DoReMi: Grounding Language Model by Detecting and Recovering from Plan-Execution Misalignment

Yanjiang Guo[1][*], **Yen-Jen Wang**[1][*], **Lihan Zha**[2][*], **Zheyuan Jiang**[1], **Jianyu Chen**[13]

[1]Institute of Interdisciplinary Information Sciences, Tsinghua University
[2]Weiyang College, Tsinghua University, [3]Shanghai Qi Zhi Institute

Abstract

Large language models encode a vast amount of semantic knowledge and possess remarkable understanding and reasoning capabilities. Previous research has explored how to ground language models in robotic tasks to ensure that the sequences generated by the language model are both logically correct and practically executable. However, low-level execution may deviate from the high-level plan due to environmental perturbations or imperfect controller design. In this paper, we propose **DoReMi**, a novel language model grounding framework that enables immediate **De**tection and **Re**covery from **Mi**salignments between plan and execution. Specifically, during low-level skill execution, we use a vision question answering (VQA) model to regularly detect plan-execution misalignments. If certain misalignment occurs, our method will call the language model to re-plan in order to recover from misalignments. Experiments on various complex tasks including robot arms and humanoid robots demonstrate that our method can lead to higher task success rates and shorter task completion times. Videos of DoReMi are available at `https://sites.google.com/view/doremi-paper`.

1 Introduction

Pretrained large language models (LLMs) encode vast amounts of semantic knowledge and exhibit remarkable reasoning ability and understanding of the world. Previous works have incorporated language models into robotic tasks to help embodied agents better understand and interact with the world to complete challenging long-horizon tasks that require complex planning and reasoning [1, 2, 3, 4].

To make the generated plan executable by embodied agents, we need to ground the language. One line of the works leverages pretrained language models in an end-to-end manner that directly maps language and image inputs into the robot's low-level action space [5, 6, 7, 8, 9]. These approaches often require large amounts of robot action data for successful end-to-end training, which is expensive to acquire [5, 6]. Moreover, these action-output models often contain large transformer-based architectures and cannot run at high frequencies. Therefore, they may not be suitable for tasks with complex dynamics (e.g., legged robots) which require high-frequency rapid response. Recently, many works adopt a hierarchical approach where high-level task planning is performed by language models, and then some low-level controllers are adopted to generate the complex robot control commands [1, 2, 4, 10]. Under this hierarchical framework, we can leverage powerful robot control methods, such as reinforcement learning, to handle complex robot dynamic control problems with high frequency.

However, these grounding methods often assume that every low-level skill can perfectly execute the high-level plan generated by the language model. In practice, low-level execution may deviate

---

[*]Equal contribution, listed alphabetically

from the high-level plan due to environmental perturbations or imperfect controller design. These misalignments between plan and execution may occur at any time during the task procedure. Previous works consider incorporating execution feedback into language prompts once the previous plan step is finished. If the step is unsuccessful, the process is repeated [10]. However, this delayed feedback can be inefficient. For instance, as illustrated in Figure 1(b), when a human is carrying a box and performing the low-level skill "Go to the gray table", if the box is accidentally dropped, it becomes futile to continue with the current skill. The human will immediately abort the current skill and call for the skill "Pick up the box". However, agents without immediate re-planning will continue going forward, and will take more time to pick up the box dropped halfway after reaching the destination.

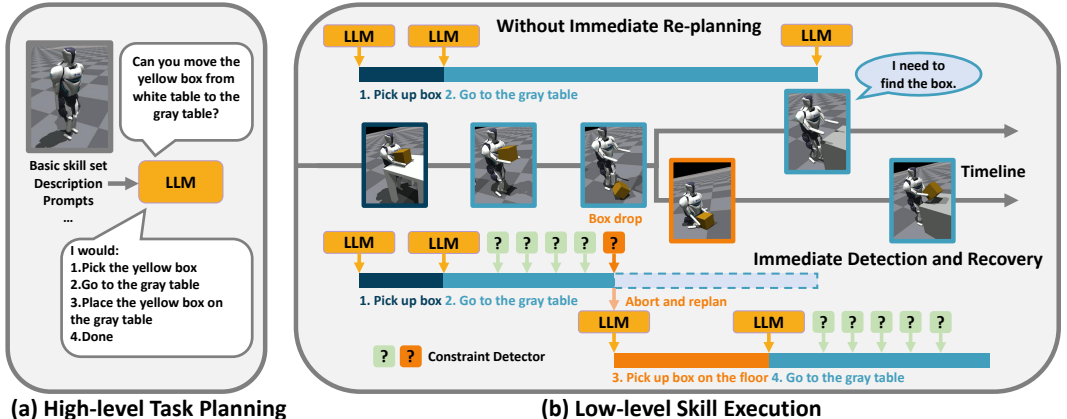(a) High-level Task Planning   (b) Low-level Skill Execution

Figure 1: Illustration of our motivation. Low-level execution may deviate from the high-level plan. DoReMi can immediately detect the misalignment between the plan and execution when the box drops accidentally and quickly recovers. Agents without immediate re-planning suffer from such misalignment.

In this paper, we propose a novel framework **DoReMi** which enables immediate **D**etecti**o**n and **Re**covery from plan-execution **Mi**salignments. Once a misalignment between plan and execution occurs, our method immediately calls the language model to re-plan, enabling timely recovery. Specifically, we maintain and update a set of constraints that indicate the alignment between the plan and execution. These constraints are provided through natural language description. During the execution of low-level skills, a visual question answering (VQA) model [11, 12] is employed as a "constraint detector" to regularly check whether the agent violates any constraint in the constraint set. If some constraints are found to be violated, which indicates that the plan and execution may be misaligned, the language model is called to re-plan the high-level task. Furthermore, under mild assumptions, we conduct a theoretical analysis to estimate how much time can be saved or how much the success rate can be improved through immediate re-planning. Experiments in physical simulations, including robot arm manipulation tasks and humanoid tasks, demonstrate that DoReMi leads to a higher success rate and shorter task execution time.

Our contributions can be summarized as follows:

- We show the importance of aligning high-level plan and low-level execution and propose the DoReMi framework, which enables immediate detection and recovery from plan-execution misalignments for complex long-horizon robotic problems.
- We explicitly maintain a constraint set through natural language description and adopt the VQA model as a general constraint detector to timely detect plan-execution misalignment and re-plan.
- Theoretical analyses and experiments on various complex robotic tasks verify the effectiveness of DoReMi, which results in less task execution time and higher success rates.

2   Related Works

Language Grounding Prior research has attempted to employ language as task abstractions and acquire control policies that are conditioned on language [13, 14, 15, 16, 17]. Furthermore, some

studies have investigated the integration of language and vision inputs within embodied tasks to directly predict the control commands [18, 19, 20]. Recent works, including [5, 19, 8, 21, 18, 22], have demonstrated significant progress in utilizing transformer-based policies to predict actions. However, these end-to-end approaches heavily depend on the scale of expert demonstrations for model training and often encounter challenges when generalizing to unseen scenarios.

Another line of work adopts a hierarchical framework and demonstrates that language models can perform high-level task planning in a zero/few-shot manner [2, 1] with appropriate grounding. Specifically, SayCan [1] incorporates an autoregressive language model and affordance functions to determine the most reasonable and executable skill from pretrained skill sets. These skill sets are trained using reinforcement learning or imitation learning [23, 24]. Huang et al. [2] leverages semantic translation to identify the skill closest to the desired output. Additionally, vision information can be integrated into prompts to assist the natural language model in the planning process [25, 26]. The above methods typically assume successful execution of each step, resulting in an open-loop system. In contrast, Inner-Monologue [10] takes into account various environmental feedback (e.g. success detectors and scene descriptions [11, 27]) upon the completion of each step. Our method emphasizes the importance of immediate detection and recovery from plan-execution misalignment, which can be highly beneficial in certain tasks.

Vision Language Model for Embodied Control. The visual language model (VLM) is trained on image-text pairs, enabling it to simultaneously understand visual and textual inputs and address a variety of downstream tasks, such as visual question answering (VQA) [12, 11], image description [28], and object detection [25]. VLM models align semantic information between visual and natural language, thereby aiding in grounding language models and facilitating embodied control. Pretrained models such as CLIP [29] have been integrated into diverse embodied tasks [30]. CLIPort [23], which incorporates Transporter [24] with CLIP, effectively combines spatial precision and semantic understanding. The PaLM-E model [31], equipped with pretrained vision transformer [32] and PaLM model [33], can leverage multimodal inputs and generate textual plans directly. Inner-Monologue [10] assumes perfect VLMs as success detectors and scene descriptors to obtain task feedback. To ensure adherence to crucial constraints, we employ the VQA model [12] as a "constraint detector", periodically verifying whether the agent satisfies specific constraints.

3   Problem Statement

Our objective is to enable the embodied agent to accomplish long-horizon tasks specified as natural language instruction i. The agent has a basic skill set Π, with each skill $\pi_j \in \Pi$ corresponding to a distinct function that can be described in natural language l_{π_j} Pretrained large language models, which have the ability to understand and reason, are used as planners [1, 2] to decompose complicated language instructions into basic skill sequences: $i \rightarrow (\pi_1, \pi_2, ..., \pi_n)$ as shown in Figure 2.
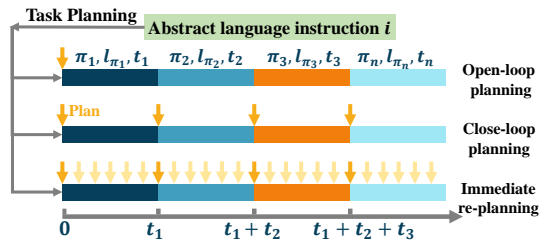

Figure 2: Comparison of the planning time.

