

Describe-Then-Act: Proactive Agent Steering via Distilled Language-Action World Models

Massimiliano Pappa^{1,2*}, Luca Romani^{1*}, Valentino Sacco^{1*},
Alessio Palma¹, Stéphane Lathuilière², Fabio Galasso¹,
Xavier Alameda-Pineda², Indro Spinelli¹

¹ Sapienza University of Rome, Italy

² Inria, Univ. Grenoble Alpes, CNRS, LJK

Abstract. Deploying safety-critical agents requires anticipating the consequences of actions before they are executed. While world models offer a paradigm for this proactive foresight, current approaches relying on visual simulation incur prohibitive latencies, often exceeding several seconds per step. In this work, we challenge the assumption that visual processing is necessary for failure prevention. We show that a trained policy’s latent state, combined with its planned actions, already encodes sufficient information to anticipate action outcomes, making visual simulation redundant for failure prevention. To this end, we introduce **DILLO** (**DI**sti**LL**ed Language-Acti**ON** World Model), a fast steering layer that shifts the paradigm from “simulate-then-act” to “describe-then-act.” DILLO is trained via cross-modal distillation, where a privileged Vision Language Model teacher annotates offline trajectories and a latent-conditioned Large Language Model student learns to predict semantic outcomes. This creates a text-only inference path, bypassing heavy visual generation entirely, achieving a $14\times$ speedup over baselines. Experiments on MetaWorld and LIBERO demonstrate that DILLO produces high-fidelity descriptions of the next state and is able to steer the policy, improving episode success rate by up to 15 pp and 9.3 pp on average across tasks. Code is available at github.com/MaxPappa/DILLO.

1 Introduction

AI-driven agents are increasingly deployed in settings that demand high reliability, from robotic manipulation to autonomous navigation. However, a critical gap remains: these agents typically operate as “black boxes”, selecting and executing actions without explicitly anticipating the consequences. In classical control theory, the standard solution is Model Predictive Control (MPC) [6]: simulate the system forward and commit only to actions whose predicted outcomes are acceptable. The same principle underlies model-based reinforcement learning [9, 26], where world models forecast future states to steer policies.

However, deploying world models in the tight control loops of dynamic agents presents a stark dilemma between foresight and latency. Current approaches generally fall into two categories. The first, post-hoc analysis [2, 21, 25], operates

* Equal contribution




Fig. 1: (*left*) Standard approaches rely on heavy visual world models to simulate future outcomes using expensive rendering. In contrast, (*right*) DILLO avoids this visual dependency entirely. By distilling the foresight of a simulator-privileged Teacher VLM into a fast Language-Action World Model, DILLO predicts future outcome descriptions and success verdicts directly from policy latents (z_t) and action chunks ($a_{t:t+k}$). This shift from “simulate-then-act” to “describe-then-act” enables fast proactive steering on consumer hardware.

as a reactive wrapper, diagnosing failures only after they occur. While computationally inexpensive, this approach is fundamentally incapable of prevention. The second paradigm, proactive visual simulation, utilizes heavyweight world models to generate high-dimensional future states (images/videos or their latent representation) [13, 16]. While effective at anticipating risks, these “visual” world models incur prohibitive computational overhead. For instance, recent steering methods report inference latencies of nearly 4 seconds per decision [31] on an enterprise RTX A6000 with 48GB of VRAM, rendering them impractical for real-time control on an embodied agent. We posit that for real-time near-future failure prevention, simulating the visual world is redundant. A policy’s internal latent representation is explicitly trained to retain task-critical features: object geometry, relative distances, and contact dynamics. If this representation already contains the information needed to predict whether an action will succeed, then why pay the cost of re-rendering pixels? We call this the *Latent Sufficiency Hypothesis*.

To empirically verify this hypothesis, we propose **DILLO**(DIstiLLed Language-ActiOn World Model), see Fig. 1, a real-time reliability layer that shifts the paradigm from “simulate-then-act” to “describe-then-act.” DILLO is trained via cross-modal knowledge distillation: a privileged Vision-Language-Model (VLM) Teacher annotates offline trajectories with semantic outcomes, having access to the simulation environment. The Student (DILLO) then learns to predict these outcomes directly from the policy’s compact latent state and candidate action chunk, without accessing raw visual observations. Architecturally, we leverage the reasoning capacity of Small Language Models, mapping projected latents directly to the language decoder. This vision-free design enables DILLO to perform fast inference, effectively decoding the agent’s internal “state of mind” into human-readable foresight.

Specifically, DILLO generates two outputs: **(i) a natural language behavior preview** (d) that describes the expected physical interaction (*e.g.*, “*The gripper goes left and forward, starting to approach the ball*”), and **(ii) a binary verdict** ($c \in \{\text{POSITIVE}, \text{NEGATIVE}\}$). A *Negative* action is one causing stagnation, erratic movement, or progression toward failure; a *Positive* action advances the task. This enables a dual-purpose mechanism: the verdict allows the agent to autonomously reject negative actions via rejection sampling, while the description provides interpretable information for human supervisors.

Our main contributions are:

- We propose the *Latent Sufficiency Hypothesis* and demonstrate empirically that a policy’s latent state coupled with the sampled actions is a sufficient statistic for predicting failure-critical outcomes, making visual simulation redundant for proactive failure prevention.
- We introduce DILLO, a Distilled Language-Action World Model that distills a VLM teacher into a latent-conditioned LLM Student, shifting the paradigm from “simulate-then-act” to “describe-then-act.” This yields a $\sim 14\times$ speedup over visual baselines [31], enabling a full correction loop in **0.26 s** on consumer hardware.
- DILLO’s latent-based descriptions match or exceed the fidelity of vision-based baselines, including an oracle with access to future observations. Furthermore, we demonstrate effective steering on single- and multi-task policies across MetaWorld and LIBERO, improving episode success rate by up to **15 pp** and **9.3 pp** on average across tasks, achieving a verdict classification accuracy of 91.4% without access to any visual observation at inference time.

2 Related Work

We propose a real-time, language-based failure detection layer within the context of three dominant paradigms: reactive failure analysis, proactive world modeling, and action-language grounding.

