Targeted Adaptation of World-Action Models via Interaction and Motion Heads

Yuze Zhu[1] Jacky Li[2] Yusen Luo[3] Ken Jiankun Zheng[1]
Chancharik Mitra[4] Trevor Darrell[1] Roei Herzig[1]
[1]University of California, Berkeley [2]University of Toronto
[3]University of Southern California [4]Massachusetts Institute of Technology

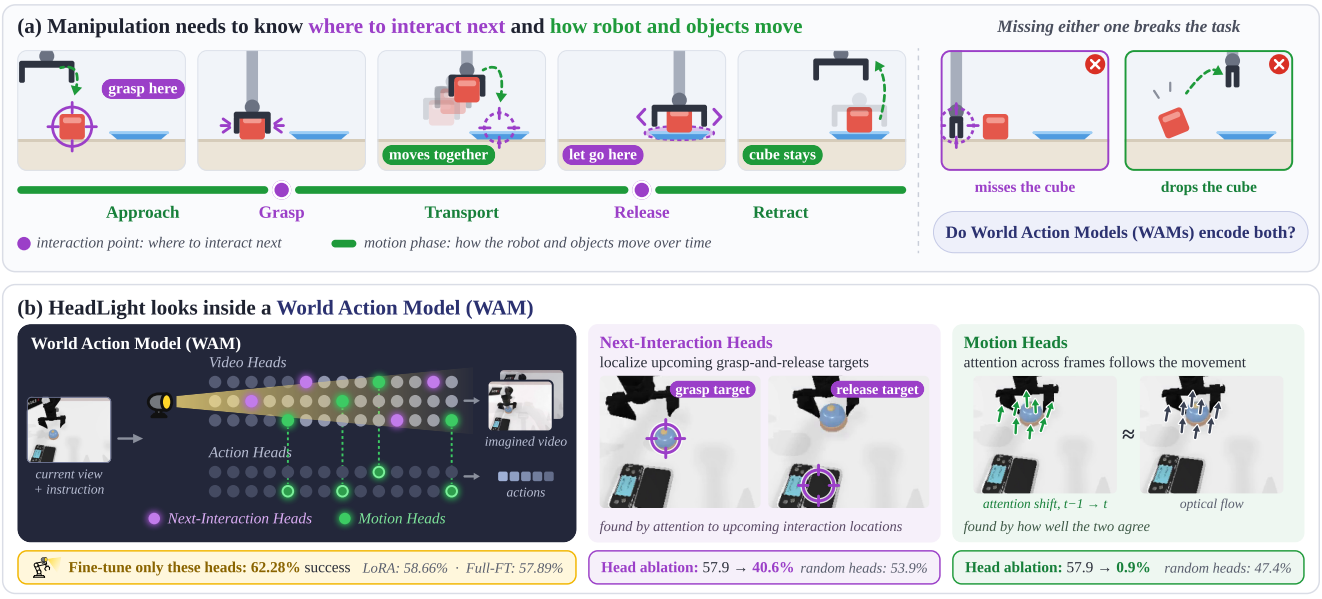


Figure 1. **HeadLight: Targeted adaptation of WAMs via specialized attention heads.** (a) Successful manipulation requires both localizing *where to interact next* and modeling *how the robot and objects move*; missing either breaks the task. (b) Within pretrained WAMs, we identify *Next-Interaction Heads* that localize upcoming grasp-and-release targets and *Motion Heads* whose cross-frame attention aligns with optical flow. Ablating either head group substantially reduces manipulation success compared with equally sized random-head controls. **HeadLight** fine-tunes only these heads, outperforming LoRA and full fine-tuning on unseen RoboTwin tasks.

Abstract

World-Action Models (WAMs) are a promising new paradigm of robot manipulation policies that jointly generate future video and actions. Successful manipulation requires both identifying where to interact and modeling how the robot and objects move. Yet standard adaptation methods, such as full fine-tuning and LoRA, do not explicitly select parameters aligned with these manipulation-specific capabilities. Motivated by this challenge, we discover two functionally specialized families of attention heads: Next-Interaction Heads, which localize upcoming grasp-and-release targets, and Motion Heads, which track movement across frames. Furthermore, we introduce HeadLight, a parameter-efficient refinement method that jointly adapts Next-Interaction Heads and Motion Heads. On RoboTwin tasks unseen during fine-tuning, HeadLight achieves 62.28% success, compared with 57.89% for full fine-tuning and 58.66% for LoRA, while updating only 1.1% of parameters. Inference-time ablations further validate the functional importance of these heads: ablating Motion Heads reduces manipulation success from 57.89% to 0.94%, while ablating Next-Interaction Heads lowers it to 40.61%, both substantially more than ablating equally many random heads. Finally, we show that our approach is effective across multiple WAM architectures and real-robot manipulation settings.

1. Introduction

Robot manipulation has been explored through various policy architectures [2, 4, 39], with World-Action Models (WAMs) emerging as a promising approach that generates both future video and actions [13, 34, 36]. Research in motor control and task-and-motion planning describes manipulation in terms of discrete contact events and continuous motion [6, 7]. These aspects translate into two complementary capabilities that WAMs need: anticipating where the next grasp or release should occur, and representing the continuous motion of the robot and manipulated objects across frames.

Current WAMs implicitly rely on large-scale video pretraining [27] to acquire these capabilities. Such pretraining allows them to generate visually plausible videos, yet *visual* plausibility alone does not ensure that the *manipulation* capabilities are actually learned [22, 26]. In practice, WAMs may fail to identify the correct interaction targets, and their generated actions may not accurately follow the motion in their predicted videos. These failures are especially pronounced when generalizing to unseen tasks. Although full fine-tuning and LoRA [11] can adapt WAMs, neither method explicitly selects parameters based on their roles in interaction targeting and motion. As a result, optimizing the prediction loss may encourage the model to fit visual details, such as appearance and texture, without necessarily improving interaction targeting and motion modeling. This raises two natural questions: Can we identify the model components associated with interaction targeting and motion? If so, can we selectively refine them to strengthen these capabilities rather than merely fitting visual details?

Previous work has shown that robot policies can be improved by adapting a small set of internal components, either through steering sparse features at inference [37] or by fine-tuning attention heads selected from demonstrations of a specific task [21]. In both cases, the identified components are selected separately for each task. We instead ask whether some heads support more fundamental capabilities shared across manipulation tasks, namely interaction targeting and motion. To answer this, we investigate attention heads in pretrained WAMs and identify two families associated with these capabilities: **Next-Interaction Heads** and **Motion Heads**. We identify the former by analyzing attention to upcoming interaction locations, and the latter by measuring the agreement between cross-frame attention and optical flow. Ablating Motion Heads reduces manipulation success from 57.89% to just 0.94%, while ablating Next-Interaction Heads reduces it to 40.61%, with both causing substantially larger declines than ablating equally many random heads. These experiments also reveal distinct failure modes: ablating Next-Interaction Heads often causes the robot to miss its intended targets, while ablating Motion Heads disrupts the motion in predicted videos. These results suggest that WAMs rely on a small number of specialized attention heads to localize interaction targets and to model motion.

Building on these findings, we introduce **HeadLight**, a parameter-efficient refinement method that uses these two head families to guide complementary adaptation pathways. We selectively fine-tune Next-Interaction Heads to improve target localization and use Motion Heads to guide the adaptation of action prediction, helping generated actions follow the predicted motion. By combining these two pathways, our method aims to improve both interaction targeting and motion execution without updating the entire model.

We evaluate HeadLight on RoboTwin [3] tasks unseen during fine-tuning. On Zero-WAM, our method achieves **62.3%** success, compared to 57.9% for full fine-tuning and 58.7% for LoRA, while updating only 1.1% of parameters. Using either adaptation pathway alone does not achieve the same improvement, supporting the importance of their combination. We further demonstrate the effectiveness of our approach on LingBot-VA [15], which uses a different WAM architecture. Finally, we validate in real-robot manipulation using a Franka robot.

To summarize, our main contributions are threefold: (i) We identify Next-Interaction Heads and Motion Heads in pretrained WAMs and validate their distinct roles in interaction targeting and motion through attention analysis and controlled ablations. (ii) We introduce HeadLight, which combines two complementary adaptation pathways guided by these heads to improve interaction targeting and motion execution while updating only a small subset of model parameters. (iii) We demonstrate improvements over standard adaptation methods on unseen manipulation tasks, with effectiveness across WAM architectures and in real-robot manipulation. Together, these findings suggest that the functional structure of WAMs can guide which parameters to adapt, improving generalization to unseen tasks.

2. Related Work

World-action models. World-action models jointly model future visual observations and robot actions for manipulation [9, 13, 15, 16, 33, 34, 38]. Among these, Zero-WAM conditions its predictions on human demonstration videos to specify unseen tasks [36]. While these models focus on learning to predict video and actions, we investigate the functional structure of their attention heads and how it can guide parameter selection for adaptation.

Functional specialization of attention heads. Prior studies have revealed functional specialization in language transformers, where attention heads capture linguistic relations [5, 30] and task representations during in-context learning [28]. Similar properties have been explored in

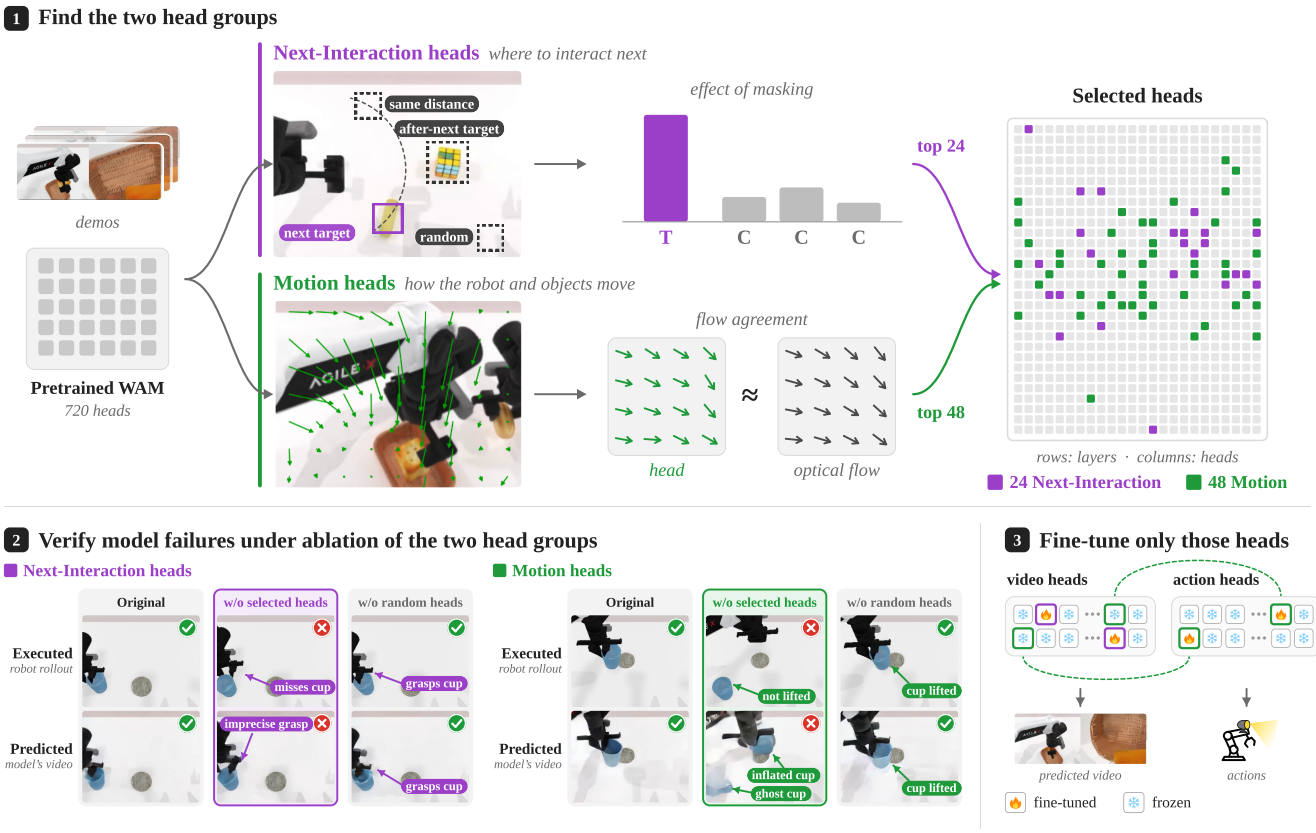


Figure 2. **Overview of HeadLight.** (1) We identify Next-Interaction Heads by their selective reliance on upcoming interaction targets over control tokens, and Motion Heads by the alignment of cross-frame attention with optical flow. (2) Ablating these heads causes distinct failures: missed grasps and disrupted object motion, respectively, whereas ablating equally many random heads largely preserves successful execution. (3) HeadLight selectively fine-tunes 24 Next-Interaction Heads in the video stream and 48 action heads paired with frozen video Motion Heads, while keeping all other heads frozen.

video diffusion models: DiTFlow extracts patch correspondences from cross-frame attention for motion transfer [25], while HALO identifies motion- and structure-sensitive heads, validates attention-derived displacements against optical flow, and leverages them for training-free motion transfer [12]. In embodied models, EWAM further reveals a layer-wise progression from semantic information to visual foresight and action [31]. Building on these observations, we identify Next-Interaction Heads and Motion Heads in WAMs and examine their functional importance for closed-loop manipulation through controlled ablations.