We denote low-level skill execution time as $(t_1, t_2, ..., t_n)$. Planning only at time $t = 0$ will lead to completely open-loop planning without the chance of re-planning, making it susceptible to environmental disturbances or intermediate failures. Previous works often consider language model planning at the time when the previous skill is finished and the next one is required [1, 10], which indicates close-loop planning at time $t \in \{0, t_1, t_1 + t_2, ..., \Sigma_{k=1}^{n} t_k\}$.

However, planning only at these switching times can be inefficient, and immediate re-planning and recovery during the skill execution process can be critical in certain situations (e.g., box drop example in Figure .1). Therefore, instead of planning only at the switching time points $t \in \{0, t_1, t_1 + t_2, ..., \Sigma_{k=1}^{n} t_k\}$, our goal is to develop a framework that enables immediate detection and re-planning throughout the entire time period $t \in [0, \Sigma_{k=1}^{n} t_k]$. This allows for timely re-planning in case of plan-execution misalignment, such as the box being dropped in the above scenario, leading to quicker recovery. A comparison of the planning time can be found in Figure. 2. In the following section, we will introduce our DoReMi framework in detail, which facilitates immediate detection and recovery from misalignment throughout the entire low-level execution time period $t \in [0, \Sigma_{k=1}^{n} t_k]$.

3

4 Method

In this section, we introduce our **DoReMi** framework which enables immediate **D**etection and **Re**covery from Plan-Execution **Mi**salignment. Our algorithm can be succinctly described in three stages depicted in Figure 3:

1. Given a set of low-level skills, scene description, and high-level task instruction, language models are utilized to generate executable high-level plans.

2. We explicitly maintain a constraint set through natural language which indicates plan-execution misalignment.

3. During low-level skill execution, we employ the VQA model [12] as a general "constraint detector" that periodically verifies the satisfaction of all constraints. If any constraints are violated, the language model is called for immediate re-planning to help recovery.
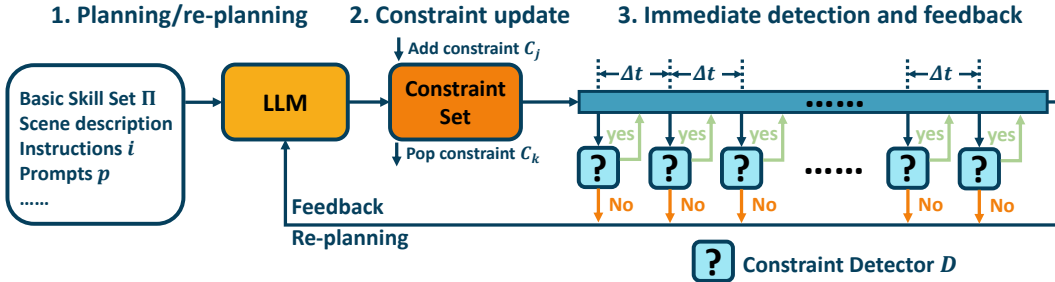

Figure 3: Our DoReMi framework mainly contains 3 components: (1) Task planning/re-planning with grounded language models; (2) Explicit maintenance of a constraint set that updates according to the planned step; (3) Utilization of VQA as general constraint detectors to identify misalignment between the plan and execution."

4.1 Language Model for Task Planning

Previous work has demonstrated that language models can plan executable skill sequences through appropriate grounding in zero/few-shot manners [26, 2], as shown in Figure 1(a). Similar to previous works [1, 10], we add instruction i, scene description, and skill language description to prompts to assist language models in high-level task planning. We try large language model GPT-4 [34] and Vicuna-13B [35]. Besides forward planning, we also use language models for re-planning if our constraint detector identifies a plan-execution misalignment. Under this scenario, we additionally include the misalignment information in prompts and then call the language model for re-plan. Through this mechanism, our agent can quickly recover from stages where the high-level plan and low-level execution are misaligned.

4.2 Constraint Set

During low-level skill execution, there are some constraints that the agent must satisfy. Violation of these constraints can indicate misalignments between the high-level plan and low-level execution. For instance, after performing the skill "pick up box", the agent must satisfy the constraint "hold box". If this constraint is violated, it may indicate an unsuccessful execution of the "pick up" skill or a dropped box. These constraints cover a variety of types. To check the constraints in a general manner without the need for specific designs or hard coding, we utilize a visual question answering (VQA) model [11, 12] as a general "constraint detector".

However, solely relying on open-ended questions such as "Is there any exception happening?" without specifying the form of the constraint can introduce ambiguity, making it challenging for the VQA model to determine plan-execution alignment. To tackle this, we propose using natural language to explicitly specify constraints for each skill. As these constraints are described in natural language, they can be easily understood and modified by humans based on requirements. For instance, some common robot skills can be categorized as shown in Table 1 [36], along with their respective constraint update processes. Navigation skills can be represented in `goto{recep}`. Manipulation motion primitives parameterized in two poses [37, 38] can be represented in `pick` and `place`. We have also explored using language models to automatically generate constraints for skills. However,

Low-level Skill	Constraint Update	VQA Question
`goto {recep}`	+no obstacle in the front	+Is there any obstacle in the front?
`pick{obj}from{recep}`	+**{agent}**holding**{obj}**	+Is the **{agent}**holding**{obj}**?
`place{obj}on{recep}`	+**{obj}**on**{recep}** -**{agent}**holding**{obj}**	+Is the **{obj}**on**{recep}**? -Is the **{agent}**holding**{recep}**?
`open{obj}`	+**{obj}** open -**{obj}**close	+Is the **{obj}** open? -Is the **{obj}**close?
`close{obj}`	-**{obj}** open +**{obj}**close	-Is the **{obj}** open? +Is the **{obj}**close?

Table 1: Constraint set update rule based on the natural language description. **{obj}** and **{recep}** correspond to objects and receptacles respectively. The symbol "**+**" indicates the addition of the constraint into the set while "**-**" means popping out this constraint. Questions are in the general structure: "Is the {constraint}?"

Algorithm 1 DoReMi (Immediate **De**tection and **Re**covery from **Mi**salignment)

Given: A high level instruction i, a skill set Π, language description l_Π for Π, language model L as task planner with initial prompt p_0, and VQA model as constraint detector D.

1: Initialize the constraint set $\mathbb{C} \leftarrow \emptyset$, the skill sequence $\pi \leftarrow \emptyset$, the number of steps $n \leftarrow 1$.
2: **while** $l_{\pi_{n-1}} \neq done$ **do**
3:     $\pi_n \leftarrow \arg\max_{\pi \in \Pi} L(l_\pi | i, p_{n-1}, l_{\pi_{n-1}}, ..l_{\pi_0})$
4:     $p_n \leftarrow p_{n-1}$
5:     Update $\mathbb{C}$ according to π_n.
6:     **while** π_n is not finished **do**
7:         Every Δt second, query agent all the constraints $c_j \in \mathbb{C}$ using the constraint detector D.
8:         **if** $\exists D(c_j) = $ **false then**
9:             Add constraint violate information into prompt p_n and **break**.
10:         **end if**
11:     **end while**
12:     Update the prompt p_n with the status of π_n, and $n \leftarrow n + 1$.
13: **end while**

---

their outputs result in large ambiguity. Therefore, in this paper, we adopt precise natural language descriptions as constraints, as shown in Table 1.

4.3 Constraint Detector and Re-planning

After the constraint set updating stage, the agent proceeds to execute the low-level skill selected by the language model. Our algorithm employs a visual question answering (VQA) model, denoted as D, which acts as a "constraint detector". It periodically checks whether the agent satisfies all constraints in the set. The visual input is captured from first-person or third-person perspective cameras. For each constraint c_j, the question input is formulated in a general structure "Is the *constraint c_j satisfied*?" The exact phrasing of the questions may vary based on grammatical considerations. For instance, for the constraint "*agent holding box*", the question would be "Is the *agent holding the box*?" Similarly, for the constraint "*green block on red block*", the question becomes "Is the *green block on red block*?". We use $D(v_t, c_j)$ to denote the answer of the VQA model D when checking c_j with vision input v_t at time t. If the constraint c_j is satisfied at time t, $D(v_t, c_j) = True$; otherwise, $D(v_t, c_j) = False$.

Every Δt second, we use the VQA model to examine whether the constraint set is fully satisfied. If so, the robot continues executing the current low-level skill; otherwise, the robot aborts the current skill, and the re-planning process is triggered. In practice, the zero-shot VQA model may make errors with a single image. Therefore, we employ an ensemble approach by considering multiple time steps and views from different cameras.

To summarize, our method has three main components: (1) We use language models as task planners. (2) We maintain a constraint set in natural language and use the VQA model to check constraint satisfaction. (3) If any constraint is violated, we trigger timely re-planning with the language model. The pseudo-code is provided in Algorithm 1.

4.4 Theoretical Analyses

We analyze the potential time savings and success rate improvements achievable through immediate detection and recovery. We denote the execution time of low-level skill with random variable t with mean $\mathbb{E}[t] = \mu$ and variance $Var(t) = \sigma^2$. Misalignment can occur at any time s within the execution time interval $[0, t]$ where $0 \le s \le t$. We define the discrete random variable M as the number of misalignment occurrences under the following assumptions: (1) Plan-execution misalignments occur independently. (2) Misalignments occur at a constant ratio λ within a small time interval: $\lim_{t \to 0} P(M = 1) = \lambda t$. (3) No two misalignments occur simultaneously: $\lim_{t \to 0} P(M = k) = 0$ for $k > 1$. Under these assumptions, the number of plan-execution misalignments follows a Poisson distribution [39]:

$$P(M = k) = \frac{(\lambda t)^k e^{-\lambda t}}{k!} \quad k = 0, 1, 2, 3... \tag{1}$$

Our analysis considers two types of misalignments: **(1) Soft misalignment**: If this occurs at time s, an agent without immediate re-planning must recover to the stage at time s, wasting time from s to t (e.g., a robot that drops an object must return to the drop location to pick it up). We assume DoReMi can detect this misalignment within Δt second and recover immediately. **(2) Critical misalignment**: If this occurs at time s, a delayed re-planning invariably results in failure; only immediate re-planning can address this misalignment (e.g., unexpected obstacles appear in front of the robot).