Reactive Post-Hoc Analysis. The most common approach to agent interpretability follows an “act-then-describe” paradigm, analyzing failures only *after* they have occurred. Methods such as Aha [2] and REFLECT [21] leverage Vision-Language Models (VLMs) to process offline trajectory data, generating human-readable summaries of past errors. Other works treat VLMs as behavioral critics [8] or as a means to ground historical robotic experiences within a linguistic framework [28]. While runtime monitors [1, 25] operate closer to the point of execution, they are primarily designed to detect anomalies or consistency violations post-occurrence. Similarly, SAFE [7] analyzes the latent states of Vision-Language-Action models (VLAs) to predict task success or failure in a task-agnostic manner. Although these methods are indispensable for data curation and debugging, they remain fundamentally reactive, offering a “failure autopsy” rather than a preemptive intervention. In physical environments where robotic actions are often irreversible, such retrospective detection cannot substitute for robust, real-time prevention.

Proactive World Models for Control. To prevent failures, an agent must anticipate them. The concept of *World Models* champions this proactive paradigm, evolving from early visual foresight [3–5] to powerful latent dynamics models like Dreamer [11–13]. While these models effectively hallucinate future rewards for policy learning, they typically output latent vectors that are opaque to human supervisors. Recently, this paradigm was extended to *semantic* policy steering by Forewarn [31], which utilizes VLMs to narrate and evaluate simulated latent states. While validating the utility of linguistic previews, such “VLM-in-the-loop” methods remain computationally prohibitive. Relying on visual encoders to process predicted latent vectors and their subsequent evaluation incurs high latency (e.g., ~ 3.7 seconds per decision [31]), rendering them impractical for the tight control loops of dynamic agents. DILLO takes a different approach: rather than performing expensive, future-state simulation, we distill semantic foresight into a fast, latent-conditioned predictor enabling real-time safety without the associated simulation and evaluation overhead.

Action-Language Grounding. Finally, we distinguish our work from rationale-driven Explainable AI (XAI) [22, 24, 30]. Methods like Embodied Chain-of-Thought (ECoT) [34] use language to expose an agent’s *internal logic* (e.g., “I need the cup, so I will move my arm”). In contrast, DILLO is outcome-driven; it does not explain the agent’s reasoning but rather predicts the external physical consequence of an action (e.g., “This action will cause the gripper to move backward, increasing its distance from the cup”). This distinction allows DILLO to function as a controller-agnostic safety module that can be layered onto any policy to provide a fast, interpretable sanity check.

3 Distilled Language Action World Model

We introduce **DILLO** (**DI**stiLLed Language-Acti**On** World Model), a framework for real-time failure prevention that decouples semantic anticipation from visual simulation. We formulate the problem as cross-modal knowledge distillation, where a lightweight student learns to approximate the semantic foresight of a privileged teacher solely from the policy’s internal representations and predicted actions, see Fig. 2.

3.1 Semantic Anticipation

Consider a robotic agent operating in a Partially Observable Markov Decision Process (POMDP). At decision step t , the agent receives an observation $o_t \in \mathcal{O}$ and processes it via the policy’s encoder $\mathcal{E}_\pi$ to extract a compact latent state:

$$z_t = \mathcal{E}_\pi(o_t) \in \mathbb{R}^{D_z} \quad (1)$$

Conditioned on this latent state, the policy proposes a sequence of k future actions, the action chunk [20] $a_{t:t+k}$. Our objective is to learn a **Language-Action World Model** modeling a distribution $P_\theta(y_{t+k} \mid z_t, a_{t:t+k})$ that predicts the semantic outcome y_{t+k} directly from these internal representations,

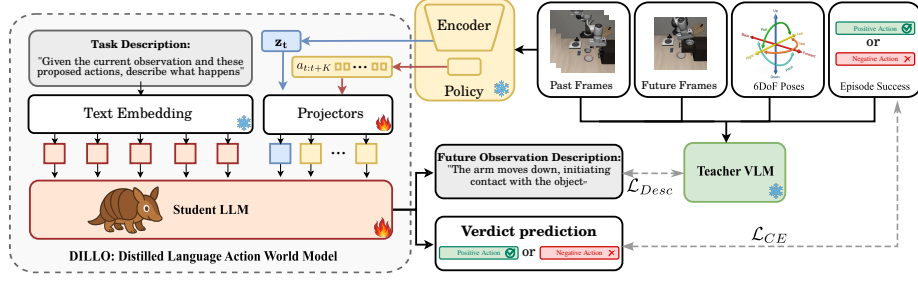


Fig. 2: DILLO Training Pipeline. We distill the foresight of a VLM teacher. The pipeline bypasses visual processing by projecting the fixed policy’s internal latents (z_t) and action chunks ($a_{t:t+k}$) directly into the LLM’s embedding space.

without accessing the raw visual observation. Standard world models use the state representation, often an RGB image, to predict high-dimensional future observations $\hat{o}_{t+k}$ (pixel space).

The semantic outcome y_{t+k} , consists of a natural language description d_{t+k} (e.g., physical dynamics, contacts) and a verdict $c_{t+k} \in \{\text{POSITIVE}, \text{NEGATIVE}\}$ indicating whether the action chunk advances or hinders the task. The distribution of this outcome is conditioned on the environment dynamics:

$$P_{env}(y_{t+k} \mid o_t, a_{t:t+k}) \quad (2)$$

This is estimated by a teacher distribution $P_{\mathcal{T}}(y_{t+k} \mid o_{t:t+k})$ that has privileged access to a world model or simulator:

$$P_{env}(o_{t+k} \mid o_{t:t+k-1}, a_{t:t+k}) \quad (3)$$

The teacher receives the full ground-truth visual history from the simulated environment, along with 6DoF poses for both the object and the gripper.

