FORGE: AN EXECUTION-FEEDBACK-GUIDED FRAMEWORK FOR BENCHMARKING GENERALIST ROBOT POLICIES

Anonymous authors
Paper under double-blind review



Figure 1: **Overview of FORGE and FORGE-Bench. (a)** Comparison with existing VLA benchmarks. **(b)** FORGE pipeline for scene construction, process-aware evaluation, and feedback-guided mutation. **(c)** VLA performance on FORGE-Bench. **(d)** FORGE-Bench statistics across 215 task-scene variants and 5 capability dimensions.

ABSTRACT

As Vision–Language–Action (VLA) models advance toward open-ended robotic manipulation, reliable evaluation of their performance across diverse capability dimensions becomes increasingly important. Yet comprehensive evaluation remains difficult due to the multimodal nature of VLA inputs and the cost of constructing executable simulation environments and reliable oracles at scale. Existing VLA benchmarks largely rely on manually specified or template-instantiated task–scene–oracle triplets constructed independently of model performance, resulting in limited input diversity, overlooked process failures, and only coarse-grained difficulty levels. In this paper, we introduce **FORGE**, an execution-feedback-guided framework that dynamically generates capability-specific mutations from validated seed tasks and scenes using observed VLA execution behavior, enabling iterative exploration of model failure modes. Using FORGE, we build **FORGE-Bench**, comprising 215 task–scene–oracle variants generated along validated mutation trajectories across five capability dimensions. Evaluation of five representative VLA models shows dimension-specific failures, terminal–process success gaps, and local performance reversals along shared mutation trajectories, providing a discriminative characterization of VLA capabilities.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

1 INTRODUCTION

As VLA models continue to advance toward general-purpose robotic manipulation (Black et al., 2024; Intelligence et al., 2025; Pertsch et al., 2025; Bu et al., 2025; Ghosh et al., 2024; Kim et al., 2024; Zitkovich et al., 2023), systematically characterizing their strengths and limitations across diverse capability dimensions, such as vision–language grounding and long-horizon reasoning, becomes increasingly important. Although existing work has made substantial progress in evaluating VLA capabilities (Liu et al., 2023; Zhang et al., 2025a; Chen et al., 2025; Fan et al., 2026; Fei et al., 2026), such multidimensional evaluation remains challenging because VLA behavior depends jointly on multimodal inputs, diverse task–environment configurations, and execution-time interactions. Consequently, existing benchmarks (Lyu et al., 2026; Li et al., 2024) typically rely on manually specified or template-instantiated task–scene–oracle triplets, limiting coverage of the task and environment space and yielding incomplete or overly coarse characterizations of VLA capabilities.

The running example in Figure 1 (a) illustrates how existing static benchmarks (Liu et al., 2023; Fei et al., 2026; Zhang et al., 2025a) can lead to insufficient evaluation. First, manually constructing evaluation scenes is costly and difficult to scale, while reusing them across fine-tuning and evaluation may introduce benchmark-specific spatial priors that inflate success rates. Second, manually constructing dense difficulty gradients is difficult to scale, often yielding only a few coarse levels that obscure gradual performance changes and capability differences. Third, existing oracles based on state predicates may overlook environmental events that fall outside the predefined predicate checks, causing incorrect or incomplete judgments of an execution. More fundamentally, VLA benchmark construction should be viewed as an exploration problem over a large task–environment space, rather than as the static enumeration of predefined tasks. This perspective leads to a central question: **How can VLA benchmarks effectively explore large task and environment spaces to discover and instantiate diagnostic tests?**

To address this challenge, we propose **FORGE**, an end-to-end framework that leverages model execution feedback to transform task specifications into executable benchmark instances, and **FORGE-Bench**, a scalable, feedback-driven benchmark constructed with FORGE across five capability dimensions: (1) **Vision–Language Grounding**, which measures whether a VLA model can identify and manipulate the intended referent by grounding linguistic descriptions in visual observations, (2) **Linguistic Robustness**, which evaluates whether the model remains robust to meaning-preserving lexical and syntactic variations in task instructions, (3) **Deployment Configuration Robustness**, which measures whether the model remains effective under perturbations to specific deployment configurations such as object mass and lighting conditions, (4) **Long-Horizon Manipulation**, which evaluates the model's ability to plan and execute extended sequences of interdependent subgoals, and (5) **Process-Level Safety**, which assesses whether the model avoids unsafe interactions and unintended state changes throughout task execution. By progressively generating capability-specific variants from validated seed scenes using execution feedback, FORGE effectively explores the complex task–environment search space and uncovers previously obscured model behaviors.

In summary, this paper makes the following contributions:

- **Execution-Feedback-Guided Benchmark Construction.** We propose **FORGE**, a scalable framework that iteratively constructs and mutates task–scene pairs using observed VLA execution outcomes, enabling traceable benchmark generation beyond static manually designed tests and facilitating the discovery of diagnostic instances across large task–scene spaces.
- **FORGE-Bench.** We introduce FORGE-Bench, a fine-grained VLA benchmark spanning five capability dimensions consisting of 215 task-scene-oracle triplets. It organizes validated variants into capability-specific mutation trajectories, enabling denser difficulty gradients and more precise characterization of performance degradation and capability differences.
- **Systematic VLA Evaluation.** We evaluate five representative VLA policies on FORGE-Bench and identify capability-specific model strengths, diverse and non-monotonic degradation patterns along mutation trajectories, and substantial gaps between terminal and process-level success, exposing previously obscured failure patterns and motivating further research on general-purpose VLA models and process-aware evaluation.

2 RELATED WORK

2.1 SIMULATION-BASED VLA BENCHMARKS

Simulation-based VLA benchmarks standardize environments, tasks, and evaluation protocols across increasingly diverse manipulation capabilities. CALVIN targets long-horizon language-conditioned execution, LIBERO studies lifelong learning and knowledge transfer, BEHAVIOR-1K scales evaluation to diverse long-horizon household activities, and VLABench emphasizes commonsense knowledge and long-horizon reasoning (Mees et al., 2022; Liu et al., 2023; Li et al., 2024; Zhang et al., 2025b). More recent benchmarks emphasize robustness and safety: LIBERO-Plus, RoboTwin 2.0, and VLA-Arena introduce controlled variations in visual conditions, object properties, language, observations, task structures, and safety constraints; SafeAgentBench studies safety-aware task planning under hazardous instructions, while ForesightSafety-VLA evaluates process-level physical risks during VLA execution. (Fei et al., 2026; Chen et al., 2025; Zhang et al., 2025a; Yin et al., 2024; Lyu et al., 2026). Most benchmarks still determine task success through symbolic state predicates, although recent safety-oriented methods incorporate simulator events, safety costs, risk exposure, or temporal specifications to capture process-level failures (Li et al., 2024; Zhang et al., 2025a; Lyu et al., 2026; Fan et al., 2026). Overall, evaluation remains largely tied to predefined task–scene configurations and oracle logic, limiting fine-grained capability characterization.

Table 1: Comparison of environment coverage, benchmark construction, oracle design, and supported evaluation dimensions. **MR**: multi-room or whole-house layouts; **Auto**: automated scene generation; **CV**: controlled test variants; **Grad**: graded difficulty; **FB**: execution-feedback-driven construction; **PO**: process-aware oracle; **FM**: failure mode diagnosis. **VIS**, **LING**, **DCR**, **LONG**, and **SAFE** denote Vision–Language Grounding, Linguistic Robustness, Deployment Configuration Robustness, Long-Horizon Manipulation, and Process Safety, respectively. ✓, △, and ✗ denote **explicit**, **partial/indirect**, and **absent** support, respectively. †**SafeAgentBench** evaluates high-level embodied planning rather than low-level VLA control.

	Environment		Construction				Evaluation		Evaluation Dimensions				
Benchmark	**Simulator**	**MR**	**Auto**	**CV**	**Grad**	**FB**	**PO**	**FM**	**VIS**	**LING**	**DCR**	**LONG**	**SAFE**
CALVIN	PyBullet	✗	✗	✗	△	✗	✗	✗	✓	△	✗	✓	✗
LIBERO	MuJoCo	✗	△	△	✗	✗	✗	✗	✓	✗	✗	✓	✗
BEHAVIOR-1K	OmniGibson	✓	△	✗	✗	✗	✗	✗	△	✗	✗	✓	✗
VLABench	MuJoCo	✗	△	△	△	✗	✗	△	✓	△	△	✓	✗
RoboTwin 2.0	SAPIEN	✗	△	△	✗	✗	✗	✗	△	✓	✗	△	✗
RoboLab	Isaac Lab	✗	△	✓	✓	✗	△	✓	✓	✓	✓	△	△
The Colosseum	CoppeliaSim	✗	✗	✓	△	✗	✗	✗	△	✗	✓	✗	✗
LIBERO-Plus	MuJoCo	✗	✗	✓	✓	✗	✗	△	✓	✓	✗	△	✗
VLA-Arena	MuJoCo	✗	△	✓	✓	✗	△	△	✓	✓	✗	✓	✓
SafeAgentBench†	AI2-THOR	✗	△	△	✗	✗	✗	✗	△	△	✗	✓	△
Safety-CHORES	AI2-THOR	✓	✓	✓	✗	✗	△	△	✓	✗	✗	✓	✓
LIBERO-Safety	MuJoCo	✗	✓	✓	✓	✗	△	△	✓	△	△	△	✓
ForesightSafety-VLA	SAPIEN	✗	✗	✓	✓	✗	△	✓	✓	✓	✗	△	✓
SafeVLA-Bench	MuJoCo	✗	✗	✗	✗	✗	△	✓	✗	✗	✗	✗	✓
FORGE-Bench	**Isaac Sim**	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

2.2 AUTOMATED BENCHMARK CONSTRUCTION METHODS

Existing work automates benchmark construction through structured task specification, generative task creation, and environment generation. BDDL formalizes task objects, initial conditions, and goals using symbolic predicates, while CBDDL extends this formulation with task customization and safety constraints (Li et al., 2024; Zhang et al., 2025a). At the task level, GenSim generates task code and curricula from prompts, RoboGen proposes skills and supervision, PARTNR combines language-model generation with simulation-in-the-loop verification, and ManiTaskGen and TEA derive feasible tasks from existing environments (Wang et al., 2024a;b; Chang et al., 2025; Dai et al., 2025; He et al., 2026). At the environment level, ProcTHOR and LUMINOUS procedurally generate diverse interactive environments, while Holodeck synthesizes 3D scenes from language descriptions. GRS reconstructs simulation scenes and task specifications, RobotArena∞ converts real-world demonstrations into reproducible simulated evaluations, and RoboTwin 2.0 supports au-

tomated task generation and structured domain randomization (Deitke et al., 2022; Zhao et al., 2021; Yang et al., 2024; Zook et al., 2025; Jangir et al., 2025; Chen et al., 2025). Most recently, RoboLab further advances scalable simulation-based evaluation through human- and LLM-enabled scene/task generation and controlled perturbation-based policy analysis (Yang et al., 2026). Overall, these approaches substantially improve the scalability and diversity of benchmark construction, while tasks, scenes, and perturbations are still typically specified independently of model behavior during evaluation.

