Disentangling Spurious Correlations in Vision-Language-Action Models via Predicting Domain-Invariant Latent Lookahead

Junghyun Kim[1,2*], Ngseo Kim[1*], ChungWoo Lee[1], Seoyeon Lee[1], Woo-Jeong Baek[5], Adam Zhou[2], Chip Huyen[2], Jun-Ki Lee[1†], Gi-Cheon Kang[3†], Byoung-Tak Zhang[1,4†]

[1]Seoul National University, [2]OpenMind, [3]Ajou University
[4]Tommoro Robotics, [5]Hyundai Motors

***Abstract*—Vision-Language-Action (VLA) models remain brittle under visual distribution shifts, often relying on spurious correlations tied to domain-specific factors rather than task-relevant structure. We propose *Domain-Invariant Latent Lookahead* (DILL), a representation-learning framework that mitigates shortcut learning in VLA policies. Our key idea is to supervise policies with domain-invariant future latents learned from domain-transformed trajectory data. A Task-Domain Encoder is trained with contrastive objectives and Gaussian disentanglement regularization to separate task-relevant structure from domain-specific visual variation. The learned encoder then provides future latents for VLA policy learning through lookahead prediction and domain disentanglement, encouraging the policy to focus on task-relevant structure rather than incidental visual factors. Under counterfactual task–view compositions, DILL reduces shortcut degree while improving OOD success. On LIBERO-Plus visual perturbations, DILL achieves the highest average success rate of 69.1%, outperforming the strongest matched-input baseline by 11.4 absolute points, while limiting the average performance drop to 12.5 points. Latent-space diagnostics further show that DILL better preserves task-consistent structure while suppressing domain-specific variation.**

***Index Terms*—Vision-Language-Action Models, Shortcut Learning, Generalist Robot Policies, Predictive Future Latents, Domain Generalization**

I. Introduction

A longstanding goal of robot learning is to build robots that generalize across a wide range of tasks and environments. Recent advances in Vision-Language-Action (VLA) models [8, 26, 21, 7] trained on large-scale robot datasets [12, 20] have established a promising path toward generalist robot policies by grounding actions in language and visual observations. Despite this progress, VLA models often remain brittle under distribution shift, particularly under incidental visual variation such as changes in viewpoint, background, lighting, or camera configuration [14, 39, 36], which should be irrelevant to task success. This brittleness undermines reliable out-of-distribution generalization and remains a major obstacle to practical deployment.

Such failures can be naturally understood through the lens of *shortcut learning* [15, 38], where policies rely on task-irrelevant visual cues that correlate with successful actions in the training distribution rather than the task-relevant semantics required for robust control. VLA models are particularly susceptible to shortcut learning for two reasons. First, robot datasets often exhibit limited within-dataset diversity and are fragmented across collection sources, causing the same task or action to repeatedly co-occur with particular visual cues [38]. Second, VLA models often inherit representations from pretrained vision-language backbones [30, 3] that are optimized for visual understanding or image-language alignment rather than control, and may therefore preserve domain-specific visual information irrelevant to action generation. When mapped to actions, these representations can allow spurious visual correlations in training data to be absorbed into the policy's decision rule [13], producing the task-substitution failure illustrated in Fig. 1: under counterfactual task–domain recomposition, the policy may execute the task associated with the observed domain.

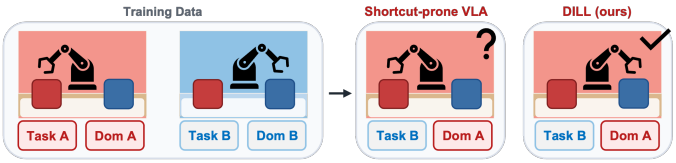

Fig. 1. **Shortcut learning under task–domain confounding.** When tasks and visual domains are spuriously correlated, a VLA may use domain-specific appearance as a shortcut for action selection. DILL reduces this failure by conditioning actions on a domain-invariant latent lookahead.

To mitigate such shortcuts, a VLA policy should learn representations that preserve task-relevant structure while discarding incidental visual variability. A natural learning signal for this purpose is future state prediction [43, 33]. Predicting how a task will evolve can encourage models to capture factors that determine future task states, such as object configuration, task progress, and action-conditioned scene changes. In this sense, future prediction provides a form of supervision that is more closely tied to control than to static visual appearance. However, future prediction alone is not sufficient. If the target is a raw future observation or an entangled visual latent, the model may still preserve domain-specific factors such as viewpoint or background that help predict future appearance but do not determine the correct action. Thus, the key question is not simply whether a VLA policy should predict the future,

but *what representation of the future it should predict*. We argue that the predictive target should be *domain-invariant*: it should retain the future task structure needed for control while suppressing domain-specific factors that can act as shortcuts.

We propose *Domain-Invariant Latent Lookahead* (DILL), a predictive representation learning framework for robust VLA control. DILL first learns a Task-Domain Encoder from domain-transformed trajectory chunks using contrastive objectives over two pair types: *task-positive pairs* that preserve trajectory chunk across domain changes, and *domain-positive pairs* that share the same domain condition across different chunks. A Gaussian disentanglement loss further separates task and domain latents: task latents remain stable across changes in viewpoint, environment appearance, camera configuration, and visual degradation, while domain latents capture such domain-specific factors. During policy learning, the pretrained Task-Domain Encoder provides future latent supervision: the VLA policy predicts a lookahead latent aligned with the future task latent, while its current representation is regularized to be disentangled from the corresponding domain latent. The predicted lookahead latent and current representation jointly condition the action head, encouraging the policy to base its actions on task-relevant future structure rather than incidental visual factors.

The central effect of DILL is to reshape the information that the policy uses to choose actions: instead of allowing domain-specific visual cues to enter the action head as reliable proxies for the task, DILL conditions actions on a predicted future latent that is stable across domain changes and informative about the commanded behavior. This makes the learned action representation less tied to where or how a scene is observed, and more tied to the information needed to produce the correct action. Empirically, we evaluate this effect through complementary behavioral and representational analyses. Under counterfactual task–view compositions, DILL substantially reduces shortcut reliance and improves OOD task success, showing that policies are less likely to substitute the commanded task with the task spuriously associated with the observed view. On LIBERO-Plus visual perturbations, DILL achieves the highest average success rate among matched-input baselines and exhibits a substantially smaller performance drop under camera, lighting, background, and sensor-noise shifts. Latent-space diagnostics further show that the final action-conditioning representation is organized more by task-consistent trajectory content than by shared visual appearance. Together, these results show that shortcut learning in VLA policies can be mitigated not only by increasing data diversity, but by explicitly shaping the predictive representation used for control.

II. Related Work

A. Spurious Correlations in Robot Learning

Spurious correlations [38] have emerged as an important challenge in robot learning, where policies may rely on incidental visual factors, such as camera viewpoint, background, and lighting, that correlate with successful actions in the training data [14, 39, 44, 36]. Prior work has addressed this issue from two perspectives: *data-centric* and *representation-centric* approaches. Data-centric approaches increase nuisance diversity through novel-view synthesis [35] or synthetic viewpoint augmentation [38]. Within representation-centric approaches, one line of work suppresses task-irrelevant factors by building invariance into visual representations, for example, by reducing sensitivity to viewpoint changes [24, 31, 27]. Another line of work injects task-relevant structure, such as spatial information [40]. Our work bridges these directions by using *predictive future latents* as task-relevant supervision while explicitly suppressing task-irrelevant visual factors through domain-invariant latent disentanglement.

B. World Models and Future Prediction in VLA Models

Recent VLA work has incorporated *future prediction* or *world modeling* signals to complement direct perception-to-action learning by anticipating the consequences of actions [43, 18, 42]. One line performs pixel-level imagination by generating future frames or subgoal observations before acting, as in the *GR* series, Ctrl-World, CoT-VLA, and SuSIE [37, 11, 16, 42, 6]. These approaches provide interpretable visual foresight, but they can be costly and prone to compounding errors. The second line moves future prediction into latent space by predicting or aligning future latent representations, including Video Prediction Policy, FLARE, VLA-JEPA, and FRAPPE [18, 43, 33, 41, 1, 2]. These methods avoid explicit pixel rollouts while encouraging representations that capture future structure useful for control. In contrast, DILL focuses on what the policy is asked to predict: rather than aligning to raw or entangled future latents, it learns future targets whose task-relevant structure is separated from domain-specific visual variation.

III. Method

A. Overview

Domain-Invariant Latent Lookahead (DILL) trains a VLA policy to predict and use future task structure without relying on visual-domain shortcuts. The central difficulty is that ordinary future prediction can preserve the same nuisance cues that make policies brittle: viewpoint, background, lighting, camera response, and image degradation may all be predictive of future appearance while remaining irrelevant to the intended action. DILL therefore separates the representation problem from the policy-learning problem.