Theorem 1 *The following equations describe the wasted time t_w under soft misalignment and the failure probability P_f under critical misalignment without immediate detection and re-planning:*

$$\mathbb{E}(t_w) = \sum_k P(M = k)\mathbb{E}(t_w|M = k) = \frac{\lambda(\mu^2 + \sigma^2)}{2} - \lambda\mu\Delta t \tag{2}$$

$$\mathbb{E}(P_f) = 1 - \mathbb{E}(e^{-\lambda t}) \approx \lambda\mu - \frac{\lambda^2(\mu^2 + \sigma^2)}{2} \tag{3}$$

The detector's reaction time, Δt, is much smaller than the average execution time μ, so $\mathbb{E}(t_w)$ is greater than 0. λ represents the misalignment occurrence ratio per second, which is very small, so $\mathbb{E}(P_f)$ is also greater than 0. Under perfect detection and re-planning conditions, DoReMi can reduce the failure probability to 0, which is significantly less than $\mathbb{E}(P_f)$. Detailed proof can be found in Appendix A.

5 Experiments

In this section, we conduct experiments involving both robotic arm manipulation tasks and humanoid robot tasks, as shown in Figure 4. These tasks incorporate various environmental disturbances and imperfect controllers, such as random dropping by the robot arms, noise in end-effector placement positions causing stacked blocks to collapse, and unexpected obstacles appearing in the robot's path.

We aim to answer the following questions: (1) Does the DoReMi framework enable immediate detection and recovery from plan-execution misalignment? (2) Does DoReMi lead to higher task success rates and shorter task execution time under environmental disturbances or imperfect controllers?

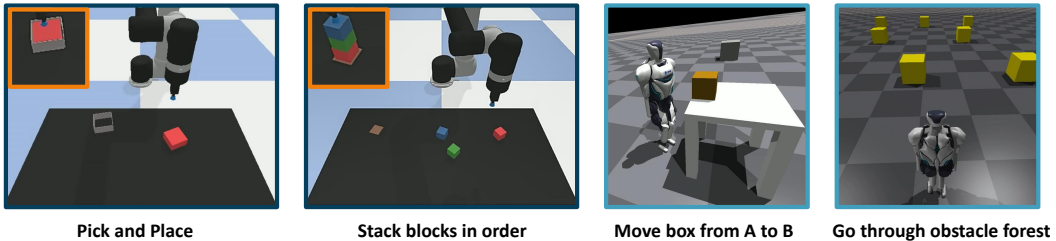| Pick and Place | Stack blocks in order | Move box from A to B | Go through obstacle forest |

Figure 4: Robot manipulation and humanoid robot tasks in our experiments. We consider various types of environmental disturbance and imperfect controllers.

5.1 Robot Arm Manipulation Tasks

Robot and Environment This environment is adapted from *Ravens* [38], a benchmark for vision-based robotic manipulation focused on pick-and-place tasks. A UR5e robot with a suction gripper works on a 0.5×1m black tabletop. Two cameras are pointing toward the workspace: one in front (facing the robot) and one in the left (facing the task area), providing 2 simulated 640×480 RGB-D images at every time step. We inject additional disturbances to the original environment and the robot controller to validate our algorithm's effectiveness.

Task 1: Pick and Place. The agent is required to pick up an $8 \times 8 \times 6$ cm cuboid block and place it in a fixture. The agent determines the pick position and place position (a_{pick}, a_{place}) based on the bird's-eye view image, and then the pick-place primitive generates the corresponding trajectory. However, we assume the block has a probability p to drop every second when sucked by the end-effector, so the agent may need to try several times to finish the task.

Task 2: Stack blocks in order. The robot is required to stack the $4 \times 4 \times 4$ cm blocks in an order given by language instructions. The agent decomposes the instruction into multiple "pick and place" sequences and executes them sequentially. We assume the controllers are not perfect by introducing uniform $[0, n]$ cm noise to the place positions. There is also a probability p that a block held by the end-effector might randomly drop every second.

5.1.1 Main Result

We maintain a constraint set to indicate the alignment between the plan and execution as described in 4.2 and use BLIP (VQA model) [12] to check constraint satisfaction using camera images. Our low-level skills are trained pick-and-place primitives conditioned on single-step instructions similar to the CLIPort [23] and Transporter Nets [38]. We compared DoReMi with 3 baselines: **(1) SayCan**: large language model (LLM) decomposes instructions into steps and executes them sequentially. However, this approach assumes the successful execution of each step without considering potential failures. **(2) CLIPort**: a multi-task CLIPort policy conditioned on the single pick-place step. It utilizes LLM to decompose instructions into steps, and our VQA is used as a success detector to check whether the current step should be repeated until success. However, CLIPort cannot perform re-planning during execution. **(3) Inner-Monologue**: similar to CLIPort, but can re-plan at the end of each step. Results are shown in Table 2.

5.1.2 Analyses

In the single-step pick-and-place task, all methods perform well when the controllers are perfect. However, with the introduction of random drop disturbances, SayCan may fail due to its lack of success detectors and re-planning capabilities. CLIPort and Inner-Monologue, while able to repeat steps until success, require more time as they cannot detect the drop immediately and only retry after completing the previous trajectory. In contrast, DoReMi can immediately detect the drop and re-plan the trajectory, resulting in lower execution times.

Stack-in-orders are multi-step tasks that require planning. SayCan consistently fails under various disturbances, including block-drop or block-collapse, due to its inability to manage plan-execution misalignments. CLIPort, armed with success detectors, can repeat the current step until success, making it more resilient to disturbances like block drops. However, in situations where re-planning is necessary, naively repeating the current step would lead to failure. For example, as depicted in Figure 5, when the robot attempts to stack the third blue block, all blocks collapse. In such cases, repeating the current step would result in failure due to the incorrect stack order. Although Inner-Monologue can re-plan at the end of each sub-task, our experiments revealed that Inner-Monologue often fails to re-plan correctly due to the imperfectness of the VQA detector in recognizing failures. There is a possibility that certain misalignments may remain undetected at the moment when the last step is finished. For example, failure to detect a block collapse in progress can directly result in an error during re-planning. In contrast, DoReMi achieves significantly better performance by continuously detecting and recovering from failures, which are more robust under imperfect VQA detectors. Furthermore, DoReMi can alleviate the high variance of single VQA detection through ensembling multi-step detection results (see Appendix C.1 for full details), which is enabled by our continuously detecting and re-planning mechanism.

Tasks with disturbance		Success Rate(%)↑				Execution Time(s)↓			
		SayCan	CLIPort	Inner Monologue	DoReMi (ours)	SayCan	CLIPort	Inner Monologue	DoReMi (ours)
Pick and place with random drop p	$p=0.0$	100($\pm$0)	100($\pm$0)	100($\pm$0)	100($\pm$0)	**2.7($\pm$0.0)**	**2.7($\pm$0.0)**	**2.7($\pm$0.0)**	**2.7($\pm$0.0)**
	$p=0.1$	96($\pm$4)	100($\pm$0)	100($\pm$0)	100($\pm$0)	3.2($\pm$0.4)	2.9($\pm$0.1)	2.9($\pm$0.1)	**2.8($\pm$0.1)**
	$p=0.2$	81($\pm$9)	100($\pm$0)	100($\pm$0)	100($\pm$0)	4.7($\pm$1.0)	3.4($\pm$0.2)	3.4($\pm$0.2)	**3.0($\pm$0.2)**
	$p=0.3$	63($\pm$9)	100($\pm$0)	100($\pm$0)	100($\pm$0)	6.6($\pm$1.3)	4.0($\pm$0.2)	4.0($\pm$0.2)	**3.3($\pm$0.2)**
Stack in order with noise n	$n=0.0$	**100($\pm$0)**	**100($\pm$0)**	**100($\pm$0)**	100($\pm$0)	**7.2($\pm$0.0)**	**7.2($\pm$0.0)**	**7.2($\pm$0.0)**	7.2($\pm$0.0)
	$n=1.0$	96($\pm$4)	96($\pm$4)	96($\pm$4)	**100($\pm$0)**	8.0($\pm$3.0)	8.0($\pm$3.0)	8.0($\pm$3.0)	**7.5($\pm$0.5)**
	$n=2.0$	63($\pm$9)	85($\pm$7)	88($\pm$7)	**96($\pm$4)**	26.1($\pm$7.7)	16.4($\pm$5.7)	15.2($\pm$5.3)	**10.2($\pm$1.7)**
	$n=3.0$	31($\pm$11)	75($\pm$10)	79($\pm$9)	**83($\pm$8)**	42.6($\pm$7.0)	23.2($\pm$6.0)	22.7($\pm$5.7)	**16.6($\pm$3.2)**
Stack in order with noise n random drop $p=0.1$	$n=0.0$	71($\pm$9)	94($\pm$7)	**98($\pm$4)**	**98($\pm$4)**	22.4($\pm$6.8)	11.0($\pm$3.6)	9.6($\pm$2.7)	**8.4($\pm$1.7)**
	$n=1.0$	71($\pm$9)	**94($\pm$7)**	**94($\pm$7)**	**94($\pm$7)**	21.4($\pm$7.1)	10.7($\pm$3.9)	10.2($\pm$3.8)	**9.6($\pm$3.2)**
	$n=2.0$	54($\pm$12)	79($\pm$9)	81($\pm$9)	**92($\pm$6)**	29.8($\pm$7.5)	19.5($\pm$6.7)	18.7($\pm$6.5)	**11.5($\pm$3.4)**
	$n=3.0$	21($\pm$9)	33($\pm$10)	42($\pm$10)	**50($\pm$10)**	48.0($\pm$6.4)	42.8($\pm$7.0)	40.2($\pm$7.0)	**26.8($\pm$4.3)**