Head selection, steering, and adaptation. Existing approaches adapt pretrained models through parameter updates or activation interventions. LoRA [11] constrains the rank of parameter updates without explicitly specifying which heads to adapt. For robot policies, activation steering modifies behavior through interventions on feedforward activations [10], while SteerAct [37] selects a sparse feature from successful and failed rollouts and modulates its activation during execution. Other approaches focus on attention heads: Sparse Attention Vectors [20] use few-shot examples to select attention-head activations for vision-language classification without fine-tuning, while Robotic Steering [21] ranks heads using nearest-neighbor action regression on task demonstrations and selectively fine-tunes them with LoRA. In contrast to task-specific head selection, HeadLight selects heads based on manipulation functions shared across tasks: selective reliance on upcoming interaction targets and alignment of cross-frame attention with optical flow.

3. The HeadLight Method

We investigate whether functional specialization exists within pretrained WAMs and can guide parameter selection, aiming to strengthen interaction targeting and motion execution rather than fitting task-irrelevant visual details. First, given a frozen WAM, we identify two functionally specialized families of video attention heads, *Next-Interaction Heads* and *Motion Heads*, by scoring each head on how

much it depends on the upcoming interaction target and how well its attention follows the motion of the scene (Sec. 3.2). Second, we introduce **HeadLight**, which uses these identified heads to establish two complementary, parameter-efficient adaptation targets: Next-Interaction Heads for localizing the upcoming interaction target, and Motion Heads for action execution (Sec. 3.3).

3.1. Preliminaries: Joint Attention in WAMs

A WAM jointly generates future videos and actions via a transformer [29] operating on both video and action tokens. At layer ℓ and head index h, we denote the attention head processing video tokens as a *video head* $(\ell, h)_v$ and its counterpart for action tokens as an *action head* $(\ell, h)_a$. The two modalities interact through *joint attention* [17]: at every layer, each action head $(\ell, h)_a$ attends to video tokens by querying the keys and values from its corresponding video head $(\ell, h)_v$. This one-to-one correspondence allows us to select target action heads through their paired video heads.

3.2. Identifying Manipulation-Relevant Heads

Given a frozen WAM, we evaluate all video heads $(\ell, h)_v$ over manipulation episodes along two axes corresponding to discrete contact events and continuous trajectories: (i) selective reliance on the upcoming interaction target, and (ii) cross-frame alignment with scene motion.

Next-Interaction Heads: Locating Where to Interact. As a manipulation task progresses, the task-relevant spatial target shifts dynamically from the object to be grasped to its intended release location. A true *Next-Interaction Head* should selectively track this *upcoming* contact region rather than merely attending to visually salient objects or static background features.

We identify such heads via causal intervention: for each head, we mask out the tokens covering the upcoming interaction target and measure how much the head's output changes. Intuitively, a head that tracks the next interaction should be disrupted far more by this removal than by removing other plausible but task-irrelevant regions. We formalize this intuition below by comparing the target-removal effect against several control sets and retaining the heads with the largest margin.

Specifically, let $\mathcal{T}$ denote the video token that covers the upcoming interaction target in frame f. We ablate $\mathcal{T}$ from head $(\ell, h)_v$ by zeroing out its attention weights and re-normalizing the remaining distribution [8], then quantify the relative change in the head's output:

$$D_{\ell,h}(\mathcal{T}) = \frac{\left\| O^{v}_{\ell,h} - O^{v,\setminus\mathcal{T}}_{\ell,h} \right\|_2^2}{\left\| O^{v}_{\ell,h} \right\|_2^2}, \tag{1}$$

where $O^{v}_{\ell,h}$ and $O^{v,\setminus\mathcal{T}}_{\ell,h}$ denote the head's outputs before and after removing $\mathcal{T}$, averaged across all video queries in frame f. A larger $D_{\ell,h}(\mathcal{T})$ indicates a stronger causal reliance on $\mathcal{T}$. However, a large $D_{\ell,h}(\mathcal{T})$ can arise for reasons other than next-interaction tracking: removing *any* token perturbs a head's output, and the amount of perturbation depends on how much attention the token receives and how different its value is from the other tokens.

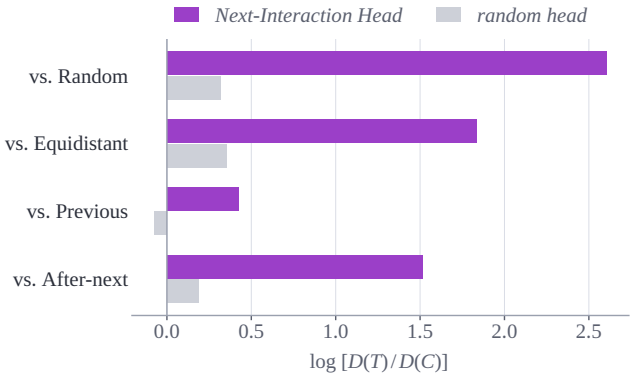


Figure 3. **Selectivity of Next-Interaction Heads.** Log-ratio margin (Eq. (2)) between masking the upcoming target $\mathcal{T}$ and masking each of four control tokens $\mathcal{C}$: Random, Equidistant (as far from the gripper as $\mathcal{T}$), Previous (the previous target), and After-next (the target after $\mathcal{T}$). A Next-Interaction Head maintains positive margins across all controls, whereas a same-layer random head stays near zero.

To rule out such alternative explanations, we apply the same removal to four control tokens $\mathcal{C}^c$ ($c \in \{1, \ldots, 4\}$) from the same frame and compute $D_{\ell,h}(\mathcal{C}^c)$ with Eq. (1). Each control isolates an alternative explanation: if removing it disrupts the head as much as removing $\mathcal{T}$, that explanation cannot be excluded. The controls are (1) a *random* token, ruling out that the head is sensitive to arbitrary removals; (2) an *equidistant* token, as far from the gripper as $\mathcal{T}$, ruling out that the head merely attends near the gripper; (3) the *previous* interaction target, ruling out that it tracks where the robot has already been; and (4) the *after-next* interaction target, ruling out that it attends to all task-relevant targets rather than only the next.

Each control yields a log-ratio margin, $\log\left(\mathbb{E}[D_{\ell,h}(\mathcal{T})]/\mathbb{E}[D_{\ell,h}(\mathcal{C}^c)]\right)$, which is positive when removing $\mathcal{T}$ disrupts the head more than removing that control. A head supports next-interaction tracking only if its margin is positive against all four controls, so we define its selectivity score as the smallest of the four margins:

$$S_{\mathrm{N}}(\ell, h) = \min_{c} \log \frac{\mathbb{E}_f\left[D_{\ell,h}(\mathcal{T})\right]}{\mathbb{E}_f\left[D_{\ell,h}(\mathcal{C}^c)\right]}, \tag{2}$$

where $\mathbb{E}_f$ averages over frames across manipulation episodes. $S_{\mathrm{N}}(\ell, h)$ is therefore high only if the head strictly relies on the upcoming target more than on all four controls. Finally, we rank all video heads by $S_{\mathrm{N}}(\ell, h)$ and select the



Figure 4. **Target-selective attention of a Next-Interaction Head.** Spatial attention maps from the same layer as the left gripper approaches the target. The Next-Interaction Head focuses sharply on the upcoming grasp target while ignoring the distractor object, whereas an unselected head attends diffusely across the scene.

top $K_{\mathrm{N}} = 24$ as the *Next-Interaction Heads* $\mathcal{H}^v_{\mathrm{N}}$, whose selectivity and attention are illustrated in Figs. 3 and 4.

Motion Heads: Capturing How the Scene Moves. While Next-Interaction Heads identify where the robot should interact, successful manipulation also requires understanding how the robot and objects move over time. We therefore seek attention heads whose internal attention patterns align with the physical motion observed in the scene. Motivated by previous observations that cross-frame attention can encode motion correspondences [12, 23, 25], we identify such heads by measuring how closely their attention-induced motion fields align with reference optical flow. This distinguishes heads whose cross-frame attention moves with the scene from heads whose attention stays at fixed image locations or is unrelated to the motion.

To measure this alignment, we first derive an attention-induced motion field for each video head (ℓ, h). For a spatial token p at coordinate $\mathbf{x}_p$ in frame f, we consider its spatial offsets relative to tokens p' in the preceding frame $f-1$, and aggregate them according to the head's cross-frame attention weights:

$$\widehat{\mathbf{u}}^f_{\ell,h}(p) = \sum_{p'} \bar{A}^{f\to f-1}_{\ell,h}(p,p')\,(\mathbf{x}_p - \mathbf{x}_{p'})\,, \qquad (3)$$

where $\bar{A}^{f\to f-1}_{\ell,h}(p,p')$ denotes the attention weight of token p attending to token p', normalized over all spatial tokens p' in frame $f-1$. As visualized in Fig. 5, this produces a dense, patch-wise displacement field that characterizes the motion correspondences encoded by each head.

We then quantify the alignment between this attention-induced field and reference optical flow $\mathbf{u}^f$ [14], using the expected cosine similarity:

$$S_{\mathrm{M}}(\ell,h) = \mathbb{E}_{f,p}\Big[\cos\big(\widehat{\mathbf{u}}^f_{\ell,h}(p), \mathbf{u}^f(p)\big)\Big]. \qquad (4)$$

We rank all video heads by $S_{\mathrm{M}}(\ell,h)$ and select the top $K_{\mathrm{M}} = 48$ as the *Motion Heads* $\mathcal{H}^v_{\mathrm{M}}$, with both selected head sets shown in Fig. 2. We use $K_{\mathrm{N}} = 24$ and $K_{\mathrm{M}} = 48$ throughout and analyze these budgets in Sec. 4.3.

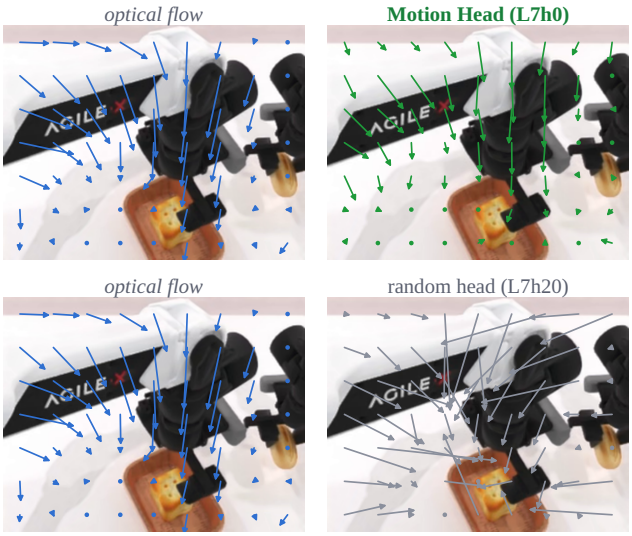


Figure 5. **Cross-frame motion alignment of Motion Heads.** Reference optical flow $\mathbf{u}^f$ (left) versus attention-induced motion fields $\widehat{\mathbf{u}}^f_{\ell,h}$ (Eq. (3), right). A Motion Head closely tracks the moving robot arm and gripper, whereas a same-layer random head yields disorganized offsets.

