

# From Foresight to Forethought: VLM-In-the-Loop Policy Steering via Latent Alignment

Yilin Wu<sup>1</sup>, Ran Tian<sup>2</sup>, Gokul Swamy<sup>1</sup>, Andrea Bajcsy<sup>1</sup>

<sup>1</sup>Carnegie Mellon University <sup>2</sup>UC Berkeley

{yilinwu, gswamy, abajcsy}@andrew.cmu.edu, rantian@berkeley.edu

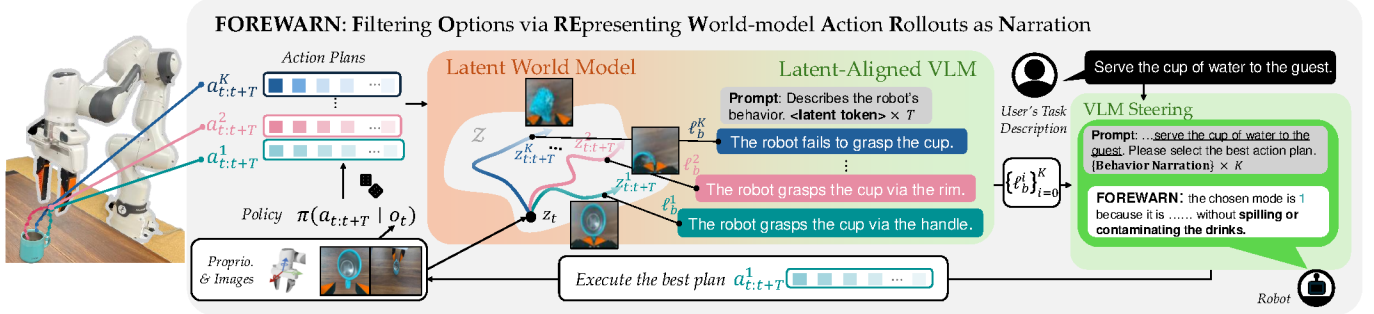


Fig. 1: We present **FOREWARN**, an VLM-in-the-loop policy steering algorithm for multi-modal generative robot policies. Our key idea is to decouple the VLM’s burden of predicting action outcomes from evaluation. By predicting action outcomes with a pre-trained latent dynamics model and aligning a VLM to reason about these latent states in text, FOREWARN can select action plans at runtime that are most appropriate for new task contexts and user needs.

**Abstract**—While generative robot policies have demonstrated significant potential in learning complex, multimodal behaviors from demonstrations, they still exhibit diverse failures at deployment-time. Policy steering offers an elegant solution to reducing the chance of failure by using an external verifier to select from low-level actions proposed by an imperfect generative policy. Here, one might hope to use a Vision Language Model (VLM) as a verifier, leveraging its open-world reasoning capabilities. However, off-the-shelf VLMs struggle to understand the consequences of low-level robot actions as they are represented fundamentally differently than the text and images the VLM was trained on. In response, we propose FOREWARN, a novel framework to unlock the potential of VLMs as open-vocabulary verifiers for runtime policy steering. Our key idea is to decouple the VLM’s burden of predicting action outcomes (*foresight*) from evaluation (*forethought*). For foresight, we leverage a latent world model to imagine future latent states given diverse low-level action plans. For forethought, we align the VLM with these predicted latent states to reason about the consequences of actions in its native representation—natural language—and effectively filter proposed plans. We validate our framework across diverse robotic manipulation tasks, demonstrating its ability to bridge representational gaps and provide robust, generalizable policy steering. Videos can be found on the project website: <https://yilinwu98.github.io/forewarn/>.

## I. INTRODUCTION

Generative imitation-based policies are an increasingly powerful way to learn low-level robot behaviors from multi-modal<sup>1</sup> expert demonstrations [6, 14, 51]. Despite their impressive ability to learn diverse behaviors directly from high-dimensional observations, these policies still degrade in nuanced and unexpected ways at runtime. For example, consider

the robot in the left of Figure 1 that must pick up a mug from the table. At training time, the generative policy learns a distribution over useful interaction modes such as grasping the cup by different parts (e.g., handle, lip and interior, etc.) shown in wrist camera photo in Figure 1.

However, at runtime, the policy exhibits a range of degradations, from complete task failures (such as the robot knocking down the cup during grasping, shown in the center of Figure 1), to inappropriate behaviors that are misaligned with the deployment context or preferences of an end-user (such as the robot placing its gripper inside of a cup of water when serving a guest shown in the middle of Figure 1). While a common mitigation strategy involves re-training the policy via more demonstrations [37] or interventions [34, 25], runtime failures are not always an indication that the base policy is fundamentally incapable of producing the desired behavior. In fact, the base policy may already contain the “right” behavior mode within its distribution (e.g., grasping the cup by the handle), but due to putting too much probability mass on an undesired mode, the robot does not reliably choose the correct action plan at runtime.

Runtime policy steering [27, 47] offers an elegant solution to this mode-selection problem. By using an external *verifier* to select candidate plans proposed by an imperfect generative policy, the robot’s behavior can be “steered” at runtime without any need for re-training. Despite the initial successes demonstrated by the policy steering paradigm, the question still remains of how to fully unlock autonomous policy steering in the open world when the robot’s environment, task context, and base policy’s performance are constantly changing.

Policy steering can be approached via the stochastic model-

<sup>1</sup>Here, multimodality only refers to the training data’s action distribution.

predictive control framework of modern control theory, which decomposes the optimal action selection (i.e. *generation*) problem of runtime policy steering into (a) *predicting* the outcomes of a given action plan and (b) *verifying* how well they align with user intent. However, this approach is only feasible when a physics-based, low-dimensional dynamics model is available for outcome prediction and a well-defined reward function can be specified for verification. In open-world environments, both of these requirements are challenging to fulfill due to the complexity of dynamics modeling and the difficulty of hand-crafting rewards to evaluate nuanced task requirements [17].

Our core idea is to leverage world models, which are well-suited for predicting action outcomes in open world settings, and VLMs, which have great potential as verifiers due to their commonsense reasoning abilities, to develop a divide-and-conquer approach to open-world policy steering. However, doing so naively is challenging as world models and VLMs operate on fundamentally different representations of reality.

To address this concern, we propose **FOREWARN: Filtering Options via REpresenting World-model Action Rollouts via Narration**. To predict challenging action outcomes (e.g., interaction dynamics of a manipulator and a deformable bag), we use state-of-the-art world models [25, 50] to predict lower-dimensional latent state representations from high-dimensional observation-action data (shown in orange in the center of Figure 1). To critique behavior outcomes under nuanced task specifications (e.g., “Serve the cup of water to the guest”), we leverage vision-language models (VLMs) [11, 29] as our open-world verifiers (shown in green in the center of Figure 1). Importantly, we demonstrate that *aligning* the VLM to reason directly about the predicted latent states from the world model enables it to understand fine-grained outcomes that it cannot directly predict zero-shot nor understand from image reconstructions. Ultimately, this alignment step enables our “VLM-in-the-loop” policy steering approach to interpret action plans as behavior narrations and select high-quality plans by reasoning over those narrations even under novel task descriptions (shown on the right of Figure 1).

We evaluate **FOREWARN** on a suite of robotic manipulation tasks, demonstrating how it can robustly filter proposed action plans to match user intent and task goals even when faced with variations not seen during training. In summary, our main contributions are:

- Formalizing runtime policy steering a stochastic model-predictive control problem, revealing the *generation-verification gap* [15, 9, 43] and where state-of-the-art models have maximal potential to shine.
- A latent space alignment strategy that enables a VLM to more reliably verify action plan outcomes predicted by a low-level, action-conditioned world model.
- A novel, fully-autonomous policy steering framework that improves a base generative imitation-based policy by over 30%, even in novel task descriptions.
- Extensive experiments in hardware showing that our latent-aligned VLM approach outperforms (by  $\sim 40\%$ ) alternative VLM approaches that do *not* decouple the

prediction of outcomes from verification.

