Dynamic Task Update Benchmark (DTU-Bench): Evaluating Runtime Task Adaptation in Robotic Manipulation

Yanran Jiang[1], Cécile Paris[1], Dana Kulić[1,2], and Leimin Tian[1]
[1]CSIRO Technology, Australia
[2]Monash University, Australia

Abstract

Vision-language-action (VLA) policies perform well on fixed-instruction robotic manipulation benchmarks, but their ability to adapt to instruction changes during execution remains understudied. We introduce the Dynamic Task Update Benchmark (DTU-Bench), evaluating 40 task revisions across seven update types and three execution stages using LIBERO, a simulated robotic manipulation benchmark. We compare performance when a new instruction is given before execution starts versus during execution, and examine whether the new task instructions are followed. Across 13 representative VLA configurations, average task success is 95.8% on the original object-manipulation tasks, compared with only 1.6% when the instruction changes which object to manipulate during execution. Controlled diagnostics show that policies often complete the original task before the updated one, and that their responses depend on the robot's physical state. We further test whether a simple scripted controller that releases the grasped object and moves the arm back helps robots respond to instruction changes. This improves success from 2% to 38% for $\pi_{0.5}$ and from 8% to 48% for DM0.5, but can cause unwanted object motion. These findings reveal limitations that fixed-instruction benchmarks and task-success metrics alone can overlook in evaluating robotic manipulation.

1. Introduction

Robots operating alongside people may need respond to changes in instructions [22, 24]. A user may change where an object should be placed, request a different object, or withdraw part of an instruction after the robot has already begun acting. Unlike starting a new task, such updates arrive while the robot is executing and may have already made some progress on the initial task. When an instruction change arrives, some actions associated with the initial task may still be useful, while others may need to be abandoned or undone [15].

Recent advances in robot learning have greatly improved language-conditioned manipulation. Generalist vision-language-action (VLA) policies such as OpenVLA [13] and π_0 [1] have demonstrated increasingly capable manipulation across tasks and environments. Benchmarks such as CALVIN [20] and LIBERO [19] have supported evaluation of VLA policies on long-horizon task execution, generalization, and knowledge transfer. However, success on tasks specified before execution does not necessarily mean a policy can adapt when the task changes during execution. A policy may successfully execute the original or the updated tasks independently, yet fail to transition between them once execution has begun.

Recent studies have addressed different challenges in dynamic robot execution. DynamicVLA [26] focuses on manipulating moving objects, while TIC-VLA [12] addresses reactive control. SwitchVLA [16] and ReSteer [6] develop training-based approaches to task switching. However, how existing VLA policies respond when different parts of a task are changed during execution remains insufficiently understood. Moreover, reaching the updated goal does not necessarily indicate that the robot has responded to the new instruction. This motivates a systematic evaluation of runtime task adaptation and post-update behavior.

We introduce Dynamic Task Update Benchmark (**DTU-Bench**) to systematically evaluate how robot policies respond to task instructions that change during execution. As illustrated in Fig. 1, DTU-Bench introduces new instructions while the robot is acting, without resetting its state. Built on LIBERO, the benchmark covers 40 task revisions across seven types of instruction changes and three execution stages. We evaluate each policy under three conditions: executing the original task, performing the revised task from the initial state of the original task, and switching to the revised task during execution. These comparisons allow us to assess whether policies can complete a revised task after making progress towards a different objective. Beyond task success, we examine whether robots maintain the requested outcome and how their behavior changes following an instruction update.

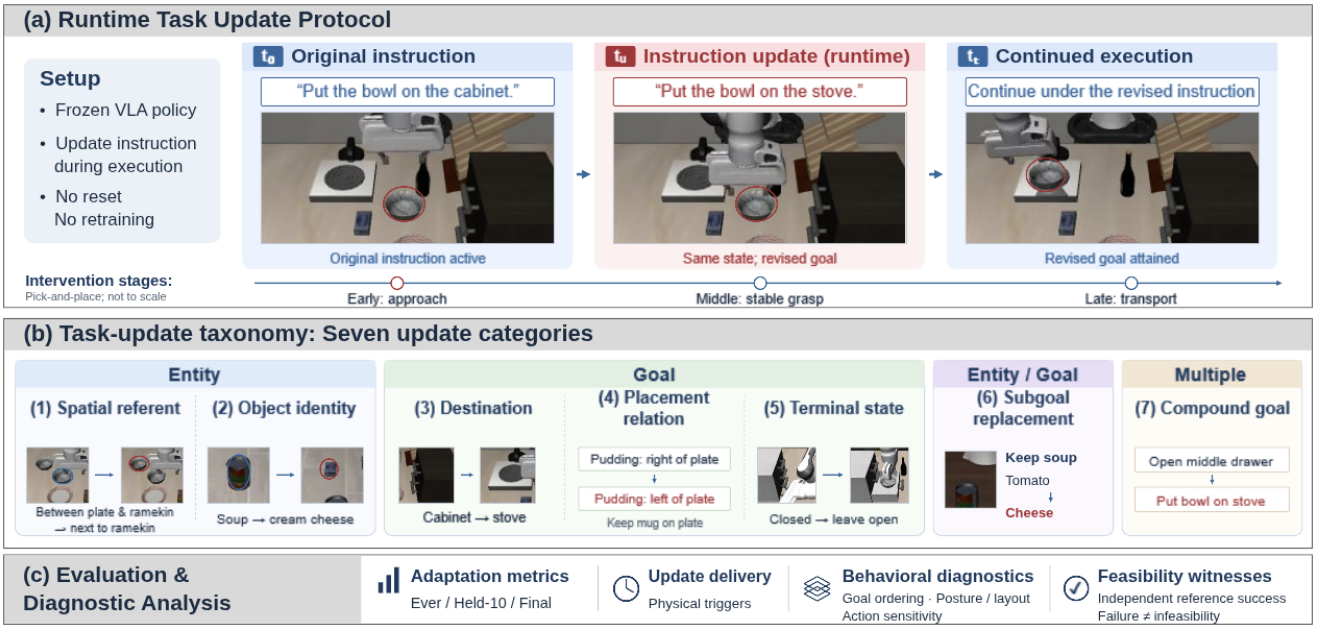


Figure 1. Overview of DTU-Bench.