3.2 The Latent Sufficiency Hypothesis

Evaluating $P_{\mathcal{T}}$ is computationally prohibitive for real-time control. Our objective is to learn a fast estimator, the student distribution $P_{\theta}(y_{t+k} \mid z_t, a_{t:t+k})$, conditioned only on the policy’s latent state and planned actions. The feasibility of this approach rests on the following hypothesis:

Hypothesis on Latent Sufficiency. *For the task of semantic world modeling, the mutual information between the policy latent-action tuple $(z_t, a_{t:t+k})$ and the future outcome y_{t+k} approximates the mutual information between the full observation history $o_{t:t+k}$ and y_{t+k} :*

$$I(o_{t:t+k}; y_{t+k}) \approx I((z_t, a_{t:t+k}); y_{t+k}). \quad (4)$$

Whether trained via Reinforcement Learning (e.g., SAC [10]) to maximize task reward, or Imitation Learning (e.g., ACT [35]) to reconstruct expert action

intent, the encoder $\mathcal{E}_\pi$ must extract and retain interaction-critical representations, such as object geometry, relative distances, and contact dynamics. We claim that $(z_t, a_{t:t+k})$, derived solely from o_t , carries as much predictive information about y_{t+k} as the full observation history seen by the teacher, making z_t a sufficient statistic for y_{t+k} and rendering generating and re-encoding raw pixels unnecessary. This sufficiency extends to failing policies. When a policy encounters an out-of-distribution state, the resulting latent z_t captures the “context of failure” (*e.g.*, uncertainty or feature mismatch), and DILLO learns to map these latents to Negative verdicts.

By minimizing the Kullback–Leibler divergence between the teacher and student distributions [14]:

$$\min_{\theta} D_{KL}(P_{\mathcal{T}}(y_{t+k} \mid o_{t:t+k}) \parallel P_{\theta}(y_{t+k} \mid z_t, a_{t:t+k})) \quad (5)$$

we distill the teacher’s future-aware foresight into a Language-Action World Model that translates the policy’s existing manifold into natural language descriptions. This reduces to a cross-entropy minimization over $(z_t, a_{t:t+k})$. DILLO thereby predicts outcomes over the horizon $t \rightarrow t+k$ with negligible computational overhead.

3.3 Models

Teacher. We instantiate the teacher $\mathcal{T}$ as a VLM with access to privileged information from the simulator state, eliminating the depth ambiguity and occlusion issues inherent to raw vision. In particular, the model has access to:

1. **Visual Context:** The sequence of RGB observations corresponding to the execution of the action chunk.
2. **Geometric Deltas:** Directional changes and 6DoF poses of the end-effector and task-relevant objects, grounding the VLM in the physics of the scene and mitigating the spatial hallucinations common in vision-only models.
3. **Episode Outcome Signal:** A binary indicator of whether the specific action chunk led to successful task completion.

In physical deployment, this privileged information can be substituted by state-of-the-art 6D pose estimation algorithms [29].

To construct the distillation dataset, the privileged teacher model processes the ground-truth state information alongside the natural language task description and the target goal configuration. We segment policy rollouts at the state-action level. For each transition, the teacher generates a natural language description $d_{\mathcal{T}}$ and assigns a binary verdict $c_{\mathcal{T}} \in \{\text{POSITIVE}, \text{NEGATIVE}\}$. A NEGATIVE verdict explicitly penalizes trajectories that exhibit stagnation, erratic actuation, or behaviors otherwise detrimental to task progression. Finally, we pair these high-level semantic labels with the base policy’s internal latent state z_t and the predicted action chunk $a_{t:t+k}$. This yields the final distillation tuple:

$$\tau = (z_t, a_{t:t+k}, d_{\mathcal{T}}, c_{\mathcal{T}})$$

Student (DILLO). The student f_θ parameterizes P_θ using the Gemma [27] family of open weights. We implement two variants to validate scalability:

- **DILLO-1B:** Uses the standard language-only Gemma-1B-it LLM backbone.
- **DILLO-4B:** Unlike standard VLM inference, which relies on a 417M-parameter SigLIP encoder, our model uses only the large language model of the Gemma-VLM-4B-it architecture. This ensures that DILLO-4B benefits from the reasoning capacity of a VLM while incurring zero visual overhead.

To map the continuous control space to the language space, we employ two learnable projectors:

- **Latent Projector** $P_z : \mathbb{R}^{D_z} \rightarrow \mathbb{R}^{D_{emb}}$, mapping the frozen policy embedding to the LLM input space.
- **Action Projector** $P_a : \mathbb{R}^{D_a} \rightarrow \mathbb{R}^{D_{emb}}$, mapping each action chunk $a_{t:t+k}$ to an action vector that is then fed to the LLM.

The input sequence to the Transformer is constructed as:

$$X_{in} = [\text{TASK PROMPT}] \oplus P_z(z_t) \oplus P_a(a_{t:t+k}) \quad (6)$$

Here, X_{in} contains no tokens derived from o_t , ensuring independence from the visual encoder during inference.

3.4 Training

End-to-end mapping of continuous control latents into discrete LLM tokens is prone to modality collapse. To stably ground the LLM in the policy’s physical manifold, we employ a progressive three-stage curriculum:

Stage 1: Projector Alignment. To bridge the dimensional gap without destabilizing pre-trained weights, we freeze the LLM backbone θ_{LLM} and optimize only the linear projectors $\{P_z, P_a\}$ [19]. We use standard autoregressive next-token prediction on the teacher’s descriptions to ensure the projected latents act as valid prompt embeddings.

Stage 2: Description Distillation. To establish a dense semantic prior, we apply Low-Rank Adaptation (LoRA) [15] to the LLM. We jointly optimize the LoRA parameters and projectors to reconstruct the teacher’s natural language rationale $d_{\mathcal{T}}$:

$$\mathcal{L}_{Desc} = -\mathbb{E}_{\mathcal{D}} \left[\sum_{i=1}^{|d|} \log p_\theta(d_i | d_{<i}, z_t, a_{t:t+k}) \right] \quad (7)$$

Stage 3: Verdict Optimization. Finally, we introduce the binary task constraint. The model generates the rationale followed by a verdict token $\hat{c} \in \{\text{POSITIVE}, \text{NEGATIVE}\}$. We minimize the joint objective:

$$\mathcal{L}_{Total} = \mathcal{L}_{Desc} + \lambda \cdot \mathcal{L}_{CE}(c_{\mathcal{T}}, \hat{c}) \quad (8)$$

where $\mathcal{L}_{CE}$ is the cross-entropy against the ground-truth verdict $c_{\mathcal{T}}$, ensuring the natural language descriptions are grounded with the task execution.