Table 2: Success rates and task execution time under different degrees of disturbances. If the task fails, the execution time is set to timeout, which is 10s for the pick-place task and 60s for the stack-in-order task. The results show the mean and standard deviation over 4 different seeds each with 12 episodes.
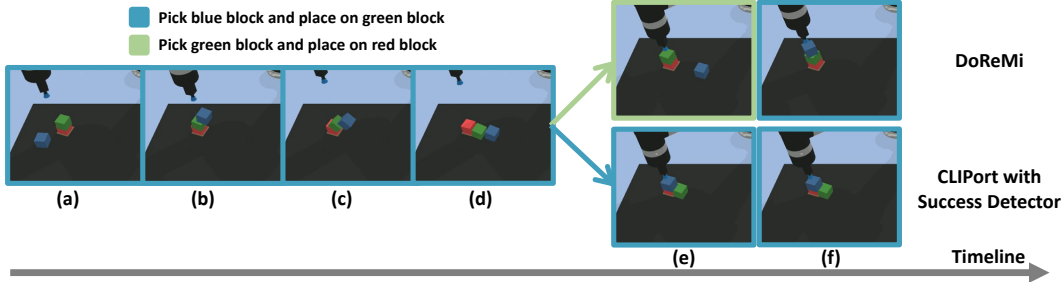

Figure 5: A comparison example. The robot arm tries to finish the step "Stack the blue block on the green block" but collapses (bcd). DoReMi detects this misalignment and replans to pick and place the green block first (e). The baseline continues to repeat the previous step (ef) and results in failure.

5.2 Humanoid Robot Tasks

Robot Description. The humanoid robot utilized in our experiments possesses 5 degrees of freedom per leg and 3 degrees of freedom per arm, totaling 16 degrees of freedom. We equip the robot with a first-view camera on its base to provide visual information. Detailed robot information can be found in Appendix B.

Task 1: Move Box from Place A to Place B. This long-horizon and challenging task involves both navigation and manipulation, commanding the robot to transport a box from one location to another. A proper solution might involve (1) Go to place A. (2) Pick up box (3) Go to place B. (4) Put down box. We introduced additional perturbations to this task by assuming that the robot has a probability p of dropping the box every second during transport.

Task 2: Go through Obstacle Forest. The robot is commanded to move forward, despite the risk of running into unexpected obstacles in its path. The robot needs to continuously detect whether there are obstacles in the front and change direction to avoid collision.

5.2.1 Train Low-level Skills

Controlling complex humanoid robots with a single policy is challenging. Thus, we train low-level skills at the category level. Following the framework in [40], we utilize reinforcement learning to train locomotion policy and use model-based methods to obtain manipulation policy. Specifically, we train 3 categories of policy: **(1) A locomotion neural network policy** conditioned on commanded linear and angular velocity, allowing the robot to execute low-level skills such as "Go forward fast", "Go forward at speed v", "Stand still", "Turn right/left", "Go to target place A", etc. **(2) A stand/squat neural network policy** conditioned on the commanded height, enabling skills like "Stand", "Squat", "Raise height with length l", etc. **(3) A hand manipulation policy** based on a linear interpolation controller, using the Deepmimic [41] algorithm with multiple motion capture data (Mocap data) to obtain policies that can produce natural behaviors. Detailed architecture and training process can be found in Appendix B.

Tasks with disturbance		Success Rate (%) ↑			Excution Time (s) ↓		
		SayCan	Inner Monologue	DoReMi (ours)	SayCan	Inner Monologue	DoReMi (ours)
Move box from A to B with random drop p	p=0.0	98($\pm$3)	98($\pm$3)	97($\pm$4)	23.8($\pm$0.7)	23.8($\pm$0.7)	24.4($\pm$1.4)
	p=0.02	61($\pm$4)	88($\pm$6)	92($\pm$5)	53.0($\pm$7.2)	44.3($\pm$7.7)	36.9($\pm$6.4)
	p=0.04	43($\pm$12)	82($\pm$8)	88($\pm$8)	67.4($\pm$3.3)	57.2($\pm$6.0)	43.2($\pm$6.7)
Go through obstacle forest with density d	d=0.0	100($\pm$0)	100($\pm$0)	100($\pm$0)	27.3($\pm$0.1)	27.3($\pm$0.1)	27.3($\pm$0.1)
	d=0.3	68($\pm$6)	68($\pm$6)	92($\pm$6)	52.4($\pm$5.0)	52.4($\pm$5.0)	38.5($\pm$5.0)
	d=0.6	40($\pm$8)	40($\pm$8)	90($\pm$4)	64.9($\pm$5.1)	64.9($\pm$5.1)	41.5($\pm$3.0)

Table 3: Success rates and task execution time under different degrees of disturbances. If tasks failed, execution time is set to timeout which is equal to 90s. The results show the mean and standard deviation over 5 different seeds each with 12 episodes.
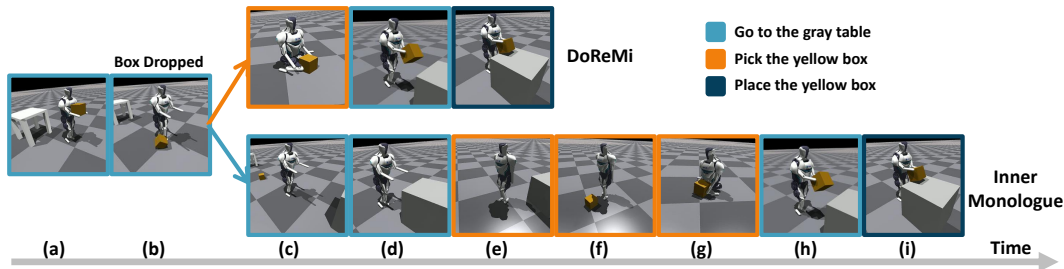

Figure 6: A comparison of DoReMi and baseline. Plan-execution misalignment occurs at the time (b). DoReMi immediately detects box drop and replans to pick up the box at the time (b). The baseline with delayed feedback keeps going forward and costs more time to pick the box at the time (c)-(g).

5.2.2 Main Result and Analyses

We follow the pipeline mentioned in the section 3. During low-level skill execution, the BLIP [12] (VQA model) consistently queries the first-view images to verify constraint satisfaction. If a constraint is violated or the current step completes, the language model is invoked for re-planning. The full prompt used can be found in Appendix D. We compare our method with **(1) SayCan** [1] which assumes every step is executed successfully, and **(2) Inner-Monologue** [10] which plans at the end of each step and uses the same VQA model as a success detector and scene descriptor.

Analyses As shown in Table 3, when the task disturbance is set to 0, indicating the perfect execution of each step, all methods exhibit high success rates as long as the language model plans correctly. However, when plan-execution misalignments occur, such as a box dropped on the floor or unexpected obstacles appearing in the path, both SayCan and Inner-Monologue struggle. SayCan fails to handle these misalignments, while Inner-Monologue's delayed re-planning leads to longer task completion time and decreased success rate. Figure 6 provides a visual comparison of the Move-box task. In this example, if the box drops (b) while the agent is executing the "Go to the gray table" skill, the agent without immediate detection continues to complete the current skill (b)-(d). Upon reaching the grey table (d), the robot re-plans but encounters difficulties in retrieving the halfway-dropped box (e)-(g). On the other hand, DoReMi immediately detects the drop and picks up the box, resulting in shorter task completion time.

6 Discussion

Limitation Our experiments indicate that the zero-shot transfer VQA model is not a perfect constraint detector. We need to employ an ensembling approach to improve detection accuracy. Our framework can benefit from a more advanced vision-language model in the future. Furthermore, our method maintains a constraint set through specified natural language to indicate plan-execution misalignments. There is potential for language models to automatically generate these constraints in future implementations.

Conclusion When employing language models for task planning in a hierarchical approach, the low-level execution might deviate from the high-level plan. We emphasised the importance of continuously aligning the plan with execution and proposed the DoReMi framework, which grounds language through immediate detection and recovery from plan-execution misalignments. Theoretical analyses and a variety of challenging tasks in disturbed environments demonstrated the effectiveness of DoReMi.

References

[1] Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, et al. Do as i can, not as i say: Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022.

[2] Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. Language models as zero-shot planners: Extracting actionable knowledge for embodied agents. In *International Conference on Machine Learning*, pages 9118–9147. PMLR, 2022.

[3] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.

[4] 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. *arXiv preprint arXiv:2209.07753*, 2022.

[5] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Jasmine Hsu, et al. Rt-1: Robotics transformer for real-world control at scale. *arXiv preprint arXiv:2212.06817*, 2022.

[6] Scott Reed, Konrad Zolna, Emilio Parisotto, Sergio Gomez Colmenarejo, Alexander Novikov, Gabriel Barth-Maron, Mai Gimenez, Yury Sulsky, Jackie Kay, Jost Tobias Springenberg, et al. A generalist agent. *arXiv preprint arXiv:2205.06175*, 2022.

[7] Eric Jang, Alex Irpan, Mohi Khansari, Daniel Kappler, Frederik Ebert, Corey Lynch, Sergey Levine, and Chelsea Finn. Bc-z: Zero-shot task generalization with robotic imitation learning. In *Conference on Robot Learning*, pages 991–1002. PMLR, 2022.