Functional validation of the identified heads. While S_{N} and S_{M} measure target selectivity and motion alignment, they do not establish whether the identified heads are necessary for closed-loop manipulation. We therefore conduct inference-time head ablations against equally sized random-head controls in Sec. 4.3, where Fig. 6 confirms their distinct functional importance for interaction targeting and motion execution.

3.3. Selective Refinement of Interaction and Motion

Once $\mathcal{H}^v_{\mathrm{N}}$ and $\mathcal{H}^v_{\mathrm{M}}$ are identified, HeadLight reuses them across downstream tasks for selective fine-tuning. On new tasks, target errors originate from predicted videos missing the right object or placement location, whereas motion errors typically arise when generated actions fail to follow a plausible video trajectory. HeadLight accordingly pairs two complementary pathways: refining $\mathcal{H}^v_{\mathrm{N}}$ in the video stream for target localization, while keeping $\mathcal{H}^v_{\mathrm{M}}$ frozen and refining its paired action heads for motion execution.

Video pathway: refining Next-Interaction Heads. To sharpen where upcoming grasps and releases appear in predicted videos, we selectively fine-tune the K_{N} Next-Interaction Heads $\mathcal{H}^v_{\mathrm{N}}$ while freezing all other video heads. Restricting video updates to $\mathcal{H}^v_{\mathrm{N}}$ focuses adaptation on localizing task-relevant contact targets, rather than fitting

Table 1. **Generalization to unseen RoboTwin [3] tasks across WAM architectures.** Success rate (%, mean±std over 3 seeds, 100 trials each) after adapting on 43 seen tasks; unadapted backbone results are from [36]. HeadLight outperforms Full FT and LoRA on LingBot-VA and Zero-WAM. Best per backbone in bold.

Method	Place object on scale	Stamp seal	Move stapler to pad	Place bread in basket	Place empty cup	Stack blocks three	Average
LingBot-VA [15]	6.17±4.87	3.67±2.49	23.33±8.22	17.33±6.18	42.33±7.85	0.00±0.00	15.47±1.30
+ Full FT	22.33±3.77	21.67±2.87	61.00±4.24	34.67±4.92	56.00±4.32	0.00±0.00	32.61±3.04
+ LoRA	20.67±3.09	21.33±2.49	64.00±2.94	37.33±1.25	61.33±4.50	0.00±0.00	34.11±1.31
+ HeadLight	**23.00**±7.07	**25.67**±2.36	**65.67**±4.50	**49.33**±2.62	**65.33**±0.47	0.00±0.00	**38.17**±1.11
Zero-WAM [36]	24.67±2.05	47.00±4.55	69.14±2.93	35.00±3.74	84.87±0.18	9.00±2.16	44.95±0.79
+ Full FT	43.33±1.25	**73.33**±6.65	85.33±0.47	42.67±2.87	91.33±2.05	**11.33**±3.09	57.89±0.77
+ LoRA	47.33±5.79	70.33±3.30	83.00±2.94	48.33±2.62	96.00±0.82	7.00±2.16	58.66±1.87
+ HeadLight	**54.00**±5.72	**73.33**±3.40	**89.00**±1.41	**52.67**±4.03	**97.00**±1.63	7.67±1.70	**62.28**±2.30

task-specific visual details such as background appearance and texture.

Action pathway: refining Motion Heads. Pretrained Motion Heads $\mathcal{H}^v_{\mathrm{M}}$ already capture general scene motion reliably, so we keep $\mathcal{H}^v_{\mathrm{M}}$ frozen to preserve these priors and instead adapt how the action stream reads them. Recall from Sec. 3.1 that each action head $(\ell, h)_a$ attends to video tokens by querying the keys and values of its paired video head $(\ell, h)_v$. We therefore select the action heads at the same layer and head indices as $\mathcal{H}^v_{\mathrm{M}}$:

$$\mathcal{H}^a_{\mathrm{M}} = \left\{ (\ell, h)_a \;\middle|\; (\ell, h)_v \in \mathcal{H}^v_{\mathrm{M}} \right\}, \tag{5}$$

which, for brevity, we also refer to as Motion Heads, and fine-tune only these K_{M} heads in the action stream. Because $\mathcal{H}^a_{\mathrm{M}}$ directly reads the motion features of $\mathcal{H}^v_{\mathrm{M}}$, adapting it helps generated actions follow predicted video trajectories without optical-flow supervision during fine-tuning.

Applicability across WAMs. HeadLight generalizes to joint-attention WAMs with either separate modality projections or a shared backbone, including human-video-conditioned Zero-WAM [36] and language-conditioned LingBot-VA [15]. It updates only a small fraction of parameters without incurring the additional inference overhead associated with activation steering methods [10, 37].

4. Experiments

In this section, we evaluate whether the identified Next-Interaction Heads and Motion Heads genuinely capture interaction targeting and motion execution, and whether parameter-efficient adaptation through these functionally specialized attention heads improves these capabilities and generalization to unseen manipulation tasks. Our experiments address three key questions: (i) Does selectively refining Next-Interaction Heads and Motion Heads outperform standard adaptation methods on unseen tasks across distinct WAM architectures? (ii) Are the identified heads causally necessary for their respective functions, with complementary benefits from their joint refinement? (iii) Do the performance improvements observed on unseen simulation tasks transfer to real-robot manipulation, where the model must generalize to unseen object–goal combinations from limited demonstrations?

4.1. Experimental Setup

Models and benchmarks. To test whether HeadLight generalizes across WAM designs, we evaluate two distinct architectures: **Zero-WAM** [36], which uses separate video and action heads conditioned on in-context human demonstration videos, and **LingBot-VA** [15], whose released model processes future video and action tokens within a shared transformer conditioned on language instructions. For simulation, we follow the cross-task protocol of Zero-WAM [36] on RoboTwin [3]. We fine-tune policies on 43 seen bimanual tasks and evaluate generalization on 6 unseen tasks, with 100 episodes × 3 seeds per task. We note that *Open Microwave* has a simulator issue and is therefore excluded from our evaluation (see Sec. E). We highlight that unseen tasks provide a direct test of manipulation-specific adaptation because they require transferring interaction targeting and motion execution to novel goals rather than memorizing training scenes.

Baselines and implementation. We compare HeadLight with the unadapted backbone, Full Fine-Tuning (Full FT), and LoRA [11]. All adaptation methods start from the same post-trained checkpoint and are trained on identical seen-task data for 2,000 steps on Zero-WAM and 10,000 steps on LingBot-VA. For LoRA, we choose rank r=77

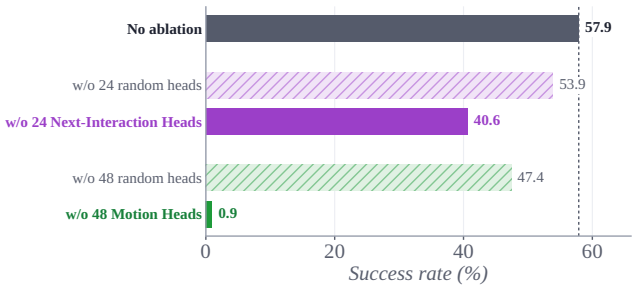


Figure 6. **Inference-time causal head ablation** on unseen RoboTwin tasks (Zero-WAM after Full FT). Lower success rates indicate stronger functional importance: ablating Next-Interaction Heads or Motion Heads substantially reduces manipulation success compared with equally sized random-head controls, confirming their critical roles in task execution.

to match HeadLight's trainable parameter budget on Zero-WAM (113.5M vs. 113.3M), applying it to the self-attention projections of all video and action heads in both WAMs. Full FT updates all model parameters. Full implementation, evaluation, and head-discovery details are provided in Secs. A and B.1.

4.2. Generalization to Unseen Tasks Across WAMs

As shown in Tab. 1, HeadLight achieves the highest overall success rate of **62.28%** on Zero-WAM, outperforming Full FT by 4.39 points and LoRA by 3.62 points. Unlike LoRA, which distributes its update capacity across attention heads without explicitly considering their functions, HeadLight adapts only the Next-Interaction and Motion Heads identified by our functional analysis. The higher success on unseen tasks suggests that selecting heads by function is an effective way to adapt WAMs.

We further evaluate HeadLight on LingBot-VA to assess its generalizability to a unified, language-conditioned WAM architecture. HeadLight consistently achieves the strongest performance, reaching **38.17%** average success and outperforming Full FT by 5.56 points and LoRA by 4.06 points. These results suggest that functionally guided head adaptation generalizes across different WAM architectures.

4.3. Ablation Studies

Functional roles of the identified heads. At inference time, we zero out each head group [19] on Zero-WAM to test their functional importance for successful manipulation. As shown in Fig. 6, ablating Next-Interaction Heads reduces success to 40.61%, while ablating Motion Heads lowers it to 0.94%, both substantially larger drops than ablating equally many random heads. The two sets are disjoint among all 720 video heads, and ablating them leads to failure modes in Fig. 2: mislocalized grasps missing the target without Next-Interaction Heads, versus ghosted object motion and stalled execution without Motion Heads.

Table 2. **Complementarity of the two adaptation pathways** on unseen RoboTwin tasks (Zero-WAM). We compare adapting Next-Interaction Heads, Motion Heads, and their combination against equally sized random-head controls. Success rates are reported as mean±std over 3 seeds, with 100 trials per seed.

Adapted heads	#Params	SR (%) ↑
None (Zero-WAM [36])	–	44.95±0.79
All (Full FT)	10.0B	57.89±0.77
Next-Interaction Heads only	37.8M	56.17±0.14
random heads control	37.8M	52.22±0.44
Motion Heads only	75.5M	56.11±1.36
random heads control	75.5M	54.33±1.65
Both (random heads control)	113.3M	57.33±1.89
Both (HeadLight)	113.3M	**62.28**±2.30

Complementarity of the two adaptation pathways. To examine whether the two adaptation pathways provide complementary benefits, we compare each pathway individually with their combination in Tab. 2. Adapting Next-Interaction Heads or Motion Heads alone outperforms adapting equally many random heads, with further gains achieved by combining both pathways in HeadLight, reaching 62.28% success, surpassing Full FT (57.89%) and outperforming either individual pathway by over six points. These results demonstrate the complementary benefits of the two adaptation pathways.

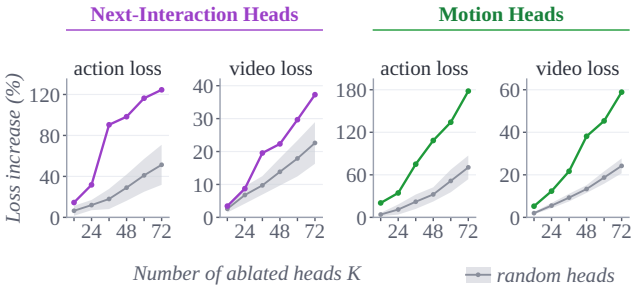


Figure 7. **Head-budget analysis via loss sensitivity.** Action- and video-loss increases from ablating top-K heads versus random-head controls (gray, mean±std) in Zero-WAM. Up to K_{N}=24, ablating Next-Interaction Heads increases action loss while leaving video loss near random-head controls; ablating Motion Heads substantially increases both losses across K.

