

GRAPPA: Generalizing and Adapting Robot Policies via Online Agentic Guidance

Author Names Omitted for Anonymous Review. Paper-ID 176.

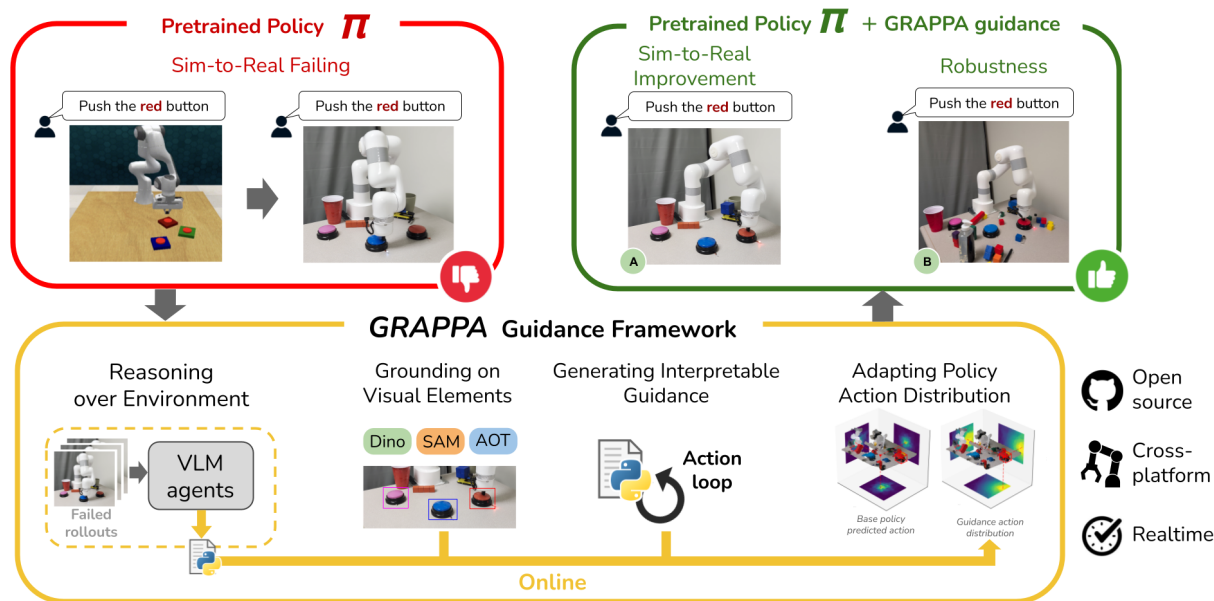


Fig. 1: **GRAPPA framework.** An overview of the proposed self-improving framework where the pretrained policy is updated by visuomotor-grounded guidance at test time. GRAPPA can guide and provide robustness to policies under two main failure modes: A) when a pretrained policy with limited to no access to real data is applied to an embodied system and B) there is a lot of noise, in the form of clutter, affecting the observations of the base policy.

Abstract—Robot learning approaches such as behavior cloning and reinforcement learning have shown great promise in synthesizing robot skills from human demonstrations in specific environments. However, these approaches often require task-specific demonstrations or designing complex simulation environments, which limits the development of generalizable and robust policies for unseen real-world settings. Recent advances in the use of foundation models for robotics (e.g., LLMs, VLMs) have shown great potential in enabling systems to understand the semantics in the world from large-scale internet data. However, it remains an open challenge to use this knowledge to enable robotic systems to understand the underlying dynamics of the world, to generalize policies across different tasks, and to adapt policies to new environments. To alleviate these limitations, we propose an agentic framework for robot self-guidance and self-improvement, which consists of a set of role-specialized conversational agents, such as a high-level advisor, a grounding agent, a monitoring agent, and a robotic agent. Our framework iteratively grounds a base robot policy to relevant objects in the environment and uses visuomotor cues to shift the action distribution of the policy to more desirable states, online, while remaining agnostic to the subjective configuration of a given robot hardware platform. We demonstrate that our approach can effectively guide manipulation policies to achieve significantly higher success rates, both in simulation and in real-world experiments, without the need for

additional human demonstrations or extensive exploration.

I. INTRODUCTION

In recent years, the advent of foundation models, such as pre-trained large language models (LLMs) and vision language models (VLMs), has enabled impressive capabilities in understanding context, scenes, and dynamics in the world. Furthermore, emergent capabilities such as in-context learning have shown great potential in the transfer of knowledge between domains, e.g., via few-shot demonstrations or zero-shot inference. However, the application of these models to robotics is still limited, given the intrinsic complexity and scarcity of human-robot interactions and the lack of large-scale datasets of human-annotated data or demonstrations [12, 40].

Approaches that leverage LLMs and VLMs alongside robotics systems often fall into one of two categories. The first category is that of using foundation models as zero-shot planners, code-generators, and task-descriptors—all of which attempt to use the foundation model to provide some high-level instruction to a low-level policy or to describe the policy’s actions in natural language [32, 34, 2, 25] or generated

plans as code [22, 35]. While these works have illustrated impressive reasoning capabilities, they still required either learning additional mappings to interface the generated natural language instructions with the low-level policy or, in the case of the code generation examples, they rely on pre-existing handcrafted primitives to compose. For the second category, in the use of foundation models for robot learning, methods may fine-tune the foundation models to improve in-domain performance by imparting strong priors for effective, albeit platform-dependent task-execution [3, 36, 19]. Here, parts of a robot policy or the entire policy is supervised to perform certain prediction tasks or to learn desired behaviors, end-to-end. While this strategy has led to impressive capabilities on a variety of tasks, methods often struggle to generalize to different object categories, tasks, and environments that are outside the data distribution that the models were trained on, or the models struggle with issues of negative transfer when trained on very diverse collections of robot data [6, 18].

In this paper, we extend the deployability of robot policies by using foundation models to generate low-level visuomotor guidance to handle out-of-distribution scenarios, such as sim-to-real differences or new tasks and robot platform variations (Figure 2). We design an agent-based framework, where a team of conversational agents works together to refine the action distribution of a robot’s base policy, via grounded visuomotor guidance (Figure 3). As illustrated in Figure 4, and upon a request from the advisor agent, the grounding agent can look for objects through a combination of detection, tracking, and high-level reasoning, to broaden or restrict the search as the context demands. When looking for a mug, for example, the agents could first look for a semantically-relevant receptacle that likely contains the mug (the cupboard) and narrow the search from there. Once the target object has been visually located, the Monitor and Advisor agents generate a guidance function that outputs a biasing guidance distribution, which, when combined with the action distribution of the Robotic agent (policy), ensures that the policy can complete task-relevant behaviors until it succeeds. In this way, the agentic framework bridges high-level reasoning with low-level control, enabling systems to reason about failure and self-guide.

In simulated and real-world experiments, we comprehensively and empirically demonstrate that our framework for **Generalizing and Adapting Robot Policies via Online Agentic Guidance (GRAPPA)** provides robustness for a variety of representative base policy classes to sim-to-real transfer paradigms and out-of-distribution settings, such as unseen objects.

In summary, our paper provides the following contributions:

- We propose GRAPPA, an agentic robot policy framework that self-improves by generating a guidance function to update base policies’ action distributions, online. Our framework is capable of learning skills from scratch after deployment and generalizes across various base policy classes, tasks, and environments.
- For robustness against cluttered and unseen environments, we propose a grounding agent that performs multi-granular object search, which enables flexible visuomotor grounding.