Fig. 3: Proactive Policy Steering via Latent Rejection Sampling. At inference time, DILLO acts as a high-speed safety filter. The policy proposes N candidate action chunks ($a^i_{t:t+K}$). Instead of re-rendering future images for each candidate, DILLO projects the policy latent (z_t) directly into the Language-Action World Model. The model predicts a semantic description and verdict for each chunk. Negative trajectories (red $\times$) are rejected, and the first verified positive plan (green $\checkmark$) is executed.

3.5 Proactive Steering via Latent Rejection Sampling

DILLO’s distilled verdict $\hat{c}$ enables it to function as a controller-agnostic, zero-overhead safety filter at inference time. Rather than executing the first proposed action chunk, the agent uses DILLO to screen a batch of N candidate plans before committing to any physical action, as illustrated in Fig. 3.

Inference Protocol. At each decision step t , the base policy encodes the current observation into a shared latent state $z_t = \mathcal{E}_\pi(o_t)$, computed *once* regardless of how many candidates are subsequently evaluated. The policy then proposes a budget of N candidate action chunks by sampling independently from its action distribution:

$$\{a^i_{t:t+K}\}_{i=1}^N \sim \pi(\cdot | z_t) \quad (9)$$

DILLO evaluates all N candidates in a single batched forward pass. For each candidate, the shared latent z_t and an encoded projection of the proposed action sequence are jointly fed into the model:

$$y^i_{t+K} = (d^i_{t+K}, c^i_{t+K}) = f_\theta(P_z(z_t), P_a(a^i_{t:t+K})), \quad i = 1, \dots, N \quad (10)$$

The agent then selects any candidate that receives a POSITIVE verdict. If no POSITIVE verdict is found within the budget, the agent falls back to executing a candidate from the initial proposal, ensuring the control loop is never stalled. In our experiments we set $N=5$. Notably, the framework naturally supports multiple resampling rounds: the policy can draw and evaluate successive batches of N candidates up to R rounds (where R is a configurable budget), allowing up to $R \cdot N$ total evaluations before fallback at the cost of added latency (see Sec. 4.4).

4 Experiments

Our evaluation validates two central claims. First, *latent sufficiency*: does DILLO, conditioned solely on latent states (z_t), produce semantic descriptions (y_{t+k}) that faithfully match ground-truth physical outcomes (o_{t+k})? Second, *utility*: does DILLO’s verdict effectively steer a base policy toward higher episode success rates? We conduct experiments in **MetaWorld** [33] (Soccer, Sweep-Into, Drawer-Open) and **LIBERO** [18] (Goal, Object, Spatial, 10, 90).

4.1 Constructing the Distillation Dataset

To train DILLO, we collect distillation datasets by rolling out pre-trained policies across two distinct algorithmic paradigms, capturing the latent transitions inherent to both specialized and generalist behaviors.

Single-Task Reinforcement Learning (MetaWorld). We roll out SAC-trained policies [10] across three manipulation tasks (Soccer, Sweep-Into, Drawer-Open). Each policy is specialized for a single task; correspondingly, a separate DILLO model is distilled per task. The collected trajectories consist of purely observational latents, $z_t = \mathcal{E}(o_t)$, where the representation is driven entirely by the RL reward signal, with no language conditioning.

Multi-Task Imitation Learning (LIBERO). We roll out action-chunking, imitation-learned policies [35] across five task suites (LIBERO-Goal, LIBERO-Object, LIBERO-Spatial, LIBERO-10, LIBERO-90). Here, a single policy must generalize over a shared set of language conditioned tasks, and the distillation dataset reflects this diversity. A single DILLO model is therefore distilled from trajectories spanning heterogeneous, language-specified goals, testing whether the latent space of a generalist policy retains sufficient structure for cross-task outcome prediction.

Failures. A reliable predictive model must identify failure modes across all levels of agent capability. A classifier trained solely on expert successes and random failures will fail to generalize to subtle execution errors. To guarantee a comprehensive distribution of both catastrophic and marginal failures, we curate our dataset using two complementary strategies:

- **Spectrum of Competency Sampling.** Rather than relying exclusively on converged experts, we sample rollouts from policy checkpoints saved at distinct performance thresholds. We incorporate low-competency models (20% success) to capture erratic exploratory behaviors, mid-competency models (50% success) to provide nuanced “near-miss” examples [1] that are difficult to classify, and high-competency models (80% success) to establish ground-truth task alignment while offering “hard negatives” through their rare failures.
- **Perception Noise Injection.** During the data collection step, we inject Gaussian noise $\epsilon \sim \mathcal{N}(0, \sigma^2)$ into the observation o_t before querying the policy $\pi(a_{t:t+k} | o_t + \epsilon)$. This represents a standard proxy for sensor uncertainty used in Sim-to-Real domain randomization. The perturbation causes

the agent to misinterpret its state, inducing trajectory drift and compounding execution errors. Crucially, this generates failures that originate from a competent action manifold rather than random initialization, systematically mimicking the covariate shift commonly observed during physical deployment.

4.2 Metrics

To support the hypothesis on “latent sufficiency”, we must quantify the alignment between the text predicted from latents (z_t) and the actual physical outcome observed in the simulator (o_{t+k}). We compare DILLO against baselines backed by Gemma-3-1B and Gemma-3-4B:

- **Zero-Shot (ZS)**: The model receives the current observation o_t and the task description.
- **Few-Shot (FS)**: The model receives o_t along with 3 in-context examples of (o_t, y_{t+k}) pairs.
- **Few-Shot Captioning (CAP)**: The model receives the current observation o_t , the *future* observation o_{t+k} , and 3 in-context examples of (o_t, o_{t+k}, y_{t+k}) tuples. It does not need to predict the future; it can caption what actually happened therefore it should be seen as an oracle.

To evaluate baselines and DILLO, we rely on two fidelity metrics:

Text-to-Observation Fidelity (T2O). To quantify low-level dynamic accuracy without the noise of linguistic ambiguity, we map both physical state transitions and generated descriptions into a shared space of canonical directional labels. A deterministic function, $f_{extract}(o_t, o_{t+k})$, computes the true displacement for the gripper and objects, as well as relative relational changes (*e.g.*, *approaching* vs. *receding*) across three orthogonal axes. Axes with negligible movement are assigned a “static” label. Simultaneously, a parser $g_{parse}(d)$ projects the natural language prediction into the same canonical set via robust synonym mapping. The final T2O score is computed as the accuracy of matched labels across all dynamic variables, serving as a direct proxy for physical grounding (see supplementary material for the definition of g_{parse}).