As shown in Fig. 2, the first stage learns a *Task-Domain Encoder* from domain-transformed trajectory chunks. This encoder decomposes a video chunk into a task latent, which should remain stable across visual-domain changes, and a domain latent, which should capture the visual condition itself. The second stage freezes this encoder and uses it to supervise VLA policy learning. From the current observation and language instruction, the policy predicts a lookahead latent aligned with the future task latent and a current representation regularized against the corresponding domain latent. The action head is conditioned on both policy latents. Thus,

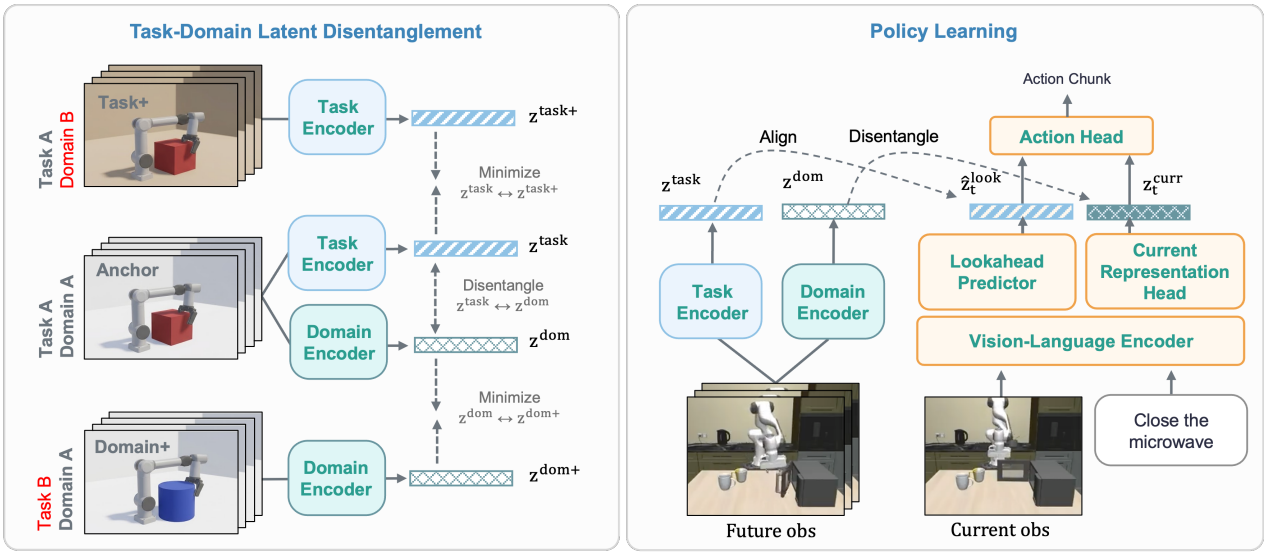

Fig. 2. **Overview of Domain-Invariant Latent Lookahead.** We first learn a Task-Domain Encoder that maps observation chunks into task and domain latents. During policy learning, a Lookahead Predictor predicts a future task latent from the current observation and instruction, while a Current Representation Head produces a domain-disentangled current representation. The two policy latents condition the action head for control.

future observations and transformed views provide training-time supervision, but test-time control still uses only the current RGB observation and instruction.

B. Learning a Task-Domain Latent Space

Let $\mathbf{o}_t^i = (o_t^i, \ldots, o_{t+T_v-1}^i)$ denote a video chunk from trajectory i. Let $\mathcal{T}_\eta$ be a domain transformation, where the condition η specifies the transformation family, parameters, and random seed. The same condition is applied consistently across all frames of a chunk, so the transformed input $x = \mathcal{T}_\eta(\mathbf{o})$ remains a coherent video observation. We pretrain the Task-Domain Encoder once on a source trajectory collection composed of ManiSkill [34], MimicGen [25], and OXE subsets [12]. These source trajectories provide the visual and temporal diversity used to define controlled interventions before any downstream LIBERO-Plus evaluation is considered. Appendix A-A provides the data and transformation details.

a) Domain interventions: The transformation families are chosen to match visual factors that often become shortcuts for robot policies. Viewpoint interventions change camera pose, crop, intrinsics, or geometric distortion. Environment-appearance interventions change backgrounds, lighting, or colors while preserving the manipulated scene content. Camera-heterogeneity interventions simulate differences in camera response, sensor noise, and image formation. Visual-degradation interventions apply blur, compression, weather effects, pixelation, or related corruptions. A sampled condition η may contain one transformation or a combination of transformations, but it is held fixed across the frames of a chunk. This gives two controlled comparisons: changing η for the same chunk changes appearance while preserving task state, whereas applying the same η to different chunks preserves the visual condition while changing task content. The transformations are therefore used to learn reusable task and domain spaces, not to train the downstream policy directly on LIBERO-Plus perturbation images.

b) Positive-pair mining: We construct two complementary positive pair types:

$$\underbrace{\big(\mathcal{T}_{\eta_1}(\mathbf{o}_t^i), \mathcal{T}_{\eta_2}(\mathbf{o}_t^i)\big)}_{\text{task-positive}}, \quad \eta_1 \neq \eta_2, \qquad \underbrace{\big(\mathcal{T}_\eta(\mathbf{o}_t^i), \mathcal{T}_\eta(\mathbf{o}_{t'}^j)\big)}_{\text{domain-positive}}, \quad i \neq j. \tag{1}$$

Task-positive pairs preserve the same trajectory content while changing the visual domain. They teach the task encoder to align observations that differ in incidental appearance but correspond to the same task state. Domain-positive pairs preserve the same domain condition while changing trajectory content. They teach the domain encoder to align observations that share visual factors but need not share task progress or action semantics. Representative task-positive transformations are shown in Fig. 3. The pair design is deliberately local. For task positives, we do not simply treat all clips from the same task or language instruction as positives, because that would allow the task branch to collapse different temporal states of the task. Instead, the positive relation preserves the underlying chunk and intervenes only on the visual domain, so invariance is tied to task state rather than task label. For domain positives, we keep the visual condition fixed while changing the trajectory content, so the domain branch is discouraged from representing action phase or object configuration as the positive signal. The two pair types therefore define complementary supervision: one asks what should remain stable under a domain intervention, and the other asks what is shared by the intervention itself. In implementation, each task-positive item samples multiple windows from the same episode and applies two sampled domain conditions, so nearby task states are not all treated as equivalent. Those additional windows become ordinary in-batch negatives after collation, which discourages

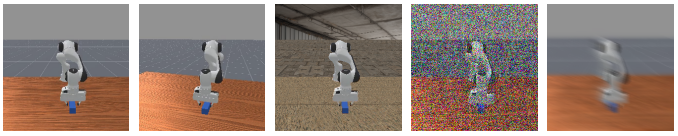

Fig. 3. **Representative task-positive transformations.** Each transformed frame preserves the same underlying trajectory chunk while changing a task-irrelevant visual domain factor.

the task branch from using only task identity. For the domain branch, each item applies several distinct domain conditions to paired clips with different trajectory content. The resulting minibatch contains anchors that share source content but differ in domain condition, and anchors that share domain condition but differ in source content. This controlled batching makes the contrastive objective test the intended factor rather than whichever correlation is easiest in the raw dataset.

c) Encoder objectives: The Task-Domain Encoder contains two video-encoder branches:

$$z^{\text{task}} = E_{\psi}^{\text{task}}(x), \qquad z^{\text{dom}} = E_{\xi}^{\text{dom}}(x). \tag{2}$$

Both branches use the same contrastive form with different positives. For a minibatch of positive pairs $\{(x_i, x_i^+)\}_{i=1}^{B}$, let z_i and z_i^+ denote the corresponding pooled latents. We use

$$\mathcal{L}_{\text{NCE}} = -\frac{1}{B}\sum_{i=1}^{B} \log \frac{\exp(\text{sim}(z_i, z_i^+)/\tau)}{\sum_{k=1}^{B} \exp(\text{sim}(z_i, z_k^+)/\tau)}, \tag{3}$$

where $\text{sim}(\cdot,\cdot)$ is cosine similarity. Applying Eq. 3 to task-positive pairs gives $\mathcal{L}_{\text{task}}$, and applying it to domain-positive pairs gives $\mathcal{L}_{\text{dom}}$.

The pair definitions establish different invariances, but they do not by themselves prevent leakage between branches. We therefore apply a Gaussian disentanglement regularizer to the concatenated task-domain latent:

$$u_i^{\text{td}} = [z_i^{\text{task}}; z_i^{\text{dom}}], \qquad \mathcal{L}_{\text{dis}}^{\text{td}} = \mathcal{R}_{\text{dis}}\left(\{u_i^{\text{td}}\}_{i=1}^{B}\right). \tag{4}$$

Under a Gaussian approximation, matching the joint latent to an isotropic Gaussian discourages cross-covariance between task and domain components. This makes the task latent a cleaner future target for policy supervision and makes the domain latent a useful nuisance factor for current-representation regularization. Appendix A-B gives the SIGReg implementation and the corresponding Gaussian argument. The Task-Domain Encoder is trained with

$$\mathcal{L}_{\text{TDD}} = \lambda_{\text{task}}\mathcal{L}_{\text{task}} + \lambda_{\text{dom}}\mathcal{L}_{\text{dom}} + \lambda_{\text{dis}}^{\text{td}}\mathcal{L}_{\text{dis}}^{\text{td}}. \tag{5}$$

C. VLA Policy Learning with Latent Lookahead

The policy-learning stage uses the pretrained Task-Domain Encoder as a frozen source of supervision. Given the current observation o_t and language instruction ℓ, the vision-language backbone [3] produces image-language tokens:

$$H_t^{\text{vlm}} = M_{\theta}(o_t, \ell). \tag{6}$$

Two policy heads map these tokens to a predicted future task latent and a current action-conditioning representation:

$$\hat{z}_t^{\text{look}} = P_{\theta}^{\text{look}}(H_t^{\text{vlm}}), \qquad z_t^{\text{curr}} = P_{\theta}^{\text{curr}}(H_t^{\text{vlm}}). \tag{7}$$

a) Frozen future targets: Let $\mathbf{o}_t^+ = (o_{t+1}, \ldots, o_{t+T_v})$ be the future observation chunk available during training. The frozen Task-Domain Encoder extracts a future task target and its corresponding domain latent:

$$z_t^{\text{task},\star} = E_{\bar{\psi}}^{\text{task}}\left(\mathbf{o}_t^+\right), \qquad z_t^{\text{dom},\star} = E_{\bar{\xi}}^{\text{dom}}\left(\mathbf{o}_t^+\right). \tag{8}$$