Our main contributions are:

- We introduce DTU-Bench, a systematic, policy-agnostic framework for evaluating runtime task updates in robotic manipulation across seven task-revision categories and three execution stages. The benchmark and evaluation code will be released.
- We evaluate 13 VLA configurations and identify limitations that fixed-instruction evaluation can overlook, including differences across update types, delayed satisfaction of updated goals, and substantial gaps between temporary and final goal satisfaction.
- We conduct controlled experiments to examine how the robot's execution state affects responses to new instructions, including whether simple scripted interventions can help recover from failed task updates.

2. Related Work

Embodied Foundation Models. Vision–language–action (VLA) models for robotic manipulation map visual observations and instructions to actions through autoregressive token generation [13, 21, 25], continuous regression [14], or flow matching [1, 2]. WorldVLA and RynnVLA-002 jointly model actions and future visual observations [4, 5], while VLA-JEPA and Light-WAM incorporate predictive representations or supervision [17, 23]. Hierarchical models such as τ_0-VLA further integrate world-model-guided planning [3]. Despite this architectural diversity, adaptation to mid-execution task changes remains insufficiently evaluated.

Dynamic Execution and Task Adaptation. DynamicVLA [26] addresses moving-object manipulation, while TIC-VLA [12] combines asynchronous reasoning and control. These methods primarily respond to changing observations rather than task objectives. More directly, SwitchVLA [16] trains a contact-aware policy for task switching, while ReSteer [6] improves policy steerability through generated training data. Rather than developing a new switching policy, DTU-Bench focuses on evaluating different types of task revisions, distinguishing revised-task competence from runtime performance and examining whether updated-goal success reflects a response to the new instruction.

Benchmarking VLA Robustness and Adaptation. LIBERO [19] provides a benchmark for language-conditioned manipulation tasks; LIBERO-Plus [10] evaluates robustness to visual, spatial, and linguistic perturbations; and LIBERO-Safety [7] examines physical and semantic safety. DTU-Bench instead focuses on changes to task requirements during execution. It investigates original tasks, updated-from-start controls, and runtime updates, assessing whether updated goals are reached and maintained, alongside diagnostics of post-update behavior.

3. DTU-Bench

DTU-Bench (Fig. 1) evaluates runtime task adaptation in robotic manipulation through structured task updates, execution-stage interventions, and measures of updated-goal attainment, persistence, and post-update behavior.

3.1. Task Update Formulation

In conventional language-conditioned manipulation, a policy executes a fixed instruction I_0 toward a goal G_0. In DTU-Bench, execution begins under I_0, and a revised instruction I_1, specifying goal G_1, is delivered at time $t = \tau$ while the system is in an intermediate state s_τ. The instruction conditioning the policy changes at the intervention time:

$$a_t \sim \begin{cases} \pi(\cdot \mid o_t, I_0), & t < \tau, \\ \pi(\cdot \mid o_t, I_1), & t \geq \tau, \end{cases} \tag{1}$$

where o_t is the current observation. Execution of I_1 starts from the physical state s_τ reached under I_0 rather than restarting the episode. This distinguishes runtime adaptation from executing the revised task from the initial state: earlier actions may still be useful, may need to be redirected, or may have created conflicts with the new goal.

3.2. What Changes: Task Update Taxonomy

We ground our taxonomy in the task representation of classical planning. Given a fixed planning domain, a planning problem can be described by a set of objects $\mathcal{O}$, an initial state s_0, and a goal specification G [11]:

$$\mathcal{P} = (\mathcal{O}, s_0, G). \tag{2}$$

Here, G specifies the conditions required to complete a task. In DTU-Bench, instruction updates change the target object, the required task outcome, or both, while preserving the objects and scene at the time of the update.

We include seven types of task updates, grouped into four broader categories (Table 1): *Entity (2 types)* changes which object the robot must manipulate (Object identity) or how the target object is identified (Spatial referent); *Goal (3 types)* changes where an object should be placed (Destination), its required relation to another object (Placement relation), or the required final state of the environment (Terminal state); *Entity/Goal (1 type)* replaces one required subgoal with another while leaving the remaining subgoals unchanged (Subgoal replacement); and *Multiple (1 type)* changes two or more independent task requirements within the same instruction (Compound goal).

Each task-update scenario is assigned to one of the seven types, although some changes may involve more than one aspect of the task. Additional experiments on task cancellation and unrelated instructions are conducted on a subset of policies, with results reported in the supplementary material.

3.3. Constructing Task-Update Scenarios from LIBERO

We construct 40 task-update scenarios using tasks from LIBERO-Spatial, Object, GOAL, and Long. Each scenario specifies an original instruction–goal pair (I_0, G_0) and an updated pair (I_1, G_1), both using objects from the same LIBERO scene. The updated instruction changes what the robot is asked to do without modifying the environment. Each scenario is evaluated using the 50 official LIBERO initial states.

Table 1. Task-update taxonomy in DTU-Bench. Categories use source tasks from the indicated LIBERO suites, with 40 task updates in total.