Choice of head budgets. To determine how many heads are needed for each adaptation pathway, we examine how head budgets affect prediction losses and unseen-task success. Specifically, we progressively ablate the top-K heads and measure the resulting increases in action and video prediction losses against equally sized random-head controls, as shown in Fig. 7. For Next-Interaction Heads, the top 24 strongly affect action prediction while leaving video predic-

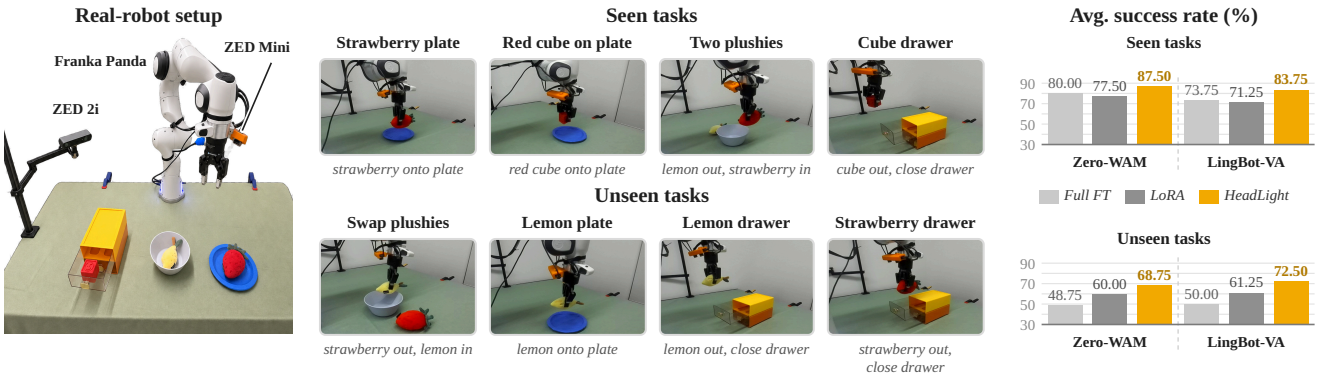


Figure 8. **Real-robot workspace, tasks, and results.** *Left:* Franka Panda workspace with third-person ZED 2i and wrist-mounted ZED Mini cameras. *Middle:* Four seen adaptation tasks (top) and four unseen tasks that recombine objects and manipulation goals (bottom), each evaluated over 20 rollouts. *Right:* Average success rate across four tasks per split. HeadLight outperforms Full FT and LoRA on both WAMs, with larger gains on unseen tasks.

tion largely intact; including more heads increasingly perturbs video prediction without improving unseen-task success (see Sec. C.2). For Motion Heads, adapting their paired action heads improves success up to $K_{\mathrm{M}}{=}48$, while ablating these 48 video heads reduces success to near zero in Fig. 6, indicating that these heads already encode scene motion. These results motivate selecting $K_{\mathrm{N}}{=}24$ for interaction targeting and $K_{\mathrm{M}}{=}48$ for motion-guided action adaptation, balancing functional specificity and adaptation effectiveness.

4.4. Real-World Manipulation

To evaluate HeadLight beyond simulation, we conduct real-world experiments on the 7-DoF Franka Panda arm shown in Fig. 8. We consider four seen tasks, each with 100 teleoperated demonstrations for adaptation, and four unseen tasks that recombine objects and manipulation goals in novel configurations. Each task is evaluated over 20 physical rollouts on both Zero-WAM and LingBot-VA, with per-task results reported in Sec. D.3.

HeadLight consistently outperforms both Full FT and LoRA across all settings on both WAMs. Beyond achieving the highest success rates on seen tasks, its gains are especially pronounced on unseen tasks, reaching 68.75% on Zero-WAM (+8.75 points over LoRA) and 72.50% on LingBot-VA (+11.25 points over LoRA). These results demonstrate that HeadLight's benefits extend beyond the training tasks and suggest that functionally guided adaptation improves the transfer of interaction targeting and motion execution to novel manipulation goals.

5. Conclusion

We show that pretrained World-Action Models (WAMs) contain two disjoint families of functionally specialized attention heads: Next-Interaction Heads, which localize upcoming interaction targets, and Motion Heads, which capture scene motion across frames. Through controlled inference-time ablations, we validate their distinct functional roles in manipulation, showing substantially larger performance drops than equally sized random-head controls. Guided by this functional structure, we introduce HeadLight, a parameter-efficient refinement method with two complementary adaptation pathways. It fine-tunes Next-Interaction Heads for interaction targeting and action heads paired with frozen video Motion Heads for motion execution. Experiments across two distinct WAM architectures, Zero-WAM and LingBot-VA, demonstrate consistent gains over full fine-tuning and LoRA on unseen RoboTwin tasks and on real-robot manipulation with a Franka Panda arm. Together, these findings suggest that functionally guided parameter selection can improve generalization by prioritizing manipulation-relevant capabilities over task-specific visual details.

6. Limitations and Future Work

Our head discovery has two requirements. First, although the identified Next-Interaction Heads can be reused across unseen tasks, their initial discovery requires datasets annotated with upcoming interaction targets. Second, identifying Motion Heads relies on attention from the current frame to the preceding frame to construct attention-induced motion fields. Extending motion-based head discovery to architectures with different temporal attention mechanisms remains an interesting direction.

Beyond discovery, HeadLight uses these functional heads to guide selective fine-tuning. Their specialized roles open up further opportunities to improve interaction targeting and motion execution. For instance, auxiliary localization objectives could sharpen interaction targeting, while task-specific cues extracted from human demonstra-

tions could provide more direct guidance to these heads. Alternatively, directly steering their activations could enable more flexible control over manipulation behavior without updating model parameters.

LLM Usage Statement

Claude Code was used to assist with code development and programmatic figure creation. ChatGPT and Claude assisted with manuscript writing and refinement. GPT Image 2.6 was used solely for background removal in the real-robot setup image. All AI-assisted content was reviewed and verified by the authors.

References

[1] 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)*, 2026. 14, 15

[2] Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Lucy Xiaoyang Shi, James Tanner, Quan Vuong, Anna Walling, Haohuan Wang, and Ury Zhilinsky. π_0: A vision-language-action flow model for general robot control. In *Proceedings of Robotics: Science and Systems (RSS)*, 2025. 2

[3] 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 *Proceedings of the International Conference on Machine Learning (ICML)*, pages 13673–13699, 2026. 2, 6, 13, 14

[4] Cheng Chi, Siyuan Feng, Yilun Du, Zhenjia Xu, Eric Cousineau, Benjamin CM Burchfiel, and Shuran Song. Diffusion Policy: Visuomotor policy learning via action diffusion. In *Proceedings of Robotics: Science and Systems (RSS)*, 2023. 2

[5] Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. What does BERT look at? An analysis of BERT's attention. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 276–286, 2019. 2

[6] J. Randall Flanagan, Miles C. Bowman, and Roland S. Johansson. Control strategies in object manipulation tasks. *Current Opinion in Neurobiology*, 16(6):650–659, 2006. 2

[7] Caelan Reed Garrett, Rohan Chitnis, Rachel Holladay, Beomjoon Kim, Tom Silver, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. Integrated task and motion planning. *Annual Review of Control, Robotics, and Autonomous Systems*, 4(1):265–293, 2021. 2

[8] Mor Geva, Jasmijn Bastings, Katja Filippova, and Amir Globerson. Dissecting recall of factual associations in auto-regressive language models. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 12216–12235, 2023. 4

[9] Yanjiang Guo, Yucheng Hu, Jianke Zhang, Yen-Jen Wang, Xiaoyu Chen, Chaochao Lu, and Jianyu Chen. Prediction with action: Visual policy learning via joint denoising process. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 112386–112410, 2024. 2

[10] Bear Häon, Kaylene Caswell Stocking, Ian Chuang, and Claire Tomlin. Mechanistic interpretability for steering vision-language-action models. In *Proceedings of the Conference on Robot Learning (CoRL)*, pages 2743–2762, 2025. 3, 6

[11] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations (ICLR)*, 2022. 2, 3, 6

[12] Sunyoung Jung, Jiwoo Park, Yoonseok Choi, Kyobin Choo, Ming-Hsuan Yang, and Seong Jae Hwang. Controlling motion transfer in diffusion transformers via attention heads. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 240–257, 2026. 3, 5

[13] 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 (ICLR)*, 2026. 2

[14] Till Kroeger, Radu Timofte, Dengxin Dai, and Luc Van Gool. Fast optical flow using dense inverse search. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 471–488, 2016. 5, 13

[15] 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 (RSS)*, 2026. 2, 6, 20

[16] Shuang Li, Yihuai Gao, Dorsa Sadigh, and Shuran Song. Unified video action model. In *Proceedings of Robotics: Science and Systems (RSS)*, 2025. 2

[17] Weixin Liang, Lili Yu, Liang Luo, Srinivasan Iyer, Ning Dong, Chunting Zhou, Gargi Ghosh, Mike Lewis, Wen-tau Yih, Luke Zettlemoyer, and Xi Victoria Lin. Mixture-of-Transformers: A sparse and scalable architecture for multi-modal foundation models. *Transactions on Machine Learning Research*, 2025. 4

[18] Bo Liu, Yifeng Zhu, Chongkai Gao, Yihao Feng, Qiang Liu, Yuke Zhu, and Peter Stone. LIBERO: Benchmarking knowledge transfer for lifelong robot learning. In *Advances in Neural Information Processing Systems (NeurIPS), Datasets and Benchmarks Track*, 2023. 14, 15

[19] Paul Michel, Omer Levy, and Graham Neubig. Are sixteen heads really better than one? In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 14014–14024, 2019. 7

[20] Chancharik Mitra, Brandon Huang, Tianning Chai, Zhiqiu Lin, Assaf Arbelle, Rogerio Feris, Leonid Karlinsky, Trevor Darrell, Deva Ramanan, and Roei Herzig. Enhancing few-shot vision-language classification with large multimodal model features. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2760–2772, 2025. 3

[21] Chancharik Mitra, Yusen Luo, Raj Saravanan, Dantong Niu, Anirudh Pai, Jesse Thomason, Trevor Darrell, Abrar Anwar, Deva Ramanan, and Roei Herzig. Mechanistic finetuning of vision-language-action models via few-shot demonstrations. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 137–155, 2026. 2, 3, 15

[22] Saman Motamed, Laura Culp, Kevin Swersky, Priyank Jaini, and Robert Geirhos. Do generative video models understand physical principles? In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 948–958, 2026. 2

[23] Jisu Nam, Soowon Son, Dahyun Chung, Jiyoung Kim, Siyoon Jin, Junhwa Hur, and Seungryong Kim. Emergent temporal correspondences from video diffusion transformers. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 123683–123727, 2025. 5

[24] NVIDIA. Cosmos 3: Omnimodal world models for physical AI. *arXiv preprint arXiv:2606.02800*, 2026. 14, 15

[25] Alexander Pondaven, Aliaksandr Siarohin, Sergey Tulyakov, Philip Torr, and Fabio Pizzati. Video motion transfer with diffusion transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 22911–22921, 2025. 3, 5

[26] Lu Qiu, Yizhuo Li, Yi Chen, Yuying Ge, Yixiao Ge, and Xihui Liu. Making foresight actionable: Repurposing representation alignment in world action models. *arXiv preprint arXiv:2606.12217*, 2026. 2

[27] 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. *arXiv preprint arXiv:2503.20314*, 2025. 2

[28] Eric Todd, Millicent L. Li, Arnab Sen Sharma, Aaron Mueller, Byron C. Wallace, and David Bau. Function vectors in large language models. In *International Conference on Learning Representations (ICLR)*, 2024. 2

[29] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 5998–6008, 2017. 4

[30] Elena Voita, David Talbot, Fedor Moiseev, Rico Sennrich, and Ivan Titov. Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 5797–5808, 2019. 2

[31] Hao Wang, Jiajun Wen, Jingzhi Liu, Shuoshuo Xue, Zhiliang Chen, Min Lin, Yicheng Chang, Xiaoyu Guo, Yukang Zhuo, Zheng Chong, Yunshuang Nie, Jian Zhang, Weijia Liufu, Qingman Wu, Heming Xu, Bingchang Song, Dantong Wu, Zhiyuan Wang, Hang Xu, Jianhua Han, Bokui Chen, Shen Zhao, Rui Li, and Xiaodan Liang. EWAM: Emergent depth-wise specialization in a unified embodied model – from semantic understanding through visual foresight to action. *arXiv preprint arXiv:2609.39973*, 2026. 3