The lookahead loss aligns $\hat{z}_t^{\text{look}}$ with $z_t^{\text{task},\star}$ using Eq. 3: predicted lookahead latents serve as anchors, stop-gradient future task latents serve as positives, and other future task latents in the minibatch serve as negatives. We denote this loss by $\mathcal{L}_{\text{look}}$. Because the target is the task branch rather than a raw visual future feature, the lookahead objective asks the policy to predict future task structure without matching the future camera condition, background, or image statistics. Freezing the encoder also prevents the target space from drifting to whatever representation most easily fits the downstream behavior-cloning loss. The policy is instead trained against a stable coordinate system learned from controlled domain interventions.

b) Current representation: The current representation is regularized against the frozen domain target from the same trajectory:

$$u_i^{\text{curr}} = \left[z_i^{\text{curr}}; \text{sg}\left(z_i^{\text{dom},\star}\right)\right], \qquad \mathcal{L}_{\text{dis}}^{\text{curr}} = \mathcal{R}_{\text{dis}}\left(\{u_i^{\text{curr}}\}_{i=1}^{B}\right). \tag{9}$$

This objective discourages the representation routed to the action head from co-varying with domain-specific appearance. At the same time, behavior cloning still requires z_t^{curr} to retain the present-state information needed for control, so DILL biases the policy away from shortcuts rather than deleting all visual information. The two policy latents therefore play different roles. The predicted lookahead latent is a future-facing task signal: it encourages the current observation and instruction to encode where the task is going in the learned task space. The current representation remains a present-facing control signal: it can encode object pose, robot state, and other immediate information needed for action prediction, but is regularized not to organize itself around the visual domain factor. Conditioning the action head on both latents gives the policy access to task progress and current control information while reducing the incentive to use incidental appearance as a proxy for the commanded behavior.

The action head receives both the predicted lookahead and the current representation:

$$\hat{a}_{t:t+K_a-1} = A_{\phi}\left([\hat{z}_t^{\text{look}}; z_t^{\text{curr}}]\right). \tag{10}$$

We supervise the action chunk with behavior cloning,

$$\mathcal{L}_{\text{act}} = \frac{1}{K_a}\sum_{k=0}^{K_a-1} \|\hat{a}_{t+k} - a_{t+k}\|_1, \tag{11}$$

and optimize

$$\mathcal{L}_{\text{policy}} = \lambda_{\text{act}}\mathcal{L}_{\text{act}} + \lambda_{\text{look}}\mathcal{L}_{\text{look}} + \lambda_{\text{dis}}^{\text{curr}}\mathcal{L}_{\text{dis}}^{\text{curr}}. \tag{12}$$

D. Training and Inference

Training proceeds in two reusable stages. First, the Task-Domain Encoder is optimized on transformed source trajectories with $\mathcal{L}_{\text{TDD}}$ and then kept fixed. Second, the VLA policy is pretrained on action-labeled source trajectories and fine-tuned on downstream demonstrations such as LIBERO [23] using $\mathcal{L}_{\text{policy}}$. Each task/domain encoder branch uses a VJEPA2 video backbone [2], branch-specific LoRA adapters [17], and an attention-pooling projection head. The policy uses a Qwen2.5-VL backbone with LoRA adaptation, a Lookahead Predictor for $\hat{z}_t^{\text{look}}$, a Current Representation Head for z_t^{curr}, and an action head that predicts a 50-step action chunk from their concatenation. At inference time, the future chunk, transformed views, and domain labels are absent; the policy only receives the current RGB observation and language instruction.

IV. Experiments

We evaluate the central claim of this paper: domain-invariant latent lookahead mitigates shortcut learning in vision-language-action policies. The main experiments focus on controlled simulation settings, where task–domain correlations and visual distribution shifts can be systematically manipulated. Specifically, we ask:

1. Does our approach reduce shortcut reliance under counterfactual task–view compositions?

2. Does our approach improve robustness under visual distribution shifts?

3. Do the learned latents successfully separate task structure from task-irrelevant domain factors?

A. LIBERO Shortcut Diagnostic Under Counterfactual Task–View Compositions

We first test shortcut mitigation in a controlled LIBERO diagnostic [38]. This diagnostic intentionally creates a spurious correlation between task identity and camera viewpoint during training, then breaks this correlation at test time to measure whether a policy follows the commanded task or the task spuriously associated with the observed view.

Protocol. The training split consists of two confounded task–view islands. A left-view task group, denoted as Task-L, is observed only from a left-view range, while a right-view task group, denoted as Task-R, is observed only from a right-view range. During evaluation, we counterfactually swap these compositions: Task-R is evaluated under the left-view boundary, and Task-L is evaluated under the right-view boundary. Additional details are provided in Appendix B-A.

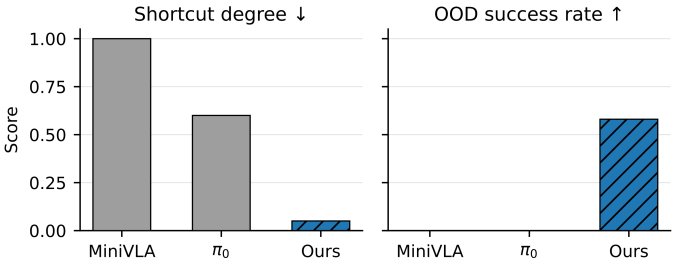

Fig. 4. **LIBERO shortcut diagnostic.** Lower shortcut degree and higher OOD success indicate less domain-induced task substitution.

Metrics. We report OOD success and shortcut degree. OOD success measures whether the commanded task is completed under the counterfactual view. Shortcut degree measures whether the policy instead executes the task group spuriously associated with the observed view during training. Lower shortcut degree indicates less shortcut reliance.

Results. Figure 4 compares our method against MiniVLA [5] and π_0 [7], using the baseline results reported in the low-diversity setting of [38]. MiniVLA and π_0 exhibit shortcut-prone behavior, with shortcut degrees of 1.00 and 0.60, respectively, and both obtain zero OOD success. In contrast, our method reduces shortcut degree to 0.05 while improving OOD success to 0.58. Importantly, shortcut degree measures view-induced task substitution rather than generic task failure. Thus, this result shows that our method does not merely improve execution under distribution shift; it specifically reduces the tendency to replace the commanded task with the task associated with the observed view. This directly supports our central hypothesis that domain-invariant latent lookahead discourages reliance on spurious task–view correlations.

Variant	Components				Metrics				
	SA	TD	CH	LA	In-dist. SR ↑	Center OOD SR ↑	Counter OOD SR ↑	Avg. SR ↑	Shortcut degree ↓
Base VLA					68%	22%	0%	30%	73%
Base VLA + SA	√				**82%**	10%	2%	31%	69%
Entangled LA	√			√	76%	24%	0%	33%	65%
DILL w/o CH	√	√		√	60%	44%	44%	49%	6%
DILL w/o LA	√	√	√		50%	**62%**	**62%**	58%	**3%**
DILL	√	√	√	√	66%	**62%**	58%	**62%**	5%

TABLE I

Ablation study. Component columns indicate source augmentation (SA), task–domain latent targets (TD), the disentangled current head (CH), and latent lookahead alignment (LA).

B. Ablation Study

We ablate DILL to identify which design choices are responsible for shortcut mitigation. The central question is not whether a model can fit the confounded training distribution, but whether it can avoid using viewpoint as a proxy for task identity when the task–view association is broken. DILL is designed to address this by reshaping the information routed to the action head: a predicted lookahead latent aligned with the future task latent, and a current representation disentangled from the corresponding domain latent. The ablation study therefore asks whether each ingredient is necessary, and whether simpler alternatives—more augmented data or generic future prediction—are sufficient.

All variants are evaluated on the LIBERO shortcut diagnostic above. *In-dist. SR* measures success on the original task–view training compositions. *Center OOD SR* evaluates each task group at an unseen midpoint viewpoint without swapping the task–view association, measuring interpolation to an unseen visual domain. *Counter OOD SR* evaluates the counterfactual task–view swaps, where the model must follow the commanded task rather than the task associated with the observed view. Finally, *shortcut degree* measures how often the model follows the task spuriously associated with the observed view; lower is better.

Compared methods. *Base VLA* is the plain behavior-cloning baseline. It uses the same VLA architecture as DILL, but directly maps the current observation-conditioned representation to actions, without source augmentation, task–domain latent targets, a disentangled current head, or latent lookahead alignment. *Base VLA + SA* keeps the same behavior-cloning objective and architecture, but trains with the augmented source data used by DILL for task–domain encoder pretraining. *Entangled LA* adds latent lookahead prediction, but uses the representation from the original video encoder [2] as the future target instead of task–domain latent targets; the predicted future latent is therefore not explicitly encouraged to discard domain-specific information. *DILL w/o CH* removes the disentangled current head while keeping task–domain latent targets and latent lookahead alignment. *DILL w/o LA* keeps the task–domain latent targets and the disentangled current head, but removes latent lookahead alignment. *DILL* is the full model, combining source augmentation, task–domain latent targets, the disentangled current head, and domain-invariant latent lookahead alignment.

