Fig. 8 presents representative executions for the five real-robot tasks, showing the temporal progression from initial observation to task completion.

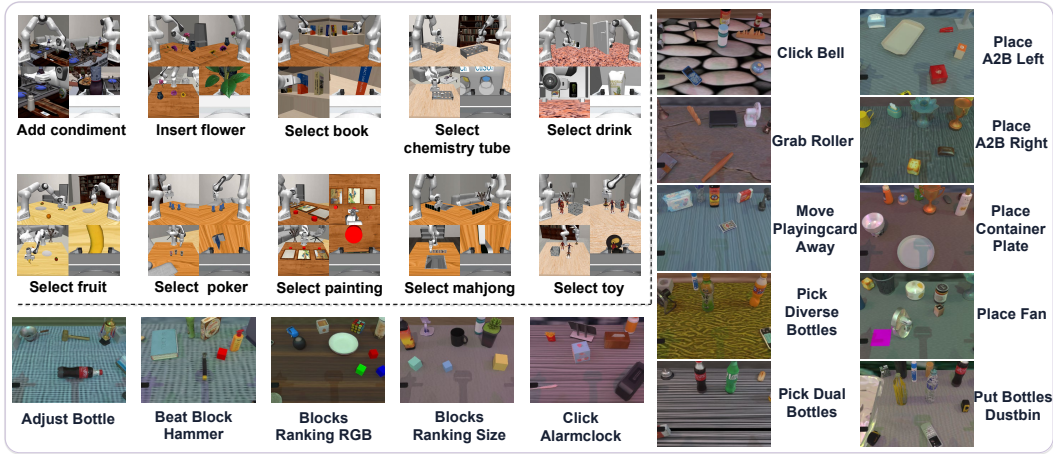


Figure 6: Task and scene visualizations for the simulation benchmarks. The figure summarizes the ten VLABench tasks and the 15-task RoboTwin 2.0 subset used in our evaluation, illustrating the diversity of object categories, workspace layouts, and manipulation goals across single-arm and bimanual settings.

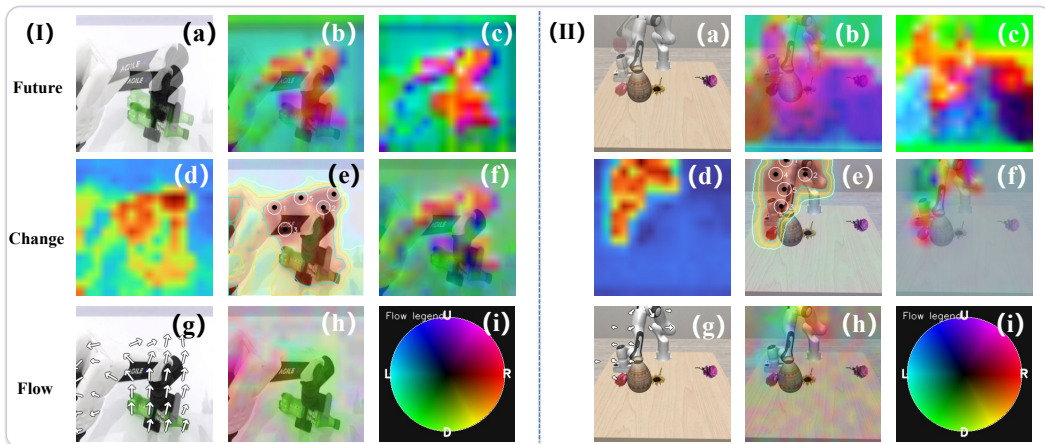


Figure 7: Qualitative world-prior visualizations on RoboTwin 2.0 (I) and VLABench (II). We show the factorized predictive signals used by BRIDGE-WA: future representations summarize the intended outcome, change maps identify action-relevant regions expected to change, and motion-flow maps indicate the local direction of scene evolution. (a) shows the future RGB target overlaid on the current observation. (b) overlays the predicted future latent representation on the scene, while (c) visualizes the future latent map. (d) presents the change map, and (e) overlays the enhanced change response on the observation; the marked points indicate high-confidence regions where action-relevant state changes are expected. (f) highlights feature-level differences between the current and future states. (g) visualizes the dense flow field on the scene, and (h) illustrates the temporal motion direction implied by the flow. (i) gives the color-wheel legend, where hue denotes motion direction.




Figure 8: Dobot real-robot execution sequences. The figure shows temporal rollouts for the five real-robot tasks, covering object selection, placement, pushing, stacking, and drawer manipulation.