Reasoning-Based Semantic Score. While T2O captures dynamics, it may miss nuanced semantic details. We employ a Judge LLM (Qwen2.5-32B-Instruct [32]) to score the semantic alignment between DILLO’s prediction d and the ground-truth reference $d_{\mathcal{T}}$. The judge decomposes the reference into a set of atomic facts $F = \{f_i\}_{i=1}^{|F|}$. Subsequently, it evaluates the extent to which the candidate description d entails each fact, assigning a support score $m_i \in \{0, 0.5, 1\}$ (contradiction, partial implication, and full support). We define the final metric as the mean fact recall:

$$\text{Score}(d_{t+k}|d_{\mathcal{T}}) = \frac{1}{|F|} \sum_{i=1}^{|F|} m_i \quad (11)$$

Table 1: Fidelity Results. We measure the alignment between the predicted description and the actual simulator state. We compare DILLO against LLM and VLM baselines (Few-Shot, Zero-Shot, and CAP) across two model scales (1B and 4B parameters) on both the MetaWorld and LIBERO suites. DILLO consistently achieves high fidelity and outperforms or matches baselines, validating that the latent state is a sufficient statistic for semantic prediction.

Environment	Task	1B Parameter Models				4B Parameter Models			
		DILLO-1B	FS	ZS	CAP	DILLO-4B	FS	ZS	CAP
MetaWorld	Soccer	0.518	0.495	0.398	0.477	0.586	0.501	0.540	0.626
	Sweep-Into	0.562	0.329	0.528	0.271	0.554	0.630	0.687	0.725
	Drawer-Open	0.706	0.626	0.419	0.612	0.705	0.555	0.553	0.609
LIBERO	Goal	0.727	0.262	0.287	0.723	0.710	0.422	0.241	0.905
	Object	0.622	0.388	0.204	0.820	0.587	0.546	0.225	0.877
	Spatial	0.668	0.546	0.341	0.794	0.656	0.563	0.194	0.902
	10	0.648	0.214	0.199	0.630	0.686	0.449	0.265	0.867
	90	0.708	0.258	0.206	0.640	0.650	0.278	0.258	0.713

Table 2: Reasoning Based LLM Score. We compare DILLO against VLM baselines (Few-Shot, Zero-Shot, and CAP) across two model scales (1B and 4B parameters). DILLO consistently outperforms standard baselines in semantic reasoning, even approaching or exceeding the CAP in the 1B regime.

Environment	Task	1B Parameter Models				4B Parameter Models			
		DILLO-1B	FS	ZS	CAP	DILLO-4B	FS	ZS	CAP
MetaWorld	Soccer	0.471	0.301	0.204	0.277	0.504	0.361	0.328	0.425
	Sweep-Into	0.592	0.372	0.284	0.388	0.607	0.391	0.381	0.364
	Drawer-Open	0.655	0.510	0.189	0.427	0.628	0.497	0.212	0.252
LIBERO	Goal	0.834	0.259	0.271	0.686	0.830	0.487	0.278	0.816
	Object	0.803	0.380	0.208	0.668	0.723	0.519	0.243	0.684
	Spatial	0.782	0.448	0.323	0.684	0.758	0.545	0.282	0.792
	10	0.822	0.227	0.205	0.611	0.797	0.493	0.282	0.777
	90	0.708	0.258	0.206	0.640	0.650	0.278	0.258	0.713

This metric offers a robust, interpretable measure of semantic coverage beyond simple lexical overlap metrics like BLEU [23] or ROUGE [17].

Episode Success Rate (ES). We measure the percentage of episodes where the agent completes the task goal. We report the success rate of the DILLO-Steered Policy versus the Base Policy (no steering). A high ES for DILLO confirms that the distilled verdict $\hat{c}$ is a reliable signal for filtering negative actions.

4.3 Validating Latent Sufficiency

We validate the Latent Sufficiency Hypothesis using the T2O fidelity and LLM-judged semantic scores reported in Tables 1 and 2. The results reveal three consistent trends across both benchmarks.

Latent beats pixels. The “blind” DILLO-4B model consistently and significantly outperforms the pixel-based Few-Shot and Zero-Shot baselines on both




Fig. 4: Qualitative validation of DILLO’s foresight for ‘Positive’ (left) and ‘Negative’ (right) outcomes. DILLO generates the predicted description (d) and verdict (check/cross) using only the policy’s latent state (z_t) and the proposed action chunk. The initial observation (o_t) and ground-truth future (o_{t+K}) are shown purely for visual comparison, confirming the high fidelity of the latent-based prediction.

MetaWorld and LIBERO, confirming that the policy latent z_t is a more informative predictor of future dynamics than the current visual observation o_t .

Latent beats the future image. Remarkably in Table 1, DILLO-4B’s T2O score also exceeds the CAP baseline on **Drawer-Open** (**0.705** vs. **0.609**), where CAP has access to the actual future observation o_{t+k} and 3 in-context examples of (o_t, o_{t+k}, y_{t+k}) , which allow for pure captioning without having to predict the future. This indicates that z_t is not merely sufficient but a *superior* predictive statistic: as a task-optimized embedding, it is less susceptible to the noise and perceptual ambiguity inherent in raw pixel data. This directly justifies the Zero-Visual-Overhead architecture.

Semantically richer descriptions. The LLM-judged scores in Table 1 show that both DILLO-1B and DILLO-4B outperform all baselines across the vast majority of tasks on both benchmarks. On MetaWorld **Drawer-Open**, DILLO-4B achieves **0.628** against **0.497** for FS and **0.252** for CAP. This strong performance extends to LIBERO, with two notable exceptions. On LIBERO-90, which spans 90 distinct tasks, performance relative to CAP suggests that larger-scale data collection is needed to handle the increased task diversity. On LIBERO-Spatial, significant variability in object placements, combined with the absence of explicit goal position conditioning in the current DILLO formulation, limits generalization to spatially complex configurations. These are targeted failure modes rather than systemic weaknesses, and both suggest clear directions for future work.