## II. RELATED WORK

**Generative Imitation-Based Policies.** With the rise of large-scale open-source datasets of expert demonstrations [2, 4, 5, 8, 13, 44], imitation learning (IL) has become a popular way to learn low-level robot control policies from data. In particular, recent advances in generative modeling have unlocked policy architectures that can model diverse, multi-modal behaviors directly from high-dimensional observations [6, 23, 33]. At the same time, generative IL policies still exhibit nuanced, hard-to-anticipate performance degradations during deployment time. These degradations range from complete task failures (e.g., inability to grasp a cup, knocking it down, or dropping it [27]) potentially due to distribution shifts or visual distractors [19, 46], to inappropriate behaviors that are misaligned with the deployment context or an end-user’s preferences (e.g., placing the gripper inside of a cup filled with water during grasping) [1]. In this work, our goal is to leverage the diverse low-level behavior distribution that the base policy has learned, but prevent these nuanced performance degradations at runtime via our novel policy steering method.

**Failure Detection, Monitoring & Prediction.** The handling of generative policy failures can be grouped into three broad categories: *posthoc* detection, *runtime* monitoring, and *failure prediction*. *Posthoc* approaches identify and explain failures present in offline robot datasets, and have recently leveraged Vision Language Models (VLMs) to accelerate this process via video captioning, highlighting critical data frames, and providing human-interpretable summaries of failures [10, 16, 26, 48]. In contrast, *runtime* monitoring aims to detect failures as they happen during robot deployment. To quickly identify nuanced failures, recent methods propose a “fast and slow” approach: a fast online detector flags unusual situations (e.g., binary anomaly classifier), while a slower VLM-based reasoner provides a deeper understanding of the event and if it is a relevant failure [1, 39]. Although these strategies can effectively identify failures, they fundamentally require the robot to start failing for the runtime monitor to activate. The final category, *failure prediction* methods, anticipate failures before they occur and unlock the potential for preemptive correction of the base policy. Here, existing approaches [22, 24, 25] often rely on out-of-distribution (OOD) detection in a latent space or dense human labels to train a binary classifier that distinguishes failures from successes. In this work, we contribute to the *predictive* category of methods. Our method anticipates future outcomes of the policy’s actions via a latent world model, and reasons about the outcomes via a VLM that is aligned with the predicted latent states.

**Policy Steering.** A traditional method to improve a base IL policy is to fine-tune it with additional intervention data [25] or recovery behaviors [37]. However, recently, runtime policy steering has become an attractive alternative to improving a generalist IL policy [27, 47] without needing any additional and expensive demonstration data. Runtime policy steering

assumes that the base policy is capable of generating the correct behaviors, but fails to select them reliably. Here, an external verifier can be used to re-rank (i.e., “steer”) the generations towards ones with good outcomes. Previous methods have explored humans-in-the-loop [47] or a  $Q$ -function learned from very large offline datasets labeled with sparse rewards. [27] as the verifier. In our framework, by using a VLM-in-the-loop as our verifier, we can perform policy steering autonomously after finetuning VLM on a small dataset and provide human-like interpretable guidance.

**Learning to Search.** From an algorithmic perspective, our approach fits within the paradigm of *learning to search* (L2S) [31]. In L2S, one learns the components required to, at test time, plan a sequence of actions, rather than merely imitating what was in the training data. L2S provides two key advantages over direct imitation algorithms like behavioral cloning. First, the agent gains the ability to reason about the consequences of its actions and potentially recover from mistakes, avoiding *compounding errors* [35, 41]. We see this manifest in our system’s ability to avoid or correct from failures the base policy would have produced otherwise. Second, *verifying* whether a plan is good is often easier to learn than *generating* a good plan in the first place [43, 28]. We see this manifest in the fact that our system only requires a limited amount of data to fine-tune our verifier VLM, rather than the larger amount of data an end-to-end approach would have required [38].

In contrast to more classical L2S approaches that require extensive *global* search in the real world [31, 52], we instead perform *local* search [3] against a learned verifier [42, 12] inside a learned world model [32]. This allows us to avoid the computational expense of global search and the potential safety violations incurred by real-world interaction. As argued by [42, 32], doing so still matches many of the guarantees of classical L2S methods if one fits the world model on a mixture of on-policy and off-policy trajectories (i.e., in a *hybrid* fashion [40, 45]), as we do consistently across all of our experiments.

### III. PROBLEM FORMULATION

We formalize the general problem of policy steering as a stochastic model-predictive control problem over the set of action plans proposed by a base action generation model.

**Setup & Notation.** Let the robot’s pre-trained multimodal imitative action generation model [6, 23] be denoted by  $\pi(\mathbf{a}_t \mid o_t)$ . We will often refer to this model as the robot’s “base policy” throughout this paper with which it performs a task for an end-user. The robot’s observations  $o \in \mathcal{O} := \mathcal{I} \times \mathcal{Q}$  combine RGB image data  $I \in \mathcal{I}$  and proprioceptive states  $q \in \mathcal{Q}$  (e.g., end-effector pose, gripper state), and  $\mathbf{a}_t := a_{t:t+T}$  denotes a robot’s  $T$  step action plan, with each action in the sequence specifying end-effector positions and rotations. Similarly,  $\mathbf{o}_t := o_{t:t+T}$  is an observation sequence. In the real world, after observing some  $o_t$ , generating some action plan  $\mathbf{a}_t \sim \pi(\cdot \mid o_t)$  and executing it open-loop, the robot observes a sequence of observations  $\mathbf{o}_t \sim \mathcal{P}(\cdot \mid o_t, \mathbf{a}_t)$ .

**Problem.** Given the robot’s current observation  $o_t$  at timestep  $t$  and  $K$  i.i.d. samples from the base policy,  $\{\mathbf{a}_t^i\}_{i=1}^K \sim \pi(\mathbf{a}_t \mid o_t)$ , the *policy steering* problem seeks to return the action plan  $\mathbf{a}_t^*$  which optimizes the following objective:

$$\mathbf{a}_t^* = \arg \max_{\mathbf{a}_t \in \{\mathbf{a}_t^i\}_{i=1}^K} \mathbb{E}_{\mathbf{o}_t \sim \mathcal{P}(\cdot \mid o_t, \mathbf{a}_t)} [R(\mathbf{o}_t; \ell)]. \quad (1)$$

Fundamentally, solving the (*behavior*) *generation* problem specified in Eq.1 requires two abilities: *prediction* and *verification*. The *prediction* problem can be clearly seen in the expectation of Eq.1, wherein we have to forward simulate the outcomes of any action plan  $\mathbf{a}_t^i$  and “imagine” the potential future observations,  $\mathbf{o}_t \sim \mathcal{P}(\mathbf{o}_t \mid o_t, \mathbf{a}_t^i)$ . The *verification* problem lies inside of the expectation and with the reward function  $R(\mathbf{o}_t; \ell)$ . Here,  $\ell \in \mathcal{L}$  is the task description, represented via a language description (e.g., “Serve the cup of water to the guest.”), and  $R$  verifies how well or how poorly the future observations align with the behavior specified by  $\ell$ .

While Eq.1 characterizes the underlying policy steering problem, it is extremely challenging to solve end-to-end because of the coupled prediction and verification steps. This is where we seek to leverage vision-language models in-the-loop to obtain a practical solution with the potential to generalize to new environments and hard-to-model steering criteria. Initially, it may be tempting use the VLM directly as a black-box solver of Eq.1 (i.e. to solve the overarching behavior generation problem) by simply passing it the  $K$  action plan options, the current observation  $o_t$ , and the task description  $\ell$ , and having it directly return  $\mathbf{a}_t^*$ . However, low-level action data is beyond the training distribution of current VLMs, which primarily focus on high-level semantic understanding and not embodied control [21]. Alternatively, we could fine-tune the VLM to directly select action plans via labeled observation-action-samples datasets (labeled with optimal action trajectories). However, this strategy is sample-inefficient, requiring extensive embodied rollouts and human annotations to generate labels. Instead, we propose tackling the problem in Eq.1 in a way that leverages the unique strengths of a VLM. We hypothesize that the right place to put VLMs into the loop is specifically in the *verification* step above, as it leverages the model’s strong reasoning abilities *given predicted outcomes*. Then, to complement the VLM, we propose that *prediction* should be handled by an embodied world model that can reason about hard-to-model outcomes directly from low-level actions. It is well known that verification is significantly easier than generation for many problems [9, 43], boding well for the sample-efficiency of our modular approach.