Fig. 2: Heatmap visualization of the guidance distribution, generated online by our proposed agentic framework, which produces code that biases a robot policy’s action distribution towards desirable behavior.

- We provide experimental results showcasing GRAPPA as a helper tool for aiding in the sim-to-real transfer of policies without the need for extensive data collection.

II. RELATED WORK

Vision-Language Models for Robot Learning: Several works explore the notion of leveraging pre-trained or fine-tuned Large Language Models (LLMs) and/or Vision-Language Models (VLMs) for high-level reasoning and planning in robotics tasks [12, 2, 22, 14, 35, 13, 10, 7, 29, 26, 21, 30]—typically decomposing high-level task specification into a series of smaller steps or action primitives, using system prompts or in-context examples to enable powerful chain-of-thought reasoning techniques. This strategy of encouraging models to reason in a stepwise manner before outputting a final answer has led to significant performance improvements across several tasks and benchmarks [12]. Despite these promising achievements, these approaches rely on handcrafted primitives [2, 14, 22, 29], struggle with low-level control, or require large datasets for retraining. Furthermore, various approaches that leverage VLMs for robot learning suffer from a granularity problem when using off-the-shelf models in a single-step/zero-shot manner [15] or are unable to perform failure correction without costly human intervention [14, 29, 23, 15]. In contrast, our framework bridges high-level reasoning with low-level control, by leveraging an agentic framework for online modification of a base policy’s action distribution at test-time, without requiring human feedback or datasets for fine-tuning. Moreover, we mitigate the granularity problem by proposing a flexible and recursive grounding mechanism that uses VLMs

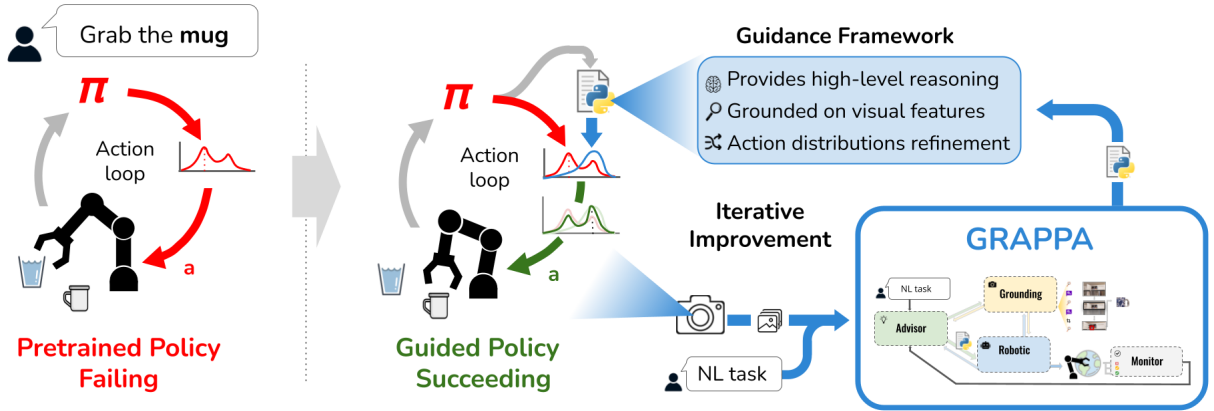


Fig. 3: An illustration of how GRAPPA intervenes in the action loop of pre-trained robotic policies in failure cases to provide visuomotor guidance generated with an agentic framework of agents to shift the action distribution for correct task execution.

to query open-vocabulary perception models.

Agent-based VLM frameworks in Robotics: Rather than using single VLMs in an end-to-end fashion, which might incur issues in generalization and robustness, various works have sought to orchestrate multiple VLM-based agents to work together in an interconnection multi-agent framework. Here, multiple agents can converse and collaborate to perform tasks, yielding improvements for the overall framework in online adaptability, cross-task generalization, and self-supervision [38, 42, 31]. These *agentic* frameworks have provided possibilities for enabling the identification of issues in task execution, providing feedback about possible improvements. Challenges remain, however, in that this feedback is often not sufficiently grounded on the spatial, visual, and dynamical properties of embodied interaction to be useful for policy adaptation; instead, the generated feedback is often too high-level or provides merely binary signals of success or failure.

Self-guided Robot Failure Recovery: Guan et al. [9] offer an analysis of frameworks for leveraging VLMs as behavior critics. Some approaches have explored integrating such pre-trained models to improve the performance of reinforcement learning (RL) algorithms. For instance, Ma et al. [26] use LLMs in a zero-shot fashion to design and improve reward functions, however this approach relies on human feedback to generate progressively human-aligned reward functions and further requires simulated retraining via RL, with high sample-complexity. On the other hand, Rocamonde et al. [33] avoid the need for explicit human feedback by directly using a VLM (CLIP) to compute the rewards to measure the proximity of a state (image) to a goal (text description), enabling gains in sample-efficiency for guiding a humanoid robot to perform various maneuvers in the MuJoCo simulator. A limitation of this approach, however, lies in the difficulty of generating rewards for long-horizon or multi-step tasks, which are characteristic of tasks involving complex agent-object interactions. Liu et al. [25] present a framework for detecting and analyzing failed executions automatically. However, their system focuses on explaining failure causes and proposing suggestions for

remediation, as opposed to also performing policy correction. In this paper, we propose a framework that directly adapts a base policy’s action distribution, during deployment, without requiring additional human feedback.

III. GROUNDING ROBOT POLICIES WITH GUIDANCE

A. Problem Formulation

We consider a pre-trained stochastic policy $\pi : O \times S \rightarrow A$ that maps observations o_t and robotic states s_t to action distributions $a_{\pi,t}$, at each time step t . Our objective is to generate a guidance distribution g_t that, when combined with this base policy, enhances overall performance during inference without requiring additional human demonstrations or extensive exploration procedures. Specifically, we aim to develop a modified policy $\pi_{\text{guided}} : O \times S \rightarrow A$ that achieves better performance on tasks where the original policy π struggles. We define this new policy as follows:

$$\pi_{\text{guided}}(a_{g,t}|o_t, s_t) = \pi(a_{\pi,t}|o_t, s_t) * G(a_{\pi,t}|o_t, s'_{t+1}), \quad (1)$$

where $G : A \times O \times S \rightarrow [0, 1]$ is a guidance function that maps observation o_t , action a_t , possible future state s'_{t+1} into a guidance score g_t . The ‘*’ operator here denotes the operation of combining both distributions conceptually, which we explore in detail in Section III-D. For the scope of this project, we assume that a dynamics model $\mathcal{D} : S \times A \rightarrow S$ is available, which can forecast possible future states of the robot $s'_{t+1} = \mathcal{D}(s_t, a_{\pi,t})$ given the current state s_t and action $a_{\pi,t}$.

Focusing on leveraging the world knowledge of Vision Language Models, while avoiding adding latency to the action loop, we choose to express these guidance functions as Python code. By integrating these code snippets into action loop of the base policies, we eliminate the need of time-consuming queries to large reasoning models. Samples of the format and content of the guidance functions generated by the framework are presented in the Appendix C.