3 FORGE: AN EXECUTION-FEEDBACK-GUIDED FRAMEWORK

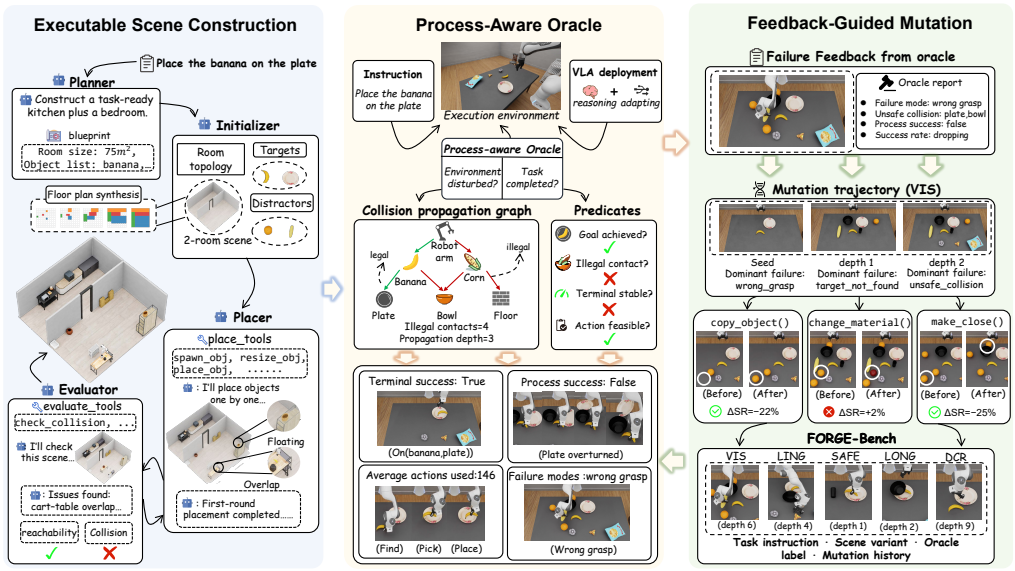

Figure 2: Overview of the FORGE framework and its three core modules.

3.1 Overview of FORGE

To systematically explore large task and environment spaces for VLA evaluation, we develop **FORGE**, a feedback-guided benchmark construction pipeline, and then use it to build **FORGE-Bench**, which organizes controlled task–scene variants across five evaluation dimensions. As shown in Figure 2, FORGE constructs the benchmark through three stages: (1) *Executable Scene Construction*, (2) *Process-Aware Oracle*, and (3) *Feedback-Guided Scene Mutation*. Starting from a fixed set of seed tasks, this pipeline progressively generates capability-specific variants that form the final benchmark.

3.2 Executable Scene Construction

FORGE constructs each executable and mutable seed scene via multi-agent collaboration among a scene planner, room initializer, asset placer, and repair evaluator, as shown in Figure 2.

Task-grounded scene planning. The scene planner converts the initial task into a structured scene blueprint. It identifies the task setting and required objects, and proposes additional distractors that may be introduced in later scene variations. The resulting blueprint defines the task environment, task-relevant objects, and a distractor pool for subsequent scene mutations.

Assets and topology initialization. Given the scene blueprint, the initializer grounds each scene-plan object in the USD asset library through a two-stage RAG-and-selection pipeline: it retrieves the top-k candidates by semantic description and required properties, then an LLM selects the final

asset based on scene context, placement, embodiment, and failure hypothesis. After grounding the required objects, FORGE further generates a floor plan to organize the scene at the room level to support mobile and whole-body VLA evaluation beyond fixed tabletop settings. Specifically, the initializer generates connected multi-room layouts using a UCB strategy that balances layout exploration, room-area balance, and geometric feasibility (see Appendix B).

Spatial and physical realization. The placer then instantiates the selected assets and translates semantic constraints, such as support, hanging, or floor placement, into simulator poses using continuously updated scene geometry and object bounds. It preserves the reachability and manipulability of task-critical entities while positioning hypothesis-relevant entities to realize the intended diagnostic condition without making the task physically unsolvable. These layouts, spatial relations, and object poses are recorded for subsequent validation and repair.

Iterative validation and repair. The evaluator node verifies whether the realized scene is a valid and executable instance of the compiled specification through semantic, physical, and task-level checks. It assesses required entities and relations, scene validity and stability, and whether the task remains reachable and feasible. When violations are detected, the evaluator returns a structured report to the placer, which revises the scene accordingly; this placement–evaluation loop continues until the scene passes validation or the retry limit is reached.

3.3 Process-aware Oracle

As summarized in the center of Figure 2, FORGE consolidates task-completion signals and broader interaction events into an execution trace for evidence-grounded verdicts and mutation feedback.

Collision propagation graph. FORGE constructs a collision propagation graph from robot-initiated contacts during execution. Each contact is linked to downstream object interactions, allowing the oracle to trace how an initial robot action propagates through the environment and to distinguish task-relevant interaction chains from unintended disturbances or illegal contacts. This graph provides process-level evidence that complements terminal spatial predicates, especially when the final task state appears successful despite invalid intermediate interactions.

Trace-grounded adjudication and feedback. FORGE combines terminal predicates with process evidence from the execution trace and collision propagation graph to adjudicate each rollout. A rollout succeeds only if the task goal is achieved without process-level violations. Otherwise, the corresponding failure mode is recorded and used to guide subsequent scene mutation.

3.4 Feedback-Guided Mutation

FORGE uses policy-execution feedback to explore the mutation space over validated seed scenes. As illustrated in Figure 2, observed execution outcomes guide subsequent scene mutations toward diagnostic variants.

Structured mutation search space. FORGE defines a structured set of capability-specific mutation operators for seed scenes (see Appendix D for detailed operator list). These operators support both parametric changes, such as appearance, pose, camera, and physical properties, and structural changes to objects, semantic relations, and environment layouts. Based on execution feedback, FORGE selects dimension-specific operators to generate the next scene variant.

Feedback-guided mutation decision. For a candidate at mutation depth d, FORGE evaluates the VLA model under the candidate configuration and summarizes its success rate and observed failure modes. This feedback guides the next mutation toward stronger evidence for the targeted failure hypothesis, without assuming monotonic difficulty progression. Candidates that cross a predefined success-rate threshold are retained as diagnostic instances, after which FORGE restarts from the seed to explore alternative mutation trajectories until the mutation budget is exhausted.

Mutation realization. We develop a custom Isaac Sim extension for realizing high-level mutations through USD layer composition. The extension translates each mutation intent into low-level

edits of USD scene attributes, including object transforms, visual properties, physical parameters, camera configurations, and scene relations. These edits are written to a new USDA override layer, allowing FORGE to modify the environment efficiently while preserving the underlying seed scene and avoiding full scene regeneration.

4 EXPERIMENTS

We organize our experiments around two main research questions:

RQ1. Can FORGE effectively construct executable and diagnostic benchmark instances?

RQ2. How do VLA models perform under variations across five dimensions?

We first describe the common experimental setup and evaluation metrics, followed by experiments addressing RQ1 and RQ2.

4.1 EXPERIMENTAL SETUP

Evaluation setup. We implement and evaluate FORGE in NVIDIA Isaac Sim 4.5.0 with PhysX (NVIDIA, 2025), using Isaac Sim APIs for scene construction, mutation, and robot control. The multi-agent workflow is implemented with LangChain (Chase, 2022) and orchestrated as a stateful graph using LangGraph (LangChain, 2024). GPT-5.6 Sol (OpenAI, 2026) serves as the common LLM backend for the scene-construction and validation agents. To mitigate evaluation bias from policies adapted to a specific simulation scene distribution, we select DROID-based models trained primarily on real-world robot data (Khazatsky et al., 2024). For the π family, we use the PolaRiS-ready joint-position checkpoints, which are further co-trained for 1k steps on a mixture of 90% DROID and 10% PolaRiS simulation data (Jain et al., 2025): **π_0-PolaRiS-DROID-joint-position** (Black et al., 2024), **π_0-FAST–PolaRiS-DROID-joint-position** (Pertsch et al., 2025), and **$\pi_{0.5}$-PolaRiS-DROID-joint-position** (Intelligence et al., 2025). We additionally evaluate **GR00T N1.7-DROID** (Bjorck et al., 2025) and **MolmoAct 2-DROID** (Fang et al., 2026). All models are deployed zero-shot in FORGE seed scenes. For benchmark construction, each task–scene pair is mutated for up to 10 steps along five evaluation dimensions, with 30 rollout episodes per variant providing success statistics and oracle-identified failure modes to guide subsequent mutations.

Baselines. We compare the scene-construction capability of FORGE with Holodeck (Yang et al., 2024) and SceneWeaver (Yang et al., 2025) on 75 tasks, including 25 from LIBERO (Liu et al., 2023), 25 from BEHAVIOR-1K (Li et al., 2024), and 25 hand-curated structurally diverse tasks beyond standard tabletop manipulation.

Metrics. We evaluate generated scenes in terms of physical plausibility, task solvability, and visual quality, using overlap, stability, task-object coverage and reachability, and LLM-based visual assessment. For VLA evaluation, we report process success rate (PSR), the terminal–process success gap (TSR–PSR), and action efficiency (Act) measured by the median number of actions among process-successful rollouts.

4.2 CAN FORGE EFFECTIVELY CONSTRUCT EXECUTABLE AND DIAGNOSTIC BENCHMARK INSTANCES?

Simulated scene construction. We first evaluate whether FORGE can reliably instantiate diverse task specifications into physically valid and executable scenes. Table 2 shows that FORGE achieves 6.2% overlap, 100.0% stability and coverage, and 90.7% reachability while maintaining competitive visual quality. Although SceneWeaver obtains higher realism and layout scores, its coverage and reachability drop to 50.6% and 20.7%, indicating that visual quality alone does not ensure executability. Additionally, iterative repair further reduces FORGE's overlap rate from 28.3% to 6.2% and improves reachability rate from 70.3% to 90.7%.

Collision-propagation-graph oracle. We validate the FORGE oracle on 396 rollouts spanning five evaluation dimensions in Table 3. Human reviewers annotate both terminal and process success, where process success additionally requires executions to avoid illegal contacts, unintended

Table 2: RQ1 scene-construction results. ↓ indicates lower is better; ↑ indicates higher is better.

Method	Physical Plausibility		Task Solvability		GPT-based Visual Quality		
	Overlap ↓	Stability ↑	Object Cov. ↑	Reach. ↑	Realism ↑	Function. ↑	Layout ↑
Expert References							
LIBERO	20.5%	100.0%	100.0%	100.0%	3.7	3.4	3.8
Generation Methods							
Holodeck	24.7%	75.4%	49.3%	12.0%	5.5	5.1	5.5
SceneWeaver	35.7%	72.4%	50.6%	20.7%	**5.7**	**5.8**	**6.5**
FORGE W/O REPAIR	28.3%	**100.0%**	**100.0%**	70.3%	4.4	4.3	4.2
FORGE	**6.2%**	**100.0%**	**100.0%**	**90.7%**	5.1	**5.8**	5.9

Table 3: Oracle validation against human process-level judgments on 396 execution rollouts. A false success is a failed execution incorrectly accepted by the oracle, whereas a false failure is a successful execution incorrectly rejected. ↑ indicates higher is better and ↓ indicates lower is better.

Oracle	Judgment quality		Error accounting			Total errors ↓
	Accuracy ↑ (%)	Failure F1 ↑ (%)	Detected failures ↑ n (rate)	False successes ↓ n (rate)	False failures ↓ n (rate)	n (rate)
Terminal predicate	79.3%	81.2%	177 (70.2%)	75 (29.8%)	**7 (4.9%)**	82 (20.7%)
FORGE Oracle	**95.7%**	**96.7%**	**248 (98.4%)**	**4 (1.6%)**	13 (9.0%)	**17 (4.3%)**
Change	+16.4 pp	+15.5 pp	+71	−71	+6	−65

interactions, and other process-level failures. Relative to the terminal-only predicate, the FORGE oracle improves accuracy from 79.3% to 95.7% and failure F1 from 81.2% to 96.7%, while decreasing false successes from 75 to 4, a 94.7% relative reduction. As a result, it makes total errors fall from 82 to 17, despite only a modest increase in false failures from 7 to 13.

Finding 1. Terminal-state success is insufficient for reliable VLA evaluation, as it can systematically overestimate capability by ignoring execution process failures.

Table 4: Mutation-search efficiency under the same 600-rollout budget. Reductions are measured relative to independently sampled root. We report search quality, sample efficiency and policy-wise maximum reduction.