Dimension	**Update category (LIBERO suite)**	**What changes?**
Entity	Spatial referent (Spatial)	Relational target description
	Object identity (Object)	Manipulated object
Goal	Destination (GOAL)	Target location
	Placement relation (Long)	Object–destination relation
	Terminal state (Long)	Required final condition
Entity / Goal	Subgoal replacement (Long)	Required subgoal object
Multiple	Compound goal (GOAL)	Multiple task requirements

3.4. When It Changes: Execution-Stage Interventions

We introduce instruction changes at three stages of execution: *early*, *middle*, and *late*. These stages are defined by task-progress events rather than fixed action counts. For example, in pick-and-place tasks, they correspond to approaching an object, grasping or lifting it, and moving it towards its destination. If the robot does not reach the required stage, the update is recorded as not delivered.

When the instruction changes, we discard any remaining actions predicted under the original instruction and generate new actions using the revised instruction and current observations.

The robot is then given 300 action steps to complete the revised task, regardless of how many steps it has already executed. We continue execution even after the revised task is completed to check whether the required outcome is maintained. We additionally evaluate selected continuations over 600 actions to examine how a longer execution horizon affects task completion and maintenance.

3.5. What Remains: Progress Compatibility and Execution State

The effect of a task update depends not only on what changes in the goal, but also on the progress made and the physical state reached during execution prior to the update. We consider two complementary aspects of compatibility with the updated goal G_1.

Task-level progress compatibility concerns which parts of the original task remain useful, which are no longer needed, and what must be done to complete the revised task.

In multi-stage tasks, completed subgoals may remain useful even when other requirements change. For example, if both tasks require opening a drawer, the robot does not need to open it again.

Execution-state compatibility concerns how the robot's physical configuration affects subsequent behavior. At the intervention state s_τ, robot posture, existing grasps, and object configurations may facilitate or constrain execution of the updated goal. For example, changing the destination of a grasped object preserves useful progress, whereas changing the target object may require releasing the grasped object and selecting another.

These two aspects motivate our analysis of post-update behavior (Sec. 4.2).

3.6. Evaluation and Diagnostic Analysis

DTU-Bench evaluates runtime task adaptation through complementary measures of task competence, update delivery, goal attainment, and post-update behavior. Table 2 summarizes the evaluation components and the evidence each provides. The benchmark protocol applies to all evaluated policies, while controlled diagnostics are performed on selected tasks and configurations.

Task competence and update delivery. For each task and initial state, we evaluate three conditions: (1) an *original control* executing I_0 from the initial state s_0; (2) an *updated-from-start control* receives I_1 before any task actions are executed and starts from the initial state s_0; and (3) *runtime updates* that begin with I_0 and switch to I_1 at the intervention state s_τ. The first two conditions evaluate performance on each task separately, while the third evaluates performance when the instruction changes during execution.

Task Success and Goal Maintenance. We evaluate the updated-goal predicate G_1 after each executed action using three metrics: *Ever*, whether the goal is satisfied at least once; *Held-10*, whether it remains satisfied for ten consecutive observations; and *Final*, whether it is satisfied at the end of the 300-step post-update execution. Held-10 measures sustained goal attainment over a short interval, while Final measures whether the goal remains satisfied at the end of execution. Their difference helps identify cases where the updated goal is temporarily attained but not maintained until the end of execution.

Matched-state behavioral diagnostics. Goal attainment alone does not establish whether a policy changes its behavior in response to an updated instruction. For selected intervention states, we restore the same recorded state and compare two continuations: one retaining the original instruction I_0 and the other receiving the updated instruction I_1. Both branches use the same updated-goal predicate G_1 and an equal action budget. We compare their goal-attainment outcomes, object interactions, and the temporal order of original- and updated-goal satisfaction. This design tests whether attaining G_1 depends on receiving the updated instruction under matched execution conditions.

Table 2. DTU-Bench evaluation components. Standard benchmark measures quantify task competence and goal satisfaction; controlled diagnostics examine post-update behavior and attainability.

Component	What it establishes
Standard benchmark evaluation	
Original control	Original-task competence
Updated-from-start	Updated-task competence
Runtime update	Updated-goal attainment from s_τ
Update delivery	Whether the intervention occurs
Held-10	Ten consecutive observations satisfying G_1
Final	Satisfaction of G_1 at execution end
Controlled diagnostic analysis	
Matched continuations	Instruction-dependent outcomes from the same state
Execution-state intervention	Sensitivity to posture, layout, and observations
Reference controller	Attainability evidence for selected states

Execution-state and attainability diagnostics. We examine how robot posture, object layout, and visual observations affect post-update behavior. A separate geometry-based reference controller uses simulator object positions and predefined waypoints to independently attempt the updated task (details in the supplementary material). Its success demonstrates attainability, while failure does not establish infeasibility.

4. Experiments

4.1. Experimental Setup

Benchmark and evaluation protocol. We evaluate DTU-Bench using all four LIBERO suites: Spatial, Object, GOAL, and Long. The full evaluation design includes 40 source tasks, 50 official initial states per task, and five execution conditions: original-from-start, updated-from-start, and runtime updates at early, middle, and late stages. The full evaluation design comprises 10,000 scheduled trials per policy configuration, with the two from-start controls shared across the three intervention stages. Each delivered update is evaluated for 300 executed actions.

We report the evaluation coverage, valid trials, and exclusions for each policy in the supplementary material. Trials with simulator errors are excluded from the success-rate calculations.