Information flow between GRAPPA VLM-agents

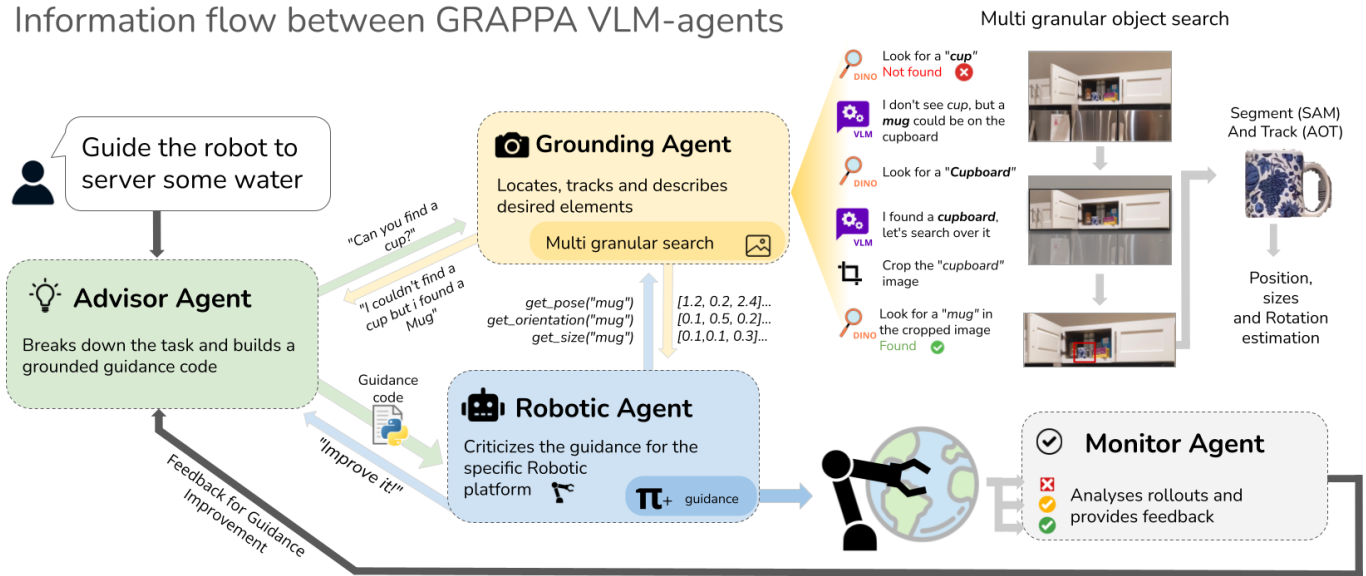


Fig. 4: Information flow between the agents to produce a guidance code. **a)** The advisor agent orchestrates guidance code generation by collaborating with other agents and using their feedback to refine the generated code. **b)** The grounding agent uses segmentation and classification models to locate objects of interest provided by the advisor, reporting findings back to the advisor. **c)** The robotic agent uses a Python interpreter to test the code for the specific robotic platform and judge the adequacy of the code. **d)** The monitor agent analyses the sequence of frames corresponding to the rollout of the guidance and give feedback on potential improvements.

B. A Multi-agent Guidance framework for Self Improvement

In order to generate the guidance function G , we leverage a group of conversational agents empowered with visual grounding capabilities and tool usage. Illustrated in Figure 4, the framework is composed of four main agents: an Advisor Agent, the Grounding Agent, the Monitor Agent, and the Robotic Agent. We provide system prompt samples in Appendix A.

Advisor Agent. A Vision Language Model is responsible for breaking down the task and communicating with the other agents to generate a sound guidance function for a given task.

Grounding Agent. A Vision Language Model that iteratively queries the free-form text segmentation models to locate [4, 24], track [5], and describe elements relevant to the task execution.

Monitor Agent. Responsible for identifying the causes of the failures in the unsuccessful rollouts, the Monitor Agent consists of a Vision Language model equipped with a key frame extractor.

Robotic Agent. Language Model equipped with descriptions of the robot platform, a robot’s dynamics model and wrapper functions for integration with the base policy. It criticizes the provided guidance functions to reinforce its relevance to the task and alignment with the robot’s capabilities.

C. Guidance Procedure

The conversational agents interact with each other through natural language and query their underlying tools to iteratively produce a guidance code tailored to the task in hand, the

environment, and the robot’s capabilities. The information flow between these agents is depicted by Figure 4.

For a given task expressed in natural language and an image of the initial state of the environment, the *Advisor Agent* uses Chain-of-Thought [37] strategy to generate a high-level plan of the steps necessary to accomplish the task. Being able to query a *Grounding Agent* and the *Robotic agent*, the Advisor is able to collect relevant information about trackable objects and elements in the environment, as well as the capabilities and limitations of the robotic platform.

For a given plan and list of relevant objects required for the task completion, the Grounding Agent uses grounding Dino [24] and the Segment Anything Model (SAM) [4] to locate the elements across multiple granularities and levels of abstraction. For instance, if an object is not immediately found, the agent will actively look for semantically similar objects or will look for higher-level elements that could encompass the missing object. For example, if the object “cup” is to be located, and it could not be immediately found, the agent could search for similar object like a “mug”. If it still struggles to locate it, the agent could search for a “shelf” and then try to find the “cup” or “mug” in the cropped image of the “shelf”. If an object is found, it is added to a tracking system [39]. This process enhances the Segment and Track Anything [4] approach with flexible multi-granular search. The object statuses are reported back to the Advisor Agent, which iterates on the action plan or proceeds with the generation of a guidance function grounded on the trackable objects located.

The Robotic agent acts as a critic to improve the guidance

function generated. Equipped with a Python interpreter and details of the base policy’s action space, the agent can evaluate the guidance function in terms of feasibility and relevance to the robot’s capabilities. Once a function suffices the system’s requirements, it is saved to be used in the action loop, in combination with the dynamics model, to provide a guidance score for possible actions sampled from the base policy.

After the execution of a rollout and the identification of failure in the task completion, the Monitor Agent is triggered to analyze the causes of the failure. By extracting key frames from the rollout video using PCA [27] and K-means clustering, the agent can feed a relevant and diverse set of images to the Visual Language Model prompted to access the failure causes. In the iterative applications of our framework, the Monitor Agent provides this feedback to the Advisor Agent, which can use this information to refine the guidance functions generated in the previous iterations.

Temporal-aware Guidance Functions. Inspired by recurrent architectures, we instruct the agents to generate guidance function conditioned on a customizable hidden state (h_t) expressed as an optional dictionary parameter as shown in the following example:

```

1 # Guidance function example in the context of
  grabbing a mug
2 def guidance_code(state,
3   hidden_state={"mug_reached": False, "mug_grabbed"
4     :False}):
5   #available grounding functions
6   #x, y, z = get_pose("mug")
7   #h, w, d = get_size("mug")
8   #rx, ry, rz = get_orientation("mug")
9   ...
  return score, new_hidden_state

```

The idea of using abstraction in a hidden state has proven to significantly improve the guidance performance, allowing the guidance functions to keep track of the task progress and adapt the guidance to longer horizon tasks. The guided policy can thus be written as:

$$\pi_{guided}(a_{g,t}|o_t, s_t, h_t) = \pi(a_{\pi,t}|o_t, s_t) * G(a_{\pi,t}|o_t, s'_{t+1}, h_t). \quad (2)$$

The complete guidance procedure is summarized in Algorithm 1. Note that we refer to the self-orchestrated conversation between the agents which yields the guidance code as the function `Generate_Guidance_Function`. For a closer look at the chain of thought employed by each agent please refer to Appendix A.

D. Guidance and policy integration

Aiming to guide a wide range of policies, our framework is designed to work both with continuous and discrete action spaces. In this section, we discuss the operation of combining the guidance function with the base policy’s action distributions. Furthermore, we discuss how deterministic regression models can be adapted to work with our framework.