Results. Table I shows that fitting the confounded training distribution is not sufficient for shortcut-free generalization. Base VLA reaches 68% in-distribution success, but obtains 0% Counter OOD success and a high shortcut degree of 73%. Adding source augmentation improves in-distribution success to 82% and slightly reduces shortcut degree to 69%, indicating that additional visual diversity and action pretraining are helpful for fitting the nominal tasks and can mildly reduce shortcut reliance. However, Counter OOD success remains only 2%, showing that data augmentation alone does not teach the policy which visual factors should be ignored when task–view correlations are counterfactually broken.

Entangled LA provides a second partial improvement. Compared to Base VLA + SA, it improves Center OOD success from 10% to 24% and further reduces shortcut degree from 69% to 65%, suggesting that future prediction can encourage representations that are somewhat more robust to unseen viewpoints. However, it still obtains 0% Counter OOD success. This failure is informative: predicting a future latent is not sufficient if the target representation remains entangled with viewpoint, background, or other domain-specific cues. The lookahead target must be tied to task-relevant structure rather than to the same visual correlations present in the training distribution.

The lower half of the table isolates the DILL-specific components. DILL w/o CH, which keeps task–domain latent targets and latent lookahead alignment but removes the disentangled current head, reaches 44% Center OOD and 44% Counter OOD success while reducing shortcut degree to 6%. This is a large jump over Entangled LA, indicating that aligning lookahead prediction with task–domain latent targets removes much of the view-induced task substitution. DILL w/o LA, which keeps the disentangled current head but removes latent lookahead alignment, achieves the strongest Counter OOD success (62%) and the lowest shortcut degree (3%), but its in-distribution success drops to 50%. This suggests that the disentangled current head is highly effective at suppressing shortcut behavior, while latent lookahead alignment helps recover a better balance between nominal task execution and counterfactual generalization.

The full DILL model provides the best overall tradeoff. It matches the best Center OOD success (62%), remains close to the best Counter OOD success (58%), keeps shortcut degree very low (5%), and achieves the highest average success across the three success metrics (62%). The key pattern is that source augmentation and generic lookahead improve some aspects of learning, but do not solve counterfactual task–view generalization. In contrast, the DILL-family variants that use task–domain latent targets sharply reduce shortcut degree, and the full model best preserves both task execution and shortcut resistance.

C. LIBERO-Plus Zero-Shot Robustness Evaluation Under Visual Distribution Shifts

We next evaluate whether shortcut mitigation translates into stronger robustness under visual distribution shifts. We use the visual perturbation categories in LIBERO-Plus [14] as a controlled zero-shot evaluation: no LIBERO-Plus perturbed images are used for training. All models are adapted on the original LIBERO training split.

Compared methods. To ensure a controlled comparison, the main table includes only methods evaluated with third-person RGB observations, excluding wrist-camera images. Under this matched-input protocol, we compare against OpenVLA [21], OpenVLA-OFT [22], WorldVLA [10], UniVLA [9], and NORA [19].

TABLE II

Zero-shot robustness evaluation on visual perturbations in LIBERO-Plus [14]. For each model, the first row shows success rate (%) on four visual perturbation categories in LIBERO-Plus. The second row shows the absolute drop relative to the original score.

Model	Original	Visual perturbations				Average
		Camera	Light	BG	Noise	
OpenVLA [21]	76.5	0.8	8.1	34.8	15.2	14.7
		↓75.7	↓68.4	↓41.7	↓61.3	↓61.8
WorldVLA [10]	79.1	0.1	43.7	17.1	10.9	18.0
		↓79.0	↓35.4	↓62.0	↓68.2	↓61.2
UniVLA [9]	95.5	1.8	69.0	81.0	21.2	43.3
		↓93.7	↓26.5	↓14.5	↓74.3	↓52.3
NORA [19]	87.9	2.2	45.7	58.6	12.8	29.8
		↓85.7	↓42.2	↓29.3	↓75.1	↓58.1
OpenVLA-OFT [22]	95.3	10.4	76.8	93.6	49.9	57.7
		↓84.9	↓18.5	↓1.7	↓45.4	↓37.6
DILL (Ours)	81.6	68.4	69.4	70.0	68.7	69.1
		↓13.2	↓12.2	↓11.6	↓12.9	↓12.5

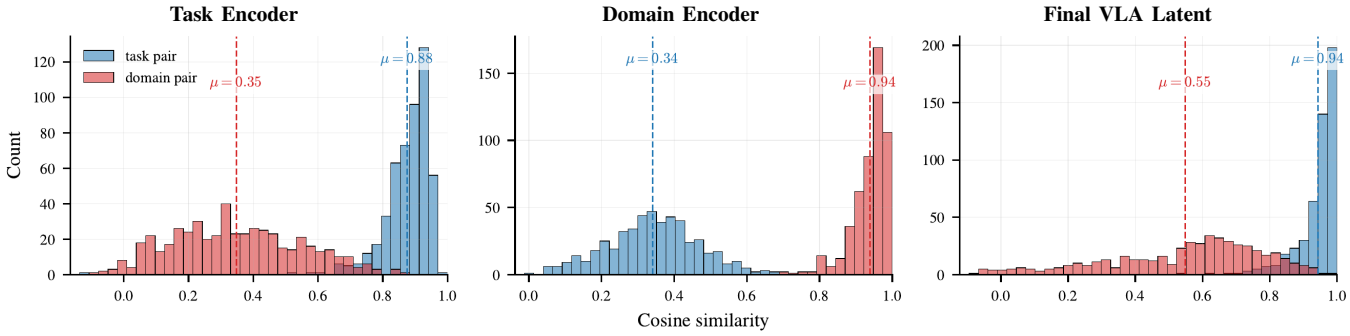

Fig. 5. **Pairwise similarity diagnostics for learned latents.** We compare cosine-similarity distributions for task pairs and domain pairs constructed from unseen task–domain combinations. From left to right, the panels show the task encoder, domain encoder, and final VLA latent. The task encoder should group task pairs, the domain encoder should group domain pairs, and the final VLA latent should preserve task-consistent structure while suppressing domain-specific variation.

Results. Table II reports the LIBERO-Plus results on visual perturbations. Our method achieves the highest average success across the four visual shift categories, improving over the strongest external baseline, OpenVLA-OFT, by 11.4 points. This gain is not explained by stronger in-distribution performance: OpenVLA-OFT and UniVLA obtain higher original LIBERO scores than our method, but suffer substantially larger drops under visual perturbations. In contrast, our method has an average drop of 12.5 points, compared to 37.6 for OpenVLA-OFT and 52.3 for UniVLA.

The improvement is especially pronounced for camera and sensor-noise perturbations, where policies must avoid relying on brittle view- or image-quality-specific cues. Our method achieves 68.4% success under camera perturbations and 68.7% under sensor noise, while all external baselines remain below 11% and 50% on these two categories, respectively. Although OpenVLA-OFT remains stronger on background and lighting perturbations, our method is substantially more balanced across visual shifts, suggesting that domain-invariant latent lookahead reduces sensitivity to spurious visual factors tied to the training distribution. The main visual-perturbation breakdown is summarized in Table II.

D. Pairwise Diagnostics of Learned Latents

Finally, we analyze whether the learned latent spaces separate task-relevant structure from domain-specific visual variation. We evaluate three representations: the task encoder output, the domain encoder output, and the final VLA latent provided to the action head. For the final VLA latent, we use the concatenated action-conditioning representation, *i.e.*, the current policy representation together with the predicted lookahead latent.

Protocol. We compute cosine similarity between normalized latents for two types of unseen task and domain pairs. *Task pairs* share the same underlying trajectory content but differ in visual domain or augmentation. *Domain pairs* share the same visual domain or augmentation pattern but differ in trajectory content. These pair combinations are not observed during training, so the diagnostic tests whether the learned representations generalize beyond memorized pairings. A task-centric representation should assign higher similarity to task pairs than to domain pairs, while a domain-centric representation should show the opposite behavior. Additional details are provided in Appendix C.

Results. Figure 5 shows that the learned factorization behaves as intended. The task encoder assigns high similarity to task pairs and substantially lower similarity to domain pairs, with mean similarities of 0.88 and 0.35, respectively.

Conversely, the domain encoder assigns high similarity to domain pairs and lower similarity to task pairs, with mean similarities of 0.94 and 0.34, indicating that the task and domain encoders capture complementary factors rather than collapsing to the same representation. The final VLA latent also remains strongly task-centric: it assigns high similarity to task pairs (μ = 0.94) while keeping domain pairs noticeably lower (μ = 0.55). Since this is the representation directly provided to the action head, the result suggests that the policy is conditioned on latent features that are stable across domain changes but still discriminative across different trajectory content. This supports the mechanism behind the robustness gains in Table II: domain-invariant latent lookahead encourages the policy to act on task-consistent structure rather than domain-specific visual cues.