VLA baselines. We evaluate a series of publicly released VLA configurations covering different action-generation mechanisms (autoregressive action tokens, direct continuous regression, and flow-based action prediction). The models also span conventional vision-language-action ar-

chitectures, action-expert policies, and approaches incorporating world-model or predictive video supervision: (1) OpenVLA [13], (2) OpenVLA-OFT [14], (3) π_0 [1], (4) $\pi_{0.5}$ [2], (5) DM0.5 [8], (6) WorldVLA [5], (7) RynnVLA-002 [4], (8) UniVLA [25], (9) VLA-JEPA [23], (10) InstructVLA [27], (11) τ_0-VLA [3], (12) MolmoAct2 [9], and (13) FlashVLA-$\pi_{0.5}$ [18].

We use released LIBERO checkpoints without additional training. Each configuration retains its native observation inputs, action representation, and inference interface while following the same task-update protocol and goal evaluation. For models with additional planning or world-model components, we evaluate only the released policy configuration supported by the tested interface, rather than assuming that the full planning system is active. Checkpoint details, execution settings, and evaluation coverage are provided in supplementary material.

Statistical analysis. We estimate pointwise 95% confidence intervals by resampling initial-state blocks within each source task, keeping all conditions and intervention stages paired. Task-cluster bootstrap and leave-one-task-out analyses assess sensitivity to task composition. Full results are reported in the supplementary material.

4.2. Controlled Diagnostic Experiments

We conduct additional diagnostics on selected LIBERO tasks using recorded intervention states and fixed policy checkpoints. These experiments examine specific benchmark behaviors rather than estimate performance across all policies or task categories.

Destination versus object updates. Using π_0 and $\pi_{0.5}$, we compare destination ($A{\to}P$ to $A{\to}Q$) and object ($A{\to}P$ to $B{\to}P$) updates in the same LIBERO-GOAL scene. We test five initial states at approach and stable-grasp stages with 300-action continuations.

Instruction and execution configuration. For $\pi_{0.5}$ and DM0.5, we vary instruction (original/updated), posture (saved/neutral), and object layout (original/swapped) across two LIBERO-Object tasks and five initial states, yielding 160 continuations. Neutral-posture interventions preserve object positions within each layout. We measure first-object contact and updated-goal success over 600 actions.

Visual-input sensitivity. We examine how external- and wrist-camera observations affect action predictions under the original and updated instructions. The experimental setup and full results are provided in the supplementary material.

Trajectory analysis. We automatically extract first finger–object contact and the times at which the original and updated task conditions are satisfied from simulator logs. Shared and obsolete task requirements are predefined for each scenario and evaluated automatically after each executed action. These records allow us to identify whether the original task is completed before the updated task, or whether the updated instructions are satisfied only temporarily. For selected 600-action trajectories, we also evaluate outcomes at 300 actions to compare different execution horizons using the same trajectory.

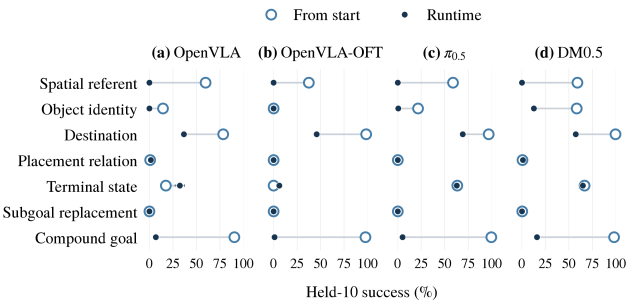


Figure 2. Updated-from-start vs. runtime Held-10 success for four representative VLA policies. Open circles indicate updated-from-start success, filled circles indicate runtime success, and connecting lines show the difference. Results for all 13 policies are provided in the supplementary material.

Matched instruction continuations. We compare original- and updated-instruction continuations from identical restored intervention states, evaluating both against the same updated-goal predicate G_1. The 300-action diagnostic contains 94 matched pairs (46 for $\pi_{0.5}$ and 48 for DM0.5) across selected LIBERO-Long, Object, and Spatial tasks. Both branches use equal continuation budgets, with pending actions discarded before fresh inference. Additional suffix comparisons reuse original branches and are not treated as independent matched pairs.

Inference-time assistance. We investigate whether simple training-free helper steps can recover failed updates. For $\pi_{0.5}$ and DM0.5, we compare three methods from the same recorded intervention states: *Direct*, which immediately executes the updated instruction; *Retreat*, which first moves the arm to a predefined joint configuration (robot default initial pose); and *Release+Retreat*, which opens the gripper before retreating when it is closed. The helpers use only robot joint and gripper information without accessing object positions, contact labels, or task-success conditions to control their actions. We evaluate all ten LIBERO-Object tasks and five LIBERO-GOAL destination-update tasks, using five initial states and two intervention stages (approach and stable grasp). All methods share the same updated instruction and 300-action budget, including assistance.

5. Results

We first compare performance across task updates and intervention stages. We then examine what robots do after an update, using matched continuations, object-selection tests, and reference trajectories.

Table 3. **Original-task and runtime updated-goal success.** Each cell reports Original Final → Runtime Final (%). Darker green indicates higher runtime success on a common 0–100 scale; bold marks the highest reported runtime score in each column.