Action-space Adaptation. We assume the availability of a dynamics model \mathcal{D} that can forecast possible future states of the robot given a possible action $a'_{\pi,t}$. In the manipulation

Algorithm 1 Guidance procedure of GRAPPA

Input

π : Base policy
 \mathcal{D} : Dynamics Model
env: Environment

```

1: for each episode do
2:    $o_t, s_t \leftarrow \text{env.init}$   $\triangleright$  observation and initial state
3:    $G \leftarrow \text{Generate\_Guidance\_Function}(o_t, s_t)$ 
4:    $h_t \leftarrow \text{Get\_Hidden\_State}(G(o_t, s_t))$ 
5:   for each time step  $t$  do
6:      $\mathcal{A}_{\pi,t} \leftarrow \{\pi(o_t, s_t)^i\}_{i=0}^n$   $\triangleright$  Sample  $n$  actions from the policy
7:      $\pi_t \leftarrow \pi(\mathcal{A}_{\pi,t}|o_t, s_t)$   $\triangleright$  Get action probabilities
8:      $\mathcal{S}_{\pi,t} \leftarrow \mathcal{D}(s_t, \mathcal{A}_{\pi,t})$   $\triangleright$  Infer possible future states
9:
10:     $G_{\pi,t} \leftarrow G(o_t, \mathcal{S}_{\pi,t}, h_t)$   $\triangleright$  Compute the guidance for the sampled possible future states
11:    Normalize  $G_{\pi,t}$ 
12:     $\pi_{guided,t} \leftarrow \pi_t * G_{\pi,t}$   $\triangleright$  Combine distributions
13:     $a_t \leftarrow \mathcal{A}_{\pi,t}[\text{argmax}(\pi_{guided,t})]$   $\triangleright$  Select the best action
14:     $o_t, s_t \leftarrow \text{env.step}(a_t)$   $\triangleright$  Execute  $a_t$ , update state  $s_{t+1}$  and observation  $o_{t+1}$ 
15:     $h_t \leftarrow \text{Get\_Hidden\_State}(G(o_t, s_t, h_t))$   $\triangleright$  Update hidden states

```

domain, a dynamics model is often available in the form of a forward kinematics model, a learned dynamics model, or a simulator. Oftentimes, the action space A of policies them-self is the same as the robot’s state S either being or joint angles of the robot or the gripper’s end-effector pose. For the last cases, where both the action and state space are expressed in $SE(3)$ integrating the guidance function with a base policy would only require a multiplication of the guidance scores with the action probabilities of the base policy. In other scenarios, adapting the robot’s action and state space to match the representation of the visual cues (position, orientation, and size) would be required.

Considering the visual grounding, the action space and the state space share the same representation ($SE(3)$), the operation to combine the guidance function with the base policy can be expressed as an element-wise weighted average:

$$\pi_{guided} = (1 - \alpha)\pi + \alpha G, \quad (3)$$

where $\alpha \in [0, 1]$ represents the percentage of guidance applied with respect to the base-policies distribution and is here denoted as *guidance factor*.

Adaptation of Regression Policies. To properly leverage the high-level guidance expressed in the guidance functions and the low-level capabilities of the base policy, it is desired that the policy’s action space be expressed as a distribution. In the case of regression policies that do not provide uncertainty estimates, several strategies can be employed to infer the action

TABLE I: Performance improvement on the RL-Bench [16] benchmark, by applying 5 iterations of guidance improvement over unsuccessful rollouts.

Model	turn tap	open drawer	sweep to dustpan of size	meat off grill	slide block to color target	push buttons	reach and drag	close jar	put item in drawer	stack blocks	Avg. Success
Act3D 25 demos/task	76	76	96	64	92	84	96	48	60	0	69.2
+1% guidance	80 (+4)	96 (+20)	96	84 (+20)	92	84	100 (+4)	84 (+36)	80 (+20)	8 (+8)	80.4 (+11.2)
+10% guidance	88 (+12)	100 (+24)	96	88 (+24)	92	84	100 (+4)	60 (+12)	80 (+20)	0	78.8 (+9.6)
Act3D 10 demos/task	32	60	84	16	60	72	68	32	44	8	47.6
+1% guidance	44 (+12)	88 (+28)	88 (+4)	24 (+8)	68 (+8)	72	76 (+8)	52 (+20)	60 (+16)	8	58 (+10.4)
+10% guidance	44 (+12)	64 (+4)	84	20 (+4)	68 (+8)	76 (+4)	76 (+8)	40 (+8)	56 (+12)	8	53.6 (+6)
Act3D 5 demos/task	24	0	84	4	8	32	8	8	12	0	18
+1% guidance	48 (+24)	16 (+16)	84	8 (+4)	12 (+4)	40 (+8)	24 (+16)	20 (+12)	20 (+8)	0	27.2 (+9.2)
+10% guidance	24	0	84	8 (+4)	12 (+4)	44 (+12)	20 (+12)	8	20 (+8)	0	22 (+4)
Diffuser actor 5 demos/task	24	64	40	28	44	68	40	24	44	0	37.6
+1% guidance	40 (+16)	92 (+28)	64 (+24)	40 (+12)	44	68	52 (+12)	24	84 (+40)	0	50.8 (+13.2)
+10% guidance	40 (+16)	84 (+20)	52 (+12)	28	52 (+8)	68	48 (+8)	32 (+8)	76 (+32)	0	48 (+10.4)

distribution. One common approach is to assume a Gaussian distribution centered at the predicted value and compute the variance using ensembles of models trained with different initialization, different data samples, or different dropout seeds or different checkpoint stages [1]. Other strategies to infer the distributions of the model include using bootstrapping, Bayesian neural networks, or using a mixture of Gaussians [28].

IV. EXPERIMENTS & RESULTS

A. Experimental Setup

Task Definitions: We demonstrate the efficacy of GRAPPA, in simulation on the RL-Bench benchmark [16] and on two challenging real-world tasks. For the real-world setup, we use the UFACTORY Lite 6 robot arm as the robotic agent and, as the end-effector, we use the included UFACTORY Gripper Lite, a simple binary gripper. The arm is mounted on a workbench. For perception, we use a calibrated RGB-D Camera, specifically the Intel RealSense Depth Camera D435i. All experiments were conducted on a desktop machine with a single NVIDIA RTX 4080 GPU, 64GB of RAM, and a AMD Ryzen 7 8700G CPU.

- **(Sim-to-real): Button-pressing:** Here we employ the diffuser policy to perform the button pressing tasks; to make the task more challenging we introduce clutter in the environment, as well as employ out-of-distribution props. The intent is to measure the ability of GRAPPA to mitigate the sim-to-real gap and remain robust to distractor objects.
- **(Sim): RL-Bench:** We consider 10 tasks on the RL-Bench benchmark [16] using a single RGBD camera input, as described in the *GNFactor* setup [41].
- **(Sim): RL-Bench, learning from scratch:** Aiming to explore the capabilities of GRAPPA on learning new skills from scratch, we selected 4 challenging tasks from the RL-Bench benchmark: *turn tap*, *push buttons*, *slide block to color target*, *reach and drag*.
- **(Real): Sim-to-real policy adaptation:** We want to evaluate the capabilities of GRAPPA on reducing the sim-to-real gap by guiding policies that were pre-trained in simulation but are deployed to real-world tasks.

- **(Real): Reach for chess piece:** Given a cluttered scene with many similar objects, we want to evaluate if the multi-granular perception framework can effectively guide the agent to identify and reach for the appropriate target. We implement this perceptual grounding and reaching task on a standard chessboard, where the agent must identify and reach for one of the chess pieces specified by natural language instruction.