Figure 4 provides qualitative validation on LIBERO. In the POSITIVE outcome (left), the predicted description $\hat{d}$ accurately articulates the arm’s successful engagement with the object. In the NEGATIVE outcome (right), DILLO correctly anticipates the impending failure, the arm disengages from the target, and issues

Table 3: Steering Success Rates and Verdict Accuracy. We report Steering Success Rates (SR) and Safety Verdict Accuracy (%). High success rates and accuracy confirm that the model effectively steers rollouts while discriminating successful dynamics from failures.

		Success Rate			Verdict Accuracy	
	Task	Base	DILLO-1B	DILLO-4B	DILLO-1B	DILLO-4B
MetaWorld	Soccer	65.0%	75.0%	80.0%	82.7%	87.4%
	Sweep-Into	80.0%	90.0%	85.0%	97.7%	92.1%
	Drawer-Open	75.0%	70.0%	90.0%	80.6%	94.7%
LIBERO	Goal	86.0%	87.5%	89.0%	80.0%	87.5%
	Object	82.0%	86.5%	88.0%	79.8%	81.2%
	Spatial	70.0%	76.0%	73.0%	77.8%	86.7%
	10	64.0%	71.5%	74.0%	88.6%	92.5%
	90	71.4%	68.6%	70.5%	86.9%	85.7%

a NEGATIVE verdict. Both the description and the verdict are generated solely from z_t and $a_{t:t+k}$, with no access to visual observations at inference time.

4.4 Proactive Policy Steering and Inference Latency

Proactive Policy Steering. We employ DILLO as a rejection sampling controller ($N=5$) to quantify steering gains over 20 episodes per task. Table 3 reports success rates across both benchmarks.

On MetaWorld, DILLO consistently improves over the base policy. On **Soccer**, DILLO-4B raises the success rate from 65.0% to **80.0%**, and even the smaller DILLO-1B achieves 75.0%. On **Sweep-Into**, DILLO-1B reaches **90.0%**, the highest score across all model variants on that task.

On LIBERO, gains are consistent across the four main suites. DILLO-4B improves LIBERO-Goal from 86.0% to **89.0%** and LIBERO-Object from 82.0% to **88.0%**, confirming that the distilled verdict generalizes effectively to multi-task, language-conditioned settings. The exception is LIBERO-90, where the base policy (71.4%) marginally outperforms both DILLO variants (68.6% and 70.5%). We analyze 900 LIBERO-90 episodes under varying rejection thresholds to isolate the source of failures. Moderate gating yields a net positive effect, improving success by +1.7% overall and by +17.5% on tasks below 70% baseline success. Heavier gating, however, over-rejects candidate actions and reduces success by -4.4%, indicating that broad multi-task deployment requires task-conditioned calibration rather than a single global rejection threshold.

Generalization across LIBERO suites. To test cross-suite transfer, we steer the LIBERO-10 ACT policy with our LIBERO-90-trained explainer. DILLO raises policy success from 64.0% to 69.5%, approaching the 71.5% of same-suite steering. Meaning that the learned semantic evaluator does not simply memorize a single suite but can generalize to similar scenarios.

Positive/Negative Classification. To understand the source of this effectiveness, we analyze the binary classification accuracy of the verdict token $\hat{c}$. Table 3 shows that DILLO-4B achieves an average accuracy of **91.4%** across tasks, without access to any visual observation at inference time. A human validation study (see supplementary material) further corroborates these results. On MetaWorld, DILLO agrees with human positive assessments on 84.1% of 60 episodes. On LIBERO, DILLO’s verdict agrees with ground truth at 72.1%, exceeding the human-vs-ground-truth agreement of 69.5%, while human evaluators rate its descriptions at 3.62/5.0, confirming interpretability to non-expert observers.

Inference Latency. Per decision step (one action chunk of 20 simulation steps), the base policy accounts for $T_{\text{sample}} \approx 10$ ms, and DILLO adds 0.26 s (1B) and 0.38 s (4B) per step. This is a $\sim 14\times$ and $\sim 10\times$ speedup over Forewarn [31] (3.70s), measured on an NVIDIA RTX 3090 against Forewarn’s more powerful RTX A6000. Additionally on LIBERO-10, DILLO-1B reduces average episode length by 36% (390 $\rightarrow$ 250 simulation steps, 19.8 $\rightarrow$ 13.0 decision steps), yielding a net episode-level overhead of only ~ 3 s over the baseline.

5 Limitations

DILLO requires re-alignment of the projectors for each new policy architecture, as the latent manifold z_t is architecture-specific. The training pipeline itself is environment- and action-representation-agnostic, enabling direct deployment on physical hardware by having the teacher annotate real trajectories. The privileged information (ground-truth 6DoF poses) is required only for the teacher during offline annotation, never at inference; in practice, it can be substituted by estimated poses using a pose estimation algorithm [29], though perception noise may degrade distillation quality.

6 Conclusion

We introduce DILLO, a Distilled Language-Action World Model that bypasses the latency bottleneck of visual simulation in proactive agent steering. By distilling the semantic foresight of a simulator-privileged teacher into a fast, latent-conditioned student, we demonstrate that a policy’s internal representations already encode the necessary dynamics to anticipate action outcomes. This Latent Sufficiency Hypothesis allows us to shift the paradigm of world modeling from computationally expensive “simulate-then-act” pipelines to a streamlined “describe-then-act” approach.

Our extensive evaluation across both single-task (MetaWorld) and multi-task (LIBERO) environments confirms that DILLO produces high-fidelity outcome descriptions and failure verdicts directly from the policy latent z_t . This architecture achieves a $\sim 14\times$ inference speedup over visual world models while effectively steering policies to high success rates. Ultimately, DILLO establishes that real-time semantic foresight does not require the expensive rendering of future states, offering a fast, interpretable, reliability layer for embodied agents.

7 Acknowledgments

I.S. gratefully acknowledges the NVIDIA Academic Grant Program for the donation of a DGX Spark. This work was further supported by Panasonic and by Sapienza grant RG123188B3EF6A80 (CENTS). We also thank CINECA for the allocation of computational resources.