[8] Mohit Shridhar, Lucas Manuelli, and Dieter Fox. Perceiver-actor: A multi-task transformer for robotic manipulation. In *Conference on Robot Learning*, pages 785–799. PMLR, 2023.

[9] Suraj Nair, Eric Mitchell, Kevin Chen, Silvio Savarese, Chelsea Finn, et al. Learning language-conditioned robot behavior from offline data and crowd-sourced annotation. In *Conference on Robot Learning*, pages 1303–1315. PMLR, 2022.

[10] 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.

[11] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433, 2015.

[12] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pages 12888–12900. PMLR, 2022.

[13] Matt MacMahon, Brian Stankiewicz, and Benjamin Kuipers. Walk the talk: Connecting language, knowledge, and action in route instructions. *Def*, 2(6):4, 2006.

[14] Devendra Singh Chaplot, Kanthashree Mysore Sathyendra, Rama Kumar Pasumarthi, Dheeraj Rajagopal, and Ruslan Salakhutdinov. Gated-attention architectures for task-oriented language grounding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.

[15] Yiding Jiang, Shixiang Shane Gu, Kevin P Murphy, and Chelsea Finn. Language as an abstraction for hierarchical deep reinforcement learning. *Advances in Neural Information Processing Systems*, 32, 2019.

[16] Dipendra Misra, John Langford, and Yoav Artzi. Mapping instructions and visual observations to actions with reinforcement learning. *arXiv preprint arXiv:1704.08795*, 2017.

[17] Hongyuan Mei, Mohit Bansal, and Matthew Walter. Listen, attend, and walk: Neural mapping of navigational instructions to action sequences. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, 2016.

[18] Andrew Silva, Nina Moorman, William Silva, Zulfiqar Zaidi, Nakul Gopalan, and Matthew Gombolay. Lancon-learn: Learning with language to enable generalization in multi-task manipulation. *IEEE Robotics and Automation Letters*, 7(2):1635–1642, 2021.

[19] Pierre-Louis Guhur, Shizhe Chen, Ricardo Garcia Pinel, Makarand Tapaswi, Ivan Laptev, and Cordelia Schmid. Instruction-driven history-aware policies for robotic manipulations. In *Conference on Robot Learning*, pages 175–187. PMLR, 2023.

[20] Prasoon Goyal, Scott Niekum, and Raymond Mooney. Pixl2r: Guiding reinforcement learning using natural language by mapping pixels to rewards. In *Conference on Robot Learning*, pages 485–497. PMLR, 2021.

[21] Yichi Zhang and Joyce Chai. Hierarchical task learning from language instructions with unified transformers and self-monitoring. *arXiv preprint arXiv:2106.03427*, 2021.

[22] Corey Lynch, Ayzaan Wahid, Jonathan Tompson, Tianli Ding, James Betker, Robert Baruch, Travis Armstrong, and Pete Florence. Interactive language: Talking to robots in real time. *arXiv preprint arXiv:2210.06407*, 2022.

[23] Mohit Shridhar, Lucas Manuelli, and Dieter Fox. Cliport: What and where pathways for robotic manipulation. In *Conference on Robot Learning*, pages 894–906. PMLR, 2022.

[24] Andy Zeng, Pete Florence, Jonathan Tompson, Stefan Welker, Jonathan Chien, Maria Attarian, Travis Armstrong, Ivan Krasin, Dan Duong, Vikas Sindhwani, et al. Transporter networks: Rearranging the visual world for robotic manipulation. In *Conference on Robot Learning*, pages 726–747. PMLR, 2021.

[25] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Open-vocabulary object detection via vision and language knowledge distillation. *arXiv preprint arXiv:2104.13921*, 2021.

[26] Andy Zeng, Adrian Wong, Stefan Welker, Krzysztof Choromanski, Federico Tombari, Aveek Purohit, Michael Ryoo, Vikas Sindhwani, Johnny Lee, Vincent Vanhoucke, et al. Socratic models: Composing zero-shot multimodal reasoning with language. *arXiv preprint arXiv:2204.00598*, 2022.

[27] Yuqing Du, Ksenia Konyushkova, Misha Denil, Akhil Raju, Jessica Landon, Felix Hill, Nando de Freitas, and Serkan Cabi. Vision-language models as success detectors. *arXiv preprint arXiv:2303.07280*, 2023.

[28] Luowei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason Corso, and Jianfeng Gao. Unified vision-language pre-training for image captioning and vqa. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 13041–13049, 2020.

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

[30] Apoorv Khandelwal, Luca Weihs, Roozbeh Mottaghi, and Aniruddha Kembhavi. Simple but effective: Clip embeddings for embodied ai. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14829–14838, 2022.

[31] Danny Driess, Fei Xia, Mehdi SM Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, et al. Palm-e: An embodied multimodal language model. *arXiv preprint arXiv:2303.03378*, 2023.

[32] Mostafa Dehghani, Josip Djolonga, Basil Mustafa, Piotr Padlewski, Jonathan Heek, Justin Gilmer, Andreas Steiner, Mathilde Caron, Robert Geirhos, Ibrahim Alabdulmohsin, et al. Scaling vision transformers to 22 billion parameters. *arXiv preprint arXiv:2302.05442*, 2023.

[33] Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*, 2022.

[34] OpenAI. Gpt-4 technical report. *arXiv*, 2023.

[35] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality, March 2023.

[36] Wei Gao and Russ Tedrake. kpam 2.0: Feedback control for category-level robotic manipulation. *IEEE Robotics and Automation Letters*, 6(2):2962–2969, 2021.

[37] Emilio Frazzoli, Munther A Dahleh, and Eric Feron. Maneuver-based motion planning for nonlinear systems with symmetries. *IEEE transactions on robotics*, 21(6):1077–1091, 2005.

[38] Andy Zeng, Pete Florence, Jonathan Tompson, Stefan Welker, Jonathan Chien, Maria Attarian, Travis Armstrong, Ivan Krasin, Dan Duong, Vikas Sindhwani, and Johnny Lee. Transporter networks: Rearranging the visual world for robotic manipulation. *Conference on Robot Learning (CoRL)*, 2020.

[39] Athanasios Papoulis and S Unnikrishna Pillai. *Probability, random variables and stochastic processes*. 2002.

[40] Yuntao Ma, Farbod Farshidian, Takahiro Miki, Joonho Lee, and Marco Hutter. Combining learning-based locomotion policy with model-based manipulation for legged mobile manipulators. *IEEE Robotics and Automation Letters*, 7(2):2377–2384, 2022.

[41] Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel Van de Panne. Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Transactions On Graphics (TOG)*, 37(4):1–14, 2018.

A   Proof for Theorem

The number of plan-execution misalignments follows a Poisson distribution [39]:

$$P(M = k) = \frac{(\lambda t)^k e^{-\lambda t}}{k!} \quad k = 0, 1, 2, 3... \tag{4}$$

Our analysis considers two types of misalignments: **(1) Soft misalignment**: If this occurs at time s, an agent without immediate re-planning must recover to the stage at time s, wasting time from s to t (e.g., a robot that drops an object must return to the drop location to pick it up). We assume DoReMi can detect this misalignment within Δt second and recover immediately. **(2) Critical misalignment**: If this occurs at time s, a delayed re-planning invariably results in failure; only immediate re-planning can address this misalignment (e.g., unexpected obstacles appear in front of the robot).

Theorem 1 *The following equations describe the wasted time t_w under soft misalignment and the failure probability P_f under critical misalignment without immediate detection and re-planning:*

$$\mathbb{E}(t_w) = \sum_k P(M = k)\mathbb{E}(t_w|M = k) = \frac{\lambda(\mu^2 + \sigma^2)}{2} - \lambda\mu\Delta t \tag{5}$$

$$\mathbb{E}(P_f) = 1 - \mathbb{E}(e^{-\lambda t}) \approx \lambda\mu - \frac{\lambda^2(\mu^2 + \sigma^2)}{2} \tag{6}$$

A.1   Lemma

Lemma 1 *Given a Poisson process which is conditional on n arrivals in the time interval $(0, t)$, the conditional pdf (probability density function) of event occurrence time $t_1, t_2..., t_n$ satisfy [39]:*

$$f(t_1, ..., t_n|M(t) = n) = \frac{n!}{t^n} \quad 0 \le t_1 \le ... \le t_n \le t \tag{7}$$

Proof Since the inter-arrival times of Poisson distribution are independent exponentially distributed, the joint pdf of the n first arrival times is:

$$\begin{aligned}
f(t_2, t_2, ..., t_n) &= f(t_1)f(t_2|t_1)...f(t_n|t_{n-1}) \\
&= \lambda e^{-\lambda t_1}\lambda e^{-\lambda(t_2 - t_1)}...\lambda e^{-\lambda(t_n - t_{n-1})} \\
&= \lambda^n e^{-\lambda t_n}
\end{aligned} \tag{8}$$

And conditional pdf can be derived:

$$\begin{aligned}
f(t_1, t_2, ..., t_n|M(t) = n) &= \frac{f(t_2, t_2, ..., t_n, M(t) = n)}{P(M(t) = n)} \\
&= \frac{f(t_1, t_2, ..., t_n)P(M(t) = n|t_1, t_2, ..., t_n)}{P(M(t) = n)} \\
&= \frac{\lambda^n e^{-\lambda t_n} e^{-\lambda(t - t_n)}}{e^{-\lambda t}(\lambda t)^n/n!} \\
&= \frac{n!}{t^n}
\end{aligned} \tag{9}$$

That is to say, each event can be considered as "placed" independently and uniformly at a given time in $[0, t]$.