Base Policies: We evaluate the effectiveness of our guidance framework using different base policies, *Act3D* [8], *3D Diffuser Actor* [17], and a *RandomPolicy*. All policies plan in a continuous space of translations, rotations, and gripper state ($SE(3) \times \mathbb{R}$), however they utilize different inference strategies:

- **Act3D** samples waypoints in the Cartesian space (\mathbb{R}^3) and predicts the orientation and gripper state for the best scoring sampled waypoint, combining a classification and regression strategies into a single policy.
- **3D Diffuser Actor**, on the other hand, uses a diffusion model to compute the target waypoints and infers the orientation and gripper state from the single forecast waypoint, thus tackling the problem as a single regression task.
- **RandomPolicy** denotes any of the former frameworks that has not been trained for a specific task, therefore the weights are randomly initialized.

The fundamentally different types of policies’ outputs make them a great use case for our policy guidance framework. Furthermore, the common representation of the action and state spaces of both policies ($SE(3) \times \mathbb{R}$) provides a straightforward integration with our grounding models.

As described in Section III-D, the regression components of the policies require an adaptation to transform the single predictions of the model into a distribution over the action space. For the sake of simplicity, we assume a Gaussian distribution over the action space, with the mean centered on the predicted values and the standard deviation fixed on a constant value. The outputs of the classification component of Act3D (waypoint positions) were directly considered as samples of a distribution over the Cartesian space (\mathbb{R}^3).

The integration of base policies with the guidance distributions was performed by applying a weighted average

parameterized by α as shown in Equation 3.

B. Experimental Evaluation

Our experimental evaluation aims to address the following questions: (1) Does GRAPPA enable policies to maintain or improve performance on tasks in out-of-distribution settings? (2) Does GRAPPA enable policies to bridge the sim-to-real gap for deployment on real-world tasks? (3) Does GRAPPA improve the performance of pre-trained base policies on specific robotics tasks and environments without additional human demonstrations? (4) Does the proposed multi-granular perception capabilities effectively guide the policy in challenging cluttered environments? (5) Does GRAPPA enable policies to learn new skills from scratch? (6) What is the effect of guidance on expert versus untrained policies?

Does GRAPPA enable policies to bridge the sim-to-real gap for deployment on real-world tasks and in out-of-distribution settings? We first train a 3D Diffuser Actor policy with 10 tasks with 100 demonstrations each with a single RGBD camera setup (GNFactor setting). Using exclusively simulation-based demonstrations, note this policy was chosen as it displayed stronger overall performance in the simulation benchmark. We roll out this policy in the real world in the button-pressing task in a cluttered environment, as depicted by Figure 2. To facilitate sim-to-real integration, we adapted the dimensions, scale, and alignment of the real-world workspace to match the simulation data, and positioned the RGBD camera (intel real sense) in a similar placement as the camera used to train the base policy. The sim-to-real setup was validated by performing simple tasks in uncluttered settings as shown in the appendix (III).

TABLE II: Real-world performance improvement in a sim-to-real setting. Diffuser actor policy trained on 10 tasks in simulation (GNFactor setup 100 demos/task) and evaluated on the tasks “Push buttons” with different guidance levels. **Red** cells refer to failures due to timeouts in task execution. **Orange** cells refer to errors in perception. **Green** cells refer to successes.

Policy	Task completion Success Rate [%]	Error Breakdown
<i>Naive Sim2Real</i>	0.0	
+15% guidance	66.7	
+50% guidance	50.0	
+75% guidance	100	
<i>Finetuned Baseline (10 real-world demos)</i>	33.0	
+15% guidance	50.0	
+50% guidance	16.7	
+75% guidance	100	

As depicted by Table II, this cluttered scenario has shown to be out-of-distribution scenarios for the base policy solely trained in simulation (Naive Sim2Real). Fine-tuning the policy with 10 real-world demonstrations of pressing buttons in uncluttered environments has proven to slightly increase the performance of the policy. These base policies combined with our guidance framework achieve a higher success rate in 5 out of the 6 scenarios with guidance. In the only exception case (finetuned baseline + 50%), the base perception models failed

to detect and keep track of the target button during the rollout leading to a detrimental guidance of the policy.

Does GRAPPA improve the performance of pre-trained base policies on specific robotics tasks and environments without additional human demonstrations? We first assess the effect of the proposed guidance on the Act3D and 3D Diffuser Actor baselines following the GNFactor [41] setup, which consists of a single RGBD camera and table-top manipulator performing 10 challenging tasks with 25 variations each. Guidance is iteratively generated for the failure cases. For the failed rollouts, our policy improvement framework ran for 5 iterations. As displayed by Table I, the framework was able to improve the success rate of the base policy on most of the tasks, with the best results achieved by using 1% guidance. The low amount of guidance has shown to be enough to bend the action distribution to the desired direction, while still preserving the low-level nuances captured by the base policy. This suggests that GRAPPA is capable of improving base policies by adding abstract understanding and grounding of the desired task, while preserving the low-level movement profiles captured by the original policies.

Does the proposed multi-granular perception capability effectively guide the policy in challenging, cluttered environments? In real-world experiments, we qualitatively demonstrate the fine-grained detection capabilities of GRAPPA by tasking it with reaching for a white knight chess piece in a cluttered chess board. Figure 5 shows the roll-out of the first guidance iteration over an untrained policy, displaying the initial and final time steps of the task. The accompanying heatmaps illustrate the distributions of the original untrained policy (Diffuser heatmaps) and the guided policy (Guidance heatmaps). Additionally, the multi-granular search results highlight the steps taken by the grounding agent to locate the target piece. After initially failing to detect the white knight directly, the agent successfully identifies the chessboard and then focuses within that region to locate the target piece. These findings demonstrate that GRAPPA effectively leverages a semantic understanding of scene components to guide the policy towards successful task execution.

Does GRAPPA enable policies to learn new skills from scratch? We evaluated the performance of the framework on learning new skills from scratch on 4 tasks of the RL-Bench benchmark: *turn tap*, *push buttons*, *slide block to color target*, *reach and drag*. In this setup, we initialized the Act3D policy with random weights and applied 100% of guidance ($\alpha = 1.0$) over the policy for the x,y, and z components. Only leveraging the waypoint sampling mechanism of Act3D and overwriting its distribution with the values queried from the guidance functions generated. The results show that the framework is capable of learning new skills from scratch, achieving a higher success rate than the base policy pre-trained on 5 demonstrations/tasks for the tasks “turn tap” and “push buttons” tasks. When utilizing only the untrained Act3D policy (without guidance) the policy achieved 0 success rate on the tasks. Figure 6 demonstrates the iterative improvement of our guidance framework. Updating the guidance code generated