[32] Yuran Wang, Siqiao Huang, Mingleyang Li, Chenhao Zhang, Jiaqi Liang, Weiyang Jin, Yue Chen, Xuemin Chi, Donghao Zhou, Qize Yu, Yu-Kai Wang, Yuhan Rui, Shenzhe Yao, Zhen Yuan, Zhenhao Shen, Kefei Zhu, Zijie Zhu, Ning Gao, Xiaowei Chi, Guanqi He, Shanghang Zhang, Hao Dong, Lin Shao, and Hang Zhao. OpenWAM: An open, modular exploration towards systematic world-action model pretraining. *arXiv preprint arXiv:2609.07398*, 2026. 14, 15

[33] Hongtao Wu, Ya Jing, Chilam Cheang, Guangzeng Chen, Jiafeng Xu, Xinghang Li, Minghuan Liu, Hang Li, and Tao Kong. Unleashing large-scale video generative pre-training for visual robot manipulation. In *International Conference on Learning Representations (ICLR)*, 2024. 2

[34] 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 Fan, and Joel Jang. World action models are zero-shot policies. *arXiv preprint arXiv:2602.15922*, 2026. 2

[35] Tianyuan Yuan, Zibin Dong, Yicheng Liu, and Hang Zhao. Fast-WAM: Do world action models need test-time future imagination? In *Advances in Neural Information Processing Systems (NeurIPS)*, 2026. 14, 15

[36] Jiaming Zhou, Qihang Zhang, Gangwei Xu, Cunxin Fan, Yujie Zhao, Ruilin Wang, Yiming Luo, Shuai Yang, Xing Zhu, Yujun Shen, Junwei Liang, and Yinghao Xu. Zero-WAM: In-context world-action modeling from human videos for open-ended task generalization. *arXiv preprint arXiv:2608.26103*, 2026. 2, 6, 7, 12, 13, 17, 18, 20

[37] Kaichen Zhou, Zhiyan Li, Chongkai Gao, Yuhan Wang, Wenlong Huang, Maggie Wang, Fangneng Zhan, Manling Li, Qianqian Wang, Mac Schwager, Zhuang Liu, Li Fei-Fei, Mengyu Wang, and Paul Pu Liang. SteerAct: Unlocking

action models through few-shot temporal steering. *arXiv preprint arXiv:2610.25632*, 2026. 2, 3, 6

[38] 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 (RSS)*, 2025. 2

[39] Brianna Zitkovich, Tianhe Yu, Sichun Xu, Peng Xu, Ted Xiao, Fei Xia, Jialin Wu, Paul Wohlhart, Stefan Welker, Ayzaan Wahid, Quan Vuong, Vincent Vanhoucke, Huong Tran, Radu Soricut, Anikait Singh, Jaspiar Singh, Pierre Sermanet, Pannag R. Sanketi, Grecia Salazar, Michael S. Ryoo, Krista Reymann, Kanishka Rao, Karl Pertsch, Igor Mordatch, Henryk Michalewski, Yao Lu, Sergey Levine, Lisa Lee, Tsang-Wei Edward Lee, Isabel Leal, Yuheng Kuang, Dmitry Kalashnikov, Ryan Julian, Nikhil J. Joshi, Alex Irpan, Brian Ichter, Jasmine Hsu, Alexander Herzog, Karol Hausman, Keerthana Gopalakrishnan, Chuyuan Fu, Pete Florence, Chelsea Finn, Kumar Avinava Dubey, Danny Driess, Tianli Ding, Krzysztof Marcin Choromanski, Xi Chen, Yevgen Chebotar, Justice Carbajal, Noah Brown, Anthony Brohan, Montserrat Gonzalez Arenas, and Kehang Han. RT-2: Vision-language-action models transfer web knowledge to robotic control. In *Proceedings of the Conference on Robot Learning (CoRL)*, pages 2165–2183, 2023. 2

Targeted Adaptation of World-Action Models via Interaction and Motion Heads

Supplementary Material

This supplementary material provides additional details and analyses supporting the main paper. Sec. A describes model architectures, training, head-level fine-tuning, and the evaluation protocol. Sec. B details the discovery procedure, reports causal ablations on Zero-WAM and four additional WAMs on RoboTwin and LIBERO, and validates head consistency. Sec. C reports per-task breakdowns, head-budget analyses, and qualitative comparisons. Sec. D details real-robot hardware, tasks, and per-task results. Sec. E documents a RoboTwin simulator artifact that leads us to exclude *Open Microwave* from evaluation.

A. Implementation and Experimental Details

A.1. Models and Training

We adapt HeadLight on Zero-WAM and LingBot-VA, which couple video and actions in different ways and both use 30 transformer layers with 24 attention heads of dimension 128 per layer. Zero-WAM employs separate video and action streams coupled through joint attention and conditions its generation on an in-context human demonstration video. In contrast, LingBot-VA processes future video and action tokens within a unified transformer with shared attention weights and conditions its generation on a language instruction. Beyond these adaptation backbones, Sec. B.3 applies head discovery and causal ablation to four more publicly released WAMs with different architectures, on RoboTwin and on LIBERO, confirming that the analysis pipeline transfers across architectures.

All adaptation methods are trained on clean demonstrations from the 43 seen tasks of RoboTwin. Zero-WAM pairs each robot episode with a human video from the Zero-WAM dataset [36] and jointly optimizes video prediction, action prediction, and multi-chunk prediction. LingBot-VA optimizes its standard video and action prediction objectives. HeadLight routes these objectives separately to its two adaptation pathways (Sec. A.2).

Starting checkpoints and optimization. For Zero-WAM, all methods start from the released RoboTwin post-trained checkpoint and train for 2,000 steps. Because the publicly released LingBot-VA RoboTwin checkpoint is trained on all 50 tasks including our held-out evaluation tasks, we first adapt the released pretrained checkpoint to the 43 seen tasks to obtain a clean starting model. All methods then train from this model for 10,000 steps. All fine-tuning runs use the AdamW optimizer with $\beta{=}(0.9, 0.95)$, zero weight decay, gradient clipping at 1.0, and a 200-step linear warm-up. On Zero-WAM, the learning rate then follows a cosine decay to $0.1\times$ the maximum; on LingBot-VA, it is held at the maximum after the warm-up and cosine-decayed to $0.1\times$ over the final $1{,}000$ steps. Both LoRA and HeadLight use a maximum learning rate of 10^{-4} with an effective batch size of 8 on Zero-WAM and LingBot-VA. Following their respective standard recipes, Full FT uses a maximum learning rate of $2{\times}10^{-5}$ with batch size 8 on Zero-WAM and 10^{-5} with batch size 16 on LingBot-VA. All experiments are conducted on 8 NVIDIA A100 or H100 GPUs.

A.2. Head-Level Selective Fine-Tuning

For each selected attention head (ℓ, h), HeadLight introduces a zero-initialized additive update to the 128 output rows of the query, key, and value projections, including their weights and biases, together with the 128 input columns of the output projection. Each selected head accounts for 1.57M trainable parameters. Adapting the 24 Next-Interaction Heads and 48 Motion Heads yields 113.3M trainable parameters in total, representing 1.1% of Zero-WAM's 10.0B parameters, while all remaining attention heads, feed-forward layers, normalizations, and embeddings stay frozen, substantially reducing training memory usage relative to Full FT.

Within each optimization step, the two adaptation pathways are updated by distinct loss signals from a single forward pass. Specifically, gradients from the video prediction objectives update only the Next-Interaction Heads, whereas gradients from the action prediction objective update only the action heads paired with the frozen Motion Heads. Because every head-level update is additive over existing linear projections, all updates can be merged directly into the base weights after training, leaving the model architecture and inference cost completely unchanged.

Adaptation on shared-weight WAMs. In the unified LingBot-VA architecture, video and action tokens share a single set of self-attention weights rather than separate modality streams. To preserve the same dual-pathway routing, HeadLight applies Next-Interaction Head updates exclusively to video tokens and Motion Head updates exclusively to action tokens, keeping separate updates per token type whenever a head belongs to both sets. Since LingBot-VA executes separate video and action forward passes at inference time, each update merges cleanly into the corresponding pass with zero inference overhead.

Baselines. Full FT updates the entire WAM backbone under the standard training objective. For LoRA, we apply

rank r=77 with scaling α=r to the query, key, value, and output projections of self-attention across all video and action heads in Zero-WAM and LingBot-VA, trained with the combined loss. This rank is chosen so that LoRA matches the trainable parameter budget of HeadLight in Zero-WAM, with 113.5M trainable parameters for LoRA versus 113.3M for HeadLight.

A.3. Simulation Evaluation Protocol

We follow the cross-task evaluation protocol of Zero-WAM [36] on RoboTwin 2.0 [3] using clean scenes and rasterized rendering. Following Zero-WAM, each unseen task is evaluated over three seeds with 100 expert-solvable scenes per seed, and all methods run on identical scenes. Each cell in Tab. 1 reports the mean and standard deviation across seeds; overall averages are computed per-seed before aggregating.

During closed-loop rollouts, the policy observes three camera views, predicts the next chunk of future video frames and end-effector actions, executes them, and replans until success or the episode limit. Zero-WAM uses one fixed human demonstration video per task with 50 denoising steps and classifier-free guidance scale 5; LingBot-VA uses the task language instruction with its released inference settings. The unadapted baselines in Tab. 1 follow the published results of [36]. We exclude *Open Microwave* due to a simulator success-check artifact. The additional WAMs of Sec. B.3 follow each model's official evaluation protocol.

B. Functional Head Discovery and Analysis

B.1. Discovery Protocol

Discovery procedure. To identify functional heads for a given WAM, we freeze the post-trained model and run a single teacher-forced forward pass over a small set of recorded demonstration episodes from its benchmark spanning diverse tasks, with only a few episodes per task. Based on an attention analysis of the visual tokens, we compute both the next-interaction selectivity score S_{N} and the motion alignment score S_{M} for every video attention head. We then rank all heads by each score and select the top K_{N} Next-Interaction Heads and the top K_{M} Motion Heads. The selected heads remain fixed and are reused across downstream tasks during adaptation. Because discovery requires no backpropagation or parameter updates, it serves as a lightweight, one-time preprocessing step before fine-tuning.

Selected heads. Fig. 9 visualizes the spatial distribution of the 24 Next-Interaction Heads and 48 Motion Heads across layers for both Zero-WAM and LingBot-VA. In both architectures, the two functional families are primarily concentrated in the middle transformer layers. They are completely disjoint in Zero-WAM and nearly disjoint in LingBot-VA.

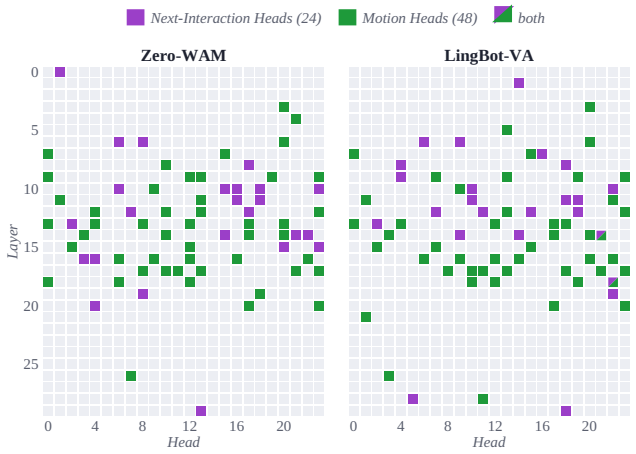


Figure 9. **Layer–head distribution of selected heads in Zero-WAM and LingBot-VA.** In both architectures, the identified Next-Interaction Heads and Motion Heads concentrate primarily in the middle transformer layers.