A.2   Proof for Theorem 1

Soft Misalignment Based on lemma A.1, each event can be considered to occur independently and uniformly at a given time in $[0, t]$, the event that occurs at s will lead to time cost $t - s$. Total time

cost without immediate replanning ($E(M) = \sum P(M = k) * k = \lambda t$):

$$\mathbb{E}(t_{delay}) = \sum_{t_i}(t - t_i) = \sum_k P(M = k)\mathbb{E}_t(t_w|M = k)$$
$$= \mathbb{E}_t[\sum_k P(M = k) * kt/2] \tag{10}$$
$$= \mathbb{E}_t[\lambda t^2/2] = \lambda(\mu^2 + \sigma^2)/2$$

Time cost without immediate replan (every event have detection time Δt):

$$\mathbb{E}(t_{doremi}) = \sum_{t_i}\Delta t = \mathbb{E}_t[M] * \Delta t = \mathbb{E}_t[\lambda t\Delta t] = \lambda\mu\Delta t \tag{11}$$

The wasted time $\mathbb{E}(t_w)$ is the difference between $\mathbb{E}(t_{delay})$ and $\mathbb{E}(t_{doremi})$:

$$\mathbb{E}(t_w) = \mathbb{E}(t_{delay}) - \mathbb{E}(t_{doremi}) = \frac{\lambda(\mu^2 + \sigma^2)}{2} - \lambda\mu\Delta t \tag{12}$$

Critical misalignment Once critical misalignment comes, a delayed replanning will lead to failure. So the failure ratio P_f equals to the probability that the misalignment occurrence number is greater than 1. We assume misalignment happen ratio λ is very small and we use second-order Tyler expansion to approximate the original equation.

$$\mathbb{E}(P_f) = \mathbb{E}_t[\sum_{k\geq 1} P(M = k)] = \mathbb{E}_t[1 - P(M = 0)] = 1 - \mathbb{E}_t(e^{-\lambda t})$$
$$= 1 - \mathbb{E}_t(1 - \lambda t + \lambda^2 t^2/2 + ...) \approx \lambda\mathbb{E}_t(t) - \lambda^2\mathbb{E}_t(t^2) \tag{13}$$
$$= \lambda\mu - \frac{\lambda^2(\mu^2 + \sigma^2)}{2}$$

B Humanoid Robot Task

B.1 Basic Humanoid Robot Information

Our robot has 17 links and 16 degrees of joint freedom(DOF), and each joint hold a corresponding motor.

Link names: "base", "left shoulder pitch", "left shoulder roll", "left elbow", "right shoulder pitch", "right shoulder roll", "right elbow", "left leg yaw", "left leg roll", "left leg pitch", "left knee", "left ankle pitch", "right leg yaw", "right leg roll", "right leg pitch", "right knee", "right ankle pitch"

DOF joint names: "left shoulder pitch joint", "left shoulder roll joint", "left elbow joint", "right shoulder pitch joint", "right shoulder roll joint", "right elbow joint", "left leg yaw joint", "left leg roll joint", "left leg pitch joint", "left knee joint", "left ankle pitch joint", "right leg yaw joint", "right leg roll joint", "right leg pitch joint", "right knee joint", "right ankle pitch joint"

B.2 Low-level Skill Training

Controlling complex humanoid robots with a single policy is challenging. Thus, we train low-level skills at the category level. Following the separate framework in [40], we utilize reinforcement learning to train locomotion policy and use model-based methods to obtain manipulation policy. In the case of our humanoid robot, there are 10 motors dedicated to the legs and 6 motors allocated to the arms. Notably, the observation of arm motors is not incorporated into the locomotion policies.

The locomotion policy and squat policy are responsible for directly controlling the 10 motors associated with the legs, leading to a 10-dimensional action space. These policies output the target position of motors and run at 50 Hz, followed by PD controller run at 1000 Hz with $k_p = 40$ and $k_d = 1$. The proprioceptive observation space of the robot includes various dimensions: 10-dimensional joint angles, 10-dimensional joint angular velocities, 10-dimensional last actions, 3-dimensional angles between the torso and gravity, 2-dimensional periodic clock signals, and reserved 3-dimensional command signals, resulting in a total of 38 basic observation spaces.

We train low-level skills with the Deepmimic algorithm based on the Legged Gym (Isaac Gym Environments for Legged Robots) environment `https://github.com/leggedrobotics/legged_gym` built with the Isaac Sim physics simulator. Motion capture data we used can be found in the poselib `https://github.com/NVIDIA-Omniverse/IsaacGymEnvs/tree/main/isaacgymenvs/tasks/amp/poselib`.

Specifically, we train 3 categories of the policy:

Locomotion Policy This neural network policy is conditioned on 3-dimensional commands which respectively represent the required velocity in x-direction, y-direction, and required yaw angular velocity. In order to obtain natural moving gaits, we use the Deepmimic algorithm with multiple Motion Capture Date(Mocap data), thus reward function has 2 parts including tracking commanded linear/angular velocity and imitating the style of Mocap data. This learned policy can help robots realize a category of sub-skills related to locomotion like: "Go forward fast", "Go forward at speed v", "Stand still", "Turn right/left", "Go to target place A", etc.

Stand and Squat Policy This neural network policy is conditioned on 2-dimensional commands which represent the demanded base height and base angular velocity. We also use Deepmimic with multiple Mocap data. The learned policy can help the robot achieve a category of policy that need changes in height like: "Stand", "Squat", "Raise height with length l", etc.

Arm Manipulation Policy Since we separate the control of arm and leg, we can use various manipulation policies including learned neural network policy or model-based policy, without influencing the leg locomotion policy. We use a linear interpolation controller to achieve the skill: "Pick up box", "Pick up box on the floor", "Put down box", etc.

Hyperparameters Deepmimic algorithm pipeline is similar to PPO. Hyperparameters of the backbone Deepmimic algorithm can be found in table 4.

Parameters	Value
Number of Environments	4096
Learning epochs	5
Steps per Environment	24
Minibatch Size	24576
Epoch length	20 seconds
Discount Factor	0.99
Generalised Advantage Estimation(GAE)	0.95
PPO clip	0.2
Entropy coefficient	0.005
Desired KL	0.01
Learning Rate	5e-4
Weight decay	0.01

Table 4: Hyperparameters of backbone PPO algorithm.

Training curves The training process for the navigation policies and stand/squat policies is illustrated in Figure 7. The navigation policy enables the robot to control its xy position within the world frame, while the height switch policy allows for adjusting the robot's z height within the world frame.

B.3   Ensembling Multi-step Detection by VQA

VQA images are the first-view camera attached to the robot. Visualization of this camera can be found in Figure 8.

VQA questions basically follow the structure mentioned in the method section 3. For the move-box environment, when the constraint refers to "the robot holding the yellow box", our question would be"Is the robot holding the yellow box"?. The box color is randomly selected as shown in Figure 9, and the question is also modified according to the box color. For an obstacle-forest environment, the question would be "Is there any obstacle in the front?" If the answer is yes, the humanoid start the collision avoidance process. In practice, we will additionally ask the question "Is the obstacle in the
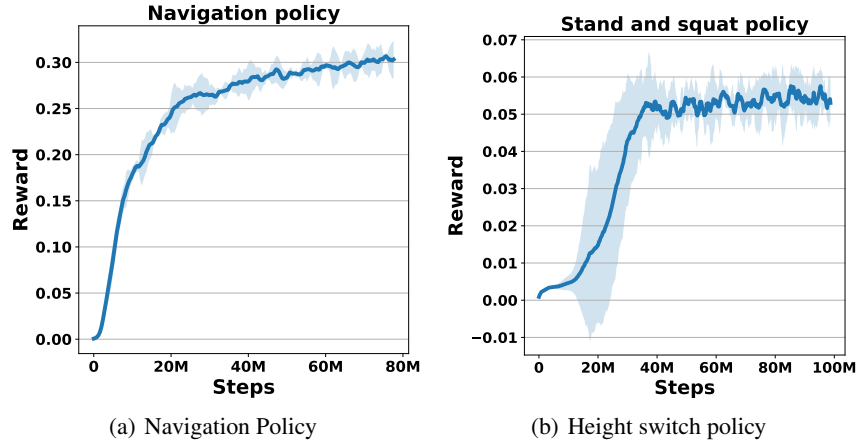
(a) Navigation Policy (b) Height switch policy

Figure 7: Traning curves for navigation policy and height switch policy by Deepmimic algorithm.

left or right?" This information is then provided to LLM and helps with replanning. The robot will turn left/right according to the language model output.

As mentioned in the paper, the VQA answer for a single vision-question pair may not always be correct. To enhance accuracy, an ensemble approach can be utilized by incorporating k consecutive frames through a neural network or a probabilistic model. However, in humanoid tasks, a simpler approach is employed. We detect constraint violations by considering 4 consecutive time-step images where VQA identifies the same constraint violation. The time step duration, denoted as Δt, is set to 0.1 seconds. Since the first-view images in these environments are quite simple, ensembled VQA answers can reach an accuracy higher than 95%.