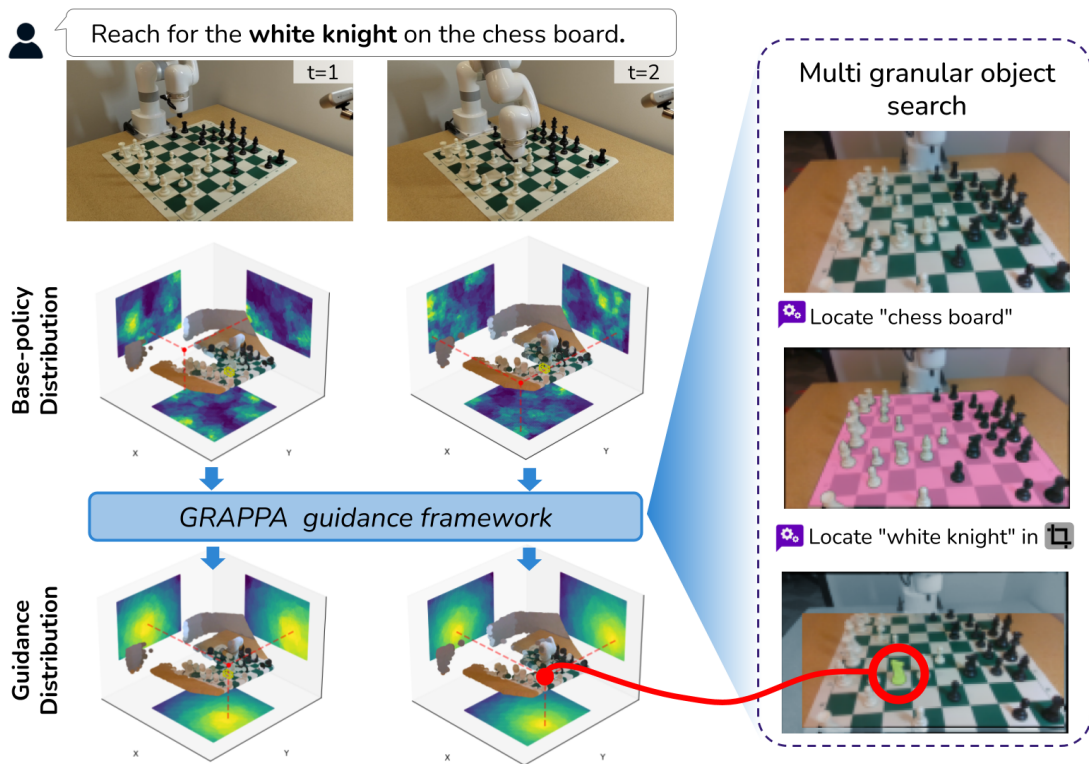


Fig. 5: Real-world results for learning skills from scratch on the Lite6 chess task. The top row shows an external view of the robot performing the tasks. The second row depicts the action heat map given by the random diffuser policy at the first and last time step. The bottom row depicts the corresponding heat maps generated after applying the guidance. We show it can successfully guide the action towards the desired object. On the right, we show a breakdown of the multi-granular search performed by other **grounding agent** to locate the white knight; we disambiguate the scene by searching in parent objects and constraining the search to semantically relevant areas.

for each failed rollouts from the previous iteration. Policy rollouts are provided in Figure 9.




Fig. 6: Performance of our framework on learning new skills from scratch (guidance over untrained policies), and iteratively improving the guidance functions generated.

It is worth mentioning that a few variations of the simulated tasks “turn tap” and “reach and drag”, which seemly would require a precise orientation control of the manipulator, could be solved by guiding only the Cartesian components of the police’s output. For these variations, a qualitative analysis shows that successful roll-outs could be achieved by tapping the end-effector on the target objects.

We perform a similar experiment in the real-world settings. Here, we run the framework only relying on the action distribution given by the guidance code ($\alpha = 1.0$), using a RandomPolicy as the base policy. We first consider the task of pressing colored buttons in a given sequence; using toy buttons made out of acrylic and paper as a prop. Figure 7 shows the roll-out corresponding to the first iteration for this task, along with heat maps depicting a projection of the output action distribution around the point of maximum. In this zero-shot scenario, GRAPPA has proven to correctly guide the robot to the desired objects preserving the prompted order; however, it struggles to capture low-level nuances of the movement, such as the appropriate pressing force and proper approach of the buttons. A simulated version of this experiment is shown in the appendix (Figure 9), demonstrating how combining the guidance with a pre-trained policy can mitigate this behavior.

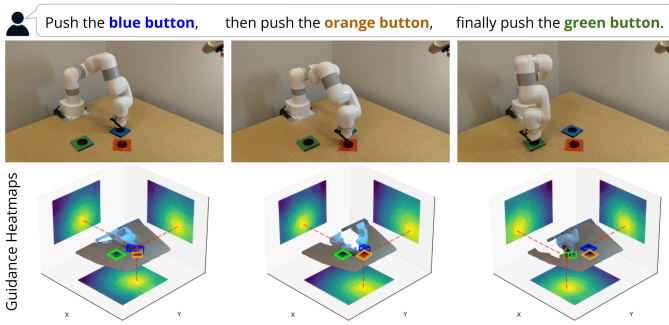


Fig. 7: Real-world results for learning skills from scratch on the Lite6 on the multiple button press task. Each column represents a keyframe in the task rollout. The first row shows a third-person view of the robot’s movement and the tabletop setting. The second shows the corresponding action distribution over the space generated by the guidance code, the red dot indicates the target waypoint for the end effector. We demonstrate that GRAPPA can guide a random base policy to successfully perform the desired task.

What is the effect of guidance on expert versus untrained policies? As discussed previously, GRAPPA can learn tasks from scratch using a random base policy with 100% of guidance (α). Given that we are not affecting the policy, the product of the iterative learning from scratch is the generated guidance script which captures a high-level understanding of the task, e.g., spatial relationships, and task completion criteria, among others. We can qualitatively see the effect of the learning process by comparing the guidance scripts for different iterations of the same task. Appendix C includes two samples of guidance code for the same task but on different iterations. We can see that the code corresponding to the second iteration is very simple and only accounts for Euclidean distance and button order; while by the fifth iteration considerations like orientation come into play.

On the other hand, applying the guidance to an existing expert policy aims to shift the action distribution to account for failure cases, like potentially out-of-distribution scenes. The goal here is to use the gained high-level understanding to aid the policy in task completion. Table I shows that adding too much weight to the guidance function can yield diminishing returns as it can overpower the nuanced low-level details from the expert policy: notice that the performance gain across the board is bigger using 1% of guidance.

V. CONCLUSION

Summary: In this work, we proposed GRAPPA, a novel framework for the self-improvement of embodied policies. Our self-guidance approach leverages the world knowledge of a group of conversational agents and grounding models to guide policies during deployment. We demonstrated the effectiveness of our approach in autonomously improving manipulation policies and learning new skills from scratch, in simulated RL-bench benchmark tasks and in two challenging real-world

tasks. Our results show that the proposed framework is especially effective in improving the following high-level task structures and key steps to solve the task. This capability can be well suited for improving pre-trained policies that struggle with long-horizon tasks or for learning new simple skills from scratch.

Limitations: From an analysis of the guided rollouts, a few of the tasks variations proved challenging for the perception models used by the grounding agent, leading to false positives detections or failure to locate specific objects. This limitation was mainly observed in simulation task, where the graphics object representations, even though simplified, do not always match the representations used to train the object detection models. This limitation could be addressed by integrating more robust object detection models or verification procedures to ensure the correct detection of objects in the scene. Moreover, occasional inaccuracies on scene understanding by the Visual Language Model (VLM) have been observed, leading to the generation of inaccurate guidance codes and unexpected behaviors. Even though recent advances in large vision-language models have shown great potential in understanding the underlying dynamics of the world from large-scale internet data, translating this knowledge into out-of-distribution domain, such as robotics, while preventing hallucinations remains an open challenge.

Future Work: Regarding future works, we think that combining the proposed framework with fine-grained exploration techniques would allow the policy to explore in a targeted manner the low-level details of the task while leveraging the high-level guidance provided by our framework. This may enrich the guidance codes with the necessary low-level details required to perform more complex tasks successfully.