Policy	Spatial referent	Object identity	Destination	Placement relation	Terminal state	Subgoal replacement	Compound goal
OpenVLA	81.2 → 0.0	85.4 → 0.0	76.8 → 36.4	24.2 → **1.8**	34.0 → **5.6**	63.5 → 0.0	67.2 → 6.8
OpenVLA-OFT	96.4 → 0.0	98.0 → 0.0	92.8 → 45.3	55.3 → 0.0	85.3 → 3.1	90.5 → 0.0	90.0 → 0.9
$\pi_{0.5}$	95.8 → 0.0	98.0 → 0.7	94.4 → **68.8**	50.3 → 0.0	74.7 → 0.7	96.0 → 0.2	90.8 → 5.3
DM0.5	98.4 → 0.1	99.8 → **13.5**	95.6 → 57.2	39.6 → 1.1	90.0 → 0.2	98.0 → 0.2	92.0 → 17.5
π_0	94.4 → 0.5	97.6 → 0.0	84.8 → 57.1	59.7 → 0.0	62.0 → 1.3	70.5 → 0.0	89.6 → 18.9
RynnVLA-002 cont.	90.2 → 0.0	94.8 → 0.1	91.6 → 42.8	61.3 → 0.0	76.0 → 0.7	94.5 → 0.0	78.0 → 2.3
VLA-JEPA	96.4 → 0.0	99.8 → 0.0	98.0 → 57.3	89.9 → 0.0	82.7 → 0.4	97.5 → 0.0	92.0 → 15.6
InstructVLA stage-1	94.6 → 0.0	99.2 → 0.1	97.6 → 53.5	70.5 → 0.0	80.0 → 1.6	88.5 → 0.0	91.2 → 2.7
τ_0-VLA low-level	97.0 → 0.1	99.2 → 0.1	95.6 → 54.0	43.0 → 0.0	90.7 → 0.0	89.5 → 0.0	87.6 → 7.9
MolmoAct2 cont.	97.0 → 0.1	99.0 → 4.9	97.2 → 58.1	47.7 → 1.6	74.7 → 0.9	95.0 → 0.0	92.8 → **19.5**
FlashVLA-$\pi_{0.5}$ sync.	97.2 → 0.3	99.0 → 1.3	94.0 → 57.6	59.7 → 0.0	83.3 → 0.0	98.0 → **1.7**	91.6 → 18.9
WorldVLA	69.6 → **1.2**	79.0 → 0.6	78.4 → 38.6	28.7 → 0.0	30.7 → 1.1	44.5 → 0.0	67.6 → 3.9
UniVLA	98.0 → 0.0	96.0 → 0.0	60.0 → 46.7	66.7 → 0.0	73.3 → 2.2	85.0 → 0.0	84.0 → 5.3

5.1. Responses Differ by Task Update Types

Table 3 compares the final task success of original-task and runtime updated-goal across seven update categories. Because these scores measure different task requirements, their difference cannot be attributed solely to the instruction change. Across the reported configurations, destination updates resulted in 36.4–68.8% runtime Final success, compared with at most 13.5% yielded by object-identity updates and 1.2% for spatial-referent updates. Destination changes often preserve manipulation of the same object, whereas object updates require selecting another target. However, scene layout and manipulation task difficulty also vary across categories, preventing attribution of the observed performance difference solely to the type of task update. We confirm this gap in a same-scene comparison (Sec. 5.5).

Finding 1: The VLAs achieved substantially higher task success rates when encountering destination updates compared to object-identity and spatial-referent updates. However, this difference may be caused by factors in addition to the update type.

5.2. Task Competence ≠ Runtime Success

A failed runtime update does not necessarily mean the system failed to respond to the new instruction. It may be due to the difficulty of the updated task itself. We compare Held-10 success across 13 VLA configurations when the updated instruction is given before execution begins versus during execution, after the robot has started the original task. Figure 2 highlights four representative policies: OpenVLA, OpenVLA-OFT, $\pi_{0.5}$, and DM0.5. Across all 13 configurations, runtime Held-10 success is often substantially lower than updated-from-start success, although the gaps vary across update categories. For example, OpenVLA-OFT shows a 96.7 pp gap on compound-goal updates (95% CI: 94.5–98.4 pp). This reveals a gap between performing a task in isolation and achieving the same goal after partial execution under a different instruction.

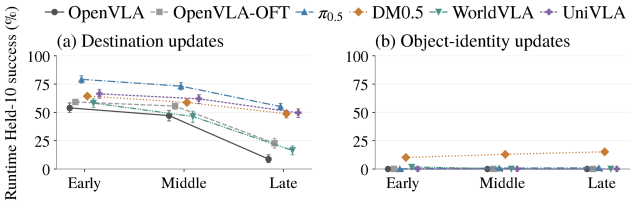


Figure 3. Effect of intervention stage on runtime success. Held-10 success for two update categories and three intervention stages for six VLA policies. Rates are computed over valid scheduled trials, including updates that were not delivered. Error bars show 95% CIs. Full 13-policy results are in the supplementary material.

Finding 2: Success on an updated task from the start does not guarantee success when the same instruction is introduced during execution.

5.3. Intervention Stage Affects Runtime Success

We examine how intervention stage affects runtime Held-10 success across all 13 evaluated policy configurations. Figure 3 shows six representative policies, illustrating different patterns across update types. For destination updates, success generally decreases from early to late intervention. For example, DM0.5 drops from 64.4% to 58.8% and 48.8%. In contrast, most policies achieve near-zero success on object-identity updates across all stages, except DM0.5, which shows a slight increase. These results suggest that the effect of intervention stage depends on the task requirements being changed, rather than following a uniform early-to-late decline.

Finding 3: Later interventions are not consistently harder; their effect on runtime success varies across task-update types.

5.4. Reaching the New Goal ≠ Adaptation Success

Runtime success measures whether the updated goal is achieved, but does not capture how or when it is reached.

Instruction-dependent goal attainment. We compare 94 matched pairs of continuations from identical intervention states, with one branch following the original instruction and the other receiving the update. Both branches are evaluated against the updated goal. Under Held-10, 37 pairs succeed with both instructions, three succeed only with the update, and 54 fail under both. All 37 both-success cases involve terminal-state updates, indicating that the updated goal can be temporarily satisfied even when following the original instruction. However, at the end of execution (Final), only five updated-instruction continuations still satisfy the updated goal, while none of the original-instruction continuations do.