Metric	Feedback-guided	Random walk	Gain
Search quality			
Maximum reduction ↑	**50.0 pp**	21.7 pp	**2.3**×
Budget-averaged best reduction ↑	**32.5 pp**	16.2 pp	**2.0**×
Worst-case maximum reduction ↑	**50.0 pp**	11.7 pp	**4.3**×
Sample efficiency			
Rollouts to ≥ 10 pp reduction ↓	120	**60**	**0.5**×
Rollouts to ≥ 20 pp reduction ↓	**120**	540	**4.5**×
Rollouts to ≥ 50 pp reduction ↓	**480**	> 600	**Not reached**
Policy-wise maximum reduction			
π_0-FAST	**53.3 pp**	20.0 pp	**2.7**×
MolmoAct 2	**56.7 pp**	23.3 pp	**2.4**×

Feedback-guided benchmark construction To evaluate feedback-guided mutation, we compare FORGE with random mutation under the same source models, mutation setup, and 600-rollout bud-

get. Table 4 shows that FORGE achieves substantially larger reductions, reaching 50.0 pp maximum and 32.5 pp budget-averaged best reduction, compared with 21.7 pp and 16.2 pp for random walk. Random walk reaches a 10 pp reduction earlier, but FORGE reaches 20 pp after only 120 rollouts and is the only method to attain 50 pp within budget. This advantage is also consistent across both source policies.

Finding 2. Execution feedback makes benchmark construction more efficient by steering mutations toward harder instances under the same rollout budget.

4.3 How do VLA models perform under variations across five dimensions?

We use π_0-FAST and MolmoAct 2 as source models to drive feedback-guided mutation, representing autoregressive action-token prediction and flow matching, respectively. This process produces 215 validated variants along mutation trajectories of up to 10 search steps. Since reporting all variants in a single table is impractical, we additionally use GR00T N1.7 as a calibration model to select three representative levels for compact presentation, while the complete results are reported in Appendix H. We further evaluate the held-out π_0 and $\pi_{0.5}$ models on these representative instances. Table 5 reports success rate and model action efficiency across the five dimensions.

Table 5: Part of five-model results on the FORGE-Bench. **PSR** is process success rate; **T–P Gap** is TSR − PSR in percentage points; **Act.** is the median number of action steps actually executed in Isaac Sim; L1–L3 are the selected mutation levels.

		Source Models						Calibration Model			Held-out Models					
		π_0-FAST			MolmoAct 2			GR00T N1.7			π_0			$\pi_{0.5}$		
Dim.	**Metric**	L1	L2	L3	L1	L2	L3	L1	L2	L3	L1	L2	L3	L1	L2	L3
Root (L0)	PSR ↑		29.3			**30.0**			6.7			26.7			15.3	
	T–P Gap ↓		14.0			**10.0**			20.0			14.0			42.7	
	Act. ↓		150.5			147.5			180.0			159.0			**146.0**	
VIS	PSR ↑	**25.3**	**21.3**	2.0	20.0	12.7	5.3	22.0	3.3	**8.0**	14.7	6.0	3.3	10.7	14.0	5.3
	T–P Gap ↓	24.0	21.3	26.0	22.0	**8.0**	10.0	31.3	12.7	10.7	**12.0**	13.3	**5.3**	38.0	28.7	18.7
	Act. ↓	158.0	**153.5**	190.0	181.5	288.0	190.5	158.0	165.0	191.0	145.0	226.0	239.0	**144.5**	155.0	**104.0**
LING	PSR ↑	21.3	7.3	8.7	**33.3**	**32.7**	10.0	20.7	26.0	**15.3**	6.0	1.3	10.0	8.7	6.0	5.3
	T–P Gap ↓	23.3	10.7	16.0	**3.3**	**6.0**	9.3	22.7	12.7	16.7	14.0	8.0	**4.7**	38.7	26.0	24.7
	Act. ↓	138.5	**138.0**	127.0	165.0	183.0	163.0	245.0	207.0	280.5	191.0	172.5	117.0	**116.0**	143.0	**116.0**
DCR	PSR ↑	**40.7**	**26.0**	7.3	34.7	24.7	13.3	2.0	6.7	5.3	22.7	14.0	**15.3**	15.3	12.0	8.7
	T–P Gap ↓	18.0	**13.3**	21.3	21.3	14.0	**7.3**	22.7	14.0	16.7	**17.3**	19.3	14.0	40.0	38.0	43.3
	Act. ↓	145.0	169.5	137.0	176.0	165.0	179.0	191.0	177.5	190.0	164.0	195.5	199.5	**122.0**	**161.0**	**126.5**
LONG	PSR ↑	16.0	1.3	0.0	**37.3**	**4.7**	0.0	12.7	2.0	0.0	12.7	0.7	**1.3**	8.7	1.3	0.0
	T–P Gap ↓	16.7	7.3	0.0	16.0	3.3	0.0	23.3	10.7	0.0	**14.0**	**2.0**	2.0	33.3	12.0	0.0
	Act. ↓	146.0	634.5	–	322.0	523.0	–	242.0	472.0	–	167.5	505.0	**377.0**	**92.0**	**208.0**	–
SAFE	PSR ↑	27.3	15.3	3.3	**38.7**	**22.7**	4.0	5.3	6.0	7.3	15.3	7.3	**12.0**	10.7	7.3	3.3
	T–P Gap ↓	30.0	32.7	27.3	**16.0**	**12.0**	**6.7**	24.7	12.7	8.7	18.0	19.3	12.0	42.7	38.7	36.0
	Act. ↓	146.0	146.0	151.0	161.5	192.5	186.0	171.0	178.5	211.5	304.0	128.5	**127.5**	**116.0**	**117.5**	130.0

Cross-dimension capability evaluation. Across the three representative levels in Table 5, MolmoAct 2 achieves the highest aggregate PSR (19.6%), while relative strengths vary across dimensions: π_0-FAST performs better on VIS (16.2%) and DCR (24.7%), whereas MolmoAct 2 leads on LING (25.3%), LONG (14.0%), and SAFE (21.8%). The complete mutation trajectories exhibit the same overall pattern, with MolmoAct 2 maintaining the highest aggregate PSR, supporting the calibrated levels as a compact summary of model capability differences. Across models, success rates also show broadly similar level-dependent changes, indicating that FORGE-Bench induces meaningful difficulty gradients while preserving model-specific capability profiles.

Finding 3. Models show dimension-specific strengths with no single model consistently leading across all evaluation dimensions, while fewer actions do not imply greater execution reliability: $\pi_{0.5}$ remains action-efficient yet exhibits a large terminal–process gap.

Finding 4. Diagnostic instances discovered by FORGE remain informative for held-out VLA models, supporting cross-model transferability of the calibrated benchmark. (see Appendix H)

Non-monotonic degradation. Figure 3(a–c) shows that mutation depth does not induce a monotone difficulty scale. LONG exhibits the clearest consistent degradation, with all three models eventually collapsing to near-zero success. Other dimensions show model-specific responses, including even local reversals. For example, VIS shows the opposite pattern from step 9 to step 10: π_0-FAST declines slightly, whereas MolmoAct 2 and GR00T rebound substantially.

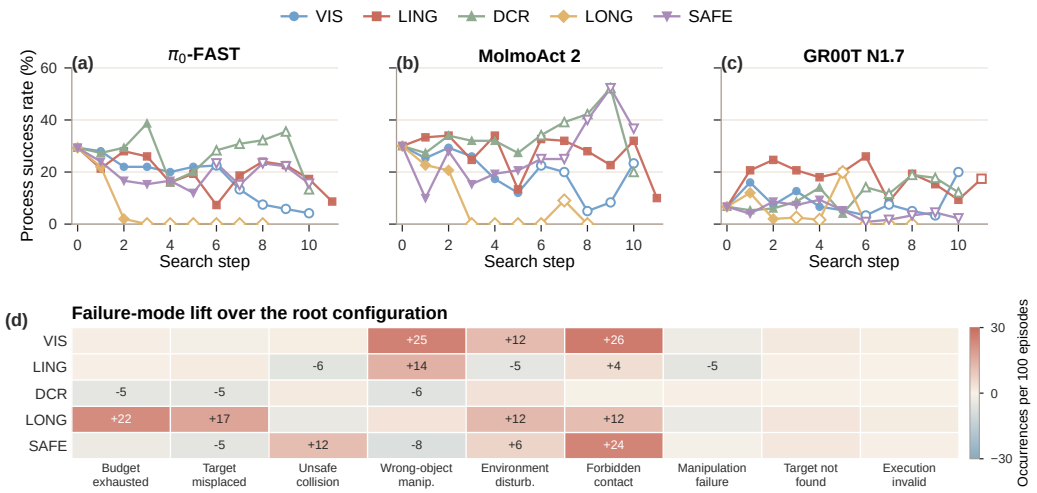

Figure 3: **Complete three-model mutation trajectories and failure signatures. (a–c)** Episode-pooled PSR over all observed search variants for the two source models and the calibration model. Open markers indicate incomplete task coverage. **(d)** Change in failure-mode incidence relative to each model's root configuration, then averaged equally over the three models.

Finding 5. At certain mutation steps discovered by FORGE, models exhibit markedly different behaviors: some show only minor fluctuations in success rate, while others undergo sharp changes. (see detailed case study in Appendix I)

Detected Failure Modes. Analysis in Figure 3(d) shows distinct diagnostic signatures: VIS mainly increases wrong-object manipulation and forbidden contact, SAFE emphasizes forbidden contact and unsafe collision, and LONG increases budget exhaustion and target misplacement.

Finding 6. Different capability dimensions expose distinct process-level failure signatures that are obscured by aggregate success rates.

5 CONCLUSION

We introduced FORGE, an execution-feedback-guided framework for constructing executable and diagnostic VLA tests, and FORGE-Bench, a benchmark spanning five capability dimensions. FORGE integrates automated scene construction, process-aware execution assessment, and feedback-guided mutation to explore complex task–scene spaces and identify diagnostic instances. Across five representative VLA models, FORGE-Bench identifies dimension-specific capability profiles, local performance reversals, terminal–process gaps, and distinct failure signatures. These results suggest that execution-feedback-guided benchmark construction can provide a more fine-grained view of VLA capabilities than aggregate success rates alone.

AI USE STATEMENT

Generative AI tools were used as research-assistance tools for literature retrieval and discovery, research ideation, feedback on research methodology and experimental design, drafting selected portions of the manuscript, and language editing and stylistic polishing. Where AI-assisted suggestions contributed to the research process, the resulting ideas, experimental choices, analyses, citations, and manuscript text were reviewed and verified by the authors before being incorporated into this work. The authors take full responsibility for the scientific claims, experimental results, citations, artifacts, and final content of the paper. Separately, large language models are used as functional components of the proposed FORGE pipeline. Their roles, configurations, and use during benchmark construction are described explicitly as part of the methodology and experimental setup, rather than as research or writing assistance.

ETHICS STATEMENT

This work studies Vision–Language–Action models in simulated robotic environments and does not involve real-world deployment on human subjects. Human annotation is used only to assess simulated robot rollouts, and no sensitive personal information is collected or analyzed. We follow the licenses and usage terms of the datasets, models, and software used in this work.

REPRODUCIBILITY STATEMENT

We provide detailed descriptions of the FORGE pipeline, scene-construction procedure, process-aware oracle, feedback-guided mutation strategy, and evaluation protocol in the main paper and appendix. We further report the simulation environment, software versions, LLM and agent frameworks, evaluated VLA checkpoints, benchmark task definitions, mutation operators, and evaluation metrics to facilitate reproduction of our results. An anonymized version of code is available at this anonymous repository to facilitate reproduction and direct comparison with our results.

REFERENCES

Johan Bjorck, Fernando Castañeda, Nikita Cherniadev, Xingye Da, Runyu Ding, Linxi Fan, Yu Fang, Dieter Fox, Fengyuan Hu, Spencer Huang, Joel Jang, Zhenyu Jiang, Jan Kautz, Kaushil Kundalia, Lawrence Lao, Zhiqi Li, Zongyu Lin, Kevin Lin, Guilin Liu, Edith LLontop, Loic Magne, Ajay Mandlekar, Avnish Narayan, Soroush Nasiriany, Scott Reed, You Liang Tan, Guanzhi Wang, Zu Wang, Jing Wang, Qi Wang, Jiannan Xiang, Yuqi Xie, Yinzhen Xu, Zhenjia Xu, Seonghyeon Ye, Zhiding Yu, Ao Zhang, Hao Zhang, Yizhou Zhao, Ruijie Zheng, and Yuke Zhu. GR00T N1: an open foundation model for generalist humanoid robots. *CoRR*, abs/2503.14734, 2025. doi: 10.48550/ARXIV.2503.14734. URL https://doi.org/10.48550/arXiv.2503.14734.

Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Lucy Xiaoyang Shi, James Tanner, Quan Vuong, Anna Walling, Haohuan Wang, and Ury Zhilinsky. π_0: A vision-language-action flow model for general robot control. *CoRR*, abs/2410.24164, 2024. doi: 10.48550/ARXIV.2410.24164. URL https://doi.org/10.48550/arXiv.2410.24164.

Qingwen Bu, Yanting Yang, Jisong Cai, Shenyuan Gao, Guanghui Ren, Maoqing Yao, Ping Luo, and Hongyang Li. Univla: Learning to act anywhere with task-centric latent actions. *CoRR*, abs/2505.06111, 2025. doi: 10.48550/ARXIV.2505.06111. URL https://doi.org/10.48550/arXiv.2505.06111.

Matthew Chang, Gunjan Chhablani, Alexander Clegg, Mikael Dallaire Cote, Ruta Desai, Michal Hlavac, Vladimir Karashchuk, Jacob Krantz, Roozbeh Mottaghi, Priyam Parashar, Siddharth Patki, Ishita Prasad, Xavier Puig, Akshara Rai, Ram Ramrakhya, Daniel Tran, Joanne Truong, John M. Turner, Eric Undersander, and Tsung-Yen Yang. PARTNR: A benchmark for planning and reasoning in embodied multi-agent tasks. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025. URL https://openreview.net/forum?id=T5QLRRHyL1.

Harrison Chase. Langchain. https://github.com/langchain-ai/langchain, 2022.

Tianxing Chen, Zanxin Chen, Baijun Chen, Zijian Cai, Yibin Liu, Qiwei Liang, Zixuan Li, Xianliang Lin, Yiheng Ge, Zhenyu Gu, Weiliang Deng, Yubin Guo, Tian Nian, Xuanbing Xie, Qiangyu Chen, Kailun Su, Tianling Xu, Guodong Liu, Mengkang Hu, Huan-ang Gao, Kaixuan Wang, Zhixuan Liang, Yusen Qin, Xiaokang Yang, Ping Luo, and Yao Mu. Robotwin 2.0: A scalable data generator and benchmark with strong domain randomization for robust bimanual robotic manipulation. *CoRR*, abs/2506.18088, 2025. doi: 10.48550/ARXIV.2506.18088. URL https://doi.org/10.48550/arXiv.2506.18088.

Liu Dai, Haina Wang, Weikang Wan, and Hao Su. Manitaskgen: A comprehensive task generator for benchmarking and improving vision-language agents on embodied decision-making. *CoRR*, abs/2505.20726, 2025. doi: 10.48550/ARXIV.2505.20726. URL https://doi.org/10.48550/arXiv.2505.20726.

Matt Deitke, Eli VanderBilt, Alvaro Herrasti, Luca Weihs, Kiana Ehsani, Jordi Salvador, Winson Han, Eric Kolve, Aniruddha Kembhavi, and Roozbeh Mottaghi. Procthor: Large-scale embodied AI using procedural generation. In Sanmi Koyejo, S. Mohamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*, 2022. URL http://papers.nips.cc/paper_files/paper/2022/hash/27c546ab1e4f1d7d638e6a8dfbad9a07-Abstract-Conference.html.

Jialiang Fan, Weizhe Xu, Oleg Sokolsky, Insup Lee, and Fanxin Kong. Safevla-bench: A benchmark for the success-safety gap in vision-language-action models. *CoRR*, abs/2606.00773, 2026. doi: 10.48550/ARXIV.2606.00773. URL https://doi.org/10.48550/arXiv.2606.00773.

Haoquan Fang, Jiafei Duan, Donovan Clay, Sam Wang, Shuo Liu, Weikai Huang, Xiang Fan, Wei-Chuan Tsai, Shirui Chen, Yi Ru Wang, Shanli Xing, Jaemin Cho, Jae Sung Park, Ainaz Eftekhar, Peter Sushko, Karen Farley, Angad Wadhwa, Cole Harrison, Winson Han, Ying-Chun Lee, Eli VanderBilt, Rose Hendrix, Suveen Ellawela, Lucas Ngoo, Joyce Chai, Zhongzheng Ren, Ali Farhadi, Dieter Fox, and Ranjay Krishna. Molmoact2: Action reasoning models for real-world deployment. *CoRR*, abs/2605.02881, 2026. doi: 10.48550/ARXIV.2605.02881. URL https://doi.org/10.48550/arXiv.2605.02881.

Senyu Fei, Siyin Wang, Junhao Shi, Zihao Dai, Jikun Cai, Pengfang Qian, Li Ji, Xinzhe He, Shiduo Zhang, Zhaoye Fei, Jinlan Fu, Jingjing Gong, and Xipeng Qiu. Libero-plus: A progressive robustness benchmark for visual-language-action models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2026, Denver, CO, USA, June 3-7, 2026*, pp. 38574–38583. Computer Vision Foundation, 2026. URL https://openaccess.thecvf.com/content/CVPR2026/html/Fei_LIBERO-Plus_A_Progressive_Robustness_Benchmark_for_Visual-Language-Action_Models_CVPR_2026_paper.html.

Dibya Ghosh, Homer Rich Walke, Karl Pertsch, Kevin Black, Oier Mees, Sudeep Dasari, Joey Hejna, Tobias Kreiman, Charles Xu, Jianlan Luo, You Liang Tan, Lawrence Yunliang Chen, Quan Vuong, Ted Xiao, Pannag R. Sanketi, Dorsa Sadigh, Chelsea Finn, and Sergey Levine. Octo: An open-source generalist robot policy. In Dana Kulic, Gentiane Venture, Kostas E. Bekris, and Enrique Coronado (eds.), *Robotics: Science and Systems XX, Delft, The Netherlands, July 15-19, 2024*, 2024. doi: 10.15607/RSS.2024.XX.090. URL https://doi.org/10.15607/RSS.2024.XX.090.

Xinyi He, Ying Yang, Chuanjian Fu, Sihan Guo, Song-Chun Zhu, Lifeng Fan, Zhenliang Zhang, and Yujia Peng. Automatic cognitive task generation for in-situ evaluation of embodied agents. *CoRR*, abs/2602.05249, 2026. doi: 10.48550/ARXIV.2602.05249. URL https://doi.org/10.48550/arXiv.2602.05249.

Physical Intelligence, Kevin Black, Noah Brown, James Darpinian, Karan Dhabalia, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Manuel Y. Galliker, Dibya Ghosh, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Devin

LeBlanc, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Allen Z. Ren, Lucy Xiaoyang Shi, Laura Smith, Jost Tobias Springenberg, Kyle Stachowicz, James Tanner, Quan Vuong, Homer Walke, Anna Walling, Haohuan Wang, Lili Yu, and Ury Zhilinsky. $\pi_{0.5}$: a vision-language-action model with open-world generalization. *CoRR*, abs/2504.16054, 2025. doi: 10.48550/ARXIV.2504.16054. URL https://doi.org/10.48550/arXiv.2504.16054.

Arhan Jain, Mingtong Zhang, Kanav Arora, William Chen, Marcel Torne, Muhammad Zubair Irshad, Sergey Zakharov, Yue Wang, Sergey Levine, Chelsea Finn, Wei-Chiu Ma, Dhruv Shah, Abhishek Gupta, and Karl Pertsch. Polaris: Scalable real-to-sim evaluations for generalist robot policies. *CoRR*, abs/2512.16881, 2025. doi: 10.48550/ARXIV.2512.16881. URL https://doi.org/10.48550/arXiv.2512.16881.

Yash Jangir, Yidi Zhang, Kashu Yamazaki, Chenyu Zhang, Kuan-Hsun Tu, Tsung-Wei Ke, Lei Ke, Yonatan Bisk, and Katerina Fragkiadaki. Robotarena ∞: Scalable robot benchmarking via real-to-sim translation. *CoRR*, abs/2510.23571, 2025. doi: 10.48550/ARXIV.2510.23571. URL https://doi.org/10.48550/arXiv.2510.23571.

Alexander Khazatsky, Karl Pertsch, Suraj Nair, Ashwin Balakrishna, Sudeep Dasari, Siddharth Karamcheti, Soroush Nasiriany, Mohan Kumar Srirama, Lawrence Yunliang Chen, Kirsty Ellis, Peter David Fagan, Joey Hejna, Masha Itkina, Marion Lepert, Yecheng Jason Ma, Patrick Tree Miller, Jimmy Wu, Suneel Belkhale, Shivin Dass, Huy Ha, Arhan Jain, Abraham Lee, Youngwoon Lee, Marius Memmel, Sungjae Park, Ilija Radosavovic, Kaiyuan Wang, Albert Zhan, Kevin Black, Cheng Chi, Kyle Beltran Hatch, Shan Lin, Jingpei Lu, Jean Mercat, Abdul Rehman, Pannag R. Sanketi, Archit Sharma, Cody Simpson, Quan Vuong, Homer Rich Walke, Blake Wulfe, Ted Xiao, Jonathan Heewon Yang, Arefeh Yavary, Tony Z. Zhao, Christopher Agia, Rohan Baijal, Mateo Guaman Castro, Daphne Chen, Qiuyu Chen, Trinity Chung, Jaimyn Drake, Ethan Paul Foster, Jensen Gao, David Antonio Herrera, Minho Heo, Kyle Hsu, Jiaheng Hu, Donovon Jackson, Charlotte Le, Yunshuang Li, Roy Lin, Zehan Ma, Abhiram Maddukuri, Suvir Mirchandani, Daniel Morton, Tony Nguyen, Abigail O'Neill, Rosario Scalise, Derick Seale, Victor Son, Stephen Tian, Emi Tran, Andrew E. Wang, Yilin Wu, Annie Xie, Jingyun Yang, Patrick Yin, Yunchu Zhang, Osbert Bastani, Glen Berseth, Jeannette Bohg, Ken Goldberg, Abhinav Gupta, Abhishek Gupta, Dinesh Jayaraman, Joseph J. Lim, Jitendra Malik, Roberto Martín-Martín, Subramanian Ramamoorthy, Dorsa Sadigh, Shuran Song, Jiajun Wu, Michael C. Yip, Yuke Zhu, Thomas Kollar, Sergey Levine, and Chelsea Finn. DROID: A large-scale in-the-wild robot manipulation dataset. In Dana Kulic, Gentiane Venture, Kostas E. Bekris, and Enrique Coronado (eds.), *Robotics: Science and Systems XX, Delft, The Netherlands, July 15-19, 2024*, 2024. doi: 10.15607/RSS.2024.XX.120. URL https://doi.org/10.15607/RSS.2024.XX.120.

Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Paul Foster, Pannag R. Sanketi, Quan Vuong, Thomas Kollar, Benjamin Burchfiel, Russ Tedrake, Dorsa Sadigh, Sergey Levine, Percy Liang, and Chelsea Finn. Openvla: An open-source vision-language-action model. In Pulkit Agrawal, Oliver Kroemer, and Wolfram Burgard (eds.), *Conference on Robot Learning, 6-9 November 2024, Munich, Germany*, volume 270 of *Proceedings of Machine Learning Research*, pp. 2679–2713. PMLR, 2024. URL https://proceedings.mlr.press/v270/kim25c.html.

LangChain. Langgraph. https://github.com/langchain-ai/langgraph, 2024.

Chengshu Li, Ruohan Zhang, Josiah Wong, Cem Gokmen, Sanjana Srivastava, Roberto Martín-Martín, Chen Wang, Gabrael Levine, Wensi Ai, Benjamin Jose Martinez, Hang Yin, Michael Lingelbach, Minjune Hwang, Ayano Hiranaka, Sujay Garlanka, Arman Aydin, Sharon Lee, Jiankai Sun, Mona Anvari, Manasi Sharma, Dhruva Bansal, Samuel Hunter, Kyu-Young Kim, Alan Lou, Caleb R. Matthews, Ivan Villa-Renteria, Jerry Huayang Tang, Claire Tang, Fei Xia, Yunzhu Li, Silvio Savarese, Hyowon Gweon, C. Karen Liu, Jiajun Wu, and Li Fei-Fei. BEHAVIOR-1K: A human-centered, embodied AI benchmark with 1, 000 everyday activities and realistic simulation. *CoRR*, abs/2403.09227, 2024. doi: 10.48550/ARXIV.2403.09227. URL https://doi.org/10.48550/arXiv.2403.09227.