Interaction targets and controls. Contact events (grasps, presses, releases) and their 3D locations are extracted from simulator states and projected onto the head-camera token grid. The upcoming interaction target $\mathcal{T}$ for frame f is the token of the earliest subsequent contact, plus the other arm's next contact token if it occurs within one frame; frames whose targets lie outside the view are skipped. The *previous* control is the most recent prior contact token, and the *after-next* control is the first contact token after $\mathcal{T}$. The *random* control samples tokens outside the 3×3 neighborhood of $\mathcal{T}$, excluding gripper tokens and previous/after-next targets; the *equidistant* control samples under the same exclusions from tokens whose gripper distance is within one token of that of $\mathcal{T}$. Both match $\mathcal{T}$'s cardinality and are averaged over 10 draws; frames where the previous or after-next target is undefined or out of view are omitted from that control. Log-ratios in Eq. (2) are computed from frame-averaged $D_{\ell,h}$, then averaged across episodes with equal weight per task, before taking the minimum across controls.

Optical flow alignment. We estimate reference optical flow $\mathbf{u}^f$ using DIS [14] between consecutive recorded head-camera frames ($f-1$ and f), and spatially average the flow within each token cell. To focus on meaningful motion, we compute the expectation in Eq. (4) over token cells with reference flow magnitudes exceeding 0.25 tokens, average across frames within each episode, and then aggregate the results equally across tasks. For visualization in Figs. 5 and 13, both the reference and attention-derived motion

Table 3. **Per-task breakdown of causal head ablations** extending Fig. 6 on Zero-WAM after Full FT, reporting mean±std over three seeds with 100 trials per seed. Across tasks, ablating $\mathcal{H}_{\mathrm{N}}^{v}$ or $\mathcal{H}_{\mathrm{M}}^{v}$ degrades success substantially more than random-head controls.

Removed heads	#Heads	Place object on scale	Stamp seal	Move stapler to pad	Place bread in basket	Place empty cup	Stack blocks three	Average
None	0	43.33±1.25	73.33±6.65	85.33±0.47	42.67±2.87	91.33±2.05	11.33±3.09	57.89±0.77
Random heads	24	44.33±0.47	52.00±3.74	84.33±3.86	40.00±4.32	92.33±1.25	10.33±2.05	53.89±1.65
Next-Interaction Heads	24	26.00±2.45	38.00±2.94	64.67±2.49	24.33±2.49	90.33±0.94	0.33±0.47	40.61±0.28
Random heads	48	37.00±1.41	55.00±2.45	72.00±2.16	30.33±3.40	80.00±3.27	10.00±2.94	47.39±1.29
Motion Heads	48	0.67±0.47	0.67±0.47	3.33±0.94	0.67±0.47	0.33±0.47	0.00±0.00	0.94±0.21

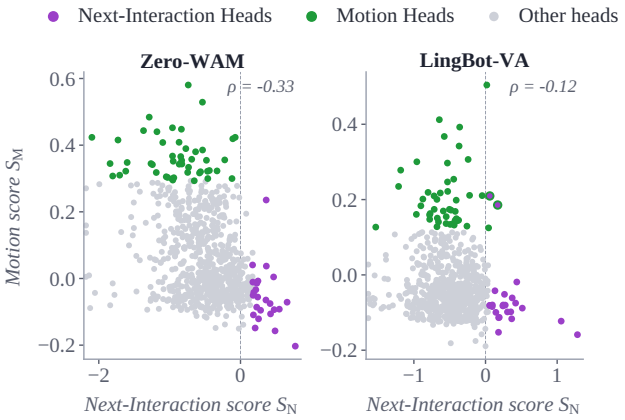


Figure 10. **Two distinct head families.** Next-Interaction score S_{N} versus Motion score S_{M} across all 720 video attention heads. In both architectures, the two families occupy opposite regions of the score space and exhibit a negative rank correlation ρ. Two LingBot-VA heads appearing in both sets are ringed.

fields are anchored at token cells in frame f, with arrows pointing to their corresponding locations in frame $f-1$.

Separation of the two functional families. Fig. 10 compares next-interaction selectivity S_{N} and motion alignment S_{M} across video attention heads. In both models, Next-Interaction Heads exhibit high target selectivity but low motion alignment, while Motion Heads show the reverse. This separation supports their distinct functional roles in interaction targeting and motion modeling.

B.2. Causal Head Ablations

Tab. 3 provides the per-task breakdown of the inference-time causal head ablations summarized in Fig. 6. In all ablations, we zero the outputs of the selected video heads on the robot's own video tokens while leaving the human demonstration conditioning and the action stream intact. Across every unseen task, removing either Next-Interaction Heads or Motion Heads degrades closed-loop manipulation substantially more than removing equally many random heads.

B.3. Causal Ablations on Other WAMs

We extend the causal ablations to four additional WAMs to examine whether the two head families play similar functional roles across architectures and benchmarks.

Models and benchmarks. We evaluate Fast-WAM [35] and Motus [1] on RoboTwin 2.0 [3], OpenWAM-α [32] and Cosmos3-Edge [24] on LIBERO [18]. The first three models use separate video and action experts, whereas Cosmos3-Edge denoises video latents and actions in a shared stream. All evaluations use released checkpoints without fine-tuning and cover tasks seen during training, testing the heads' causal roles in distribution.

Protocol. We identify both head families using the procedure in Sec. B.1 with a few demonstrations per task. Head budgets are guided by offline loss comparisons against layer-matched random ablations, following Fig. 7. The resulting $K_{\mathrm{N}}/K_{\mathrm{M}}$ are $30/48$ for Fast-WAM, $24/48$ for Motus, $48/18$ for OpenWAM-α, and $30/19$ for Cosmos3-Edge. At inference time, we zero the selected heads' outputs on video tokens and compare against equally many uniformly sampled random heads on identical episodes (5,000 each for Fast-WAM and Motus, and 2,000 each for OpenWAM-α and Cosmos3-Edge), following each model's official evaluation protocol. The unablated success rates in Tab. 4 are published results.

Two architectures require minor adjustments. Since Fast-WAM encodes only the current frame at inference time, we identify its Motion Heads using training-time video prediction and ablate them during current-frame encoding. For Cosmos3-Edge, we measure motion alignment from the current to the last predicted latent frame and restrict ablation to video-latent tokens in the shared stream.

Results. As shown in Tab. 4, ablating the selected heads generally leads to larger success drops than ablating equally

Table 4. **Causal head ablation on four additional public WAMs.** On each released checkpoint, we rediscover the two head families (Sec. B.1) and zero them on the video tokens at inference time, without fine-tuning, against the same number of random video heads on the same episodes. Original SR: published success rate (%) of the unablated checkpoint over all tasks of the benchmark; the ablation columns give the success rate (%) after zeroing the random or the selected heads.

(a) Ablating Next-Interaction Heads

Model	Benchmark	Original SR	Random head ablation	Next-Interaction head ablation
Fast-WAM [35]	RoboTwin 2.0	91.9	90.0	85.6
Motus [1]	RoboTwin 2.0	88.7	85.9	85.5
OpenWAM-α [32]	LIBERO [18]	99.3	98.6	91.9
Cosmos3-Edge [24]	LIBERO [18]	98.0	81.9	51.7

(b) Ablating Motion Heads

Model	Benchmark	Original SR	Random head ablation	Motion head ablation
Fast-WAM [35]	RoboTwin 2.0	91.9	91.8	90.5
Motus [1]	RoboTwin 2.0	88.7	85.0	75.8
OpenWAM-α [32]	LIBERO [18]	99.3	99.4	98.5
Cosmos3-Edge [24]	LIBERO [18]	98.0	96.6	37.5

many random heads across the four additional WAMs. The magnitude of these effects varies depending on how each model uses video information to predict action.

Next-Interaction Heads have a greater impact when actions rely on current-frame video tokens. Ablating them produces larger drops in Fast-WAM, OpenWAM-α, and Cosmos3-Edge, where the current image reaches the actions only through the video stream. The drop is much smaller in Motus, where action tokens also attend to an understanding branch built on a pretrained vision-language model. This branch encodes the current image and the instruction, giving action tokens a direct path to where they should interact, bypassing the video expert.

Motion Heads matter most when actions depend on predicted future frames. Their impact is therefore limited in Fast-WAM and OpenWAM-α. Fast-WAM does not generate future frames at inference: actions are denoised from features of the current frame, computed once by the video expert, so there are no consecutive frames for *Motion Heads* to relate. OpenWAM-α does denoise future frames jointly with the actions, but its action tokens also attend directly to the clean current-frame tokens. It also replans from a new observation after executing only 10 of its 32 predicted actions, which limits how far errors in the predicted motion can accumulate. In contrast, Motus and Cosmos3-Edge jointly denoise future frames and actions and execute full action chunks, leaving predicted-motion errors uncorrected within a chunk, which yields larger ablation drops. The effect is strongest in Zero-WAM, which generates future video before decoding actions. As a result, Motion Head ablation reduces success to nearly zero (Fig. 6).

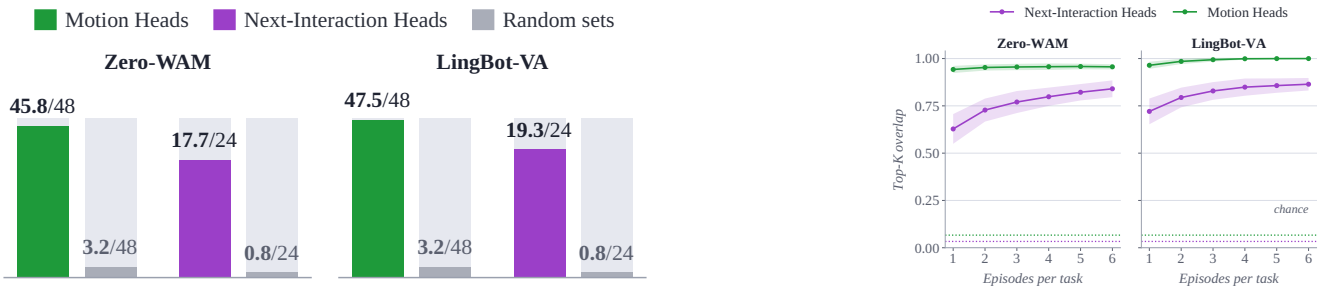
B.4. Consistency and Generality of Selected Heads

We examine the stability of the two head families in Zero-WAM and LingBot-VA, following the discovery procedure in Sec. 3.2.

Stability across demonstration subsets and data sizes. Following [21], we repeat head selection on disjoint demonstration subsets and measure the overlap between the selected sets. As shown in Fig. 11a, both families exhibit substantially greater overlap than random selections. Motion Heads remain stable even with a single demonstration per task, while Next-Interaction Heads become increasingly consistent as more demonstrations are used (Fig. 11b).

B.5. Visualizing Functional Specialization

Across additional tasks, the same Next-Interaction Head shifts its attention from grasp targets to subsequent release targets as interactions unfold (Fig. 12), while the same Motion Head produces attention-induced motion fields that align with reference optical flow (Fig. 13). Both figures show the same patterns in LingBot-VA, providing qualitative evidence of both specializations in a unified, language-conditioned WAM.



(a) Consistency of head selection.

(b) Discovery data size.

Figure 11. **Stability of head selection.** (a) Mean number of heads shared by the top-K sets selected from two disjoint subsets of the discovery episodes, over 200 random partitions into five subsets. Gray bars show the number shared by two random sets of K heads, and light bars mark K. (b) Overlap between top-K head sets selected from two disjoint demonstration sets with the same number of episodes per task, reported as mean±std over 200 draws with dotted lines showing random-set chance levels. Motion Heads stabilize from a single episode per task, while Next-Interaction Heads converge steadily as more episodes are added.