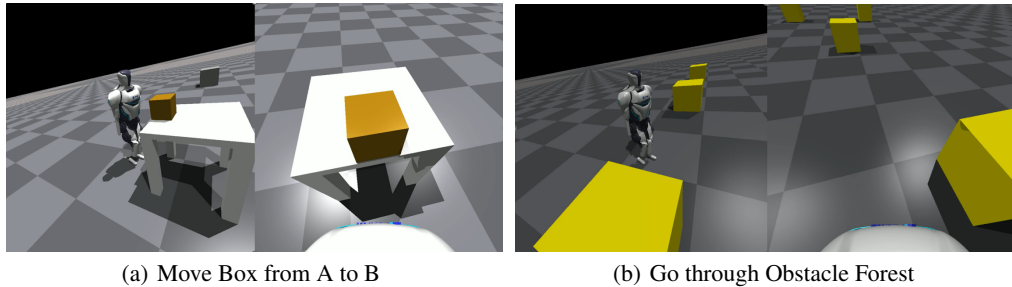
(a) Move Box from A to B (b) Go through Obstacle Forest

Figure 8: Robot and tasks. For every task, the photo on the Left is recorded from a third-view camera and the photo on the right is from the first-view camera. We use first-view images as inputs to VQA models.
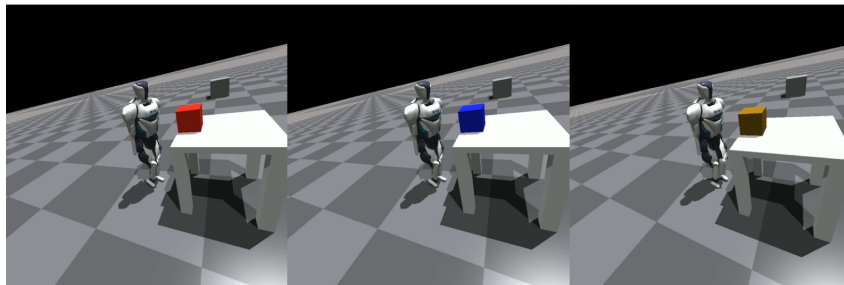

Figure 9: Different box colors.

16

B.4 Constraint update details

See Table 5.

Low-level Skill	Constraint Update	VQA Question
goto **{recep}**/go forward	+no obstacle in the front	+Is there any obstacle in the front?
pick**{obj}**from**{recep}**	**{agent}**holding**{obj}**+	Is the **{agent}**holding**{obj}**?+
place**{obj}**on**{recep}**	**{obj}**on**{recep}**+ **{agent}**holding**{obj}**-	Is the **{obj}**on**{recep}**?+ Is the **{agent}**holding**{recep}**?-

Table 5: Constraint set update rule based on the natural language template. **{obj}** and **{recep}** correspond to objects and receptacles. "**+**" in the front means adding this constraint into the set before skill execution while "**+**" in the end stands for adding this constraint into the set after skill is finished. Similarly, "**-**" means pop this constraint. Questions are in the general structure: "Is the {constraint}?"

B.5 Case study

We visualize different scenarios in Figure 10. First-view images are used to query the VQA. We utilize a green border to indicate situations where the VQA system believes no misalignment has occurred, while a red border is used to represent situations where the VQA system believes that certain constraints have been violated.
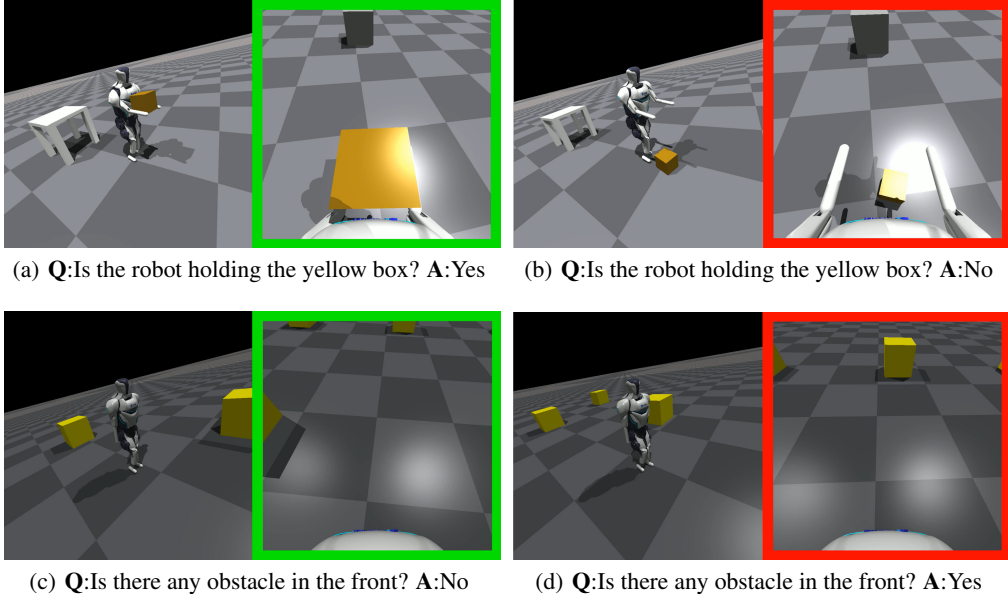

(a) **Q**:Is the robot holding the yellow box? **A**:Yes

(b) **Q**:Is the robot holding the yellow box? **A**:No

(c) **Q**:Is there any obstacle in the front? **A**:No

(d) **Q**:Is there any obstacle in the front? **A**:Yes

Figure 10: Case visualization.

C Robot Arm Manipulation Task

C.1 Implementation Details

Low-level Policy The low-level policy is similar to CLIPort[23] and Transporter Network [38]. Detailed references for its implementation can be found at `https://github.com/google-research/ravens`. This policy has been trained to perform single-step pick-and-place tasks based on language descriptions, and its performance is nearing perfection. However, for the purpose of our study, we presume this original policy to be perfect and introduce additional perturbations to the location placement.

Ensembling Multi-step Detection by VQA Similar to B.3, we also incorporate four consecutive frames to perform ensembled detection. For the pick-and-place task, the constraint is "the robot holding the red box", and the corresponding question is generated as "Is the robot holding the red box?". For the stack-block-in-order task, additional constraints are "the blue on the green box", "the green box on the red box" and "the red box on the brown fixture". Randomly changing the stack order has no effect on the task performance, as shown in Figure 11. We also clip the image fed to the VQA to enable faster inference.
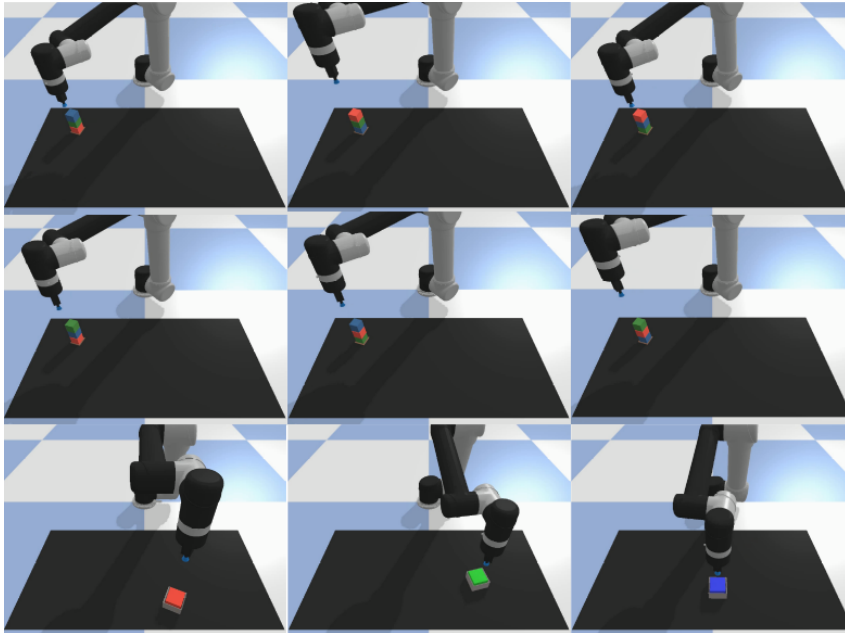

Figure 11: Our method is agnostic to different stack orders.

Baseline To adapt to our tasks, we slightly modify the original implementation of three baselines: (1) SayCan: the implementation is based on `https://github.com/google-research/google-research/tree/master/saycan`. The low-level policy is adopted the same as ours. (2) CLIPort: the implementation is based on `https://github.com/google-research/ravens`. The oracle success detector is replaced with our VQA detector. (3) Inner Monologue: we reproduce the implementation based on [10]. Both the low-level policy and LLM planner are the same as ours. The original success detector is also replaced with our VQA detector.

C.2 Ablation Study

We evaluated the robustness of DoReMi in various environmental conditions by testing our method under distinct levels of perturbations. We observed that when the positional noise level n of the end-effector exceeds 0.03 cm, the frequency of constraint violations escalates, leading to an almost zero success rate and dramatically increased execution time, despite accurate detection of constraint violations. This observation aligns with the theoretical expectation derived from a simple computation of block placement probabilities.

Our primary interest lies in the response of DoReMi to a spectrum of drop perturbation levels, in addition to the $p = 0.1$ presented in Table 2. As illustrated in Table 6, DoReMi demonstrates admirable performance under a variety of scenarios, outperforming the best-performing baseline. We attribute this largely to DoReMi's robust detection mechanism and its ability to swiftly recover from misalignment between plan and execution.

We are also curious to explore how the ensembling of multi-step detection would perform under various environmental settings. As indicated in Table 6, our findings suggest that ensembling can markedly enhance DoReMi's effectiveness in scenarios where strong perturbations exist and a single detection error could potentially result in a complete episode failure.