Furthermore, the guidance function generation could be further improved by composing and adapting from a repository of successful guidance functions from previous experiences. This could be achieved by incorporating Retrieval Augmented Generation (RAG) [20] into our multi agent framework. This modification could allow the guidance system to learn new simple skills from scratch by interacting with the environment and leveraging this collected knowledge to guide the policy more effectively.

Aiming to incorporate the knowledge captured by the guidance functions into the base policy, an experience replay and finetuning mechanisms could be incorporated into our current system. This modification could allow the framework to use past guided experiences to improve the base policy in a sample efficient manner. This could be achieved in a targeted manner by leveraging Low Rank Adaptation (LoRA) [11].

Reproducibility Discussion: Intending to encourage other researchers to build upon the introduced framework, we take steps to ensure the usability and reproducibility of our work. The source code for GRAPPA will be open-sourced and linked to on a project website. We will provide dockerized scripts to facilitate the setup across different development environments. Additionally, in Section A we include the prompts used to configure each agent. The temperature of the model was set

as zero to reduce variations in runs, as using fixed seeds for the experiments. More hyperparameter details will be available in the open-sourced repository.

REFERENCES

- [1] Moloud Abdar, Farhad Pourpanah, Sadiq Hussain, Dana Rezazadegan, Li Liu, Mohammad Ghavamzadeh, Paul Fieguth, Xiaochun Cao, Abbas Khosravi, U Rajendra Acharya, et al. A review of uncertainty quantification in deep learning: Techniques, applications and challenges. *Information fusion*, 76:243–297, 2021.
- [2] Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, et al. Do as i can, not as i say: Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022.
- [3] 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, Alexander 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, 2023.
- [4] Yangming Cheng, Liulei Li, Yuanyou Xu, Xiaodi Li, Zongxin Yang, Wenguan Wang, and Yi Yang. Segment and track anything. *arXiv preprint arXiv:2305.06558*, 2023.
- [5] Yangming Cheng, Liulei Li, Yuanyou Xu, Xiaodi Li, Zongxin Yang, Wenguan Wang, and Yi Yang. Segment and track anything, 2023. URL <https://arxiv.org/abs/2305.06558>.
- [6] Open X-Embodiment Collaboration, Abhishek Padalkar, Acorn Pooley, Ajinkya Jain, Alex Bewley, Alex Herzog, Alex Irpan, Alexander Khazatsky, Anant Rai, Anikait Singh, Anthony Brohan, Antonin Raffin, Ayzaan Wahid, Ben Burgess-Limerick, Beomjoon Kim, Bernhard Schölkopf, Brian Ichter, Cewu Lu, Charles Xu, Chelsea Finn, Chenfeng Xu, Cheng Chi, Chenguang Huang, Christine Chan, Chuer Pan, Chuyuan Fu, Coline Devin, Danny Driess, Deepak Pathak, Dhruv Shah, Dieter Buechler, Dmitry Kalashnikov, Dorsa Sadigh, Edward Johns, Federico Ceola, Fei Xia, Freek Stulp, Gaoyue Zhou, Gaurav S. Sukhatme, Gautam Salhotra, Ge Yan, Giulio Schiavi, Hao Su, Hao-Shu Fang, Haochen Shi, Heni Ben Amor, Henrik I Christensen, Hiroki Furuta, Homer Walke, Hongjie Fang, Igor Mordatch, Ilija Radosavovic, Isabel Leal, Jacky Liang, Jaehyung Kim, Jan Schneider, Jasmine Hsu, Jeannette Bohg, Jeffrey Bingham, Jiajun Wu, Jialin Wu, Jianlan Luo, Jiayuan Gu, Jie Tan, Jihoon Oh, Jitendra Malik, Jonathan Tompson, Jonathan Yang, Joseph J. Lim, João Silvério, Junhyek Han, Kanishka Rao, Karl Pertsch, Karol Hausman, Keegan Go, Keerthana Gopalakrishnan, Ken Goldberg, Kendra Byrne, Kenneth Oslund, Kento Kawaharazuka, Kevin Zhang, Keyvan Majd, Krishan Rana, Krishnan Srinivasan, Lawrence Yunliang Chen, Lerrel Pinto, Liam Tan, Lionel Ott, Lisa Lee, Masayoshi Tomizuka, Maximilian Du, Michael Ahn, Mingtong Zhang, Mingyu Ding, Mohan Kumar Srirama, Mohit Sharma, Moo Jin Kim, Naoaki Kanazawa, Nicklas Hansen, Nicolas Heess, Nikhil J Joshi, Niko Suenderhauf, Norman Di Palo, Nur Muhammad Mahi Shafullah, Oier Mees, Oliver Kroemer, Pannag R Sanketi, Paul Wohlhart, Peng Xu, Pierre Sermanet, Priya Sundaresan, Quan Vuong, Rafael Rafailov, Ran Tian, Ria Doshi, Roberto Martín-Martín, Russell Mendonca, Rutav Shah, Ryan Hoque, Ryan Julian, Samuel Bustamante, Sean Kirmani, Sergey Levine, Sherry Moore, Shikhar Bahl, Shivin Dass, Shuran Song, Sichun Xu, Siddhant Halder, Simeon Adebola, Simon Guist, Soroush Nasiriany, Stefan Schaal, Stefan Welker, Stephen Tian, Sudeep Dasari, Suneel Belkhal, Takayuki Osa, Tatsuya Harada, Tatsuya Matsushima, Ted Xiao, Tianhe Yu, Tianli Ding, Todor Davchev, Tony Z. Zhao, Travis Armstrong, Trevor Darrell, Vidhi Jain, Vincent Vanhoucke, Wei Zhan, Wenxuan Zhou, Wolfram Burgard, Xi Chen, Xiaolong Wang, Xinghao Zhu, Xuanlin Li, Yao Lu, Yevgen Chebotar, Yifan Zhou, Yifeng Zhu, Ying Xu, Yixuan Wang, Yonatan Bisk, Yoonyoung Cho, Youngwoon Lee, Yuchen Cui, Yueh hua Wu, Yujin Tang, Yuke Zhu, Yunzhu Li, Yusuke Iwasawa, Yutaka Matsuo, Zhuo Xu, and Zichen Jeff Cui. Open X-Embodiment: Robotic learning datasets and RT-X models. <https://arxiv.org/abs/2310.08864>, 2023.
- [7] Yan Ding, Xiaohan Zhang, Chris Paxton, and Shiqi Zhang. Task and motion planning with large language models for object rearrangement. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2086–2092. IEEE, 2023.
- [8] Theophile Gervet, Zhou Xian, Nikolaos Gkanatsios, and Katerina Fragkiadaki. Act3d: 3d feature field transformers for multi-task robotic manipulation. In *7th Annual Conference on Robot Learning*, 2023.
- [9] 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. *arXiv preprint arXiv:2402.04210*, 2024.
- [10] Huy Ha, Pete Florence, and Shuran Song. Scaling up and distilling down: Language-guided robot skill acquisition. In *Conference on Robot Learning*, pages 3766–3777. PMLR, 2023.