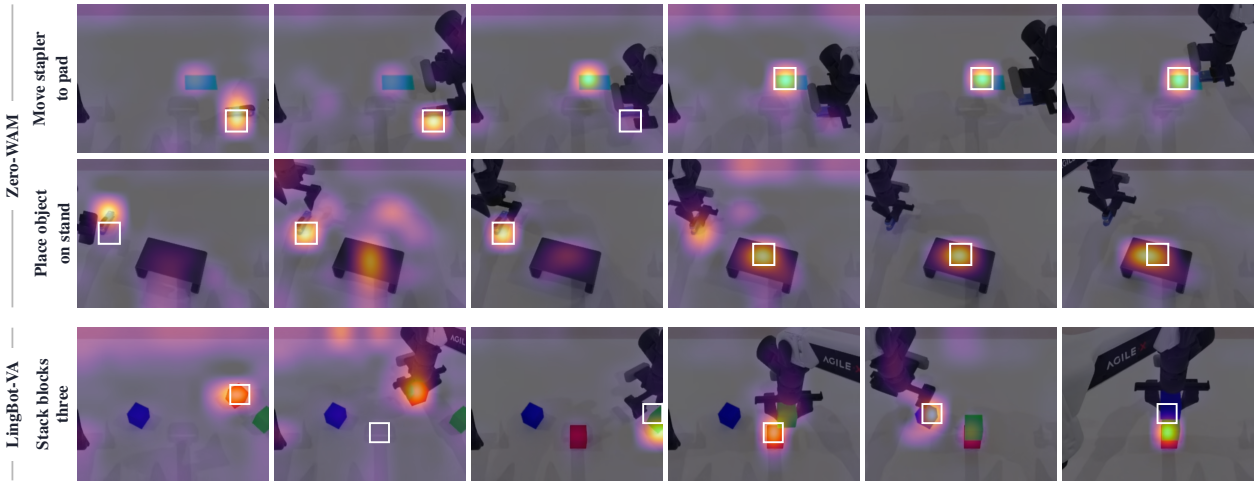


Figure 12. **Next-Interaction Heads over an episode.** Attention maps of a Next-Interaction Head in Zero-WAM (L10h23, six consecutive frames spaced 0.32 s apart) and in LingBot-VA (L11h10, every fifth frame of the episode), with each map scaled to its own maximum as in Fig. 4. The white square marks the upcoming interaction target $\mathcal{T}$. In Zero-WAM, attention focuses on the object to be grasped and, once it is grasped, shifts toward the upcoming release location; in LingBot-VA, it likewise concentrates on the upcoming target at key transitions.

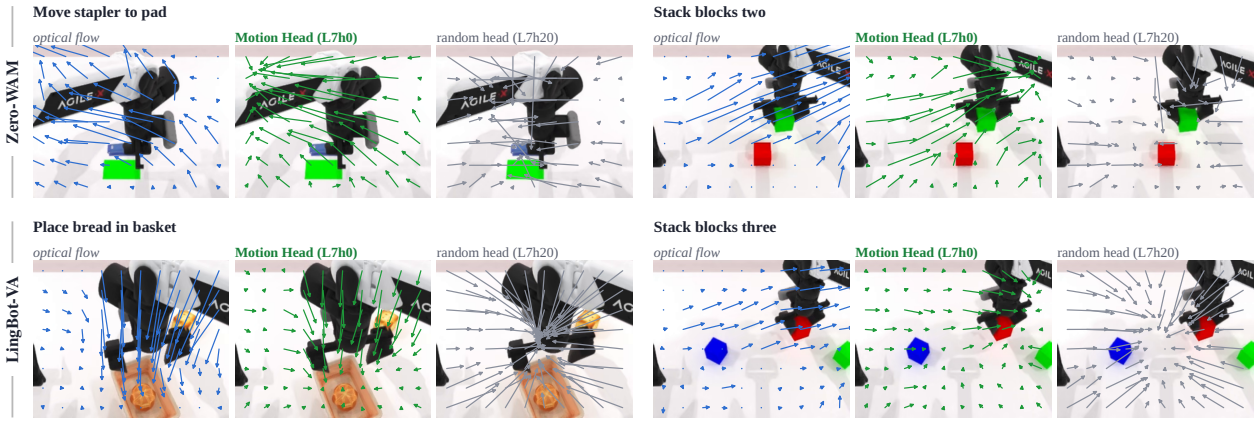


Figure 13. **Motion Heads on additional tasks.** Reference optical flow $\mathbf{u}^f$ and attention-induced motion fields $\widehat{\mathbf{u}}^f_{\ell,h}$ of the Motion Head L7h0 and the random head L7h20 from Fig. 5 on two tasks per model; L7h0 is a selected Motion Head in both models, L7h20 in neither. The Motion Head closely tracks the moving arm and manipulated objects in every example, whereas the random head does not.

Table 5. **Per-task breakdown of pathway complementarity** corresponding to Tab. 2 on the six unseen RoboTwin tasks with Zero-WAM, reporting mean±std over three seeds with 100 trials per seed. Combining both pathways in HeadLight achieves the best or tied-best performance on five of six tasks. Best per column in bold.

Adapted heads	Place object on scale	Stamp seal	Move stapler to pad	Place bread in basket	Place empty cup	Stack blocks three	Average
None (Zero-WAM [36])	24.67±2.05	47.00±4.55	69.14±2.93	35.00±3.74	84.87±0.18	9.00±2.16	44.95±0.79
All (Full FT)	43.33±1.25	**73.33**±6.65	85.33±0.47	42.67±2.87	91.33±2.05	**11.33**±3.09	57.89±0.77
Next-Interaction Heads only	44.33±1.70	59.00±4.55	86.00±0.82	48.33±2.87	91.67±3.09	7.67±0.47	56.17±0.14
random heads control	38.33±7.72	61.00±2.94	80.00±2.94	40.00±5.35	86.67±2.62	7.33±0.47	52.22±0.44
Motion Heads only	41.67±5.56	62.67±6.65	86.00±2.16	45.33±4.50	92.67±0.47	8.33±3.09	56.11±1.36
random heads control	40.00±1.63	59.33±8.73	84.00±0.82	44.67±2.87	91.33±1.70	6.67±0.47	54.33±1.65
Both (random heads control)	45.33±3.30	68.67±5.56	84.33±2.49	45.67±3.40	92.67±2.05	7.33±1.25	57.33±1.89
Both (HeadLight)	**54.00**±5.72	**73.33**±3.40	**89.00**±1.41	**52.67**±4.03	**97.00**±1.63	7.67±1.70	**62.28**±2.30

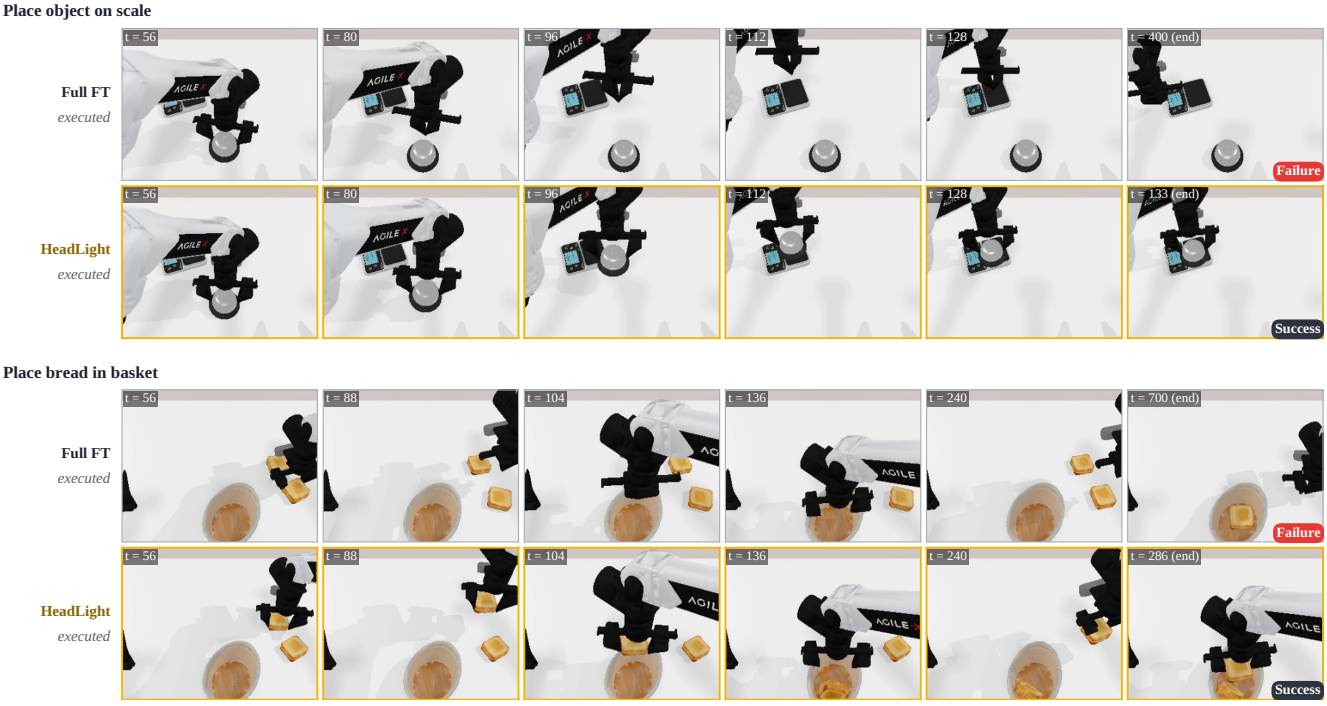


Figure 14. **Same-scene rollouts on unseen RoboTwin tasks with Zero-WAM.** Executed head-camera frames at matching action steps t, with the final column showing the end of each episode. Full FT closes its gripper on or beside the object but fails to lift it. HeadLight successfully grasps, carries, and places the object in both scenes.

C. Additional Experimental Results

C.1. Per-Task Pathway Complementarity

Tab. 5 breaks down Tab. 2 per task. Neither pathway alone matches the joint combination: HeadLight achieves the highest or tied-highest success on five of six unseen tasks.

C.2. Head Budgets

Tab. 6 reports unseen-task success when varying K within each pathway, complementing Fig. 7. Next-Interaction Head success peaks at K_{N}=24 and does not improve at 36 or 48, consistent with larger budgets increasingly perturbing video prediction. Motion Head success improves steadily up to K_{M}=48, supporting the default budgets.

Table 6. **Online head-budget ablation per task** on unseen RoboTwin tasks with Zero-WAM. The first row is the unadapted model. Each other row adapts only the top-K Next-Interaction Heads in the video stream or only the action heads paired with the top-K Motion Heads, reporting mean±std over three seeds with 100 trials per seed. Success peaks at K_{N}=24 and rises up to K_{M}=48.

Adapted heads	K	#Params	Place object on scale	Stamp seal	Move stapler to pad	Place bread in basket	Place empty cup	Stack blocks three	Average
None (Zero-WAM [36])	0	–	24.67±2.05	47.00±4.55	69.14±2.93	35.00±3.74	84.87±0.18	9.00±2.16	44.95±0.79
Next-Interaction Heads	12	18.9M	39.00±2.16	53.33±5.31	77.67±3.77	38.33±2.36	82.33±0.94	9.67±2.05	50.06±1.85
	24	37.8M	44.33±1.70	59.00±4.55	86.00±0.82	48.33±2.87	91.67±3.09	7.67±0.47	56.17±0.14
	36	56.6M	43.67±4.71	52.67±3.68	84.33±2.49	38.67±4.92	92.00±1.63	5.67±1.25	52.83±0.49
	48	75.5M	42.67±6.94	56.33±5.44	84.67±1.89	44.00±2.16	94.67±2.87	8.00±0.82	55.06±1.98
Motion Heads	12	18.9M	34.00±2.16	57.67±2.49	80.33±0.47	42.67±3.77	91.00±3.56	3.00±0.82	51.44±1.09
	24	37.8M	38.67±1.70	60.67±5.73	80.33±1.70	47.00±5.72	93.33±1.25	4.67±1.25	54.11±1.10
	36	56.6M	37.67±2.62	64.00±4.32	82.33±0.47	46.33±2.49	92.33±0.94	6.67±0.47	54.89±0.67
	48	75.5M	41.67±5.56	62.67±6.65	86.00±2.16	45.33±4.50	92.67±0.47	8.33±3.09	56.11±1.36

Table 7. **Real-robot tasks.** Each unseen task recombines the objects and manipulation steps of a seen task into a novel goal. A rollout succeeds when the final scene satisfies the specified success condition.