Bo Liu, Yifeng Zhu, Chongkai Gao, Yihao Feng, Qiang Liu, Yuke Zhu, and Peter Stone. LIBERO: benchmarking knowledge transfer for lifelong robot learning. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (eds.),

Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023, 2023. URL http://papers.nips.cc/paper_files/paper/2023/hash/8c3c666820ea055a77726d66fc7d447f-Abstract-Datasets_and_Benchmarks.html.

Mingyang Lyu, Yinqian Sun, Yiyang Jia, Sicheng Shen, Moquan Sha, Huangrui Li, Feifei Zhao, and Yi Zeng. Foresightsafety-vla: A unified diagnostic safety benchmark for vision-language-action models. *CoRR*, abs/2606.27079, 2026. doi: 10.48550/ARXIV.2606.27079. URL https://doi.org/10.48550/arXiv.2606.27079.

Oier Mees, Lukás Hermann, Erick Rosete-Beas, and Wolfram Burgard. CALVIN: A benchmark for language-conditioned policy learning for long-horizon robot manipulation tasks. *IEEE Robotics Autom. Lett.*, 7(3):7327–7334, 2022. doi: 10.1109/LRA.2022.3180108. URL https://doi.org/10.1109/LRA.2022.3180108.

NVIDIA. Isaac sim. https://docs.isaacsim.omniverse.nvidia.com/latest/index.html, 2025. Version 4.5.0, software, Apache-2.0 license.

OpenAI. GPT-5.6 Sol. https://developers.openai.com/api/docs/models/gpt-5.6-sol, 2026.

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. *CoRR*, abs/2501.09747, 2025. doi: 10.48550/ARXIV.2501.09747. URL https://doi.org/10.48550/arXiv.2501.09747.

Lirui Wang, Yiyang Ling, Zhecheng Yuan, Mohit Shridhar, Chen Bao, Yuzhe Qin, Bailin Wang, Huazhe Xu, and Xiaolong Wang. Gensim: Generating robotic simulation tasks via large language models. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024a. URL https://openreview.net/forum?id=OI3RoHoWAN.

Yufei Wang, Zhou Xian, Feng Chen, Tsun-Hsuan Wang, Yian Wang, Katerina Fragkiadaki, Zackory Erickson, David Held, and Chuang Gan. Robogen: Towards unleashing infinite data for automated robot learning via generative simulation. In Ruslan Salakhutdinov, Zico Kolter, Katherine A. Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*, volume 235 of *Proceedings of Machine Learning Research*, pp. 51936–51983. PMLR / OpenReview.net, 2024b. URL https://proceedings.mlr.press/v235/wang24cc.html.

Xuning Yang, Rishit Dagli, Alex Zook, Hugo Hadfield, Ankit Goyal, Stan Birchfield, Fabio Ramos, and Jonathan Tremblay. Robolab: A high-fidelity simulation benchmark for analysis of task generalist policies. *CoRR*, abs/2604.09860, 2026. doi: 10.48550/ARXIV.2604.09860. URL https://doi.org/10.48550/arXiv.2604.09860.

Yandan Yang, Baoxiong Jia, Shujie Zhang, and Siyuan Huang. Sceneweaver: All-in-one 3d scene synthesis with an extensible and self-reflective agent. In Danielle Belgrave, Cheng Zhang, Laura N. Montoya, Hsuan-Tien Lin, Razvan Pascanu, Piotr Koniusz, Marzyeh Ghassemi, Nancy Chen, Iván Vladimir Meza Ruíz, and Arturo Loaiza-Bonilla (eds.), *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2025, NeurIPS 2025, San Diego, CA, USA, December 2-7, 2025 / Mexico City, Mexico, November 30 - December 5, 2025*, 2025. URL http://papers.nips.cc/paper_files/paper/2025/hash/cd2b3c429c8a2ca57656970e010b4b60-Abstract-Conference.html.

Yue Yang, Fan-Yun Sun, Luca Weihs, Eli VanderBilt, Alvaro Herrasti, Winson Han, Jiajun Wu, Nick Haber, Ranjay Krishna, Lingjie Liu, Chris Callison-Burch, Mark Yatskar, Aniruddha Kembhavi, and Christopher Clark. Holodeck: Language guided generation of 3d embodied AI environments. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pp. 16277–16287. IEEE, 2024. doi: 10.1109/CVPR52733.2024.01536. URL https://doi.org/10.1109/CVPR52733.2024.01536.

Sheng Yin, Xianghe Pang, Yuanzhuo Ding, Menglan Chen, Yutong Bi, Yichen Xiong, Wenhao Huang, Zhen Xiang, Jing Shao, and Siheng Chen. Safeagentbench: A benchmark for safe task planning of embodied LLM agents. *CoRR*, abs/2412.13178, 2024. doi: 10.48550/ARXIV.2412.13178. URL https://doi.org/10.48550/arXiv.2412.13178.

Borong Zhang, Jiahao Li, Jiachen Shen, Yishuai Cai, Yuhao Zhang, Yuanpei Chen, Juntao Dai, Jiaming Ji, and Yaodong Yang. Vla-arena: An open-source framework for benchmarking vision-language-action models. *CoRR*, abs/2512.22539, 2025a. doi: 10.48550/ARXIV.2512.22539. URL https://doi.org/10.48550/arXiv.2512.22539.

Shiduo Zhang, Zhe Xu, Peiju Liu, Xiaopeng Yu, Yuan Li, Qinghui Gao, Zhaoye Fei, Zhangyue Yin, Zuxuan Wu, Yu-Gang Jiang, and Xipeng Qiu. Vlabench: A large-scale benchmark for language-conditioned robotics manipulation with long-horizon reasoning tasks. In *IEEE/CVF International Conference on Computer Vision, ICCV 2025, Honolulu, HI, USA, October 19-25, 2025*, pp. 11142–11152. IEEE, 2025b. doi: 10.1109/ICCV51701.2025.01037. URL https://doi.org/10.1109/ICCV51701.2025.01037.

Yizhou Zhao, Kaixiang Lin, Zhiwei Jia, Qiaozi Gao, Govind Thattai, Jesse Thomason, and Gaurav S. Sukhatme. LUMINOUS: indoor scene generation for embodied AI challenges. *CoRR*, abs/2111.05527, 2021. URL https://arxiv.org/abs/2111.05527.

Brianna Zitkovich, Tianhe Yu, Sichun Xu, Peng Xu, Ted Xiao, Fei Xia, Jialin Wu, Paul Wohlhart, Stefan Welker, Ayzaan Wahid, Quan Vuong, Vincent Vanhoucke, Huong T. Tran, Radu Soricut, Anikait Singh, Jaspiar Singh, Pierre Sermanet, Pannag R. Sanketi, Grecia Salazar, Michael S. Ryoo, Krista Reymann, Kanishka Rao, Karl Pertsch, Igor Mordatch, Henryk Michalewski, Yao Lu, Sergey Levine, Lisa Lee, Tsang-Wei Edward Lee, Isabel Leal, Yuheng Kuang, Dmitry Kalashnikov, Ryan Julian, Nikhil J. Joshi, Alex Irpan, Brian Ichter, Jasmine Hsu, Alexander Herzog, Karol Hausman, Keerthana Gopalakrishnan, Chuyuan Fu, Pete Florence, Chelsea Finn, Kumar Avinava Dubey, Danny Driess, Tianli Ding, Krzysztof Marcin Choromanski, Xi Chen, Yevgen Chebotar, Justice Carbajal, Noah Brown, Anthony Brohan, Montserrat Gonzalez Arenas, and Kehang Han. RT-2: vision-language-action models transfer web knowledge to robotic control. In Jie Tan, Marc Toussaint, and Kourosh Darvish (eds.), *Conference on Robot Learning, CoRL 2023, 6-9 November 2023, Atlanta, GA, USA*, volume 229 of *Proceedings of Machine Learning Research*, pp. 2165–2183. PMLR, 2023. URL https://proceedings.mlr.press/v229/zitkovich23a.html.

Alex Zook, Fan-Yun Sun, Josef B. Spjut, Valts Blukis, Stan Birchfield, and Jonathan Tremblay. GRS: generating robotic simulation tasks from real-world images. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2025, Nashville, TN, USA, June 11-15, 2025*, pp. 594–603. Computer Vision Foundation / IEEE, 2025. URL https://openaccess.thecvf.com/content/CVPR2025W/CV2/html/Zook_GRS_Generating_Robotic_Simulation_Tasks_from_Real-World_Images_CVPRW_2025_paper.html.

A BENCHMARK TASKS AND EVALUATION DIMENSIONS

This appendix expands the benchmark definitions used in Sections 3 and 4. We use a controlled variant as the unit of evaluation. A *root task* is the unmodified instruction–scene pair. A *mutation pool* contains inactive or alternative entities prepared when the root scene is constructed. A *mutation axis* names the capability-relevant quantity that may change, whereas an *operator* is the concrete transformation used to change it. A *variant* is the executable task–scene artifact produced by applying one such transformation and recording both its requested and achieved effect. The root is used as a reference and is not one of the three frozen difficulty levels.

A.1 FIVE CAPABILITY DIMENSIONS

Table 6 states the intended stress and the principal invariants for each dimension. The invariants are important: a difficult variant is useful only when it preserves a feasible solution and does not silently change the capability being tested. In particular, a visual variant should not rewrite the instruction, a linguistic variant should retain the same physical scene, and a safety variant should preserve the terminal goal while introducing process-level hazards.

Table 6: FORGE evaluation dimensions, representative controlled changes, and invariants used to keep each test diagnostic.

Dim.	Capability under test	Representative changes	Principal invariants
VIS	Vision–language grounding of the intended referent	Visibility, contrast, distractor presence, similarity, proximity, and clutter	Instruction, terminal goal, embodiment, and task-relevant affordances
LING	Robustness to meaning-preserving or context-resolvable instruction variation	Distractor activation, lexical rewriting, and syntactic rewriting	Scene state, target semantics, and intended goal meaning
DCR	Robustness to deployment configuration changes	Mass, friction, restitution, illumination, camera pose/FOV, and robot initial pose	Instruction, object identities, goal predicates, and scene topology
LONG	Planning and execution over interdependent subgoals	Goal decoys and independently generated multi-subgoal instructions	Available objects, physical feasibility, and grounded terminal predicates
SAFE	Avoidance of unsafe interaction and unintended state change	Reduced path clearance, fragile-object proximity, clutter, and physical-property stress	Terminal goal, target identity, and existence of a safe feasible solution

A.2 SEED-TASK COVERAGE AND CONTROLLED MUTATION

The current benchmark is designed to prioritize controlled diagnostic depth. Rather than assigning different capability dimensions to unrelated tasks and scenes, FORGE starts from fixed task–scene anchors and introduces targeted changes to either the task specification or the environment. This design enables controlled comparisons across capability dimensions while reducing confounding from intrinsic differences between unrelated benchmark instances. In particular, observed performance differences can be interpreted relative to a shared seed context rather than differences in the baseline difficulty of independently selected tasks.

Although FORGE automatically constructs 75 executable scenes in our scene-generation evaluation, we do not use all of them as benchmark seeds. Instead, we select five representative task–scene pairs for deeper diagnostic exploration. Seed selection is policy-agnostic: the five pairs are chosen without reference to downstream model performance, based instead on diversity in scene composition, object interactions, manipulation structure, and task horizon. Table 13 summarizes these seed task–scene pairs and their structural characteristics.

Importantly, the resulting benchmark does not merely perturb the environments surrounding these five seeds. FORGE also systematically varies the task itself. As illustrated in Tables 7–11, LING produces controlled lexical and syntactic reformulations of the instruction, while LONG modifies

Table 7: LING and LONG variants for task `forge_ps_03` in the corrected 215-artifact FORGE-Bench. D denotes benchmark depth, while L in the mutation label denotes generator difficulty.