### IV. OUR APPROACH: FOREWARN

Our key idea is to adopt a divide-and-conquer strategy, explicitly decoupling the VLM’s burden of predicting action outcomes from evaluation during policy steering. Specifically, we take advantage of recent advances in latent dynamics models [25, 50] which can learn lower-dimensional latent state representations from high-dimensional observation-action data collected on a robot. By passing possible action generations

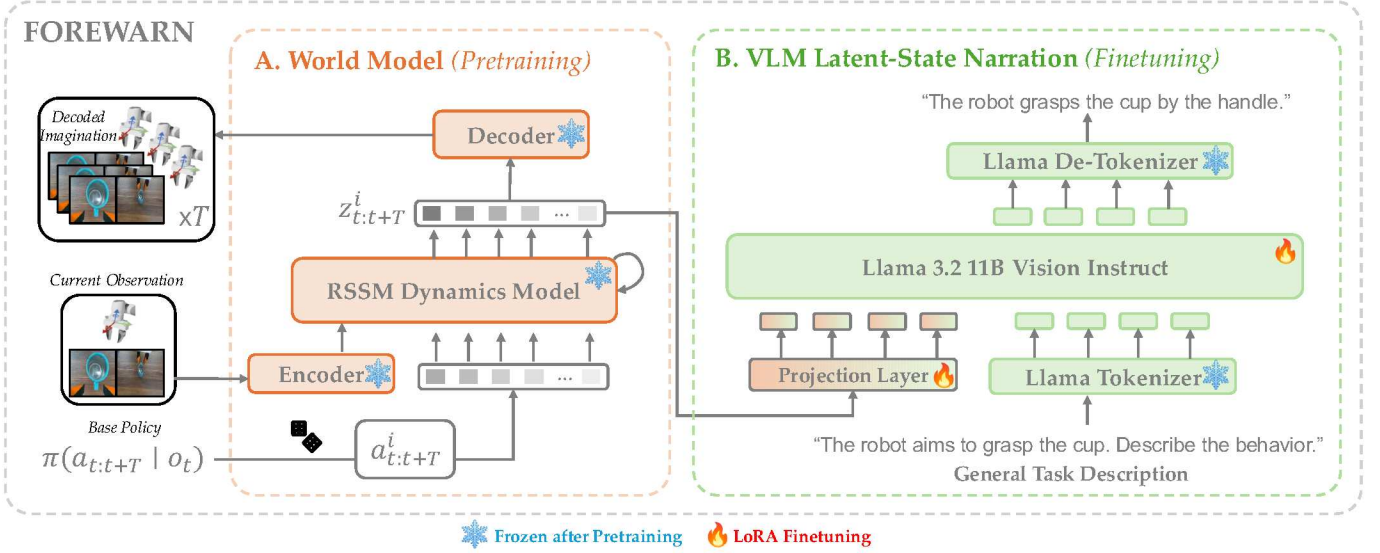


Fig. 2: **Training FOREWARN.** In part A (Sec. IV-A), a Recurrent State Space Model (RSSM) is pretrained to learn good latent embeddings of the dynamics conditioned on the observations and actions. In part B (Sec. IV-B), the sequence of learned latent embeddings is projected through a linear layer to the text embedding space, similar to the original vision token processing in the Llama-3.2 Model. The projection layer and Llama model are finetuned together using LoRA [20], but the world model remains frozen.

to the latent dynamics model, we can efficiently predict diverse future outcomes that would hard to model otherwise. Importantly, even though the world model learns a latent state that is an approximate sufficient statistic of the dynamics, one still needs to teach the VLM to evaluate outcomes *in the latent representation of the world model*, rather than from the raw image observations, as we will address in detail below.

At the highest level, our mathematical formulation of model-predictive “VLM-in-the-loop” policy steering is:

$$\begin{aligned} \mathbf{a}_t^* &= \arg \max_{\mathbf{a}_t \in \{\mathbf{a}_t^j\}_{j=1}^K} \mathbb{E}_{\mathbf{z}_t \sim f_\phi(\mathbf{z}_t, \mathbf{a}_t)} [R_\psi^{\text{VLM}}(\mathbf{z}_t; \ell)], \\ \text{s.t.} \quad & \mathbf{z}_t = \mathcal{E}_\phi(o_t), \end{aligned} \quad (2)$$

where  $\mathbf{z}_{t+1} \sim f_\phi(\mathbf{z}_t, \mathbf{a}_t)$  is a probabilistic latent dynamics model,  $\mathbf{z}_t := \mathbf{z}_{t:t+T}$ ,  $\mathcal{E}_\phi : \mathcal{O} \rightarrow \mathcal{Z}$  is an observation encoder that maps the robot’s current observations  $o_t$  into a latent space  $\mathcal{Z}$ , and  $R_\psi^{\text{VLM}}(\mathbf{z}_t; \ell)$  represents an implicit reward function embedded in our VLM (parameterized by  $\psi$ ) which evaluates the predicted outcomes given the task description. Both  $f_\phi$  and  $\mathcal{E}_\phi$  are part of our world model and parameterized by  $\phi$ . In short, to realize our idealized policy steering objective (1), we approximate the expectation over future outcomes with world model rollouts and leverage a fine-tuned VLM as a latent-outcome-conditioned verifier, leading to (2). We call our overall policy steering approach **FOREWARN: Filtering Options via REpresenting World-model Action Rollouts as Narration**, visualize it in Figure 1, and now discuss the details.

#### A. Foresight: Predicting Outcomes via Latent World Models

In this work, we adopt the Dreamerv3 world model [18] as our mechanism for predicting future outcomes of ac-

tion plans. This world model (visualized on the left in Fig. 2) is pre-trained via an offline dataset that contains trajectories of the robot’s observations and actions:  $\mathbb{D}_{\text{WM}} = \{ \{ (o_\tau^j, a_\tau^j) \}_{\tau=0}^{H-1} \}_{j=1}^M$ , where  $M$  is the total number of trajectories with each trajectory representing the robot’s outcome  $o_{0:H}^j$  induced by an action plan  $a_{0:H}^j$ , and  $H$  denotes the length of these trajectories. The training data consists of both successful and failed rollouts from the base policy  $\pi(\mathbf{a} | o)$  and additional demonstration data. This allows the world model to accurately predict the outcomes of both good and bad actions plans, a sufficient condition for successful behavior generation [32].

The world model training loss incentivizes the latent state to be informative for high-quality image decoding as well as highly predictive of the next latent state given an action (see [18] and Appendix A2 for more details). After training, the world model is frozen, and we utilize the trained  $\mathcal{E}_\phi$  and  $f_\phi$  throughout the rest of our policy steering algorithm.

#### B. Forethought: Latent-Text Alignment for Outcome Reasoning and Policy Steering

Once the world model encoder  $\mathcal{E}_\phi$  learns an effective low-dimensional latent state representation  $\mathbf{z}_t$  and the latent dynamics model  $f_\phi$  learns to predict future latent states conditioned on the robot’s actions, we can shift our focus to evaluating the imagined outcomes of any candidate low-level action plan. We hypothesize that the VLM’s open-world reasoning capabilities could allow it to be an effective verifier here, enabling better decision-making downstream.

However, before we can unlock the potential of the VLM as a verifier, we need to first enable the VLM to reason directly about the predicted latent states (rather than in terms of raw

image observations). We propose approaching this problem as a *latent-text alignment* problem: by mapping latent states generated by the world model to a textual representation, we can enable the VLM to more easily evaluate plans due to its strong natural language understanding abilities. We then further frame this alignment problem as a *visual question-answering (VQA)* task, where we prompt the VLM to narrate the real-world behavior of the robot conditioned on a sequence of latent states produced by the world model (visualized on the right of Figure 2). We now discuss this process in detail.

**VLM Backbone.** We use the Llama-3.2-11B-Vision-Instruct model [11] as our VLM backbone. The original Llama model utilizes a Vision Transformer (ViT) to tokenize visual inputs into latent tokens, which are then processed by the LLM backbone to generate text outputs. In our setting, we adapt the Llama model by replacing its observation (image) tokenization module with our world model’s encoder  $\mathcal{E}_\phi$  and latent dynamics model  $f_\phi$ . To align the latent dynamics embedding space with the text embedding space, we introduce a linear layer as an adapter. Specifically,  $\mathcal{E}_\phi$  encodes the current robot observation, and the forward dynamics model  $f_\phi$  predicts future latent states based on a candidate action plan (left side of Figure 2). These latent states are passed through the linear layer to produce a sequence of latent tokens, which serve as input to the LLM backbone (right side of Figure 2).