Delayed goal completion. In a LIBERO-Object task where the instruction changes from placing alphabet soup to placing cream cheese, DM0.5's Held-10 success increases from 4/15 at 300 actions to 15/15 at 600 actions. All 15 trajectories complete the original goal before the updated goal (Fig. 6). This behavior extends across multiple VLAs: among successful Object updates, the original goal is completed first in 95.8% of DM0.5 cases (184/192), 90.8% for MolmoAct2 (59/65), 94.7% for FlashVLA (18/19), and 90.0% for $\pi_{0.5}$ (9/10). These results show that achieving the updated goal does not necessarily indicate a prompt response to the instruction change.

Temporary goal satisfaction. Terminal-state updates provide another example of why goal attainment alone can be misleading. For terminal-state updates, DM0.5 achieves 64.9% Held-10 success but only 0.22% Final success, while $\pi_{0.5}$ achieves 62.9% and 0.67%, respectively. For example, a robot may temporarily leave a drawer open before continuing to close it. Such outcomes distinguish temporary goal attainment from sustained satisfaction.

Preserving useful progress. Runtime updates do not always require abandoning previous work. In a LIBERO-Long example, the instruction changes from placing cream cheese and butter in a basket to placing cream cheese and soup in the basket. DM0.5 preserves the already completed cheese placement and adds the newly requested soup without placing the obsolete butter.

Finding 4: Reaching the updated goal does not necessarily mean the robot changed its behavior promptly or maintained the goal afterward.

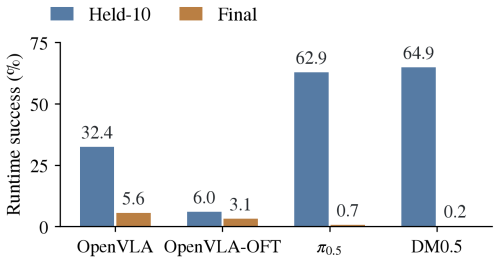


Figure 4. Held-10 and Final success for LIBERO-Long terminal-state updates. Each policy has 450 scheduled trials (3 tasks × 50 states × 3 intervention stages). Held-10 requires the revised goal to hold for 10 consecutive executed-action observations; Final requires it to hold at the last observation. Full 13-policy results are in the supplementary material.

5.5. Task Context Affects Post-Update Behavior

Updated-goal success alone does not reveal how policies respond to instruction changes from different execution states. We therefore conduct controlled experiments to examine how prior task progress and robot posture affect subsequent behavior.

Destination versus object updates. To examine whether partial execution affects different types of updates, we compare destination and object changes from matched intervention stages using π_0 and $\pi_{0.5}$. Both policies succeed in all 10 destination-update continuations, including after grasping the original object. In contrast, neither succeeds in the five object-update continuations after stable grasp, although all updated-from-start controls succeed. Before grasping, object-update success is 0/5 for π_0 and 3/5 for $\pi_{0.5}$. These results show that the gap between destination and object updates (Sec. 5.1) persists when scene and task difficulty are held fixed: progress made under the original instruction remains useful for destination changes but can hinder updates requiring a different object.

Instruction, posture, and object layout. We next investigate how the robot's current configuration affects object selection. In 160 matched continuations with $\pi_{0.5}$ and DM0.5, we vary the instruction, robot posture, and object layout across two Object source tasks. Figure 5 illustrates how DM0.5 selects different objects under the same updated instruction when the arm posture changes. At the saved intervention posture, both policies first contact the nearer candidate object in all 40 tested conditions per policy, while only 20/40 contacts match the specified instruction. Swapping the object layout changes which object is contacted first. At the neutral posture, instruction-consistent first contact increases to 40/40 for DM0.5 and 29/40 for $\pi_{0.5}$. The results therefore differ between policies. Moreover, even when DM0.5 contacts the updated target first in all 20 relevant neutral-posture continuations, only 18/20

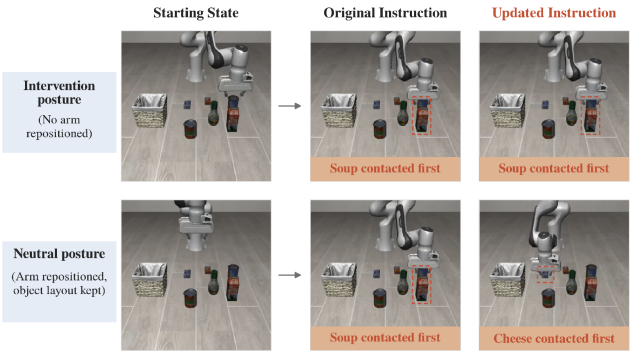


Figure 5. Object selection under different robot postures and instructions. DM0.5 continuations under original (soup) and updated (cheese) instructions at the saved intervention posture (top) and neutral posture (bottom), with object layout unchanged. Red dashed boxes mark the first contacted object. Results are consistent across five tested states per condition.

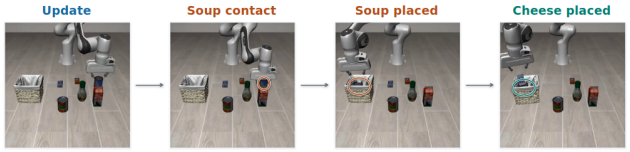


Figure 6. Example DM0.5 trajectory in LIBERO-Object, where the instruction changes from alphabet soup to cream cheese. The frames show the robot contacting and placing the original object before completing the revised goal.

achieve the updated goal at the final evaluation step. These results show that selecting the requested object and completing the updated task are different outcomes.