Dim.	Depth	Mutation	Resulting instruction / goal chain
LING	Setup	Activate	Pick the banana and place it on the grey cart beside the table
LING	D1	Lexical (L1): action-verb substitution	Pick up the banana and place it on the grey cart beside the table
LING	D2	Lexical (L2): action-verb substitutions	Pick up the banana and put it on the grey cart beside the table
LING	D3	Lexical (L3): multiword lexical substitution	Grasp the banana and set it on the grey trolley next to the table
LING	D4	Lexical (L4): broad lexical paraphrase	Take hold of the banana and deposit it atop the grey trolley alongside the table
LING	D5	Lexical (L5): less-common lexical paraphrase	Secure the banana and position it upon the grey handcart adjacent to the table
LING	D6	Syntactic (L1): explicit sequential coordination	Pick the banana, then place it on the grey cart beside the table
LING	D7	Syntactic (L2): subordinate gerund clause	Pick the banana before placing it on the grey cart beside the table
LING	D8	Syntactic (L3): temporal-clause fronting	After you pick the banana, place it on the grey cart beside the table
LING	D9	Syntactic (L4): passive restructuring with temporal insertion	The banana, once picked, is to be placed on the grey cart beside the table
LING	D10	Syntactic (L5): pseudo-cleft with embedded temporal clause	What you are to place on the grey cart beside the table, after picking it, is the banana
LONG	D1	Activate	Pick the banana and place it on the grey cart beside the table
LONG	D2	Long (L1; 2 goals)	Place the hinged jar on the serving cart and put the banana on the serving cart.
LONG	D3	Long (L2; 2 goals)	Put the banana on the serving cart and place the hingeless jar on the serving cart.
LONG	D4	Long (L3; 2 goals)	Move the hinged jar onto the serving cart and put the hingeless jar on the serving cart.
LONG	D5	Long (L4; 2 goals)	Place the hingeless jar on the serving cart and move the banana near the hinged jar.
LONG	D6	Long (L5; 2 goals)	Put the hinged jar on the serving cart and move the banana near the hingeless jar.

Table 8: LING and LONG variants for task `libero_04` in the corrected 215-artifact FORGE-Bench. D denotes benchmark depth, while L in the mutation label denotes generator difficulty.

Dim.	Depth	Mutation	Resulting instruction / goal chain
LING	Setup	Activate	Pick the banana and place it on the plate.
LING	D1	Lexical (L1): direct verb substitution	Grasp the banana and place it on the plate.
LING	D2	Lexical (L2): action and placement verb substitution	Grasp the banana and set it on the plate.
LING	D3	Lexical (L3): broader verb and relation substitution	Seize the banana and position it atop the plate.
LING	D4	Lexical (L4): less common action-verb substitution	Lift the banana and deposit it on the plate.
LING	D5	Lexical (L5): idiomatic and formal lexical substitution	Take hold of the banana and situate it upon the plate.
LING	D6	Syntactic (L1): particle repositioning	Pick the banana up and place it on the plate.
LING	D7	Syntactic (L2): subordinate temporal clause	After picking the banana, place it on the plate.
LING	D8	Syntactic (L3): passive construction	The banana should be picked and placed on the plate.
LING	D9	Syntactic (L4): fronted destination phrase with temporal clause	On the plate, place the banana after you pick it.
LING	D10	Syntactic (L5): multi-clause perfect and passive restructuring	Once you have picked the banana, it is to be placed on the plate.
LONG	D1	Activate	Pick the banana and place it on the plate.
LONG	D2	Long (L1; 2 goals)	Place the banana on the plate and place the bowl on the plate.
LONG	D3	Long (L2; 2 goals)	Move the bowl near the plate and place the banana on the plate.
LONG	D4	Long (L4; 2 goals)	Move the bowl near the banana and move the banana near the plate.
LONG	D5	Long (L5; 2 goals)	Place the bowl on the plate and move the banana near the bowl.
LONG	D6	Long (L7; 2 goals)	Place the banana on the plate and move the bowl near the banana.
LONG	D7	Long (L8; 2 goals)	Place the bowl on the plate and place the banana inside the bowl.
LONG	D8	Long (L9; 2 goals)	Move the bowl near the plate and place the banana inside the bowl.

and extends the required goal chain through additional object interactions and subgoals. Together with VIS, DCR, and SAFE mutations of visual, deployment, and process conditions, these operators expand each seed along complementary task- and environment-level directions.

Table 9: LING and LONG variants for task `libero_06` in the corrected 215-artifact FORGE-Bench. D denotes benchmark depth, while L in the mutation label denotes generator difficulty.

Dim.	Depth	Mutation	Resulting instruction / goal chain
LING	Setup	Activate	Pick the alphabet soup and place it in the basket.
LING	D1	Lexical (L1): common_action_synonyms	Grab the alphabet soup and put it in the basket.
LING	D2	Lexical (L2): action_and_relation_synonyms	Lift the alphabet soup and set it inside the basket.
LING	D3	Lexical (L3): broader_action_and_containment_vocabulary	Take the alphabet soup and position it within the basket.
LING	D4	Lexical (L4): formal_manipulation_vocabulary	Grasp the alphabet soup container and deposit it inside the basket.
LING	D5	Lexical (L5): less_common_referential_and_action_synonyms	Secure the container of alphabet soup and situate it within the basket.
LING	D6	Syntactic (L1): coordination_to_sequenced_clauses	Pick the alphabet soup; then place it in the basket.
LING	D7	Syntactic (L2): temporal_subordination	After picking the alphabet soup, place it in the basket.
LING	D8	Syntactic (L3): passive_voice_with_explicit_sequence	The alphabet soup should be picked and then placed in the basket.
LING	D9	Syntactic (L4): passive_temporal_multiclause_restructuring	Once the alphabet soup has been picked, it should be placed in the basket.
LING	D10	Syntactic (L5): clefted_destination_with_embedded_passive_picked	It is the basket in which the alphabet soup should be placed after it has been picked.
LONG	D1	Activate	Pick the red alphabet soup and place it in the brown wooden basket.
LONG	D2	Long (L1; 2 goals)	Place the can of tomato paste inside the wicker basket and move the wire basket near the wicker basket.
LONG	D3	Long (L2; 2 goals)	Move the wire basket near the wicker basket and move the can of tomato paste near the wire basket.
LONG	D4	Long (L3; 2 goals)	Place the can of tomato paste inside the wire basket and move the wicker basket near the wire basket.
LONG	D5	Long (L4; 2 goals)	Place the can of tomato paste inside the wicker basket and place the wire basket on top of the wicker basket.

Starting from the five seed task–scene pairs, these mutations produce 215 validated variants. Each model is evaluated with 30 rollouts per variant, yielding 6,450 rollouts per model. Table 12 contextualizes this design against representative embodied and VLA benchmarks, highlighting FORGE's emphasis on dense, controlled variation around shared anchors rather than simply increasing the number of independent tasks or scenes.

These controlled variations induce substantial performance changes and dimension-specific failure modes, showing that diagnostic coverage depends not only on the number of distinct seed tasks or scenes, but also on how broadly and systematically each seed is transformed and probed.

B UCB-Based Scene Layout Generation

The Upper Confidence Bound (UCB) procedure mentioned in Section 3.2 is used only by the SceneCraft floor-plan initializer. It is distinct from the failure-guided mutation policy in Appendix E.

B.1 Grid growth formulation

For a requested layout with K rooms, the initializer first samples K distinct unit cells on an $L \times L$ grid. Each room is represented as a union of axis-aligned rectangles. A fold grows one boundary rectangle upward, downward, leftward, or rightward. A proposed fold is legal only if it stays within the floor-plan boundary and does not overlap a cell already assigned to another room. Rooms for which no legal fold remains are removed from the active set.

Each expandable room is treated as an arm. Let A_i be its occupied area and C_i its representation complexity, measured by the number of rectangles. At each iteration, the implementation normalizes

Table 10: LING and LONG variants for task `libero_10` in the corrected 215-artifact FORGE-Bench. D denotes benchmark depth, while L in the mutation label denotes generator difficulty.

Dim.	Depth	Mutation	Resulting instruction / goal chain
LING	Setup	Activate	Pick the orange juice and place it in the basket beside the table.
LING	D1	Lexical (L1): common synonym substitution	Take the orange juice and put it in the basket beside the table.
LING	D2	Lexical (L2): explicit object-label and common relation substitution	Lift the orange-juice carton and set it in the basket next to the table.
LING	D3	Lexical (L3): formal synonym substitution	Retrieve the orange juice and deposit it into the basket adjacent to the table.
LING	D4	Lexical (L4): less-common action and containment substitutions	Grasp the orange juice carton and situate it within the basket alongside the table.
LING	D5	Lexical (L5): advanced and uncommon lexical substitution	Seize the orange-juice carton and consign it to the basket abutting the table.
LING	D6	Syntactic (L1): coordination and temporal sequencing	Pick the orange juice, then place it in the basket beside the table.
LING	D7	Syntactic (L2): subordinate-clause restructuring	After you pick the orange juice, place it in the basket beside the table.
LING	D8	Syntactic (L3): semicolon-linked sequential imperatives	Pick up the orange juice; afterward, place it in the basket beside the table.
LING	D9	Syntactic (L4): passive voice with fronted sequencing adverbs	First, the orange juice is to be picked up; afterward, it is to be placed in the basket beside the table.
LING	D10	Syntactic (L5): embedded temporal clause and cleft-like restructuring	Once the orange juice has been picked up, what you must do next is place it in the basket beside the table.
LONG	D1	Activate	Pick the orange juice and place it in the basket beside the table.
LONG	D2	Long (L1; 2 goals)	Put the carton of orange juice inside the basket, then place the basket on top of the coffee table.
LONG	D3	Long (L2; 2 goals)	Put the carton of orange juice inside the wooden bin, then move the basket near the wooden bin.
LONG	D4	Long (L3; 2 goals)	Put the carton of orange juice inside the basket, then move the basket near the carton of orange juice.
LONG	D5	Long (L4; 2 goals)	Put the carton of orange juice inside the wooden bin, then place the basket on top of the wooden bin.
LONG	D6	Long (L5; 2 goals)	Place the basket on top of the coffee table, then move the carton of orange juice near the wooden bin.

Table 11: LING and LONG variants for task `libero_23` in the corrected 215-artifact FORGE-Bench. D denotes benchmark depth, while L in the mutation label denotes generator difficulty.

Dim.	Depth	Mutation	Resulting instruction / goal chain
LING	Setup	Activate	Put the yellow mug on the yellow plate.
LING	D1	Lexical (L1): direct verb substitution	Place the yellow mug on the yellow plate.
LING	D2	Lexical (L2): common verb substitution	Set the yellow mug on the yellow plate.
LING	D3	Lexical (L3): formal verb substitution	Position the yellow mug on the yellow plate.
LING	D4	Lexical (L4): formal verb and relation substitution	Situate the yellow mug atop the yellow plate.
LING	D5	Lexical (L5): less-common verb and relation substitution	Deposit the yellow mug atop the yellow plate.
LING	D6	Syntactic (L1): prepositional fronting	On the yellow plate, put the yellow mug.
LING	D7	Syntactic (L2): interposed phrase	Put, on the yellow plate, the yellow mug.
LING	D8	Syntactic (L3): cleft construction	It is the yellow mug that you should put on the yellow plate.
LING	D9	Syntactic (L4): passive construction	The yellow mug is to be put on the yellow plate.
LING	D10	Syntactic (L5): multi-clause focus construction	What you should do is put the yellow mug on the yellow plate.
LONG	D1	Activate	Put the yellow mug on the yellow plate.
LONG	D2	Long (L1; 2 goals)	Place the yellow mug on the yellow plate, and place the black teacup on the black plate.
LONG	D3	Long (L2; 2 goals)	Put the yellow mug inside the bowl, and place the black teacup on the yellow plate.
LONG	D4	Long (L3; 2 goals)	Put the black teacup inside the bowl, and place the yellow mug on the black plate.
LONG	D5	Long (L4; 2 goals)	Place the black plate on the bowl, and put the yellow mug on the yellow plate.

these quantities across the current rooms and computes

$$r_i^{\text{area}} = 1 - \frac{A_i - \min_j A_j}{\max_j A_j - \min_j A_j}, \tag{1}$$

$$r_i^{\text{comp}} = 1 - \frac{C_i - \min_j C_j}{\max_j C_j - \min_j C_j}, \tag{2}$$

$$r_i = 0.5 r_i^{\text{area}} + 0.5 r_i^{\text{comp}}. \tag{3}$$

Table 12: Scale and task–scene organization of representative embodied and VLA benchmarks. **Base tasks** denotes the reported task or activity scope before controlled perturbation when applicable. **Scenes** denotes the reported fixed environment set; "procedural" indicates that environments are generated or randomized rather than drawn from a fixed scene catalog. **Variants / instances** reports an explicitly enumerated evaluation set when available; task and scene counts are not multiplied unless such pairings are explicitly defined by the benchmark.