**VLM Finetuning.** Our objective is to enable the VLM to narrate the real-world behavior of the robot, denoted as  $\ell_b^i$ , based on a sequence of generated latent states,  $z_{t:t+T}^i$  from the world model. These behavior narrations highlight fine-grained motion details (e.g., “the robot grasps the cup by the handle”) rather than the high-level semantics of the motion (e.g., “the robot grasps the cup”). This design encourages the model to capture nuanced behaviors and identify potential failures in the robot’s predicted outcomes. We construct a VQA dataset to finetune the VLM to achieve this objective (described in Sec. V and App. A3).

**VLM-In-the-Loop Policy Steering.** Our fine-tuned VLM can now describe the outcomes of candidate action plan in natural language narration. However, our ultimate goal is to query the VLM to identify the best action sequence. To accomplish this, we propose querying the same VLM again (i.e. using it as a verifier), leveraging its open-world reasoning capabilities and natural language understanding to select the best action plan. Mathematically, policy steering can be expressed as:

$$\begin{aligned} \mathbf{a}_t^* &= \arg \max_{\mathbf{a}_t \in \{\mathbf{a}_t^i\}_{i=1}^K} \mathbb{E}_{\mathbf{z}_t \sim f_\phi(\mathbf{z}_t, \mathbf{a}_t)} \left[ R_\psi^{\text{VLM}}(\mathcal{T}_\psi^{\text{VLM}}(\mathbf{z}_t; \ell); \ell) \right], \\ \text{s.t.} \quad & z_t = \mathcal{E}_\phi(o_t). \end{aligned} \quad (3)$$

where  $\mathcal{T}_\psi^{\text{VLM}} : \mathcal{Z}^T \rightarrow \mathcal{L}_b$  denotes the inference process that translates a candidate action plan into the associated robot behavior narration. Instead of instructing the VLM to explicitly assign the reward for each predicted action plan outcome, we leverage its implicit knowledge about the reward and directly query the VLM for the best action plan among the

$K$  candidates, conditioned on the task description  $\ell$  (see the example in right side of Figure 1), which can be seen as a form of multiple-choice question answering (MCQA).

In summary, by integrating the predictions from the world model with the VLM’s open-vocabulary behavior narration generation and commonsense reasoning, *FOREWARN* guides the robot towards action plans that are aligned with the deployment context and task goals and enables the robot to proactively prevent failures.

## V. EXPERIMENTS

In this section, we instantiate several real-world manipulation tasks to study our method. We first investigate how effectively we can translate low-level actions into high-level behavior descriptions (Sec. V-A). Then we evaluate the closed-loop policy steering performance as well as our method’s robustness to novel task descriptions,  $\ell$  (Sec. V-B).

**Real Robot Setup.** We use a Franka Emika robotic manipulator equipped with a 3D printed gripper from [7] in our experiments. The base generative policy controls the low-level robot actions (which are the robot’s end-effector pose and gripper opening) controlled at 15Hz. The robot uses RGB images from a wrist-mounted camera and a third-person camera mounted in front of the robot. See App. B1 for more details.

**Manipulation Tasks.** We consider three real-world robot manipulation tasks that exhibit underlying multi-modal behaviors, hard-to-model outcomes, and nuanced failures. In the **Cup** task, the robot must grasp a cup placed in front of it and in the **Bag** task, the robot must pick up a bag of chips from the table. Fundamentally, both tasks can be accomplished in diverse ways: for example, for the **Cup** task, the robot can pick up the cup by the handle or by placing its fingers inside the cup by grasping the lip; for the **Bag** task, the robot can grab the bag by any edge, or by squeezing the center. Furthermore, the **Bag** task has more challenging dynamics and potential outcomes because the bag is deformable. We use this task to study how our framework performs when faced with harder-to-predict interaction outcomes and nuanced failures (e.g., crushing the chips inside the bag).

We also introduce a third task that features longer horizon and more complex interactions **Fork-to-Bowl Transfer**. In this task, the robot must pick up a fork from the table and place it inside a bowl. This task is considerably more challenging than the other two tasks for three reasons: 1) it requires longer-horizon planning to effectively navigate distinct phases; 2) requires reasoning about interactions between objects (the fork and the bowl); and 3) it introduces different multi-modal motion choices, including selecting the optimal picking location (tines or handle) and determining the appropriate placement strategy (inside versus outside the bowl, low versus high release). In the top row of Figure 3, we visualize the camera observations of the robot interacting with these items for Bag and Cup tasks.



**Base Policy.** For our multi-modal imitative action generation model,  $\pi(\mathbf{a}_t | o_t)$ , we use a Diffusion Policy [6] trained on 100 teleoperated demonstrations per task. The policy takes inputs from wrist and third-person cameras, along with proprioceptive states, to predict a distribution over  $T$ -step action plan, where  $T = 64$ . Please see App. A1 for more details.

**World Model Training.** We collected 250 real-world trajectories per task, including both successful and failed rollouts from the base policy, along with additional 100 demonstrations used in base policy, for training and evaluating the world model (300 for training and 50 for testing). The world model is trained to predict  $T = 64$  future latent states given the current  $o_t$  and an action plan  $\mathbf{a}_t$ .

**VLM Fine-tuning.** We construct our VQA dataset for fine-tuning from the same offline dataset,  $\mathbb{D}_{\text{WM}}$ , used to train the world model. For each  $T$ -step trajectory snippet  $\{(o_\tau^j, a_\tau^j)\}_{\tau=t}^{t+T}$  from the dataset, the encoder  $\mathcal{E}_\phi$  processes the initial observation  $o_t^j$  at timestep  $t$ , and the forward dynamics model  $f_\phi$  predicts latent states  $z_{t:t+T}^j$  based on the action plan  $a_{t:t+T}^j$ . We note that in our specific implementation, although  $f_\phi$  is a stochastic dynamics function, we use only the most likely prediction as the outcome associated with the action plan. To avoid “semantically repetitive” latent states between adjacent low-level actions, we downsample the  $T$ -step latent states to  $T/4$ . These downsampled latent states, along with the task description  $\ell$ , are provided as input to the VLM, and we manually annotate the corresponding behavior narrations,  $\ell_b \in \mathcal{L}_b$ , for the associated observations  $o_{t:t+T}^j$ . We fine-tune the model using the Low-Rank Adaptation (LoRA) technique [20], keeping both the encoder  $\mathcal{E}_\phi$  and the latent dynamics model  $f_\phi$  frozen during the fine-tuning process.

**VLM-In-the-Loop Policy Steering.** When deploying **FOREWARN** for run-time policy steering, we begin by sampling 100 action plans from the base policy and aggregating them into  $K = 6$  modes using the non-maximum suppression (NMS) scheme from [36]. These 6 aggregated action plans are then passed to **FOREWARN** for interpretation and evaluation. For each candidate action plan, we use only the most likely future latent state predictions from  $f_\phi$  as input to the VLM for reasoning.

#### A. From Action Rollouts to Behavior Narration

As discussed in Sec. IV-B, if we want to use the VLM as an open world verifier  $R_\psi^{\text{VLM}}(\cdot; \ell)$ , we need to enable the model to understand the underlying textual representation of low-level action outcomes. In this section, we study if our latent-aligned VLM can accurately describe the outcomes of low-level actions. We also compare our approach with several baselines to investigate the advantages of using an explicit world model for predicting action outcomes and decoding a robot’s action plans into behavior narrations.

**Baselines.** We compare our approach, **FOREWARN**, against four baselines (more implementation details provided in App. A4). (1) **FOREWARN-Oracle**, is an upper-bound on

our method’s performance assuming that we had access to *ground-truth* future observations (instead of relying on the latent dynamics  $f_\phi$  to predict future outcomes). This method uses the encoder  $\mathcal{E}_\phi$  on ground-truth future observations to get privileged (posterior) future latent states  $z_{t:t+T}$  as input for the VLM. (2) **VLM-Act**, which directly fine-tunes the original Llama-3.2-11B-Vision-Instruct model to generate behavior narrations end-to-end from the current observation  $o_t$  and an action plan  $a_{t:t+T}$  (represented as text), without explicitly predicting outcomes. (3) **VLM-Img**, which utilizes the *decoded* world model’s predictions (i.e., the predicted future visual observations) given a robot’s planned actions. We use GPT-4o [29] to process the predicted visual observations and generate behavior narrations in a zero-shot manner. (4) **VLM-Img-Oracle**, which is similar to **VLM-Img** but is an upper bound on this method by using ground-truth visual observations instead of predicted ones.