Attainability of failed updates. We use a separate geometry-based controller (Sec. 3.6) that uses simulator object positions and predefined waypoints to complete the updated task without a VLA policy. Across 20 Object-task tests, it succeeds in 10 cases, nine of which fail under the evaluated VLAs. Conversely, DM0.5 succeeds in three cases where the reference controller fails. Thus, some VLA failures occur despite achievable updated tasks, while controller failure does not establish infeasibility.

Finding 5: Prior progress and robot state can facilitate or hinder responses to task updates.

5.6. Assistance to Improve Task Update Success

The preceding experiments suggest that the robot's physical state affects its response to instruction changes. We therefore test whether simple training-free assistance can improve task completion without retraining the policy (Fig. 7). For each policy and intervention stage, we evaluate 50 Object trials (10 tasks $\times$ 5 states) and 25 Destination trials (5 tasks $\times$ 5 states).

Object updates. Before grasping, fixed-joint retreat

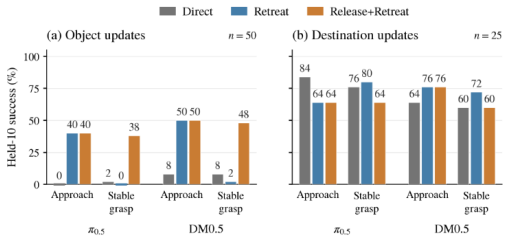


Figure 7. Training-free assistance for task updates.

(i.e., return to the robot original pose) increases Held-10 success from 0/50 to 20/50 for $\pi_{0.5}$ and from 4/50 to 25/50 for DM0.5. After stable grasp, however, retreat alone provides little benefit (1/50 to 0/50 and 4/50 to 1/50, respectively). Releasing the object before retreating instead increases stable-grasp success to 19/50 for $\pi_{0.5}$ and 24/50 for DM0.5. Relative to direct updates, the paired improvements are 36 percentage points (95% task-cluster bootstrap CI: [12, 62]) and 40 percentage points ([18, 64]), respectively. The improvements remain positive when any single Object task is excluded.

Destination updates. The same assistance does not consistently improve tasks that retain the original target object. For $\pi_{0.5}$, retreat decreases approach-stage success from 21/25 to 16/25, while release followed by retreat decreases stable-grasp success from 19/25 to 16/25. Results for DM0.5 are mixed. These outcomes suggest that unnecessary physical interventions can disrupt progress that remains useful for the updated task.

Physical side effects. Although release followed by retreat improves completion of some object updates, unsupported downward object motion occurs in 49/50 stable-grasp trials for each policy. Thus, improved task success does not necessarily imply desirable physical execution. The interventions are performed by external scripted controllers rather than learned recovery behavior of the VLA policies.

Finding 6: Simple physical assistance can mitigate some task-update failures without policy retraining, although its benefits and physical consequences depend on the update type and execution stage.

6. Conclusion

We introduced DTU-Bench to evaluate how VLA policies respond to instruction changes during execution. Across 13 configurations, we found substantial differences across update types and execution stages, while success on fixed tasks did not guarantee runtime success. Behavioral diagnostics revealed delayed completion and temporary satisfaction of updated tasks. Training-free assistance improved some failures but could also disrupt useful progress and introduce undesirable object motion. These findings highlight the im-

portance of evaluating task success alongside post-update behavior and physical execution.

LLM Usage Statement

ChatGPT and OpenAI Codex (OpenAI) were used to assist with manuscript wording, experimental code development and debugging, statistical analysis, and figure formatting and layout adjustments. The authors are responsible for the experimental methodology, results, and scientific conclusions.

References

[1] Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, et al. π_0: A vision-language-action flow model for general robot control. *arXiv preprint arXiv:2410.24164*, 2024. 1, 2, 5

[2] Kevin Black, Noah Brown, James Darpinian, Karan Dhabalia, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, and Manuel Y Galliker. $\pi_{0.5}$:: A vision-language-action model with open-world generalization. In *9th Annual Conference on Robot Learning*, 2025. 2, 5

[3] Xiaowei Cai, Yunuo Cai, Bingao Chen, Jingxiao Chen, Zhi Chen, Siyuan Feng, Tengyu Hou, Jingshun Huang, Han Jiang, Runkun Ju, et al. τ_0-VLA: a hierarchical robot foundation model with world-model-guided test-time computation. *arXiv preprint arXiv:2608.16885*, 2026. 2, 5

[4] Jun Cen, Siteng Huang, Yuqian Yuan, Kehan Li, Hangjie Yuan, Chaohui Yu, Bohan Hou, Yuming Jiang, Jiayan Guo, Xin Li, et al. Rynnvla-002: A unified vision-language-action and world model. *arXiv preprint arXiv:2511.17502*, 2025. 2, 5

[5] Jun Cen, Chaohui Yu, Hangjie Yuan, Yuming Jiang, Siteng Huang, Jiayan Guo, Xin Li, Yibing Song, Hao Luo, Fan Wang, et al. Worldvla: Towards autoregressive action world model. *arXiv preprint arXiv:2506.21539*, 2025. 2, 5

[6] Zhenyang Chen, Alan Tian, Liquan Wang, Benjamin Joffe, Yingyan Celine Lin, Yuxiao Chen, Siddharth Karamcheti, and Danfei Xu. Resteer: Quantifying and refining the steerability of multitask robot policies. *arXiv preprint arXiv:2603.17300*, 2026. 1, 2

[7] Rongxu Cui, Zongzheng Zhang, Jingrui Pang, Haohan Chi, Jinbang Guo, Saining Zhang, Shaoxuan Xie, Xin Jin, Yao Mu, Jiaolong Yang, et al. Libero-safety: A comprehensive benchmark for physical and semantic safety in vision-language-action models. In *European Conference on Computer Vision*, pages 512–530. Springer, 2026. 2