Comparison to Base VLA. We further compare DILL against an architecture-matched Base VLA. Base VLA uses the same VLA architecture, observation and language inputs, and action head as our method, but removes domain-invariant latent lookahead and its associated representation objectives; it predicts actions directly from the current observation-conditioned representation using behavior cloning. We therefore compare the actual representation routed into the action head: the standard Base VLA latent versus DILL's concatenation of the current representation and predicted lookahead latent. In this comparison, task pairs share trajectory content under different visual domains, while domain pairs share visual domain but differ in trajectory content. Table III summarizes the result. Base VLA exhibits the undesirable ordering for shortcut-robust control: its domain-pair similarity (0.70) is higher than its task-pair similarity (0.52), yielding a negative task-centric gap of −0.18. DILL reverses this ordering: task-pair similarity rises to 0.94, domain-pair similarity drops to 0.55, and the gap increases to 0.39. This indicates that the action-conditioning latent becomes stable across domain changes while remaining discriminative across different trajectory content. The full distribution visualization is provided in Appendix C-B.

Representation	Task-pair sim. ↑	Domain-pair sim. ↓	Gap Δ_{task} ↑
Base VLA final latent	0.52	0.70	−0.18
DILL final latent	0.94	0.55	0.39

TABLE III
Base VLA vs. DILL final latent similarity. The gap is task-pair similarity minus domain-pair similarity.

V. Limitations

DILL is designed to reduce shortcut reliance induced by visual-domain factors. In our experiments, these factors include changes in viewpoint, scene appearance, camera configuration, and visual degradation. While this captures a common source of brittleness in VLA policies, it does not cover all possible spurious correlations. For example, biases in initial states, object layouts, language templates, task frequencies, or demonstrator styles may also influence action prediction without being part of the intended task semantics. Addressing such non-visual shortcuts would require defining additional nuisance factors, obtaining corresponding paired interventions, or developing objectives that can discover them more automatically.

Another important consideration is the trade-off between robustness and in-distribution performance. In the training distribution, domain-dependent cues can be highly predictive of demonstrated actions, even when they are not causally tied to the intended task semantics. Moreover, the same visual factor may play different roles across contexts: a background pattern may be a spurious cue in one setting, but part of the task-relevant scene configuration in another. Overly strong invariance may therefore suppress information that is useful for control in some situations, reducing peak in-distribution performance even as it improves robustness when those cues become misleading. This suggests that robustness should not be pursued by uniformly removing domain-specific information in all cases. Instead, future work should explore adaptive objectives that preserve action-relevant factors and suppress action-irrelevant shortcuts in a context-dependent manner.

VI. Conclusion

A central obstacle to robust VLA control is that policies can treat domain-specific appearance cues as action-relevant when they are only reliable within the training distribution. Domain-Invariant Latent Lookahead addresses this shortcut learning problem by supervising VLA representations with future latents in a space where domain-specific visual variation has been separated from task-relevant structure. By aligning lookahead prediction with this future task latent and regularizing the current representation against the domain latent, the policy is encouraged to condition actions on stable task progress rather than shortcut cues such as background, viewpoint, or lighting.

Acknowledgments

This work was supported by grants from the IITP (RS-2021-II211343-GSAI/10%, RS-2022-II220951-LBA/15%, RS-2022-II220953-PICA/15%), the NRF (RS-2024-00353991-SPARC/15%, RS-2023-00274280-HEI/15%), the KEIT (RS-2025-25453780/15%), and the KIAT (RS-2025-25460896/15%), funded by the Korean government. Part of this work was conducted while the first author was serving as a Research Fellow at OpenMind.

References

[1] Mahmoud Assran, Quentin Duval, Ishan Misra, Piotr Bojanowski, Pascal Vincent, Michael Rabbat, Yann LeCun, and Nicolas Ballas. Self-supervised learning from images with a joint-embedding predictive architecture. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15619–15629, 2023.

[2] Mido Assran, Adrien Bardes, David Fan, Quentin Garrido, Russell Howes, Matthew Muckley, Ammar Rizvi, Claire Roberts, Koustuv Sinha, Artem Zholus, et al. V-jepa 2: Self-supervised video models enable understanding, prediction and planning. *arXiv preprint arXiv:2506.09985*, 2025.

[3] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. Qwen2.5-vl technical report, 2025. URL https://arxiv.org/abs/2502.13923.

[4] Randall Balestriero and Yann LeCun. Lejepa: Provable and scalable self-supervised learning without the heuristics, 2025. URL https://arxiv.org/abs/2511.08544.

[5] Suneel Belkhale and Dorsa Sadigh. Minivla: A better vla with a smaller footprint, 2024. URL https://github.com/Stanford-ILIAD/openvla-mini.

[6] Kevin Black, Mitsuhiko Nakamoto, Pranav Atreya, Homer Walke, Chelsea Finn, Aviral Kumar, and Sergey Levine. Zero-shot robotic manipulation with pre-trained image-editing diffusion models. *arXiv preprint arXiv:2310.10639*, 2023.

[7] Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, et al. π_0: A vision-language-action flow model for general robot control. *arXiv preprint arXiv:2410.24164*, 2024.

[8] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. *arXiv preprint arXiv:2307.15818*, 2023.

[9] Qingwen Bu, Yanting Yang, Jisong Cai, Shenyuan Gao, Guanghui Ren, Maoqing Yao, Ping Luo, and Hongyang Li. UniVLA: Learning to act anywhere with task-centric latent actions, 2025. URL https://arxiv.org/abs/2505.06111.

[10] Jun Cen, Chaohui Yu, Hangjie Yuan, Yuming Jiang, Siteng Huang, Jiayan Guo, Xin Li, Yibing Song, Hao Luo, Fan Wang, Deli Zhao, and Hao Chen. WorldVLA: Towards autoregressive action world model, 2025. URL https://arxiv.org/abs/2506.21539.

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

[12] Open X-Embodiment Collaboration, Abby O'Neill, Abdul Rehman, Abhinav Gupta, Abhiram Maddukuri, Abhishek Gupta, Abhishek Padalkar, et al. Open X-Embodiment: Robotic learning datasets and RT-X models. https://arxiv.org/abs/2310.08864, 2023.

[13] Pim De Haan, Dinesh Jayaraman, and Sergey Levine. Causal confusion in imitation learning. *Advances in neural information processing systems*, 32, 2019.

[14] Senyu Fei, Siyin Wang, Junhao Shi, Zihao Dai, Jikun Cai, Pengfang Qian, Li Ji, Xinzhe He, Shiduo Zhang, Zhaoye Fei, et al. Libero-plus: In-depth robustness analysis of vision-language-action models. *arXiv preprint arXiv:2510.13626*, 2025.

[15] Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2 (11):665–673, 2020.

[16] Yanjiang Guo, Lucy Xiaoyang Shi, Jianyu Chen, and Chelsea Finn. Ctrl-world: A controllable generative world model for robot manipulation. *arXiv preprint arXiv:2510.10125*, 2025.

[17] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Liang Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *Iclr*, 1(2):3, 2022.

[18] Yucheng Hu, Yanjiang Guo, Pengchao Wang, Xiaoyu Chen, Yen-Jen Wang, Jianke Zhang, Koushil Sreenath, Chaochao Lu, and Jianyu Chen. Video prediction policy: A generalist robot policy with predictive visual representations. *arXiv preprint arXiv:2412.14803*, 2024.

[19] Chia-Yu Hung, Qi Sun, Pengfei Hong, Amir Zadeh, Chuan Li, U-Xuan Tan, Navonil Majumder, and Soujanya Poria. NORA: A small open-sourced generalist vision language action model for embodied tasks, 2025. URL https://arxiv.org/abs/2504.19854.

[20] Alexander Khazatsky, Karl Pertsch, Suraj Nair, Ashwin Balakrishna, Sudeep Dasari, Siddharth Karamcheti, Soroush Nasiriany, Mohan Kumar Srirama, Lawrence Yunliang Chen, Kirsty Ellis, et al. Droid: A large-scale in-the-wild robot manipulation dataset. *arXiv preprint arXiv:2403.12945*, 2024.

[21] Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, Quan Vuong, Thomas Kollar, Benjamin Burchfiel, Russ Tedrake, Dorsa Sadigh, Sergey Levine, Percy Liang, and Chelsea Finn. Openvla: An open-source vision-language-action model. *arXiv preprint arXiv:2406.09246*, 2024.

[22] Moo Jin Kim, Chelsea Finn, and Percy Liang. Fine-tuning vision-language-action models: Optimizing speed and success. *arXiv preprint arXiv:2502.19645*, 2025.

[23] Bo Liu, Yifeng Zhu, Chongkai Gao, Yihao Feng, Qiang Liu, Yuke Zhu, and Peter Stone. Libero: Benchmarking

knowledge transfer for lifelong robot learning. *Advances in Neural Information Processing Systems*, 36:44776–44791, 2023.

[24] Fanfan Liu, Feng Yan, Liming Zheng, Chengjian Feng, Yiyang Huang, and Lin Ma. Robouniview: Visual-language model with unified view representation for robotic manipulation. *arXiv preprint arXiv:2406.18977*, 2024.

[25] Ajay Mandlekar, Soroush Nasiriany, Bowen Wen, Iretiayo Akinola, Yashraj Narang, Linxi Fan, Yuke Zhu, and Dieter Fox. Mimicgen: A data generation system for scalable robot learning using human demonstrations. In *7th Annual Conference on Robot Learning*, 2023.

[26] Octo Model Team, Dibya Ghosh, Homer Walke, Karl Pertsch, Kevin Black, Oier Mees, Sudeep Dasari, Joey Hejna, Charles Xu, Jianlan Luo, Tobias Kreiman, You Liang Tan, Lawrence Yunliang Chen, Pannag Sanketi, Quan Vuong, Ted Xiao, Dorsa Sadigh, Chelsea Finn, and Sergey Levine. Octo: An open-source generalist robot policy. In *Proceedings of Robotics: Science and Systems*, Delft, Netherlands, 2024.