**Metrics.** We adopt the metrics from [10] to evaluate the alignment between predicted behavior narrations and ground-truth narrations: (1) **LLM Score**: A similarity score (ranging from 0 to 1) determined by the GPT-4o model. (2) **GT Accuracy**: A binary score (0 or 1) indicating whether the predictions match the ground-truth narrations, as determined by a human labeler (in this case, the authors). For further details on the motivation behind using these metrics for evaluation, please refer to App. B3.

	GT Accuracy $\uparrow$			LLM Score $\uparrow$		
	Cup	Bag	Average	Cup	Bag	Average
<b>FOREWARN-Oracle</b>	0.92 $\pm$ 0.02	0.77 $\pm$ 0.03	0.85 $\pm$ 0.03	0.86 $\pm$ 0.01	0.76 $\pm$ 0.03	0.81 $\pm$ 0.02
<b>FOREWARN (Ours)</b>	<b>0.87<math>\pm</math>0.02</b>	<b>0.75<math>\pm</math>0.03</b>	<b>0.82<math>\pm</math>0.03</b>	<b>0.82<math>\pm</math>0.02</b>	<b>0.72<math>\pm</math>0.02</b>	<b>0.76<math>\pm</math>0.02</b>
<b>VLM-Act</b>	0.37 $\pm$ 0.03	0.36 $\pm$ 0.06	0.37 $\pm$ 0.05	0.52 $\pm$ 0.05	0.50 $\pm$ 0.07	0.51 $\pm$ 0.06
<b>VLM-Img-Oracle</b>	0.61 $\pm$ 0.04	0.43 $\pm$ 0.03	0.52 $\pm$ 0.04	0.65 $\pm$ 0.02	0.62 $\pm$ 0.02	0.64 $\pm$ 0.02
<b>VLM-Img</b>	0.36 $\pm$ 0.05	0.33 $\pm$ 0.04	0.35 $\pm$ 0.05	0.56 $\pm$ 0.03	0.60 $\pm$ 0.04	0.58 $\pm$ 0.04

TABLE I: Alignment Between Predicted Behavior Narrations and Ground-Truth Narrations. **FOREWARN** outperforms all baselines across both tasks and achieves performance comparable to **FOREWARN-Oracle**, which has access to ground-truth action outcomes and represents the upper bound for our approach. We use 50 rollouts to evaluate the performance. For **FOREWARN**, **FOREWARN-Oracle** and **VLM-Act**, the mean and standard deviation are reported by running 3 seeds for the finetuning experiments while **VLM-Img** and **VLM-Img-Oracle**, report 3 queries of GPT-4o.

**Results: On the Value of Explicit Action Outcome Prediction.** Table I presents the GT Accuracy and LLM Score for our approach and each baseline. The results are averaged across 30 test rollouts for each task. The results show that **VLM-Act** performs poorly, achieving less than 50% GT Accuracy across all tasks. This underperformance is due to its inability to interpret low-level actions without the grounding provided by a world model’s future outcome predictions. In contrast, **FOREWARN**, which leverages an explicit world model, outperforms **VLM-Act** by over 50% on every task, despite both being fine-tuned on the same dataset. These results demonstrate that decoupling the VLM’s burden of predicting action outcomes enables the model to produce

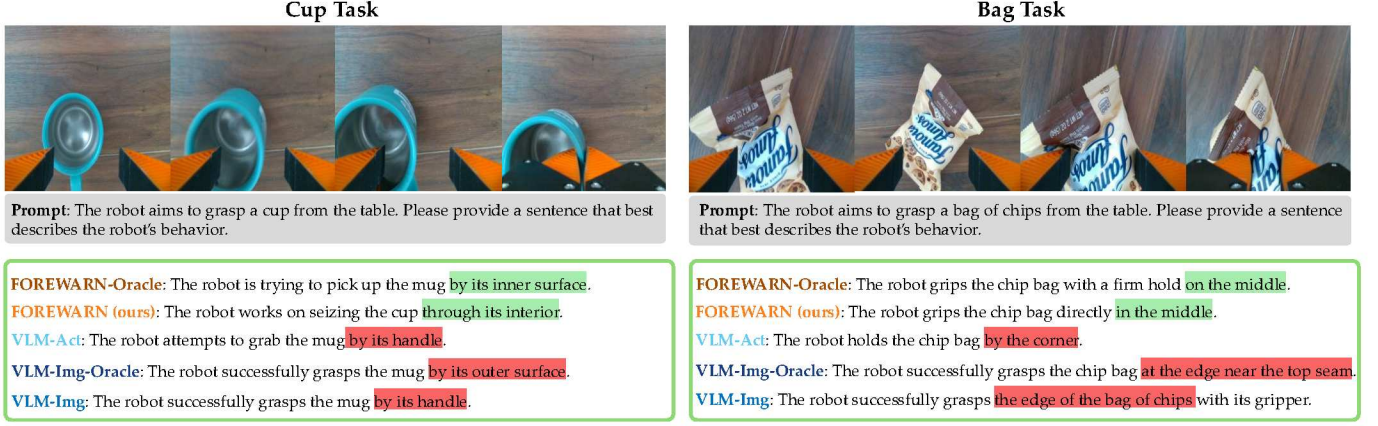


Fig. 3: **Examples of Behavior Narrations Predicted by Each Approach.** The top row displays the ground-truth robot observations and the prompt used for querying VLMs. Only **FOREWARN** and **FOREWARN-Oracle** consistently produce accurate outcome narrations, effectively capturing nuanced motion details. In contrast, the baselines frequently hallucinate or fail to capture critical contact details between the gripper and objects. For instance, in the **Bag** task, **VLM-Act**, **VLM-Img**, and **VLM-Img-Oracle** all hallucinate that the robot is grasping the edge of the bag, whereas it is actually grasping the middle.

more accurate outcome narrations than directly training the VLM to both predict outcomes and generate narrations end-to-end.

**Results: On the Value of VLM Fine-Tuning.** Interestingly, despite being among the most advanced VLMs, both **VLM-Img** and **VLM-Img-Oracle** struggle to accurately interpret robot behaviors directly from visual observations, even with access to ground-truth observations. As shown in Table I, these methods fall behind **FOREWARN** by at least 30% in **GT Accuracy** and 16% in **LLM Score**. These results show that existing state-of-the-art VLMs struggle to decode fine-grained motion details from video observations, underscoring the importance of fine-tuning for improved performance in such tasks.

**Results: Qualitative Examples.** In Figure 3, we visualize behavior narrations generated by our approach and the baselines. **FOREWARN** consistently produces more accurate outcome narrations, effectively capturing nuanced motion details. In contrast, the baselines often hallucinate or fail to capture critical contact details between the gripper and objects.

### B. Policy Steering for Open-World Alignment

The results from the previous section demonstrate our approach’s effectiveness in decoding predicted latent states into nuanced behavior narrations by explicitly using a world model for outcome prediction. In this section, we compare our approach against several baselines to evaluate its system-level policy steering performance and robustness under **novel** task descriptions and specifications. Our experiments focus on evaluating the impact of leveraging the VLM as both an interpreter and evaluator of predicted action outcomes.