References

1. Agia, C., Sinha, R., Yang, J., Cao, Z., Antonova, R., Pavone, M., Bohg, J.: Unpacking failure modes of generative policies: Runtime monitoring of consistency and progress. In: Agrawal, P., Kroemer, O., Burgard, W. (eds.) Conference on Robot Learning, 6-9 November 2024, Munich, Germany. Proceedings of Machine Learning Research, vol. 270, pp. 689–723. PMLR (2024), <https://proceedings.mlr.press/v270/agia25a.html>
2. Duan, J., Pumacay, W., Kumar, N., Wang, Y.R., Tian, S., Yuan, W., Krishna, R., Fox, D., Mandlekar, A., Guo, Y.: AHA: A vision-language-model for detecting and reasoning over failures in robotic manipulation. In: The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025 (2025), <https://openreview.net/forum?id=JVkdSi7Ekg>
3. Ebert, F., Finn, C., Dasari, S., Xie, A., Lee, A.X., Levine, S.: Visual foresight: Model-based deep reinforcement learning for vision-based robotic control. CoRR **abs/1812.00568** (2018), <http://arxiv.org/abs/1812.00568>
4. Finn, C., Goodfellow, I.J., Levine, S.: Unsupervised learning for physical interaction through video prediction. In: Lee, D.D., Sugiyama, M., von Luxburg, U., Guyon, I., Garnett, R. (eds.) Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain. pp. 64–72 (2016), <https://proceedings.neurips.cc/paper/2016/hash/d9d4f495e875a2e075a1a4a6e1b9770f-Abstract.html>
5. Finn, C., Levine, S.: Deep visual foresight for planning robot motion. In: 2017 IEEE International Conference on Robotics and Automation, ICRA 2017, Singapore, Singapore, May 29 - June 3, 2017. pp. 2786–2793. IEEE (2017). <https://doi.org/10.1109/ICRA.2017.7989324>, <https://doi.org/10.1109/ICRA.2017.7989324>
6. García, C.E., Prett, D.M., Morari, M.: Model predictive control: Theory and practice—a survey. *Automatica* **25**(3), 335–348 (1989). [https://doi.org/https://doi.org/10.1016/0005-1098\(89\)90002-2](https://doi.org/https://doi.org/10.1016/0005-1098(89)90002-2), <https://www.sciencedirect.com/science/article/pii/0005109889900022>
7. Gu, Q., Ju, Y., Sun, S., Gilitschenski, I., Nishimura, H., Itkina, M., Shkurti, F.: SAFE: Multitask failure detection for vision-language-action models. In: The Thirty-ninth Annual Conference on Neural Information Processing Systems (2025), <https://openreview.net/forum?id=XPyAukgsFf>
8. Guan, L., Zhou, Y., Liu, D., Zha, Y., Amor, H.B., Kambhampati, S.: "task success" is not enough: Investigating the use of video-language models as behavior critics for catching undesirable agent behaviors. CoRR **abs/2402.04210** (2024). <https://doi.org/10.48550/ARXIV.2402.04210>, <https://doi.org/10.48550/arXiv.2402.04210>
9. Ha, D., Schmidhuber, J.: Recurrent world models facilitate policy evolution. In: Bengio, S., Wallach, H.M., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett,

- R. (eds.) Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada. pp. 2455–2467 (2018), <https://proceedings.neurips.cc/paper/2018/hash/2de5d16682c3c35007e4e92982f1a2ba-Abstract.html>
10. Haarnoja, T., Zhou, A., Abbeel, P., Levine, S.: Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In: Dy, J.G., Krause, A. (eds.) Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018. Proceedings of Machine Learning Research, vol. 80, pp. 1856–1865. PMLR (2018), <http://proceedings.mlr.press/v80/haarnoja18b.html>
11. Hafner, D., Lillicrap, T.P., Ba, J., Norouzi, M.: Dream to control: Learning behaviors by latent imagination. In: 8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020. OpenReview.net (2020), <https://openreview.net/forum?id=S1l0TC4tDS>
12. Hafner, D., Lillicrap, T.P., Norouzi, M., Ba, J.: Mastering atari with discrete world models. In: 9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021. OpenReview.net (2021), <https://openreview.net/forum?id=0oabwyZb0u>
13. Hafner, D., Pasukonis, J., Ba, J., Lillicrap, T.P.: Mastering diverse domains through world models. CoRR **abs/2301.04104** (2023). <https://doi.org/10.48550/ARXIV.2301.04104>, <https://doi.org/10.48550/arXiv.2301.04104>
14. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network (2015), <https://arxiv.org/abs/1503.02531>
15. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: Lora: Low-rank adaptation of large language models. In: The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022. OpenReview.net (2022), <https://openreview.net/forum?id=nZeVKeeFYf9>
16. Li, C., Krause, A., Hutter, M.: Robotic world model: A neural network simulator for robust policy optimization in robotics. CoRR **abs/2501.10100** (2025). <https://doi.org/10.48550/ARXIV.2501.10100>, <https://doi.org/10.48550/arXiv.2501.10100>
17. Lin, C.Y.: ROUGE: A package for automatic evaluation of summaries. In: Text Summarization Branches Out. pp. 74–81. Association for Computational Linguistics, Barcelona, Spain (Jul 2004), <https://aclanthology.org/W04-1013/>
18. Liu, B., Zhu, Y., Gao, C., Feng, Y., Liu, Q., Zhu, Y., Stone, P.: Libero: Benchmarking knowledge transfer for lifelong robot learning. Advances in Neural Information Processing Systems **36**, 44776–44791 (2023)
19. Liu, H., Li, C., Wu, Q., Lee, Y.J.: Visual instruction tuning. In: Oh, A., Naumann, T., Globerson, A., Saenko, K., Hardt, M., Levine, S. (eds.) Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023 (2023), http://papers.nips.cc/paper_files/paper/2023/hash/6dcf277ea32ce3288914faf369fe6de0-Abstract-Conference.html
20. Liu, Y., Hamid, J.I., Xie, A., Lee, Y., Du, M., Finn, C.: Bidirectional decoding: Improving action chunking via closed-loop resampling. CoRR **abs/2408.17355** (2024). <https://doi.org/10.48550/ARXIV.2408.17355>, <https://doi.org/10.48550/arXiv.2408.17355>
21. Liu, Z., Bahety, A., Song, S.: REFLECT: summarizing robot experiences for failure explanation and correction. In: Tan, J., Toussaint, M., Darvish, K. (eds.) Conference on Robot Learning, CoRL 2023, 6-9 November 2023, Atlanta, GA, USA.