[8] Dexmal Team. DM0.5: An open-world foundation model for general-purpose embodied intelligence. https://www.dexmal.com/blog/dm0.5/index_en.html, 2026. 5

[9] Haoquan Fang, Jiafei Duan, Donovan Clay, Sam Wang, Shuo Liu, Weikai Huang, Xiang Fan, Wei-Chuan Tsai, Shirui Chen, Yi Ru Wang, et al. Molmoact2: Action reasoning models for real-world deployment. *arXiv preprint arXiv:2605.02881*, 2026. 5

[10] Senyu Fei, Siyin Wang, Junhao Shi, Zihao Dai, Jikun Cai, Pengfang Qian, Li Ji, Xinzhe He, Shiduo Zhang, Zhaoye Fei, et al. Libero-plus: A progressive robustness benchmark for visual-language-action models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 38574–38583, 2026. 2

[11] Maria Fox and Derek Long. Pddl2. 1: An extension to pddl for expressing temporal planning domains. *Journal of artificial intelligence research*, 20:61–124, 2003. 3

[12] Zhiyu Huang, Yun Zhang, Johnson Liu, Rui Song, Chen Tang, and Jiaqi Ma. Tic-vla: A think-in-control vision-language-action model for robot navigation in dynamic environments. *arXiv preprint arXiv:2602.02459*, 2026. 1, 2

[13] 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. 1, 2, 5

[14] Moo Jin Kim, Chelsea Finn, and Percy Liang. Fine-tuning vision-language-action models: Optimizing speed and success. *arXiv preprint arXiv:2502.19645*, 2025. 2, 5

[15] Boyi Li, Philipp Wu, Pieter Abbeel, and Jitendra Malik. Interactive task planning with language models. *arXiv preprint arXiv:2310.10645*, 2023. 1

[16] Meng Li, Zhen Zhao, Zhengping Che, Fei Liao, Kun Wu, Zhiyuan Xu, Pei Ren, Zhao Jin, Ning Liu, and Jian Tang. Switchvla: Execution-aware task switching for vision-language-action models. *arXiv preprint arXiv:2506.03574*, 2025. 1, 2

[17] Ziang Li, Dongzhou Cheng, Yibin Wang, Shiyue Wang, Xiaoyang Xu, Lingxuan Weng, Juan Wang, and Jiaqi Wang. Light-wam: Efficient world action models with state-fusion action decoding. *arXiv preprint arXiv:2606.08242*, 2026. 2

[18] Zekai Li, Jiaming Tang, and Zhijian Liu. Flashvla: Streaming action decoding for fast and asynchronous vla inference. *arXiv preprint arXiv:2608.27384*, 2026. 5

[19] Bo Liu, Yifeng Zhu, Chongkai Gao, Yihao Feng, Qiang Liu, Yuke Zhu, and Peter Stone. LIBERO: Benchmarking knowledge transfer for lifelong robot learning. In *NeurIPS*, pages 44776–44791, 2023. 1, 2

[20] Oier Mees, Lukas Hermann, Erick Rosete-Beas, and Wolfram Burgard. Calvin: A benchmark for language-conditioned policy learning for long-horizon robot manipulation tasks. *IEEE Robotics and Automation Letters*, 7(3): 7327–7334, 2022. 1

[21] Karl Pertsch, Kyle Stachowicz, Brian Ichter, Danny Driess, Suraj Nair, Quan Vuong, Oier Mees, Chelsea Finn, and Sergey Levine. Fast: Efficient action tokenization for vision-language-action models. *arXiv preprint arXiv:2501.09747*, 2025. 2

[22] Pratyusha Sharma, Balakumar Sundaralingam, Valts Blukis, Chris Paxton, Tucker Hermans, Antonio Torralba, Jacob Andreas, and Dieter Fox. Correcting robot plans with natural language feedback. *arXiv preprint arXiv:2204.05186*, 2022. 1

[23] Jingwen Sun, Wenyao Zhang, Zekun Qi, Shaojie Ren, Zezhi Liu, Hanxin Zhu, Guangzhong Sun, Xin Jin, and Zhibo

Chen. Vla-jepa: Enhancing vision-language-action model with latent world model. In *European Conference on Computer Vision*, pages 478–497. Springer, 2026. 2, 5

[24] Christopher Thierauf, Ravenna Thielstrom, Bradley Oosterveld, Will Becker, and Matthias Scheutz. "do this instead"—robots that adequately respond to corrected instructions. *ACM Transactions on Human-Robot Interaction*, 13 (3):1–23, 2024. 1

[25] Yuqi Wang, Xinghang Li, Wenxuan Wang, Junbo Zhang, Yingyan Li, Yuntao Chen, Xinlong Wang, and Zhaoxiang Zhang. Unified vision-language-action model. In *International Conference on Learning Representations*, pages 80929–80944, 2026. 2, 5

[26] Haozhe Xie, Beichen Wen, Jiarui Zheng, Zhaoxi Chen, Fangzhou Hong, Haiwen Diao, and Ziwei Liu. Dynamicvla: A vision-language-action model for dynamic object manipulation. *arXiv preprint arXiv:2601.22153*, 2026. 1, 2

[27] Shuai Yang, Hao Li, Bin Wang, Yilun Chen, Yang Tian, Tai Wang, Hanqing Wang, Feng Zhao, Yiyi Liao, and Jiangmiao Pang. Vision-language-action instruction tuning: From understanding to manipulation. In *International Conference on Learning Representations*, pages 152896–152943, 2026. 5