**Baselines.** We first keep the VLM’s role as the verifier unchanged and ablate the effect of using an explicit world model to predict action outcomes. Specifically, we compare

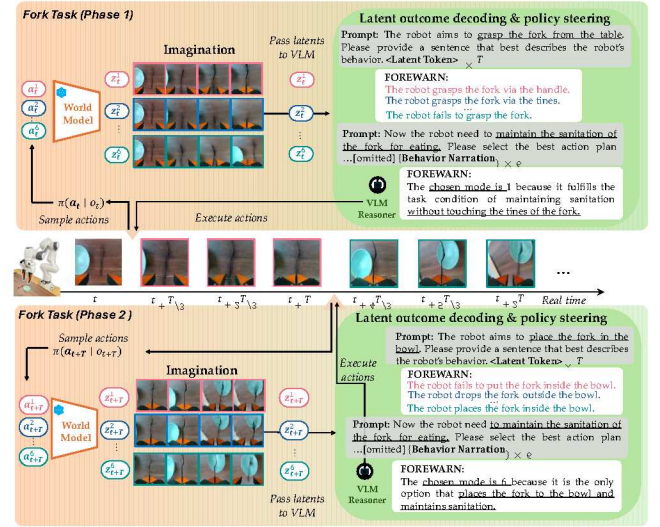


Fig. 4: **Policy Steering: Fork Task.** We visualize the steering process for the **Fork** task including two phases (Pick and Place). For each phase, we visualize the imagined  $T$ -step rollouts decoded from the world model for the 3 out of 6 action plans sampled from the base policy on the left. On the right, we show the behavior narrations generated from our finetuned VLM  $\mathcal{T}_{\psi}^{\text{VLM}}$  and the VLM’s reasoning  $R_{\psi}^{\text{VLM}}(\cdot; \ell)$  about the outcomes based on the task description  $\ell$  and behavior narrations to select the best action plan to execute. The time axis shows the real-world execution of the selected behavior from the perspective of the wrist-camera.

**FOREWARN** to **VLM-Act** (described in subsection V-A), which directly fine-tunes the original Llama model to predict action outcome narrations without utilizing a world model. Similar to our approach, **VLM-Act** then queries the VLM to select the best motion plan based on the task description  $\ell$ .



Next, we keep the world model unchanged and ablate the effect of alignment: mapping latent states generated by the world model to their underlying textual representation that the VLM can easily reason about.

We compare our approach against two more baselines: (3) **VLM-DynLat-Category**, which also leverages a world model for outcome predictions and a VLM for action plan selection. However, instead of decoding the predicted outcomes into text representations and leveraging the VLM’s open-world knowledge for policy steering, it directly fine-tunes the VLM to predict a set of indices of the successful candidate action plans and randomly selects one index from the set to execute the corresponding action plan. (4) **Classifier-Dyn-Latent**, which is similar to **VLM-DynLat-Category**, but instead of relying on a VLM, it directly takes the predicted latent embeddings  $\hat{z}_{t:t+T}$  as input and trains a transformer-based binary classifier (commonly used in prior failure-prediction work [25]) to predict success or failure for each candidate action plan and randomly selects one predicted to succeed.

**Metrics.** We evaluate each policy steering method using the **Success Rate**. For each method, we conduct 20 trials with randomly initialized task configurations and report the average success rate across these trials. A trial is considered successful if the robot successfully completes the task while aligning with the end-user’s preferences.

Method	Success Rate $\uparrow$					
	Training Task Description			Novel Task Description		
	Cup	Bag	Fork	Cup	Bag	Fork
<b>Base Policy</b>	0.30 $\pm$ 0.14	0.20 $\pm$ 0.13	0.10 $\pm$ 0.09	0.50 $\pm$ 0.16	0.40 $\pm$ 0.15	0.30 $\pm$ 0.14
<b>FOREWARN (Ours)</b>	<b>0.80<math>\pm</math>0.13</b>	<b>0.70<math>\pm</math>0.14</b>	<b>0.70<math>\pm</math>0.14</b>	<b>0.80<math>\pm</math>0.13</b>	<b>0.70<math>\pm</math>0.14</b>	<b>0.60<math>\pm</math>0.15</b>
<b>VLM-DynLat-Category</b>	<b>0.80<math>\pm</math>0.13</b>	0.40 $\pm$ 0.15	0.50 $\pm$ 0.16	0.30 $\pm$ 0.14	0.40 $\pm$ 0.15	0.30 $\pm$ 0.14
<b>Classifier-Dyn-Latent</b>	<b>0.80<math>\pm</math>0.13</b>	<b>0.70<math>\pm</math>0.14</b>	<b>0.70<math>\pm</math>0.14</b>	0.00 $\pm$ 0.00	0.10 $\pm$ 0.09	0.20 $\pm$ 0.13
<b>VLM-Act</b>	0.40 $\pm$ 0.15	0.20 $\pm$ 0.13	0.20 $\pm$ 0.13	0.30 $\pm$ 0.14	0.50 $\pm$ 0.16	0.20 $\pm$ 0.13

TABLE II: **Policy Steering.** The success rate is reported by averaging over 20 different rollouts. **FOREWARN** outperforms all the baselines in both training and novel task contexts.

**Results: Policy Steering Performance.** We first evaluate our approach and the baselines under task descriptions  $\ell$  that fall within the training distribution. Table II shows the success rates of the base robot policy (without policy steering), our approach (**FOREWARN**), and the baselines for each task. Our results demonstrate that **FOREWARN** can effectively steer the policy towards safe and aligned behavior modes by leveraging the VLM as an interpreter and evaluator of predicted latent action outcomes.

Interestingly, **Classifier-Dyn-Latent** achieves comparable performance to our approach in all tasks under this setting while **VLM-DynLat-Category** fails to match **FOREWARN** in more complicated **Bag** and **Fork** tasks. **VLM-Act** has similar performance as base policy because the behavior narrations, directly generated from low-level action sequences, fail to capture accurate motion details and thus provide no useful signal for VLM to steer the policy. Next, we assess their generalization abilities under novel task descriptions to evaluate the robustness of our approach.

**Results: Robustness to Novel Task Descriptions.** We evaluate the success rate of each approach when the task description is altered to introduce novel scenarios. Specifically, in the **Cup** task, we modify the original task description from “Please grasp a cup from the table and serve the cup of water to the guest” to a novel description: “Please grasp a cup from the table, but note that the handle is covered with oil.” This change makes behaviors where the robot grasps the cup by the handle unsafe, while grasping the cup by the rim becomes the desired behavior. In the **Bag** task, we modify the original task description from “Please pick up a bag of chips from the table and minimize the contact region to avoid crushing contents inside” to a novel description: “Please pick up a bag of chips from the table and maximize stability without dropping the bag.” This change makes behaviors where the robot picks up the bag by the corner less preferred, as the bag may slip from the gripper, while squeezing the middle of the bag to secure it becomes the desired behavior. In the **Fork** task, we modify the original prompt from “Please pick up the fork and place it in the bowl while maintaining the sanitation for eating.” to a novel description: “Please pick up the fork and place it in the bowl while maximizing the contact region during grasping.” This change makes the behaviors where the robot picks up the handle and drop the fork high less preferred as the fork handle is more narrow than tines and dropping high can make fork fall out of the bowl while grasping by tines is more secure.

Interestingly, while both **VLM-DynLat-Category** and **Classifier-Dyn-Latent** improve task success rates when task descriptions fall within the training distribution, their performance deteriorates significantly with novel task descriptions. This indicates that the VLM struggles to reason directly about predicted action outcomes from the world model’s latent states and essentially degrades to a traditional end-to-end model. To fully leverage the VLM’s open-world reasoning capabilities for generalized policy steering, it is essential to enable the model to interpret predicted action outcomes through textual representations.

**Results: Qualitative Examples.** In Figure 4, we present examples of runtime policy steering using our approach for the **Fork** task and additional examples for **Cup** and **Bag** tasks are included in Appendix B2. **FOREWARN** consistently demonstrates superior performance by selecting motion plans that align with task descriptions and user preferences, even in scenarios with novel task specifications. In contrast, the baselines either fail to interpret action outcomes effectively, resulting in unsafe behaviors, or experience severe performance degradation in novel task specifications.

**Results: Out-of-distribution Generalization.** Table II has shown our system’s robustness to novel task descriptions. Moreover, we further test our system’s generalization capability by adding variations in the environment. Specifically, we study the variations in object appearances including colors and sizes and variations in background. Among the six tested scenarios shown in Fig 5, our method can generalize to those variations with small performance drop.






Fig. 5: **Generalization to Environmental Changes.** For each task, we test our method against similar objects of different colors and sizes and also change the table cover to test the system against background variations.

**Results: Policy Steering Speed.** Our system queries the VLM twice to first generate behavior narrations and then select the best action plan. The overall inference time is 3.7 seconds among which the generation of behavior narrations for 6 candidate action plans takes 1.3 seconds. In comparison, **VLM-Act** takes 22.0 seconds in total for VLM inference and behavior narration runs 19.7 seconds, much slower than **FOREWARN**, partially due to the usage of token per patch in images rather than a single token per state (image) like **FOREWARN**. Additionally, our world model and VLM communicate directly in latent space, avoiding image decoding from world model and encoding from VLM, which could further speeds up inference.