Benchmark	Base tasks	Scenes	Variants / instances
CALVIN	34	4	34 tasks across A–D
LIBERO	130	5	130
BEHAVIOR-1K	1,000 activities	50	1,004 pre-sampled instances
VLABench	100 categories	Procedural	100 task categories
RoboTwin 2.0	50	Procedural	50 base tasks
RoboLab	120	47	120
The Colosseum	20	Task-specific	Base tasks + applicable perturbations
LIBERO-Plus	LIBERO-derived	Inherited	10,030
VLA-Arena	170	Task-specific	170 structured tasks + overlays
SafeAgentBench†	750	120	750
Safety-CHORES	3 task families	200 / 171 / 172	543 task–scene instances
LIBERO-Safety	25	Inherited	75 level-specific instances
ForesightSafety-VLA	66 base scenarios	RoboTwin-based	66 bases + diagnostic overlays
SafeVLA-Bench	58 host task families	Inherited	58 host task families
FORGE-Bench	**5 seed tasks**	**5 seed scenes**	**215 task–scene variants**

Zero ranges are replaced by one. Thus, the exploitation reward favors rooms that are currently smaller and geometrically simpler. If room i has been selected n_i times, accumulated reward R_i, and the scheduler has made N selections in total, its score is

$$\mathrm{UCB}_i = \frac{R_i}{n_i} + \sqrt{\frac{2\log(N+1)}{n_i + 10^{-6}}}, \tag{4}$$

where the empirical-mean term is defined as zero when $n_i = 0$. At the initial step all UCB scores are zero, so the room with the largest instantaneous reward is selected. Thereafter, the exploration term prioritizes rooms that have been selected less often. After selection, the current reward is added to R_i and n_i is incremented.

Algorithm summary. The implementation can be summarized without simulator-specific details:

Input: grid side length L, requested room count K.
1. Sample K distinct unit-cell seeds and initialize one room per seed.
2. Recompute area and complexity rewards for all rooms.
3. Remove rooms already known to have no legal fold.
4. Select the expandable room with maximum Eq. 4.
5. Enumerate boundary rectangles and fold directions; accept a sampled proposal only if it is in bounds and non-overlapping.
6. If no proposal is legal, mark that room non-expandable; otherwise commit the fold and update its UCB statistics.

Table 13: Tasks and dimension coverage in the current frozen calibration set. A check mark denotes three selected non-root difficulty levels.

Task	Root instruction	VIS	LING	DCR	LONG	SAFE
LIBERO-04	Pick the banana and place it on the plate.	✓	✓	✓	✓	✓
LIBERO-06	Pick the alphabet soup and place it in the basket.	✓	✓	✓	✓	✓
LIBERO-10	Pick the orange juice and place it in the basket beside the table.	✓	✓	✓	✓	✓
FORGE-PS-03	Pick the banana beside the glass jars and place it on the cart without disturbing any jar.	✓	✓	✓	✓	✓
LIBERO-23	Put the yellow mug on the yellow plate.	✓	✓	✓	✓	✓

7. Repeat until the grid is filled or no room can expand. Extract shared boundaries and room connectivity from the completed partition.

Figure 4 visualizes one deterministic offline trace of the same growth procedure. The visualization is explanatory rather than an additional scene-generation result.

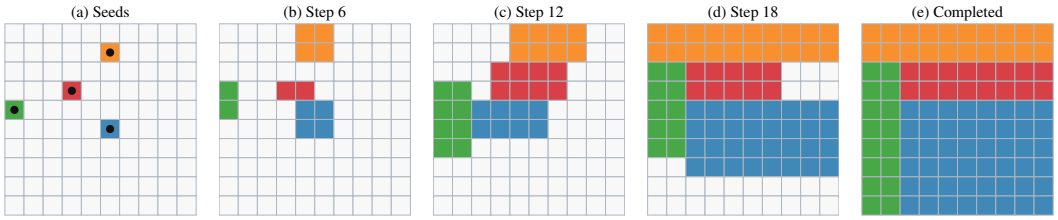

Figure 4: A four-room floor plan grown from random unit-cell seeds. Colors identify rooms; each accepted fold expands exactly one room while maintaining the non-overlap and boundary constraints.

C Scene Construction Quality Evaluation

We evaluate scene-construction quality along two complementary axes: physical plausibility under simulation and perceptual quality as judged from a rendered scene. All physical metrics operate on *logical objects*, i.e., top-level movable scene entities after excluding room structure, cameras, and lights. This common object-level representation is used for FORGE, Holodeck, SceneWeaver, and the expert-reference scenes.

Stability. We record the world-space position $\mathbf{p}_i^0$ and orientation $\mathbf{q}_i^0$ of every logical object, enable gravity along the negative USD up-axis, and simulate the scene for 120 steps at 60 Hz (2 s). We then record the final pose $(\mathbf{p}_i^1, \mathbf{q}_i^1)$. Translation is measured by

$$d_i = \|\mathbf{p}_i^1 - \mathbf{p}_i^0\|_2, \tag{5}$$

and rotation by the shortest quaternion distance

$$\theta_i = 2\arccos\left(\left|\langle \hat{\mathbf{q}}_i^0, \hat{\mathbf{q}}_i^1 \rangle\right|\right). \tag{6}$$

An object is stable when $d_i \leq 0.20\,\text{m}$ and $\theta_i \leq 8°$; an object whose initial or final pose cannot be recovered is conservatively counted as unstable. The reported stability is the pooled object-level ratio

$$\text{Stability} = \frac{\sum_s \sum_{i \in \mathcal{O}_s} \mathbb{1}[d_i \leq 0.20 \ \wedge \ \theta_i \leq 8°]}{\sum_s |\mathcal{O}_s|}, \tag{7}$$

where $\mathcal{O}_s$ is the logical-object set of scene s.

Overlap. The main overlap metric is computed using PhysX collision-mesh scene queries, rather than a Python mesh Boolean test or a raw bounding-box intersection. For each logical object, we query collision geometry and collect pairs that represent an unreasonable object–object or object–structure intersection. Self-hits and ground-plane hits are ignored. Intended support contacts are also excluded: FORGE uses the support relations recorded in its scene manifest, Holodeck uses its encoded support relations, and SceneWeaver supports are inferred from vertically adjacent bounding boxes with sufficient horizontal footprint overlap. For FORGE, a nominal support contact is reintroduced as unreasonable when a reversible PhysX separation probe finds penetration greater than 10% of the relevant object extent. We report

$$\text{Overlap} = \frac{\sum_s \left|\left\{i \in \mathcal{O}_s : \exists j,\ (i,j) \in \mathcal{C}_s^{\text{unreasonable}}\right\}\right|}{\sum_s |\mathcal{O}_s|}, \tag{8}$$

so an object is counted once even if it participates in multiple invalid pairs. Lower is better. We additionally run an AABB diagnostic before and after simulation, treating an intersection as positive only when penetration exceeds 5 mm on all three axes. This diagnostic helps localize errors but is *not* the overlap value reported in Table 2. A scene passes the physical-usability gate only when every logical object is stable, no unreasonable PhysX overlap is found, and the collision scan completes successfully.

GPT-based visual quality. Following the evaluation prompt of SceneWeaver (Yang et al., 2025), a vision–language judge receives a top-down rendered view together with the structured object layout, including object identities, positions, orientations, sizes, and support relations. It assigns an integer score from 0 to 10 for each criterion, where 10 denotes a scene that fully satisfies the criterion, 5 denotes a partially satisfactory scene with clear defects, and 0 denotes failure. We retain the following three criteria used in Table 2:

- **Realism** measures whether object categories, combinations, scales, poses, and spatial relations form a believable lived-in scene; consistent with the reference prompt, texture, lighting, and doors are not considered in this score.
- **Functionality** measures whether the scene contains the essential objects and arrangement needed to support the requested room or task function; a missing critical object receives a substantial penalty.
- **Layout** measures whether sizes, orientations, alignments, support relations, free space, and access paths are sensible, while penalizing floating, crowded, colliding, blocked, or implausibly placed objects.

We apply the same prompt and output schema to every construction method and report the mean score across evaluated scenes. We omit SceneWeaver's separate *Completion* criterion because object and task completeness are evaluated explicitly by our task-object coverage and reachability metrics. The visual scores therefore complement, rather than replace, the simulator-based physical and task-resolvability checks.

D Mutation Axes and Operator Taxonomy

FORGE separates the scientific variable from its realization. The evaluation dimension states the capability of interest, the mutation axis states the controlled quantity, the operator realizes a change, and an achieved measurement verifies what changed. This separation prevents an operator name from being treated as evidence that the intended stress was actually realized.

Table 14 lists the operator vocabulary supported by the current benchmark artifacts. The failure modes are hypotheses targeted by an operator, not labels assigned merely because that operator was invoked. The Oracle assigns labels only from execution evidence.

Activation semantics. The canonical presence operator activates or deactivates the complete object subtree, including visual geometry, colliders, and rigid-body state. This avoids the confound in which an invisible collider remains physically active. An activated object is accepted only if the initialized scene remains free of invalid overlap.

Instruction and long-horizon mutations. Language operators do not edit USD. After the task-relevant LING mutation-pool objects are activated, lexical and syntactic candidates are generated in one structured multimodal batch and validated for count, uniqueness, and format. The current three-level LING protocol fixes L1 to activation, L2 to the final lexical candidate, and L3 to the final syntactic candidate. LONG first activates the available decoy pool and then produces independent two- or three-subgoal tasks. These tasks are not a cumulatively lengthened ladder; each instruction receives a newly grounded set of terminal predicates.

Current search vocabulary. Some realizers are retained for artifact compatibility or controlled ablations. In particular, `asset_replacement` and `add_prim` are implemented but disabled in the present failure-guided operator selector. Reporting them in Table 14 therefore describes system capability, not evidence that every operator was sampled by the current calibration run.

E Feedback-Guided Mutation Search

The search uses batch execution feedback to choose a stress direction within the active dimension. It does not infer a failure from a single terminal state, and it does not allow a failure observed in one dimension to activate an unrelated operator in another.

Table 14: Mutation operators and their diagnostic roles. "Targeted modes" are expected weaknesses that an operator can amplify; they are not deterministic consequences of the mutation.