- Proceedings of Machine Learning Research, vol. 229, pp. 3468–3484. PMLR (2023), <https://proceedings.mlr.press/v229/liu23g.html>
22. Milani, S., Topin, N., Veloso, M., Fang, F.: Explainable reinforcement learning: A survey and comparative review. *ACM Comput. Surv.* **56**(7), 168:1–168:36 (2024). <https://doi.org/10.1145/3616864>, <https://doi.org/10.1145/3616864>
 23. Papineni, K., Roukos, S., Ward, T., Zhu, W.: Bleu: a method for automatic evaluation of machine translation. In: Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, July 6–12, 2002, Philadelphia, PA, USA. pp. 311–318. ACL (2002). <https://doi.org/10.3115/1073083.1073135>, <https://aclanthology.org/P02-1040/>
 24. Sammani, F., Mukherjee, T., Deligiannis, N.: NLX-GPT: A model for natural language explanations in vision and vision-language tasks. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18–24, 2022. pp. 8312–8322. IEEE (2022). <https://doi.org/10.1109/CVPR52688.2022.00814>, <https://doi.org/10.1109/CVPR52688.2022.00814>
 25. Sinha, R., Elhafi, A., Agia, C., Foutter, M., Schmerling, E., Pavone, M.: Real-time anomaly detection and reactive planning with large language models. In: Kulić, D., Venture, G., Bekris, K.E., Coronado, E. (eds.) *Robotics: Science and Systems XX*, Delft, The Netherlands, July 15–19, 2024 (2024). <https://doi.org/10.15607/RSS.2024.XX.114>, <https://doi.org/10.15607/RSS.2024.XX.114>
 26. Sutton, R.S.: Dyna, an integrated architecture for learning, planning, and reacting. *SIGART Bull.* **2**(4), 160–163 (1991). <https://doi.org/10.1145/122344.122377>, <https://doi.org/10.1145/122344.122377>
 27. Team, G.: Gemma 3 technical report. CoRR **abs/2503.19786** (2025). <https://doi.org/10.48550/ARXIV.2503.19786>, <https://doi.org/10.48550/arXiv.2503.19786>
 28. Wang, Z., Liang, B., Dhat, V., Brumbaugh, Z., Walker, N., Krishna, R., Cakmak, M.: I can tell what I am doing: Toward real-world natural language grounding of robot experiences. In: Agrawal, P., Kroemer, O., Burgard, W. (eds.) *Conference on Robot Learning*, 6–9 November 2024, Munich, Germany. Proceedings of Machine Learning Research, vol. 270, pp. 1863–1890. PMLR (2024), <https://proceedings.mlr.press/v270/wang25g.html>
 29. Wen, B., Yang, W., Kautz, J., Birchfield, S.: Foundationpose: Unified 6d pose estimation and tracking of novel objects. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16–22, 2024. pp. 17868–17879. IEEE (2024). <https://doi.org/10.1109/CVPR52733.2024.01692>, <https://doi.org/10.1109/CVPR52733.2024.01692>
 30. Wojciechowski, A., Lango, M., Dusek, O.: Faithful and plausible natural language explanations for image classification: A pipeline approach. In: Al-Onaizan, Y., Bansal, M., Chen, Y. (eds.) *Findings of the Association for Computational Linguistics: EMNLP 2024*, Miami, Florida, USA, November 12–16, 2024. pp. 2340–2351. Association for Computational Linguistics (2024). <https://doi.org/10.18653/v1/2024.FINDINGS-EMNLP.130>, <https://doi.org/10.18653/v1/2024.findings-emnlp.130>
 31. Wu, Y., Tian, R., Swamy, G., Bajcsy, A.: From foresight to forethought: Vlm-in-the-loop policy steering via latent alignment. In: *ICLR 2025 Workshop on World Models: Understanding, Modelling and Scaling*
 32. Yang, A., Yang, B., Zhang, B., Hui, B., Zheng, B., Yu, B., Li, C., Liu, D., Huang, F., Wei, H., Lin, H., Yang, J., Tu, J., Zhang, J., Yang, J., Yang, J., Zhou, J., Lin, J., Dang, K., Lu, K., Bao, K., Yang, K., Yu, L., Li, M., Xue, M., Zhang, P., Zhu,

- Q., Men, R., Lin, R., Li, T., Xia, T., Ren, X., Ren, X., Fan, Y., Su, Y., Zhang, Y., Wan, Y., Liu, Y., Cui, Z., Zhang, Z., Qiu, Z.: Qwen2.5 technical report (2024). <https://doi.org/10.48550/ARXIV.2412.15115>, <https://doi.org/10.48550/arXiv.2412.15115>
33. Yu, T., Quillen, D., He, Z., Julian, R., Hausman, K., Finn, C., Levine, S.: Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In: Kaelbling, L.P., Kragic, D., Sugiura, K. (eds.) 3rd Annual Conference on Robot Learning, CoRL 2019, Osaka, Japan, October 30 - November 1, 2019, Proceedings. Proceedings of Machine Learning Research, vol. 100, pp. 1094–1100. PMLR (2019), <http://proceedings.mlr.press/v100/yu20a.html>
 34. Zawalski, M., Chen, W., Pertsch, K., Mees, O., Finn, C., Levine, S.: Robotic control via embodied chain-of-thought reasoning. In: Agrawal, P., Kroemer, O., Burgard, W. (eds.) Conference on Robot Learning, 6-9 November 2024, Munich, Germany. Proceedings of Machine Learning Research, vol. 270, pp. 3157–3181. PMLR (2024), <https://proceedings.mlr.press/v270/zawalski25a.html>
 35. Zhao, T.Z., Kumar, V., Levine, S., Finn, C.: Learning fine-grained bimanual manipulation with low-cost hardware. In: Proceedings of Robotics: Science and Systems (RSS) (2023)