Method	Time (Seconds) ↓			
	World Model Prediction	Behavior Narration	Evaluation	Total
<b>FOREWARN</b>	0.1	1.3	2.3	3.7
<b>VLM-Act</b>	-	19.7		22.0

TABLE III: **Inference time for the Policy Steering System.** Inference time for each component in the system (averaged across 3 runs) shows that **FOREWARN** greatly reduces the time to generate behavior narrations from our modified VLM compared to **VLM-Act**.

**Supplementary Experiments & Analyses.** To provide an in-depth analysis of our approach, we performed additional studies on the impact of each component on the overall system and a detailed comparison with baselines. We also showcase an additional application of our system as a runtime monitor which returns if a single action plan is good or bad, opening up potential future directions for soliciting supervisor assistance. All these experiments can be found in App. B2.

## VI. LIMITATIONS

While **FOREWARN** exhibits strong policy steering across diverse task settings, it is not without limitations. First, it assumes base policy is sufficiently competent—i.e., already

containing the correct behavior. Future work should investigate how to detect if none of the policy’s generated action plans are suitable for the deployment context, and how to improve the base policy via targeted fine-tuning data.

A second limitation lies in modeling real-world interaction dynamics. Our component-level analysis in App. B2 revealed that our system’s primary failures stem from the world model’s imprecise “imagination”, exacerbated by our limited training data. More advanced visual features (e.g., DINO features [30]) might improve the world model’s robustness to visual distractors. Ongoing research on large-scale, generalizable world models [49] for manipulation may also inform future extension of this framework.

Another bottleneck of the current system is the inference overhead, a common challenge for large autoregressive models like LLMs and world models. Some approaches to this include: 1) hierarchical decomposition of **FOREWARN** running high-level reasoning at low frequency while maintaining low-level control at high frequency like the common practice of hierarchical Vision Language Action Models (VLAs), 2) using advancement in quantization and caching techniques, 3) distilling our latent-aligned VLM into a model with smaller size, 4) strategically allocate inference time by identifying keypoints when it is “worth” reasoning for longer or deliberately.

Finally, our VLM is built upon an open-source Llama-3.2-11B-Vision-Instruct model, whose visual reasoning capabilities lag behind its language-based commonsense reasoning. A stronger backbone could yield higher-quality behavior descriptions, especially if it excels at video captioning tasks.

## VII. CONCLUSION

In this work, we investigate the problem of policy steering for multi-modal generative policies. We propose a novel framework **FOREWARN** that unlocks Vision Language Models (VLMs) to serve as open-vocabulary verifiers for run-time policy steering. Our method decouples the steering problem as future prediction (foresight) and outcome assessment (forethought). Through an explicit world model and a latent-space alignment strategy, we enable VLMs to reason about sensorimotor data using natural language. Our experiments across diverse manipulation tasks confirm that **FOREWARN** not only provides interpretable and reliable failure detection, but also significantly enhances policy success rates through flexible, generalizable steering. These results highlight the promise of combining learned dynamics models with language-based reasoning to improve test-time performance of robot policies.

## ACKNOWLEDGMENTS

This work is funded in part by a Google Research Scholar Award. Gokul Swamy is supported by STTR grant. We also want to thank Junwon Seo for providing feedback for the paper and Michelle Zhao, Pranay Gupta, Kensuke Nakamura, Lasse Peters for the helpful discussions of ideas.

## REFERENCES

- [1] Christopher Agia, Rohan Sinha, Jingyun Yang, Ziang Cao, Rika Antonova, Marco Pavone, and Jeannette Bohg. Unpacking failure modes of generative policies: Runtime monitoring of consistency and progress. In *8th Annual Conference on Robot Learning*, 2024.
- [2] anonymous authors. Anonymous title. In *Robotics: Science and Systems*, 2024.
- [3] James Bagnell, Sham M Kakade, Jeff Schneider, and Andrew Ng. Policy search by dynamic programming. *Advances in neural information processing systems*, 16, 2003.
- [4] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Jasmine Hsu, Julian Ibarz, Brian Ichter, Alex Irpan, Tomas Jackson, Sally Jesmonth, Nikhil Joshi, Ryan Julian, Dmitry Kalashnikov, Yuheng Kuang, Isabel Leal, Kuang-Huei Lee, Sergey Levine, Yao Lu, Utsav Malla, Deeksha Manjunath, Igor Mordatch, Ofir Nachum, Carolina Parada, Jodilyn Peralta, Emily Perez, Karl Pertsch, Jor-nell Quiambao, Kanishka Rao, Michael Ryoo, Grecia Salazar, Pannag Sanketi, Kevin Sayed, Jaspiar Singh, Sumedh Sontakke, Austin Stone, Clayton Tan, Huong Tran, Vincent Vanhoucke, Steve Vega, Quan Vuong, Fei Xia, Ted Xiao, Peng Xu, Sichun Xu, Tianhe Yu, and Brianna Zitkovich. Rt-1: Robotics transformer for real-world control at scale. In *arXiv preprint arXiv:2212.06817*, 2022.
- [5] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, Pete Florence, Chuyuan Fu, Montse Gonzalez Arenas, Keerthana Gopalakrishnan, Kehang Han, Karol Hausman, Alex Herzog, Jasmine Hsu, Brian Ichter, Alex Irpan, Nikhil Joshi, Ryan Julian, Dmitry Kalashnikov, Yuheng Kuang, Isabel Leal, Lisa Lee, Tsang-Wei Edward Lee, Sergey Levine, Yao Lu, Henryk Michalewski, Igor Mordatch, Karl Pertsch, Kanishka Rao, Krista Reymann, Michael Ryoo, Grecia Salazar, Pannag Sanketi, Pierre Sermanet, Jaspiar Singh, Anikait Singh, Radu Soricut, Huong Tran, Vincent Vanhoucke, Quan Vuong, Ayzaan Wahid, Stefan Welker, Paul Wohlhart, Jialin Wu, Fei Xia, Ted Xiao, Peng Xu, Sichun Xu, Tianhe Yu, and Brianna Zitkovich. Rt-2: Vision-language-action models transfer web knowledge to robotic control. In *arXiv preprint arXiv:2307.15818*, 2023.
- [6] Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research*, 2024.
- [7] Cheng Chi, Zhenjia Xu, Chuer Pan, Eric Cousineau, Benjamin Burchfiel, Siyuan Feng, Russ Tedrake, and Shuran Song. Universal manipulation interface: In-the-wild robot teaching without in-the-wild robots. In *Proceedings of Robotics: Science and Systems (RSS)*, 2024.
- [8] AgiBot World Colosseum contributors. Agibot world colosseum. <https://github.com/OpenDriveLab/Agibot-World>, 2024.
- [9] Stephen A Cook. The complexity of theorem-proving procedures. In *Logic, automata, and computational complexity: The works of Stephen A. Cook*, pages 143–152. 2023.
- [10] Jiafei Duan, Wilbert Pumacay, Nishanth Kumar, Yi Ru Wang, Shulin Tian, Wentao Yuan, Ranjay Krishna, Dieter Fox, Ajay Mandlekar, and Yijie Guo. Aha: A vision-language-model for detecting and reasoning over failures in robotic manipulation. In *International Conference on Learning Representations*, 2025.
- [11] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- [12] Nicolas Espinosa-Dice, Sanjiban Choudhury, Wen Sun, and Gokul Swamy. Efficient imitation under misspecification. *arXiv preprint arXiv:2503.13162*, 2025.
- [13] Hao-Shu Fang, Hongjie Fang, Zhenyu Tang, Jirong Liu, Junbo Wang, Haoyi Zhu, and Cewu Lu. Rh20t: A robotic dataset for learning diverse skills in one-shot. In *RSS 2023 Workshop on Learning for Task and Motion Planning*, 2023.
- [14] Zipeng Fu, Tony Z. Zhao, and Chelsea Finn. Mobile aloha: Learning bimanual mobile manipulation with low-cost whole-body teleoperation. In *Conference on Robot Learning (CoRL)*, 2024.
- [15] Kurt Godel. Letter to john von neumann, 1956. URL <https://ecommons.cornell.edu/server/api/core/bitstreams/46ae9c4-288b-457d-ab3e-bb6cb1a4b88e/content>.
- [16] Lin Guan, Yifan Zhou, Denis Liu, Yantian Zha, Heni Ben Amor, and Subbarao Kambhampati. Task success is not enough: Investigating the use of video-language models as behavior critics for catching undesirable agent behaviors. In *First Conference on Language Modeling*, 2024. URL <https://openreview.net/forum?id=otKo4zFKmH>.
- [17] Dylan Hadfield-Menell, Smitha Milli, Pieter Abbeel, Stuart J Russell, and Anca Dragan. Inverse reward design. *Advances in neural information processing systems*, 30, 2017.
- [18] Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. Mastering diverse domains through world models. *arXiv preprint arXiv:2301.04104*, 2023.
- [19] Asher J Hancock, Allen Z Ren, and Anirudha Majumdar. Run-time observation interventions make vision-language-action models more visually robust. *arXiv preprint arXiv:2410.01971*, 2024.
- [20] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language