Tasks with disturbance		Success Rate(%) ↑			Execution Time(s) ↓		
		Inner Monologue	DoReMi (w/o ensembling)	DoReMi (ours)	Inner Monologue	DoReMi (w/o ensembling)	DoReMi (ours)
Stack in order with noise n random drop $p=0.05$	$n=0.0$	**100**($\pm$**0**)	**100**($\pm$**0**)	**100**($\pm$**0**)	7.9($\pm$0.7)	7.5($\pm$0.5)	**7.4**($\pm$**0.5**)
	$n=1.0$	94($\pm$7)	96($\pm$4)	**98**($\pm$**4**)	9.3($\pm$3.3)	8.6($\pm$2.9)	**8.1**($\pm$**1.0**)
	$n=2.0$	83($\pm$8)	88($\pm$7)	**94**($\pm$**7**)	17.3($\pm$5.8)	12.1($\pm$2.9)	**10.8**($\pm$**2.7**)
	$n=3.0$	63($\pm$9)	67($\pm$10)	**73**($\pm$**11**)	36.3($\pm$7.2)	25.8($\pm$7.1)	**19.9**($\pm$**3.9**)
Stack in order with noise n random drop $p=0.15$	$n=0.0$	92($\pm$6)	92($\pm$6)	**94**($\pm$**7**)	10.6($\pm$4.3)	9.7($\pm$3.2)	**8.9**($\pm$**2.2**)
	$n=1.0$	88($\pm$7)	90($\pm$7)	**92**($\pm$**6**)	14.8($\pm$5.1)	12.5($\pm$4.2)	**10.3**($\pm$**3.2**)
	$n=2.0$	73($\pm$11)	79($\pm$9)	**85**($\pm$**7**)	25.2($\pm$6.3)	21.3($\pm$5.7)	**14.0**($\pm$**3.7**)
	$n=3.0$	23($\pm$9)	33($\pm$10)	**44**($\pm$**11**)	47.8($\pm$6.5)	40.6($\pm$6.9)	**29.3**($\pm$**4.1**)

Table 6: Ablation study over different degrees of perturbations and whether adopt ensembling or not. The results show the mean and standard deviation over 4 different seeds each with 12 episodes.

C.3 VQA Detector's Accuracy Analysis

To analyze the accuracy of the VQA detector, we categorize all the detection results into True Positives (TP), True Negatives (TN), False Positives (FP), and False Negatives (FN), using these to calculate relevant accuracy metrics as outlined in Table 7.

A True Positive (TP) refers to the VQA correctly identifying that no constraint violation has occurred, whereas a True Negative (TN) signifies a successful detection of a constraint violation. A False Positive (FP) is when the VQA fails to recognize a constraint violation, and a False Negative (FN) is when the VQA incorrectly identifies a normal condition as a violation.

Utilizing the count in each of these categories, we compute the True Positive Rate (TPR) as $TPR = \frac{TP}{TP+FN}$. TPR reflects the accuracy with which the VQA identifies normal conditions. Similarly, the True Negative Rate (TNR) is calculated as $TNR = \frac{TN}{TN+FP}$, representing the accuracy of VQA in detecting constraint violations.

Further, we determine the Positive Prediction Value (PPV) as $PPV = \frac{TP}{TP+FP}$, and the Negative Prediction Value (NPV) as $NPV = \frac{TN}{TN+FN}$. These metrics correspond to the precision of the VQA detector in identifying normal conditions and constraint violations, respectively.

As per Table 7, both TPR and PPV maintain high values across various settings, which suggests that the VQA detector excels at identifying normal conditions. However, TNR is typically lower, particularly under conditions of low perturbations, indicating that the VQA detector may not be adept at detecting all constraint violations. The fluctuating detections become particularly pronounced when violations are infrequent. Similarly, NPV also trends lower across settings, signifying that our VQA detector might misidentify normal conditions as constraint violations at times, leading to redundant re-planning efforts.

C.4 Case study

Similar to B.5, we use a green border around the image to indicate that our VQA detector determines there is no constraint violation in the image, and a red border to indicate that VQA believes there is a constraint violation in the image.
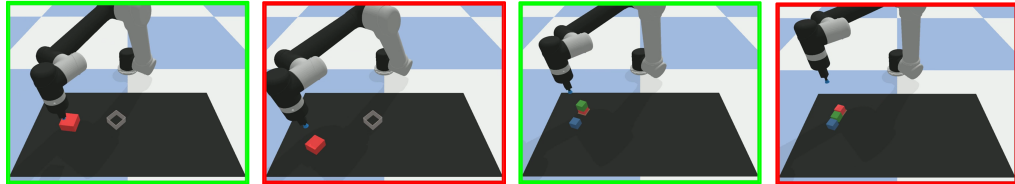
(a) **Q**:Is the robot holding the red block? **A**:Yes
(b) **Q**:Is the robot holding the red block? **A**:No
(c) **Q**:Is the green block on the red block? **A**:Yes
(d) **Q**:Is the green block on the red block? **A**:No

Figure 12: Case visualization for robot arm experiment.

In the table manipulation tasks, the constraint is updated as shown in Table 8.

Stack-block-in-order		TP	TN	FP	FN	TPR	TNR	PPV	NPV
	n=0	150	0	0	0	1.00	N/A	1.00	N/A
p=0	n=1	188	0	0	4	0.98	N/A	1.00	0.00
	n=2	249	27	3	6	0.98	0.90	0.99	0.82
	n=3	272	81	1	2	0.99	0.99	1.00	0.98
	n=0	173	5	0	0	1.00	1.00	1.00	1.00
p=0.05	n=1	204	7	1	1	1.00	0.88	1.00	0.88
	n=2	196	23	3	4	0.98	0.88	0.98	0.85
	n=3	253	92	2	6	0.98	0.98	0.99	0.94
	n=0	202	13	1	1	1.00	0.93	1.00	0.93
p=0.1	n=1	180	8	0	1	0.99	1.00	1.00	0.89
	n=2	212	14	3	1	1.00	0.82	0.99	0.93
	n=3	231	100	4	7	0.97	0.96	0.98	0.93
	n=0	182	24	1	0	1.00	0.96	0.99	1.00
p=0.15	n=1	178	12	2	0	1.00	0.86	0.97	1.00
	n=2	175	26	1	3	0.98	0.96	0.99	0.90
	n=3	208	128	6	9	0.96	0.96	0.97	0.93

Table 7: Statistics of VQA detection. The number of results of TP, TN, FP, FN are summed over 4 different seeds each with 12 episodes.

Low-level Skill	Constraint Update	VQA Question
pick {obj}	{agent}holding{obj}+	Is the {agent}holding{obj}?+
place{obj} on {recep}	{obj}on{recep}+	Is {obj}on{recep}?+

Table 8: A concrete constraint set update example based on the natural language description for the table manipulation tasks. {obj} and {recep} correspond to objects and receptacles respectively. "+" in the rear means adding this constraint into the set after skill execution. Questions are in the general structure: "Is the {constraint}?"

D   Prompts

Listing 1: Zero-shot prompt in DoMeRi on Move Box from A to B.

```
================
Human:
You are a robot, and your valid actions are
1. "pick up" something,
2. "go" somewhere,
3. "put down" something.

You see a white table, a gray table,
and a yellow box on the white table.

You know your current position is far away from the white table
and the gray table.

The white table is also far away from the gray table.

Your task is to move a yellow box from the white table to the
gray table.

What do you do?

Please list your plan step by step and do not output redundant
messages except steps.
```

```
LLM:
1. Go to the white table.
2. Pick up the yellow box.
3. Go to the gray table.
4. Put down the yellow box on the gray table.

================

Robot action: Go to the white table.

Robot action: Pick up the yellow box.

Robot action: Go to the gray table.

Detector: The robot is not holding the yellow box.

Robot: The robot is not holding the yellow box now,
and this box is next to the robot. What can I do?
Please list the new plan step by step and do not output redundant
messages except steps.

LLM:
1. Pick up the yellow box.
2. Go to the gray table.
3. Put down the yellow box on the gray table.

================

Robot action: Pick up the yellow box again.

Robot action: Go to the gray table.

Robot action: Put down the yellow box on the gray table.

Robot action: Done.

================
```

Listing 2: Few-shot in context learning prompt in table manipulation tasks.

```
Task: Pick the red block and place it in the gray base.
Robot action: Pick the red block and place it in the gray base.
Robot action: Done.
Task: Pick the red block and place it in the gray base.
Robot action: Pick the red block and place it in the gray base.
Detector: The robot is not holding the red block.
Robot action: Pick the red block and place it in the gray base.
Robot action: Done.
Task: Stack the blocks in the order of red, green, and blue.
Robot action: Pick the red block on top of the base.
Robot action: Pick the green block and place it on top of the
red block.
Robot action: Pick the blue block and place it on top of the
green block.
Robot action: Done.
Task: Stack the blocks in the order of red, green, and blue.
Robot action: Pick the red block on the base.
```

```
Robot action: Pick the green block and place it on top of the
red block.
Detector: The green block is not on the red block.
Robot action: Pick the green block and place it on top of the
red block.
Robot action: Pick the blue block and place it on top of the
green block.
Robot action: Done.
Task: Stack the blocks in the order of red, green, and blue.
Robot action: Pick the red block on the base.
Robot action: Pick the green block and place it on top of the
red block.
Robot action: Pick the blue block and place it on top of the
green block.
Detector: The green block is not on the red block. The blue
block is not on the green block.
Robot action: Pick the green block and place it on top of the
red block.
Robot action: Pick the blue block and place it on top of the
green block.
Robot action: Done.
```

Listing 3: Few-shot in context learning prompt in humanoid tasks.

```
Task: Move the yellow box from the white table to the gray
table.
1. Go to the white table
2. Pick up the yellow box
3. Go to the gray table
4. Put the yellow box on the gray table.
5. Done
Task: Move the yellow box from the white table to the gray
table.
1. Go to the white table
2. Pick up the yellow box
3. Go to the gray table [Detector: The robot is not holding the
box]
4. Pick up the yellow box
5. Go to the gray table
6. Put the yellow box on the gray table.
7. Done
Task: Go forward
1. Go forward
2. Done
Task: Go forward
1. Go forward [Detector: obstacle in the left]
2. Turn right
3. Go forward
4. Done
```