[27] Jing-Cheng Pang, Nan Tang, Kaiyuan Li, Yuting Tang, Xin-Qiang Cai, Zhen-Yu Zhang, Gang Niu, Masashi Sugiyama, and Yang Yu. Learning view-invariant world models for visual robotic manipulation. In *The Thirteenth International Conference on Learning Representations*, 2025.

[28] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. SDXL: Improving latent diffusion models for high-resolution image synthesis. In *International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=di52zR8xgf.

[29] Poly Haven. Poly Haven: A free 3D asset library. https://polyhaven.com/, 2024. Textures licensed under CC0 1.0 Universal Public Domain Dedication.

[30] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.

[31] Younggyo Seo, Junsu Kim, Stephen James, Kimin Lee, Jinwoo Shin, and Pieter Abbeel. Multi-view masked world models for visual robotic manipulation. In *International Conference on Machine Learning*, pages 30613–30632. PMLR, 2023.

[32] Stability AI. Stable Diffusion XL Base 1.0. https://huggingface.co/stabilityai/stable-diffusion-xl-base-1.0, 2023. Hugging Face model card.

[33] Jingwen Sun, Wenyao Zhang, Zekun Qi, Shaojie Ren, Zezhi Liu, Hanxin Zhu, Guangzhong Sun, Xin Jin, and Zhibo Chen. Vla-jepa: Enhancing vision-language-action model with latent world model. *arXiv preprint arXiv:2602.10098*, 2026.

[34] Stone Tao, Fanbo Xiang, Arth Shukla, Yuzhe Qin, Xander Hinrichsen, Xiaodi Yuan, Chen Bao, Xinsong Lin, Yulin Liu, Tse kai Chan, Yuan Gao, Xuanlin Li, Tongzhou Mu, Nan Xiao, Arnav Gurha, Viswesh Nagaswamy Rajesh, Yong Woo Choi, Yen-Ru Chen, Zhiao Huang, Roberto Calandra, Rui Chen, Shan Luo, and Hao Su. Maniskill3: Gpu parallelized robotics simulation and rendering for generalizable embodied ai. *Robotics: Science and Systems*, 2025.

[35] Stephen Tian, Blake Wulfe, Kyle Sargent, Katherine Liu, Sergey Zakharov, Vitor Guizilini, and Jiajun Wu. View-invariant policy learning via zero-shot novel view synthesis. *arXiv preprint arXiv:2409.03685*, 2024.

[36] Guodong Wang, Chenkai Zhang, Qingjie Liu, Jinjin Zhang, Jiancheng Cai, Junjie Liu, and Xinmin Liu. Libero-x: Robustness litmus for vision-language-action models. *arXiv preprint arXiv:2602.06556*, 2026.

[37] 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. *arXiv preprint arXiv:2312.13139*, 2023.

[38] Youguang Xing, Xu Luo, Junlin Xie, Lianli Gao, Hengtao Shen, and Jingkuan Song. Shortcut learning in generalist robot policies: The role of dataset diversity and fragmentation. *arXiv preprint arXiv:2508.06426*, 2025.

[39] Borong Zhang, Jiahao Li, Jiachen Shen, Yishuai Cai, Yuhao Zhang, Yuanpei Chen, Juntao Dai, Jiaming Ji, and Yaodong Yang. Vla-arena: An open-source framework for benchmarking vision-language-action models. *arXiv preprint arXiv:2512.22539*, 2025.

[40] Ji Zhang, Shihan Wu, Xu Luo, Hao Wu, Lianli Gao, Heng Tao Shen, and Jingkuan Song. Inspire: Vision-language-action models with intrinsic spatial reasoning. *arXiv preprint arXiv:2505.13888*, 2025.

[41] Han Zhao, Jingbo Wang, Wenxuan Song, Shuai Chen, Yang Liu, Yan Wang, Haoang Li, and Donglin Wang. Frappe: Infusing world modeling into generalist policies via multiple future representation alignment. *arXiv preprint arXiv:2602.17259*, 2026.

[42] Qingqing Zhao, Yao Lu, Moo Jin Kim, Zipeng Fu, Zhuoyang Zhang, Yecheng Wu, Zhaoshuo Li, Qianli Ma, Song Han, Chelsea Finn, et al. Cot-vla: Visual chain-of-thought reasoning for vision-language-action models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 1702–1713, 2025.

[43] Ruijie Zheng, Jing Wang, Scott Reed, Johan Bjorck, Yu Fang, Fengyuan Hu, Joel Jang, Kaushil Kundalia, Zongyu Lin, Loic Magne, et al. Flare: Robot learning with implicit world modeling. *arXiv preprint arXiv:2505.15659*, 2025.

[44] Xueyang Zhou, Yangming Xu, Guiyao Tie, Yongchao Chen, Guowen Zhang, Duanfeng Chu, Pan Zhou, and Lichao Sun. Libero-pro: Towards robust and fair evaluation of vision-language-action models beyond memorization. *arXiv preprint arXiv:2510.03827*, 2025.

Appendix Overview

This appendix provides additional details, analyses, and supplementary experiments for DILL. Appendix A describes the method implementation and training objectives. Appendix B gives simulation benchmark details beyond the main text, including the full LIBERO-Plus breakdown used only as supplementary context. Appendix C provides additional representation diagnostics.

Contents.

- Appendix A: Method details and training objectives.
- Appendix B: Simulation benchmark details.
- Appendix C: Representation diagnostics.

Appendix A
Additional Method Details

A. Domain Transformations and Positive Pair Mining

This subsection supports Sec. III-B with source-data statistics and implementation details for the domain interventions and positive-pair mining used to train the Task-Domain Encoder. The main text describes the conceptual role of task-positive and domain-positive pairs; here we record the exact data mixture, transformation families, and batching procedure.

a) Data sources: Task-domain encoder pretraining uses a source trajectory collection composed of MimicGen [25], ManiSkill [34], and subsets of OXE [12]. The OXE subset includes Bridge and Fractal/RT-1 trajectories. Table IV summarizes the pretraining data. In our main experiments, the Task-Domain Encoder is pretrained once on this source trajectory collection and then kept fixed during VLA policy pretraining and downstream policy learning. This design decouples Task-Domain Encoder pretraining from downstream policy learning: downstream users can train only the policy-side modules unless they choose to further adapt the Task-Domain Encoder with additional domain transformations.

TABLE IV
Source trajectory collection used for Task-Domain Encoder pretraining. The MimicGen row aggregates 16 task datasets, while OXE aggregates Bridge and Fractal/RT-1 subsets.

Source	Frames	Trajectories	Tasks
MimicGen, 16 datasets	1,010,618	3,200	16
ManiSkill, 2 datasets	1,632,291	14,057	17
OXE, Bridge + Fractal/RT-1	5,785,810	140,404	20,553
Total	8,428,719	157,661	20,586

For simulation data, controllable rendering enables viewpoint changes and, when segmentation masks are available, mask-based background replacement. We additionally apply image-level transformations directly to frames; these transformations are used for both simulation data and recorded trajectories such as the OXE subsets.

b) Transformation families: We group domain transformations into four families that reflect common visual domain shifts: viewpoint changes, environment appearance variation, camera heterogeneity, and visual degradation. A sampled domain condition can include a single transformation or a combination of transformations from these families. Table V summarizes the transformation types and the shared domain condition used for domain-positive pairs. The main text visualizes representative single-family examples; Table V records the broader transformation set sampled during training.

TABLE V
Domain transformation families used for positive pair mining. A domain condition η consists of a transformation family and its sampled parameters or random seed. The last column lists the condition shared by domain-positive pairs.

Family	Transformations	Shared condition for domain positives
Viewpoint changes	Rendered camera view, random crop, intrinsics change, radial distortion, warping	Camera ID or camera parameters; crop seed; focal scale; distortion and warp parameters
Environment appearance	Background replacement, lighting changes, color changes	Background texture IDs; mask regions; replacement mode; lighting or color parameters
Camera heterogeneity	Camera unprocessing, sensor noise, camera color response changes	Camera-pipeline seed; color matrices; gamma; ISO level; shot/read noise; channel gains
Visual degradation	Blur, weather effects, compression, pixelation, image corruptions	Corruption type, severity, and random seed

c) Background replacement: For ManiSkill trajectories [34] with available segmentation masks, we perform mask-based background replacement for table, floor, and wall regions. Table and floor textures are sampled from Poly Haven [29], while wall and background assets are generated with Stable Diffusion XL using text-to-image prompts for robot-relevant environments such as factories, kitchens, laboratories, workspaces, loading docks, and server rooms [28, 32]. The resulting texture pool contains 92 table textures, 163 floor textures, and 604 wall/background textures, for a total of 859 region-specific assets. At training time, textures are sampled by region and composited into the corresponding segmentation masks. For domain-positive pairs, the same texture IDs, mask regions, and replacement mode define the shared domain condition. Representative samples from the texture pool are shown in Fig. 6.

d) Pair mining details: For the task encoder, each task-positive pair is constructed from the same episode and the same temporal window. Each training item samples two non-overlapping windows from the same episode, denoted by $\mathbf{o}^i_{t_a}$ and $\mathbf{o}^i_{t_b}$, and samples two domain conditions η_1 and η_2. We then form two task-positive pairs:

$$\big(\mathcal{T}_{\eta_1}(\mathbf{o}^i_{t_a}), \mathcal{T}_{\eta_2}(\mathbf{o}^i_{t_a})\big), \qquad \big(\mathcal{T}_{\eta_1}(\mathbf{o}^i_{t_b}), \mathcal{T}_{\eta_2}(\mathbf{o}^i_{t_b})\big). \tag{13}$$

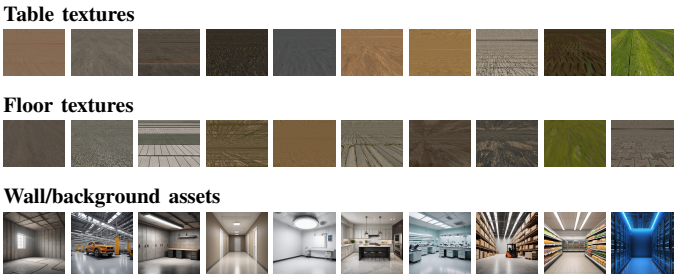

Fig. 6. Representative assets for mask-based background replacement. We sample table and floor textures from Poly Haven and generate wall and background assets with Stable Diffusion XL.