- models. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=nZeVKeeFYf9>.
- [21] Yafei Hu, Quanting Xie, Vidhi Jain, Jonathan Francis, Jay Patrikar, Nikhil Keetha, Seungchan Kim, Yaqi Xie, Tianyi Zhang, Hao-Shu Fang, et al. Toward general-purpose robots via foundation models: A survey and meta-analysis. *arXiv preprint arXiv:2312.08782*, 2023.
- [22] Motonari Kambara and Komei Sugiura. Future success prediction in open-vocabulary object manipulation tasks based on end-effector trajectories, 2025. URL <https://arxiv.org/abs/2412.19112>.
- [23] Seungjae Lee, Yibin Wang, Haritheja Etukuru, H Jin Kim, Nur Muhammad Mahi Shafiullah, and Lerrel Pinto. Behavior generation with latent actions. In *Forty-first International Conference on Machine Learning*, 2024.
- [24] Huihan Liu, Shivin Dass, Roberto Martín-Martín, and Yuke Zhu. Model-based runtime monitoring with interactive imitation learning. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.
- [25] Huihan Liu, Yu Zhang, Vaarij Betala, Evan Zhang, James Liu, Crystal Ding, and Yuke Zhu. Multi-task interactive robot fleet learning with visual world models. In *8th Annual Conference on Robot Learning*, 2024.
- [26] Zeyi Liu, Arpit Bahety, and Shuran Song. Reflect: Summarizing robot experiences for failure explanation and correction. In *Conference on Robot Learning*, pages 3468–3484. PMLR, 2023.
- [27] Mitsuhiko Nakamoto, Oier Mees, Aviral Kumar, and Sergey Levine. Steering your generalists: Improving robotic foundation models via value guidance. *Conference on Robot Learning (CoRL)*, 2024.
- [28] Andrew Y Ng et al. Algorithms for inverse reinforcement learning.
- [29] OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rameesh Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Madeline Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan



- Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. Gpt-4 technical report, 2024. URL <https://arxiv.org/abs/2303.08774>.
- [30] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy V Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel HAZIZA, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *Transactions on Machine Learning Research*, 2023.
  - [31] Nathan D Ratliff, David Silver, and J Andrew Bagnell. Learning to search: Functional gradient techniques for imitation learning. *Autonomous Robots*, 27:25–53, 2009.
  - [32] Juntao Ren, Gokul Swamy, Zhiwei Steven Wu, J Andrew Bagnell, and Sanjiban Choudhury. Hybrid inverse reinforcement learning. *arXiv preprint arXiv:2402.08848*, 2024.
  - [33] Moritz Reuss, Ömer Erdiñç Yağmurlu, Fabian Wenzel, and Rudolf Lioutikov. Multimodal diffusion transformer: Learning versatile behavior from multimodal goals. In *Robotics: Science and Systems*, 2024.
  - [34] Stéphane Ross and Drew Bagnell. Efficient reductions for imitation learning. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 661–668. JMLR Workshop and Conference Proceedings, 2010.
  - [35] Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635. JMLR Workshop and Conference Proceedings, 2011.
  - [36] Ari Seff, Brian Cera, Dian Chen, Mason Ng, Aurick Zhou, Nigamaa Nayakanti, Khaled S Refaat, Rami Al-Rfou, and Benjamin Sapp. Motionlm: Multi-agent motion forecasting as language modeling. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8579–8590, 2023.
  - [37] Nur Muhammad Mahi. Shafiullah, Siyuan. Feng, Lerrel. Pinto, and Russ. Tedrake. Supervised policy learning for real robots, July 2024. URL <https://supervised-robot-learning.github.io>. Tutorial presented at the Robotics: Science and Systems (RSS), Delft.
  - [38] Shai Shalev-Shwartz and Amnon Shashua. On the sample complexity of end-to-end training vs. semantic abstraction training. *arXiv preprint arXiv:1604.06915*, 2016.
  - [39] Rohan Sinha, Amine Elhafsi, Christopher Agia, Matthew Foutter, Ed Schmerling, and Marco Pavone. Real-time anomaly detection and reactive planning with large language models. In *Robotics: Science and Systems*, 2024.
  - [40] Yuda Song, Yifei Zhou, Ayush Sekhari, J Andrew Bagnell, Akshay Krishnamurthy, and Wen Sun. Hybrid rl: Using both offline and online data can make rl efficient. *arXiv preprint arXiv:2210.06718*, 2022.
  - [41] Gokul Swamy, Sanjiban Choudhury, J Andrew Bagnell, and Steven Wu. Of moments and matching: A game-theoretic framework for closing the imitation gap. In *International Conference on Machine Learning*, pages 10022–10032. PMLR, 2021.
  - [42] Gokul Swamy, David Wu, Sanjiban Choudhury, Drew Bagnell, and Steven Wu. Inverse reinforcement learning without reinforcement learning. In *International Conference on Machine Learning*, pages 33299–33318. PMLR, 2023.
  - [43] Gokul Swamy, Sanjiban Choudhury, Wen Sun, Zhiwei Steven Wu, and J Andrew Bagnell. All roads lead to likelihood: The value of reinforcement learning in fine-tuning. *arXiv preprint arXiv:2503.01067*, 2025.
  - [44] Open X-Embodiment Team. Open X-Embodiment: Robotic learning datasets and RT-X models. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.
  - [45] Anirudh Vemula, Yuda Song, Aarti Singh, Drew Bagnell, and Sanjiban Choudhury. The virtues of laziness in model-based rl: A unified objective and algorithms. In *International Conference on Machine Learning*, pages 34978–35005. PMLR, 2023.
  - [46] Joseph A. Vincent, Haruki Nishimura, Masha Itkina, Paarth Shah, Mac Schwager, and Thomas Kollar. How generalizable is my behavior cloning policy? a statistical approach to trustworthy performance evaluation. *IEEE Robotics and Automation Letters*, 9(10):8619–8626, 2024. doi: 10.1109/LRA.2024.3445635.
  - [47] Yanwei Wang, Lirui Wang, Yilun Du, Balakumar Sundaralingam, Xuning Yang, Yu-Wei Chao, Claudia Perez-D’Arpino, Dieter Fox, and Julie Shah. Inference-time policy steering through human interactions. *arXiv preprint arXiv:2411.16627*, 2024.
  - [48] Zihan Wang, Brian Liang, Varad Dhat, Zander Brumbaugh, Nick Walker, Ranjay Krishna, and Maya Cakmak. I can tell what i am doing: Toward real-world natural language grounding of robot experiences. In *8th Annual Conference on Robot Learning*, 2024.
  - [49] Jialong Wu, Shaofeng Yin, Ningya Feng, Xu He, Dong Li, Jianye Hao, and Mingsheng Long. ivideopt: Interactive videopts are scalable world models. In *Advances in Neural Information Processing Systems*, 2024.
  - [50] Philipp Wu, Alejandro Escontrela, Danijar Hafner, Pieter Abbeel, and Ken Goldberg. Daydreamer: World models for physical robot learning. In *Conference on robot learning*, pages 2226–2240. PMLR, 2023.
  - [51] Tony Z. Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual manipulation with low-cost hardware, 2023. URL <https://arxiv.org/abs/2304.13705>.
  - [52] Brian D Ziebart, J Andrew Bagnell, Anind K Dey, et al. Maximum entropy inverse reinforcement learning. 2008.