Split	Task	Success condition
Seen	Strawberry plate	strawberry plush on the blue plate
	Red cube on plate	red cube on the blue plate
	Two plushies	lemon plush out of the bowl, strawberry plush in the bowl
	Cube drawer	red cube out of the drawer on the table, drawer closed
Unseen	Swap plushies	strawberry plush out of the bowl, lemon plush in the bowl
	Lemon plate	lemon plush on the blue plate
	Lemon drawer	lemon plush out of the drawer on the table, drawer closed
	Strawberry drawer	strawberry plush out of the drawer on the table, drawer closed

C.3. Qualitative Comparison

Fig. 14 compares closed-loop rollouts of Full FT and HeadLight on identical unseen scenes. In both examples, the Full FT gripper closes without lifting the object and reaches the target empty-handed, despite its predicted video completing the grasp and transport. HeadLight reliably grasps, transports, and places the objects in both scenes.

D. Real-Robot Experiments

D.1. Hardware and Data Collection

The real-robot experiments in Sec. 4.4 are carried out on a 7-DoF Franka Panda arm equipped with a parallel-jaw gripper, observed by a third-person ZED 2i camera and a wrist-mounted ZED Mini camera as shown in Fig. 8. All eight tasks share a single tabletop workspace and a common set of objects: a red cube, a strawberry plush, a lemon plush, a blue plate, a bowl, and a drawer box. We collect 100 teleoperated demonstrations for each of the four seen tasks, and every method is adapted on the same 400 demonstrations. HeadLight reuses the heads discovered on RoboTwin for each WAM. Consistent with the simulation setup in Sec. 4.1, LingBot-VA is conditioned on task language instruction and Zero-WAM is conditioned on a human demonstration video of the task.

At control time, policies output absolute end-effector poses relative to the initial episode pose along with a binary gripper command at 13.5 Hz. Zero-WAM processes both camera views at 192×352 resolution and predicts action chunks spanning 16 actions and 8 video frames, corresponding to 1.2 s per chunk. Human demonstration videos for Zero-WAM are recorded from the third-person camera and resampled to match the frame rate of the robot observations.

D.2. Tasks and Evaluation Protocol

Tab. 7 defines the eight real-robot tasks, and Fig. 15 displays keyframes of a successful rollout for each task together with its language instruction. The four seen tasks provide the teleoperated demonstrations used for adaptation, whereas the four unseen tasks have no robot demonstrations and instead recombine the same objects and manipulation prim-

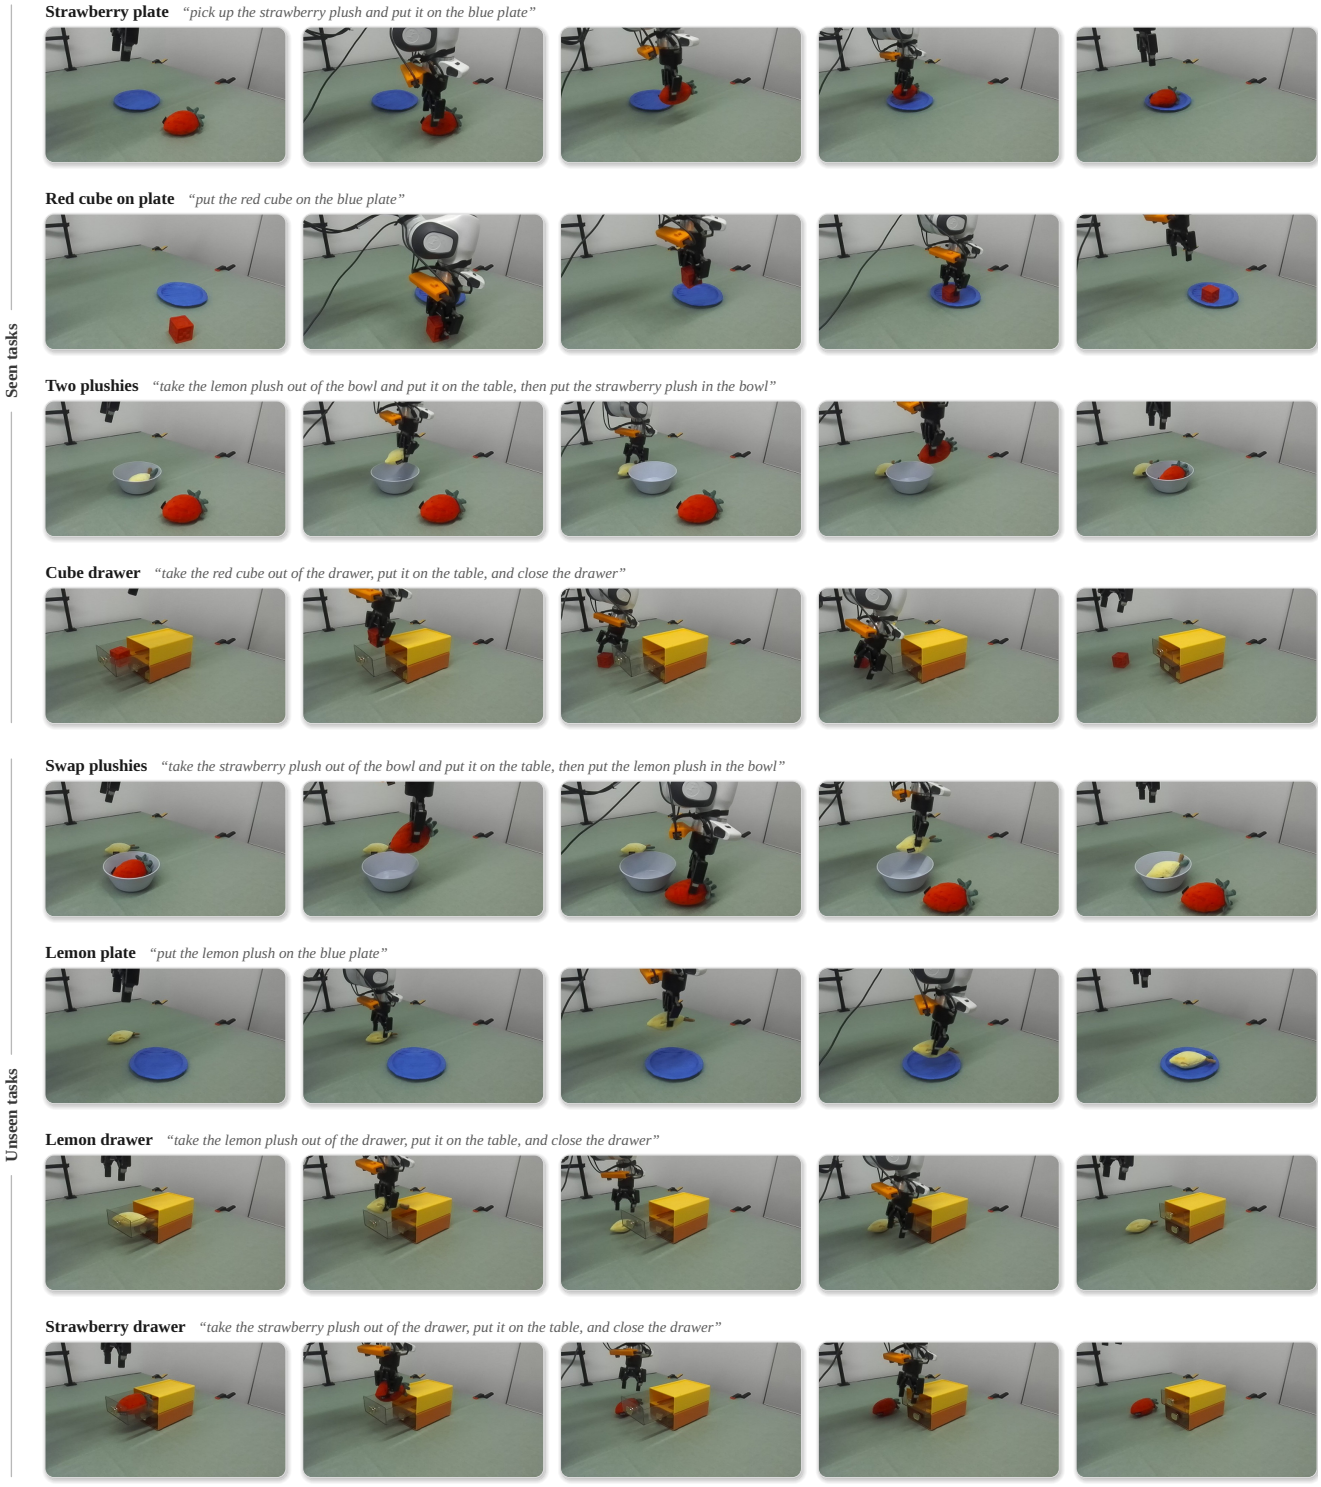


Figure 15. **Real-robot task rollouts.** Third-person keyframes ordered over time from left to right together with language instructions for all eight tasks in Tab. 8. Seen tasks provide teleoperated adaptation demonstrations, whereas unseen tasks recombine the same tabletop objects into novel manipulation goals.

Table 8. **Real-robot manipulation on a Franka Panda arm.** Policies adapted on four seen tasks with 100 demonstrations each are evaluated over 20 rollouts per task on both seen tasks and four unseen object–goal recombinations, matching the averages in Fig. 8. HeadLight consistently outperforms Full FT and LoRA across both WAMs, with the largest gains on unseen tasks. Best per backbone in bold.

	Seen tasks					Unseen tasks				
Method	Strawberry plate	Red cube on plate	Two plushies	Cube drawer	Avg. (%)	Swap plushies	Lemon plate	Lemon drawer	Strawberry drawer	Avg. (%)
Zero-WAM [36] + Full FT	18/20	17/20	**16/20**	13/20	80.00	7/20	13/20	10/20	9/20	48.75
Zero-WAM [36] + LoRA	17/20	17/20	15/20	13/20	77.50	8/20	15/20	**13/20**	12/20	60.00
Zero-WAM [36] + HeadLight	**19/20**	**20/20**	15/20	**16/20**	**87.50**	**10/20**	**18/20**	**13/20**	**14/20**	**68.75**
LingBot-VA [15] + Full FT	18/20	16/20	11/20	14/20	73.75	10/20	11/20	9/20	10/20	50.00
LingBot-VA [15] + LoRA	17/20	15/20	10/20	**15/20**	71.25	11/20	14/20	11/20	**13/20**	61.25
LingBot-VA [15] + HeadLight	**20/20**	**19/20**	**13/20**	**15/20**	**83.75**	**14/20**	**17/20**	**14/20**	**13/20**	**72.50**

itives into novel goals. A rollout is counted as successful if, once the robot stops, every object named in the instruction rests at its designated target location on the plate, in the bowl, or on the table, and for drawer tasks the drawer is pushed closed. Each method is evaluated over 20 physical rollouts per task.

D.3. Results

Tab. 8 reports the full success counts per-task of the real-robot corresponding to the averages shown in Fig. 8. HeadLight achieves the highest average success rate on both seen and unseen splits with each of Zero-WAM and LingBot-VA. Its advantages are most pronounced on unseen object–goal recombinations, where HeadLight outperforms LoRA by 8.75 percentage points on Zero-WAM and by 11.25 percentage points on LingBot-VA, confirming that functionally guided head adaptation transfers reliably to physical manipulation.

E. Notes on the Simulation Benchmark

Exclusion of Open Microwave. Zero-WAM [36] originally evaluated seven unseen RoboTwin tasks, including *Open Microwave*. We exclude it because the simulator judges success solely by the hinge angle of the door: the simulated door is lightweight and undamped, so incidental contact near the hinge or the top edge can swing it open without grasping the handle (Fig. 16). Replaying logged rollouts confirms that roughly 90% of successes across adapted policies exploit this artifact. The unadapted Zero-WAM average in Tab. 1 is recomputed on the remaining six tasks from the published per-task scores (44.95% vs. 46.95% across all seven).



Figure 16. **Open Microwave counts successes without the intended interaction.** *Top:* the RoboTwin scripted expert grasps the free edge of the door and pulls it open. *Bottom:* a Full FT rollout on the same scene touches the hinge-side corner of the door; the lightweight, undamped door swings from 0° to 50° within a single action and triggers the success check, highlighted in red.