Within each pair, the two clips share the same underlying temporal window but differ in domain condition. Across the two pairs, the temporal windows are distinct while the task identity and domain conditions are shared. After collation, the two task-positive pairs per item are flattened from $[B, 2, \cdots]$ to $[2B, \cdots]$ before applying InfoNCE. Thus, examples from the same episode but different temporal windows are included as ordinary in-batch negatives. This discourages the task encoder from solving the contrastive objective using task identity alone, while avoiding false negatives from the same underlying temporal segment.

For the domain encoder, a domain-positive pair is constructed by applying the same domain condition to two different clips. Each training item first samples a source pair $(\mathbf{o}_a, \mathbf{o}_b)$, either from two non-overlapping windows of the same episode or from cross-task clips within the same camera bucket. It then samples four distinct domain-transformation types without replacement, with corresponding domain conditions $\eta_1, \ldots, \eta_4$, and forms four domain-positive pairs:

$$\big(\mathcal{T}_{\eta_m}(\mathbf{o}_a), \mathcal{T}_{\eta_m}(\mathbf{o}_b)\big), \qquad m = 1, \ldots, 4. \tag{14}$$

Within each pair, the clips differ in trajectory content but share the same domain condition. Across the four pairs in the same item, the source clips are fixed while the domain-transformation type changes. After collation, the four domain-positive pairs per item are flattened from $[B, 4, \cdots]$ to $[4B, \cdots]$ before applying InfoNCE. Therefore, the domain encoder receives controlled in-batch comparisons: for an anchor under domain condition η_m, other clips from the same item but with $\eta_{m'} \neq \eta_m$ appear among the standard in-batch negatives. No separate hard-negative loss is used; the dataloader simply ensures that informative non-matching domain conditions are present in the minibatch.

All transformations are applied consistently across the frames of a chunk so that each video chunk corresponds to a coherent visual domain.

B. Gaussian Disentanglement Objective

The main text defines the task-positive and domain-positive contrastive losses. This appendix gives the implementation of the Gaussian disentanglement regularizer used to reduce leakage between the two latent branches. The task and domain encoders are trained with different positive pairs, but without an additional constraint, the resulting task and domain latents can still encode overlapping information. We therefore regularize concatenated latents toward an isotropic Gaussian distribution.

For a minibatch, we construct

$$u_i^{\text{td}} = [z_i^{\text{task}}; z_i^{\text{dom}}], \qquad i = 1, \ldots, B, \tag{15}$$

where $[\cdot;\cdot]$ denotes concatenation. Here $z_i^{\text{task}}, z_i^{\text{dom}} \in \mathbb{R}^{D_z}$, and therefore $u_i^{\text{td}} \in \mathbb{R}^D$ with $D = 2D_z$. We implement the disentanglement regularizer $\mathcal{R}_{\text{dis}}$ using SIGReg [4]. Given a batch of vectors $U = \{u_i\}_{i=1}^B \subset \mathbb{R}^D$, we sample M random unit directions $r_m \sim \text{Unif}(\mathbb{S}^{D-1})$, project $s_{i,m} = r_m^\top u_i$, and match each projected distribution to a standard Gaussian. Using a Gaussian kernel with bandwidth σ, let $\text{EP}_{\sigma,m} = \text{EP}_\sigma(\{s_{i,m}\}_{i=1}^B)$ for the projected samples along direction r_m. We compute

$$\begin{aligned}\text{EP}_{\sigma,m} = \frac{1}{B^2}\sum_{i=1}^{B}\sum_{j=1}^{B}\exp\left(-\frac{(s_{i,m}-s_{j,m})^2}{2\sigma^2}\right) \\ - \frac{2\sigma}{B\sqrt{\sigma^2+1}}\sum_{i=1}^{B}\exp\left(-\frac{s_{i,m}^2}{2(\sigma^2+1)}\right) + \frac{\sigma}{\sqrt{\sigma^2+2}}.\end{aligned} \tag{16}$$

The regularizer averages this statistic over random projection directions:

$$\mathcal{R}_{\text{dis}}(U) = \frac{1}{M}\sum_{m=1}^{M}\text{EP}_\sigma\left(\{r_m^\top u_i\}_{i=1}^B\right). \tag{17}$$

The task-domain disentanglement loss is

$$\mathcal{L}_{\text{dis}}^{\text{td}} = \mathcal{R}_{\text{dis}}\left(\{u_i^{\text{td}}\}_{i=1}^B\right). \tag{18}$$

a) Why Gaussian matching encourages disentanglement: The Gaussian disentanglement loss does not guarantee exact independence for arbitrary distributions. Its motivation is clearest under a joint Gaussian approximation. Let

$$u = \begin{bmatrix} z^a \\ z^b \end{bmatrix}$$

be a jointly Gaussian random vector with covariance

$$\Sigma = \begin{bmatrix} \Sigma_a & C \\ C^\top & \Sigma_b \end{bmatrix}. \tag{19}$$

If the concatenated latent is isotropic Gaussian, then $\Sigma = I$, which implies

$$\Sigma_a = I, \qquad \Sigma_b = I, \qquad C = 0. \tag{20}$$

For jointly Gaussian variables, zero cross-covariance is sufficient for independence, so

$$p(z^a, z^b) = p(z^a)p(z^b). \tag{21}$$

Thus, driving the concatenated latent $[z^a, z^b]$ toward an isotropic Gaussian encourages the two components to be disentangled.

We use this argument twice. For task-domain latent disentanglement, $z^a = z^{\text{task}}$ and $z^b = z^{\text{dom}}$. For policy learning,

$z^a = z^{\text{curr}}$ and $z^b = z^{\text{dom},\star}$. In practice, this should be interpreted as an approximate disentanglement regularizer rather than a guarantee of exact independence.

C. Architecture and Implementation Details

a) Task and domain encoders: In the main text, E_ψ^{task} and E_ξ^{dom} denote the full task and domain encoder branches that output pooled latent vectors. Each branch consists of a VJEPA2 video backbone [2], branch-specific LoRA adapters [17], and an attention-pooling projection head.

b) Multi-query attention pooling: Let $H \in \mathbb{R}^{N \times D_v}$ be a sequence of input tokens and let $Q \in \mathbb{R}^{R \times D_q}$ be R learned query vectors. We project queries, keys, and values using $W_q \in \mathbb{R}^{D_q \times d}$, $W_k \in \mathbb{R}^{D_v \times d}$, and $W_v \in \mathbb{R}^{D_v \times D_p}$. A multi-query attention pooler computes

$$A = \text{softmax}\left(\frac{(QW_q)(HW_k)^\top}{\sqrt{d}}\right), \tag{22}$$

$$U = AHW_v \in \mathbb{R}^{R \times D_p}. \tag{23}$$

The pooled tokens U are flattened and further projected to the output latent dimension D_z:

$$\begin{aligned} \text{Head}(H) &= W_o \,\text{vec}(U) + b_o, \\ W_o &\in \mathbb{R}^{D_z \times (RD_p)}, \\ b_o &\in \mathbb{R}^{D_z}. \end{aligned} \tag{24}$$

The projection heads for the task encoder, domain encoder, Lookahead Predictor, and Current Representation Head all use this attention-pooling-plus-linear architecture, with separate parameters.

c) Policy-side heads: The policy uses a Qwen2.5-VL [3] backbone with LoRA adaptation. On top of the VLM tokens, we train two AttentiveLatentHead modules and one ResNetActionHead. The Lookahead Predictor maps VLM tokens to $\hat{z}_t^{\text{look}}$, and the Current Representation Head maps the same tokens to z_t^{curr}. Each latent head uses 8 learned query tokens, a two-layer attentive pooler with 16 attention heads, and a linear projection to a 4096-dimensional latent.

d) Action head: The action head is an MLP-ResNet that maps the concatenated policy latents to an action chunk. In the dual-head setting, the action head input dimension is $4096 + 4096 = 8192$. We use hidden dimension 2048, two residual MLP blocks, and output a chunk size of $K_a = 50$.

TABLE VI
POLICY-SIDE HEAD BUDGET FOR THE VLA IMPLEMENTATION.

Module	Parameters
Lookahead Predictor, AttentiveLatentHead	167.85M
Current Representation Head, AttentiveLatentHead	167.85M
ResNetActionHead, dual-head input	28.48M
Total policy-side heads	364.18M

D. Training Stage Summary

Table VII summarizes which modules are updated in each training stage. In the main experiments, the Task-Domain Encoder is pretrained once on the source trajectory collection and then kept fixed. This decouples reusable Task-Domain Encoder training from downstream VLA policy training.