- [11] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.
- [12] Yafei Hu, Quanting Xie, Vidhi Jain, Jonathan Francis, Jay Patrikar, Nikhil Keetha, Seungchan Kim, Yaqi Xie, Tianyi Zhang, Zhibo Zhao, et al. Toward general-purpose robots via foundation models: A survey and meta-analysis. *arXiv preprint arXiv:2312.08782*, 2023.
- [13] Chenguang Huang, Oier Mees, Andy Zeng, and Wolfram Burgard. Visual language maps for robot navigation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 10608–10615. IEEE, 2023.
- [14] Wenlong Huang, Fei Xia, Ted Xiao, Harris Chan, Jacky Liang, Pete Florence, Andy Zeng, Jonathan Tompson, Igor Mordatch, Yevgen Chebotar, et al. Inner monologue: Embodied reasoning through planning with language models. *arXiv preprint arXiv:2207.05608*, 2022.
- [15] Wenlong Huang, Chen Wang, Ruohan Zhang, Yunzhu Li, Jiajun Wu, and Li Fei-Fei. Voxposer: Composable 3d value maps for robotic manipulation with language models. *arXiv preprint arXiv:2307.05973*, 2023.
- [16] Stephen James, Zicong Ma, David Rovick Arrojo, and Andrew J Davison. Rlbench: The robot learning benchmark & learning environment. *IEEE Robotics and Automation Letters*, 5(2):3019–3026, 2020.
- [17] Tsung-Wei Ke, Nikolaos Gkanatsios, and Katerina Fragkiadaki. 3d diffuser actor: Policy diffusion with 3d scene representations. *arXiv preprint arXiv:2402.10885*, 2024.
- [18] Alexander Khazatsky, Karl Pertsch, Suraj Nair, Ashwin Balakrishna, Sudeep Dasari, Siddharth Karamcheti, Soroush Nasiriany, Mohan Kumar Srirama, Lawrence Yunliang Chen, Kirsty Ellis, et al. Droid: A large-scale in-the-wild robot manipulation dataset. *arXiv preprint arXiv:2403.12945*, 2024.
- [19] Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, et al. Openvla: An open-source vision-language-action model. *arXiv preprint arXiv:2406.09246*, 2024.
- [20] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. Retrieval-augmented generation for knowledge-intensive nlp tasks, 2021. URL <https://arxiv.org/abs/2005.11401>.
- [21] Xiaoqi Li, Mingxu Zhang, Yiran Geng, Haoran Geng, Yuxing Long, Yan Shen, Renrui Zhang, Jiaming Liu, and Hao Dong. Manipllm: Embodied multimodal large language model for object-centric robotic manipulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18061–18070, 2024.
- [22] Jacky Liang, Wenlong Huang, Fei Xia, Peng Xu, Karol Hausman, Brian Ichter, Pete Florence, and Andy Zeng. Code as policies: Language model programs for embodied control. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9493–9500. IEEE, 2023.
- [23] Jacky Liang, Fei Xia, Wenhao Yu, Andy Zeng, Montserrat Gonzalez Arenas, Maria Attarian, Maria Bauza, Matthew Bennice, Alex Bewley, Adil Dostmohamed, et al. Learning to learn faster from human feedback with language model predictive control. *arXiv preprint arXiv:2402.11450*, 2024.
- [24] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*, 2023.
- [25] Zeyi Liu, Arpit Bahety, and Shuran Song. Reflect: Summarizing robot experiences for failure explanation and correction. *arXiv preprint arXiv:2306.15724*, 2023.
- [26] Yecheng Jason Ma, William Liang, Guanzhi Wang, De-An Huang, Osbert Bastani, Dinesh Jayaraman, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Eureka: Human-level reward design via coding large language models. *arXiv preprint arXiv:2310.12931*, 2023.
- [27] Andrzej Maćkiewicz and Waldemar Ratajczak. Principal components analysis (pca). *Computers & Geosciences*, 19(3):303–342, 1993.
- [28] José Mena, Oriol Pujol, and Jordi Vitrià. A survey on uncertainty estimation in deep learning classification systems from a bayesian perspective. *ACM Computing Surveys (CSUR)*, 54(9):1–35, 2021.
- [29] Zawalski Michał, Chen William, Pertsch Karl, Mees Oier, Finn Chelsea, and Levine Sergey. Robotic control via embodied chain-of-thought reasoning. *arXiv preprint arXiv:2407.08693*, 2024.
- [30] Yao Mu, Qinglong Zhang, Mengkang Hu, Wenhai Wang, Mingyu Ding, Jun Jin, Bin Wang, Jifeng Dai, Yu Qiao, and Ping Luo. Embodiedgpt: Vision-language pre-training via embodied chain of thought. *Advances in Neural Information Processing Systems*, 36, 2024.
- [31] Meenal Parakh, Alisha Fong, Anthony Simeonov, Tao Chen, Abhishek Gupta, and Pulkit Agrawal. Lifelong robot learning with human assisted language planners. *arXiv e-prints*, pages arXiv–2309, 2023.
- [32] Krishan Rana, Jesse Haviland, Sourav Garg, Jad Abou-Chakra, Ian Reid, and Niko Suenderhauf. Sayplan: Grounding large language models using 3d scene graphs for scalable robot task planning. In *7th Annual Conference on Robot Learning*, 2023.
- [33] Juan Rocamonde, Victoriano Montesinos, Elvis Nava, Ethan Perez, and David Lindner. Vision-language models are zero-shot reward models for reinforcement learning. *arXiv preprint arXiv:2310.12921*, 2023.
- [34] Saumya Saxena, Blake Buchanan, Chris Paxton, Bingqing Chen, Narunas Vaskevicius, Luigi Palmieri, Jonathan Francis, and Oliver Kroemer. Grapheqa: Using 3d semantic scene graphs for real-time embodied ques-

- tion answering. *arXiv preprint arXiv:2412.14480*, 2024.
- [35] Ishika Singh, Valts Blukis, Arsalan Mousavian, Ankit Goyal, Danfei Xu, Jonathan Tremblay, Dieter Fox, Jesse Thomason, and Animesh Garg. Progprompt: Generating situated robot task plans using large language models. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11523–11530. IEEE, 2023.
- [36] Octo Model Team, Dibya Ghosh, Homer Walke, Karl Pertsch, Kevin Black, Oier Mees, Sudeep Dasari, Joey Hejna, Tobias Kreiman, Charles Xu, et al. Octo: An open-source generalist robot policy. *arXiv preprint arXiv:2405.12213*, 2024.
- [37] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models, 2023. URL <https://arxiv.org/abs/2201.11903>.
- [38] Mengdi Xu, Peide Huang, Wenhao Yu, Shiqi Liu, Xilun Zhang, Yaru Niu, Tingnan Zhang, Fei Xia, Jie Tan, and Ding Zhao. Creative robot tool use with large language models. *arXiv preprint arXiv:2310.13065*, 2023.
- [39] Zongxin Yang and Yi Yang. Decoupling features in hierarchical propagation for video object segmentation. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [40] Sriram Yenamandra, Arun Ramachandran, Mukul Khanna, Karmesh Yadav, Jay Vakil, Andrew Melnik, Michael Büttner, Leon Harz, Lyon Brown, Gora Chand Nandi, et al. Towards open-world mobile manipulation in homes: Lessons from the neurips 2023 homerobot open vocabulary mobile manipulation challenge. *arXiv preprint arXiv:2407.06939*, 2024.
- [41] Yanjie Ze, Ge Yan, Yueh-Hua Wu, Annabella Macaluso, Yuying Ge, Jianglong Ye, Nicklas Hansen, Li Erran Li, and Xiaolong Wang. Gnfactor: Multi-task real robot learning with generalizable neural feature fields, 2024. URL <https://arxiv.org/abs/2308.16891>.
- [42] Ceng Zhang, Xin Meng, Dongchen Qi, and Gregory S Chirikjian. Rail: Robot affordance imagination with large language models. *arXiv preprint arXiv:2403.19369*, 2024.