TABLE VII
TRAINING STAGE SUMMARY FOR DOMAIN-INVARIANT LATENT LOOKAHEAD.

Stage	Data	Updated modules	Objective
Task-domain encoder pretraining	Source trajectory collection	Task/domain encoder LoRA and projection heads	$\mathcal{L}_{\text{TDD}}$
VLA policy pretraining	Source trajectory collection	VLM LoRA, lookahead/current heads, action head	$\mathcal{L}_{\text{policy}}$
Downstream policy tuning	Downstream action-labeled demos (e.g., LIBERO)	VLM LoRA, lookahead/current heads, action head	$\mathcal{L}_{\text{policy}}$

a) Optional Task-Domain Encoder adaptation: Although the main experiments keep the pretrained Task-Domain Encoder fixed for downstream policy tuning, the same task-domain disentanglement objective could be used to adapt the task and domain encoders on additional downstream demonstration datasets. When controllable rendering or segmentation masks are unavailable, such optional adaptation would rely on image-level transformations such as cropping, warping, lighting and color changes, camera-pipeline perturbations, and visual corruptions.

APPENDIX B
SIMULATION BENCHMARK EXPERIMENTS

A. LIBERO Shortcut Diagnostic

We use the LIBERO shortcut diagnostic of [38] to test whether a policy follows the commanded task or instead executes the task spuriously associated with the observed view. Unlike standard robustness evaluation, this diagnostic explicitly separates shortcut-driven task substitution from general execution failure.

a) Task–view confounding: The benchmark constructs two confounded task–view islands. We refer to the left-view task group as Task-L and the right-view task group as Task-R. Task-L contains LIBERO task IDs $\{0, 1, 3, 5, 8\}$ and is observed only in the left-view range (10°–25°) during training. Task-R contains task IDs $\{2, 4, 6, 7, 9\}$ and is observed only in the right-view range (55°–70°). At test time, we evaluate counterfactual task–view compositions by swapping these associations: Task-R is evaluated at the left viewpoint (10°), and Task-L is evaluated at the right viewpoint (70°). A task-faithful policy should follow the language instruction under these swapped views, whereas a shortcut-prone policy will execute the task spuriously associated with the observed view.

Model	Orig.	LIBERO-Plus perturbations							Total
		Camera	Robot	Lang.	Light	BG	Noise	Layout	
OpenVLA [21]	76.5	0.8 ↓75.7	3.5 ↓73.0	23.0 ↓53.5	8.1 ↓68.4	34.8 ↓41.7	15.2 ↓61.3	28.5 ↓48.0	15.6 ↓60.9
WorldVLA [10]	79.1	0.1 ↓79.0	27.9 ↓51.2	41.6 ↓37.5	43.7 ↓35.4	17.1 ↓62.0	10.9 ↓68.2	38.0 ↓41.1	25.0 ↓54.1
UniVLA [9]	95.5	1.8 ↓93.7	46.2 ↓49.3	69.6 ↓25.9	69.0 ↓26.5	81.0 ↓14.5	21.2 ↓74.3	31.9 ↓63.6	43.9 ↓51.6
NORA [19]	87.9	2.2 ↓85.7	37.0 ↓50.9	65.1 ↓22.8	45.7 ↓42.2	58.6 ↓29.3	12.8 ↓75.1	62.1 ↓25.8	39.0 ↓48.9
OpenVLA-OFT [22]	95.3	10.4 ↓84.9	38.7 ↓56.6	70.5 ↓24.8	76.8 ↓18.5	93.6 ↓1.7	49.9 ↓45.4	69.9 ↓25.4	55.8 ↓39.5
DILL (Ours)	81.6	68.4 ↓13.2	21.4 ↓60.2	2.9 ↓78.7	69.4 ↓12.2	70.0 ↓11.6	68.7 ↓12.9	55.1 ↓26.5	50.8 ↓30.8

TABLE VIII

Supplementary full LIBERO-Plus breakdown. Rows report success rate (%) and absolute drop from the original LIBERO score. The main text focuses on visual perturbations because they match the visual-shortcut mechanism targeted by DILL.

b) Metrics: We report two complementary metrics. *OOD success rate* measures whether the commanded task is completed under the counterfactual view. *Shortcut degree* measures whether the policy instead executes the task spuriously associated with the observed view during training, thereby isolating view-induced task substitution.

B. LIBERO-Plus Evaluation Details

The main text reports the visual perturbation categories because they are most directly aligned with our claim about visual shortcut mitigation. Table VIII provides the full seven-category LIBERO-Plus breakdown as supplementary context. The non-visual categories should be interpreted separately from the visual robustness results: DILL targets domain-specific visual shortcuts, while language and robot-state robustness are separate failure modes.

a) Scope of the full breakdown: DILL is strongest on perturbations that directly change visual appearance or viewpoint, including Camera, Light, BG, and Noise. The low Language score reflects instruction-generalization failures, and the Robot score reflects sensitivity to robot-state variation. These axes require different invariances than the visual-domain invariance targeted by DILL, so the full table is included for completeness rather than as the central comparison.

Appendix C
Representation Diagnostics

A. Pairwise latent similarity protocol

We use pairwise latent similarity diagnostics to test whether the action-conditioning representation is organized around task content rather than domain-specific visual factors. For each representation, we compute cosine similarity between normalized latent vectors for two types of unseen task–domain pairs.

Task pairs. A task pair consists of two observations that share the same underlying trajectory content but differ in visual domain condition. A task-centric latent should assign high similarity to these pairs because the task-relevant trajectory content is preserved across the domain change.

Domain pairs. A domain pair consists of two observations that share the same visual domain condition but differ in trajectory content. A task-centric latent should assign lower similarity to these pairs because the underlying behavior is different. Conversely, a domain-centric latent would assign high similarity to domain pairs, even when the trajectory content changes.

Metric. For a task-centric action representation, we summarize the diagnostic using the task–domain separation gap,

$$\begin{aligned}\Delta_{\text{task}} = \mathbb{E}_{(i,j)\in\mathcal{P}_{\text{task}}}[\text{sim}(z_i, z_j)] \\ - \mathbb{E}_{(i,j)\in\mathcal{P}_{\text{dom}}}[\text{sim}(z_i, z_j)].\end{aligned} \tag{25}$$

where $\mathcal{P}_{\text{task}}$ denotes task pairs and $\mathcal{P}_{\text{dom}}$ denotes domain pairs. A larger positive gap indicates that the representation is more aligned with task content and less dominated by shared visual domain.

The main text reports the DILL factorization diagnostics and summarizes the direct Base VLA vs. DILL final-latent comparison in prose. This appendix provides the pair construction, metric, full visualization, and compact numerical summary behind those diagnostics.

B. Base VLA vs. DILL Final Latent Visualization

Figure 8 visualizes the final action-conditioning latent comparison. For Base VLA, the final latent assigns higher similarity to domain pairs (0.70) than to task pairs (0.52), indicating organization around shared visual appearance. DILL reverses this ordering: task-pair similarity rises to 0.94 while domain-pair similarity is lower at 0.55. This supports the interpretation that DILL changes the representation used for action prediction rather than only changing downstream success rates.

Task pair: same trajectory, different domain **Domain pair:** same domain, different trajectory

Fig. 7. **Pair construction for latent similarity diagnostics.** Task pairs test whether a latent remains stable across visual-domain changes when the underlying trajectory content is preserved. Domain pairs test whether a latent collapses observations that share visual-domain cues despite different trajectory content.

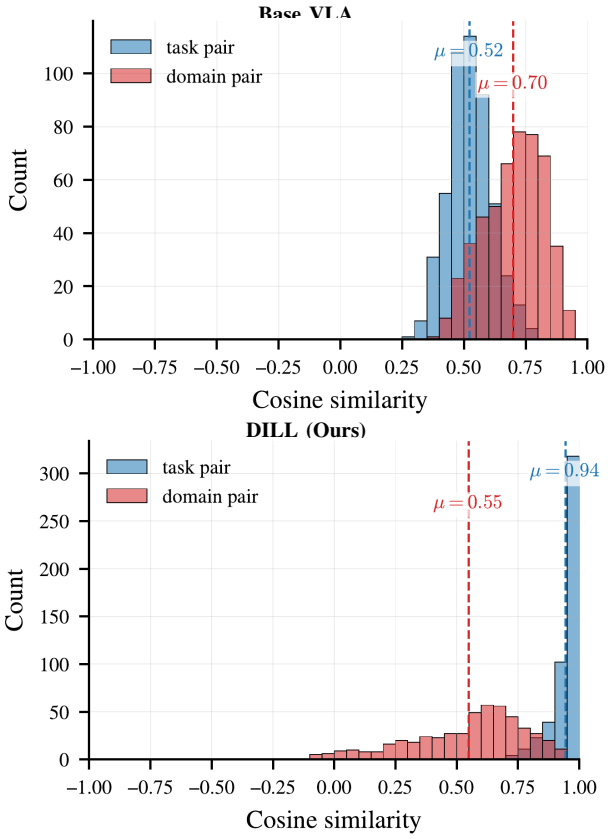

Fig. 8. **Base VLA vs. DILL final action-conditioning latent similarity.** Task pairs share trajectory content under different visual domains, while domain pairs share visual domain but differ in trajectory content.