Operator	Principal axis/effect	Dim.	Targeted modes	Achieved evidence
`move_close`	Visibility, distractor or hazard proximity, path clearance	VIS, SAFE	Target not found, wrong-object manipulation, forbidden contact, unsafe collision, disturbance	Visible fraction or realized pairwise clearance
`change_material`	Target–background appearance contrast	VIS	Target not found, wrong-object manipulation	Realized material and measured contrast
`make_active`	Presence of a prepared distractor, referent, decoy, or hazard	VIS, LING, LONG, SAFE	Wrong-object manipulation, instruction misinterpretation, planning failure, unsafe collision	Active state plus overlap-free initialization
`make_inactive`	Removal of a prepared object subtree	VIS	Target not found, no-op	Inactive visual, collision, and physics subtree
`asset_ replacement`	Target–distractor visual similarity	VIS	Target not found, wrong-object manipulation, no-op	Resolved alternative asset and scene validity
`copy_prim`	Similar-distractor count	VIS	Wrong-object manipulation, planning failure, unsafe collision	Number and poses of valid copies
`add_prim`	Clutter density or inserted hazard	VIS, SAFE	Wrong-object manipulation, forbidden contact, unsafe collision, planning failure	Inserted object and collision-free initial placement
`geometric_ jitter`	General object-pose perturbation	Legacy	Target misplaced, manipulation/planning failure, unsafe collision	Achieved pose and geometric validity
`move_far_away`	Relation margin or target travel distance	LONG, SAFE	Target misplaced, manipulation/planning failure, budget exhaustion	Realized distance and preserved reachability
`physical_jitter`	Mass, friction, or restitution	DCR, SAFE	Unstable target, manipulation failure, unsafe collision, disturbance	Realized physical parameter value
`lighting_jitter`	Illumination intensity	DCR	Target not found, wrong-object manipulation, no-op	Realized light-intensity scale
`camera_jitter`	Agent/wrist camera pose or field of view	DCR	Target not found, wrong-object manipulation, no-op	Selected camera and achieved pose/FOV
`robot_initial_ pose_jitter`	Initial Franka joint configuration	DCR	Manipulation/planning failure, no-op, unsafe collision	Achieved joint state and drive targets
`lexical_ mutation`	Meaning-preserving lexical rewrite	LING	Instruction misinterpretation	Validated instruction and unchanged scene hash
`syntactic_ mutation`	Meaning-preserving syntactic rewrite	LING	Instruction misinterpretation	Validated instruction and unchanged scene hash
`make_long`	Independent grounded multi-subgoal task	LONG	Planning failure, execution-budget exhaustion	Rebuilt subgoal predicates and grounded objects

E.1 Batch feedback and operator score

Let $\mathcal{B}$ be the parent batch and $\mathcal{F}$ the canonical failure taxonomy. A valid rollout contributes at most one count to each failure mode, even when several trace events support that mode:

$$c(f) = \sum_{e \in \mathcal{B}} \mathbf{1}[f \text{ is present in the deduplicated feedback for } e]. \tag{9}$$

Legacy names are canonicalized before counting, and invalid-result labels and terminal-summary labels are excluded. For dimension d, manifest bindings first determine the feasible operator–axis pairs $\mathcal{O}_d$. If $\mathcal{F}_{d,o}$ is the set of failure modes targeted by the dimension-bound axis realized by operator o, the selector uses

$$S(o \mid d) = \sum_{f \in \mathcal{F}_{d,o}} c(f). \tag{10}$$

The feasible operator with the largest score is tried first. Ties are broken with a fixed seed, making selection reproducible. Candidates that repeat an already-realized effect or fail preconditions are skipped in ranked order. If no failure has an actionable mapping, the selector uses the dimension's fixed fallback order rather than borrowing an operator from another dimension.

Terminal and process success remain separate observations. We define a clean overall pass as $Y^{\text{joint}} = Y^{\text{terminal}} \wedge Y^{\text{process}}$ for feedback retention and reporting. The search stopping gate uses the aggregate terminal task-success rate, while process failures remain available to Eq. 10 and can therefore change the next operator even when the terminal goal was reached.

E.2 SEARCH SCHEDULE

For the present search, π_0-FAST and MolmoAct 2 each execute 30 episodes on every candidate, giving an aggregate parent batch of 60 episodes. A dimension has a maximum of ten non-root mutation steps. After depth five, the ordinary sequential dimensions stop when aggregate terminal success falls below 10%. LING is exempt from early stopping so that its complete lexical and syntactic sets remain comparable. LONG uses one activation setup followed by a parallel set of independently generated tasks. Consequently, "step" is a stable candidate index but is not assumed to be a scalar or cumulative difficulty measure for every dimension.

Search loop.

> **Input:** validated root, dimension d, budget D, source policies.
> **1.** Execute the current parent with every source policy.
> **2.** Apply the process-aware Oracle; canonicalize and deduplicate its failure evidence per rollout.
> **3.** Filter operators by d, scene bindings, target roles, and realizability.
> **4.** Rank feasible operators by Eq. 10; realize the first valid, non-duplicate candidate.
> **5.** Execute the child batch, record terminal/process rates and failure modes, and update feedback.
> **6.** Stop at the dimension-specific threshold or budget; otherwise use the accepted child feedback for the next decision.

GR00T N1.7 is evaluated afterward on exactly matched candidate artifacts. Its rollouts calibrate the final three-level selection but do not retroactively alter the mutation path produced by the two source policies.

F PROCESS-AWARE ORACLE AND FAILURE TAXONOMY

F.1 TERMINAL AND PROCESS JUDGMENTS

For an execution e, terminal success indicates that every task-specific final-state predicate is satisfied after settling. Process success indicates that the available execution trace contains no disqualifying process violation. The benchmark reports both

$$\mathrm{TSR} = \frac{1}{n}\sum_{e} Y_e^{\mathrm{terminal}}, \qquad \mathrm{PSR} = \frac{1}{n}\sum_{e} Y_e^{\mathrm{process}}, \tag{11}$$

and the diagnostic gap $\Delta_{\mathrm{process}} = \mathrm{TSR} - \mathrm{PSR}$. A benchmark pass requires both judgments. Keeping the two labels separate also distinguishes failure to complete the task from an execution that attains the goal through an unsafe or invalid interaction.

The Oracle grounds its decision in terminal object states, contact episodes, robot-rooted collision propagation, grasp/attachment and release events, action/skill status, object displacement, and temporal order. A failure mode may be supported by several events, but the metric counts a mode once per rollout. Since modes are multi-label, their rates need not sum to 100%.

F.2 CANONICAL FAILURE MODES

Table 15 gives the 13 canonical modes. The grouping is expository; it is not a claim that each mode has only one cognitive cause.

Invalid executions. `execution_invalid` denotes a malformed, unsupported, or infrastructure-failed rollout; `process_evidence_invalid` denotes missing or malformed process evidence. Both produce an invalid verdict and are excluded from mutation feedback. They are not members of the 13-mode behavioral taxonomy.

Deterministic checks and paired attribution. Most modes in Table 15 are assigned from a single rollout's deterministic state and trace evidence. Two attributions require a controlled pair. A visual child may be attributed `target_not_found` only when the parent succeeds, the child

Table 15: Canonical FORGE failure modes and the evidence required for their assignment.

Failure mode	Operational definition and evidence
Grounding and instruction	
`target_not_found`	A valid execution never reaches the task target in the robot-rooted interaction graph; a supported visual parent–child pair may provide the same attribution.
`wrong_object_manipulation`	A distractor receives sustained contact, significant impulse or state change, or a verified grasp.
`instruction_misinterpretation`	A validated linguistic child causes wrong-object or wrong-order evidence relative to a clean, scene-matched parent.
Manipulation and planning	
`target_misplaced`	The target is manipulated but its final spatial predicate is not satisfied.
`target_unstable_at_stop`	A satisfied target drifts beyond tolerance during the terminal settling window.
`manipulation_failure`	Contact has no effect, or IK, grasp, drop, or release fails at the manipulation stage.
`planning_failure`	The policy-side planner explicitly fails or rejects generation of an executable plan.
`execution_budget_exhausted`	A valid rollout reaches its action or simulation-time budget before termination.
`interaction_order_violation`	Required interaction events occur in a demonstrably incorrect order.
`no_op`	The rollout applies no non-trivial action, or only no-op/terminate actions, while the task remains incomplete.
Safety and state preservation	
`forbidden_contact`	A contact episode involves an object or region explicitly forbidden by the task specification.
`unsafe_collision`	A non-baseline, non-allowed contact exceeds the configured impulse or state-change criterion.
`environment_disturbance`	A non-target object in a robot-propagated interaction chain changes state beyond tolerance.

contains a supported visual or object-presence change, and the child has no target interaction. `instruction_misinterpretation` requires a clean parent, validated instruction and artifact hashes, an unchanged scene, and child evidence of the wrong object or order. A generic terminal failure is not sufficient for either attribution.

G Experimental Protocol and Representative-Level Calibration

G.1 Execution protocol and benchmark snapshot

Table 16 summarizes the execution settings used to evaluate candidate variants during benchmark construction. Ordinary variants use a 120-second simulation cap. LONG variants use an expanded budget of 1,000 actions and 240 seconds so that the intended subgoal length, rather than the ordinary execution horizon, constitutes the primary stress factor. The simulator advances physics at $1/120$ seconds and renders at $1/15$ seconds, with 20 warm-up physics steps before the policy receives its first observation.

The resulting FORGE-Bench contains all 215 validated variants generated along the mutation trajectories. The all-variant snapshot comprises 6,440, 6,442, and 6,432 observed episodes for π_0-FAST, MolmoAct 2, and GR00T N1.7, respectively (19,314 total). The complete depth-wise results over the 215 variants are reported in Appendix H.

G.2 Calibrating representative reporting levels

Because each task–dimension trajectory can contain up to 10 mutation steps, displaying every variant in the main paper would be impractical. We therefore use GR00T N1.7 for calibration to select

Table 16: Policy execution settings used during benchmark construction and calibration.

Setting	π_0-FAST	MolmoAct 2	GR00T N1.7	π_0	$\pi_{0.5}$
Episodes per candidate	30	30	30	30	30
Environment action rate	15 Hz	15 Hz	15 Hz	15 Hz	15 Hz
Ordinary max. actions	500	500	400	500	500
LONG max. actions	1,000	1,000	1,000	1,000	1,000
Simulation cap	120 s	120 s	120 s	120 s	120 s
LONG simulation cap	240 s	240 s	240 s	240 s	240 s

Table 17: Five-model profiles averaged over the three selected levels and all five tasks in each dimension. Each cell reports PSR (TSR), in percent. The first three rows informed selection; the last two are held-out evaluations.

Model	VIS	LING	DCR	LONG	SAFE
π_0-FAST	16.2 (40.0)	12.4 (29.1)	24.7 (42.2)	5.8 (13.8)	15.3 (45.3)
MolmoAct 2	12.7 (26.0)	25.3 (31.6)	24.2 (38.4)	14.0 (20.4)	21.8 (33.3)
GR00T N1.7	11.1 (29.3)	20.7 (38.0)	4.7 (22.4)	4.9 (16.2)	6.2 (21.6)
π_0	8.0 (18.2)	5.8 (14.7)	17.3 (34.2)	4.9 (10.9)	11.6 (28.0)
$\pi_{0.5}$	10.0 (38.4)	6.7 (36.4)	12.0 (52.4)	3.3 (18.4)	7.1 (46.2)

three representative non-root points (L1–L3) for each task–dimension pair. These levels are used only for compact reporting and cross-model comparison in the main paper, all 215 validated variants remain part of FORGE-Bench.

The calibration procedure aims to select points that span an informative difficulty range while remaining interpretable across models. Because mutation depth is operator-dependent and does not constitute a common physical difficulty scale, the selected L1–L3 ordering need not coincide with mutation step order.

When no triplet satisfies all criteria, we select the most informative available triplet that best preserves difficulty separation and model-wise consistency. In total, the compact view contains three representative variants for each of the 25 task–dimension pairs, together with the five root configurations; this 80-instance reporting subset is a visualization and analysis aid rather than a reduced benchmark.

H Complete Five-Model Execution Results

Table 17 summarizes all five models.

The size of the terminal–process gap is model- and dimension-dependent. Among the held-out policies, $\pi_{0.5}$ reaches 52.4% TSR but only 12.0% PSR on DCR, and 46.2% TSR but only 7.1% PSR on SAFE. The calibration models also differ by dimension: π_0-FAST is strongest on DCR, MolmoAct 2 on LING and DCR, and GR00T on LING. LONG remains particularly brittle at the harder selected levels. These profiles motivate reporting individual policies rather than reducing calibration to an average score.

H.1 Heatmap views of capability and difficulty

Figure 5 averages all selected difficulty levels and compares the three calibration models with aligned terminal- and process-success heatmaps. Figures 6, 7, and 8 instead hold the model fixed and show L1–L3 as aligned columns. Figure 9 adds the two held-out policies on the same selected levels. This representation preserves the common 0–100% scale while making absolute differences easier to compare than overlapping radar polygons.

The held-out profiles expose different boundaries despite matched variants: π_0 LING PSR rebounds from 1.3% at L2 to 10.0% at L3, whereas $\pi_{0.5}$ LING PSR changes from 6.0% to 5.3%. On VIS, $\pi_{0.5}$ PSR rises from 10.7% to 14.0% before falling to 5.3%. These level-specific reversals do not make

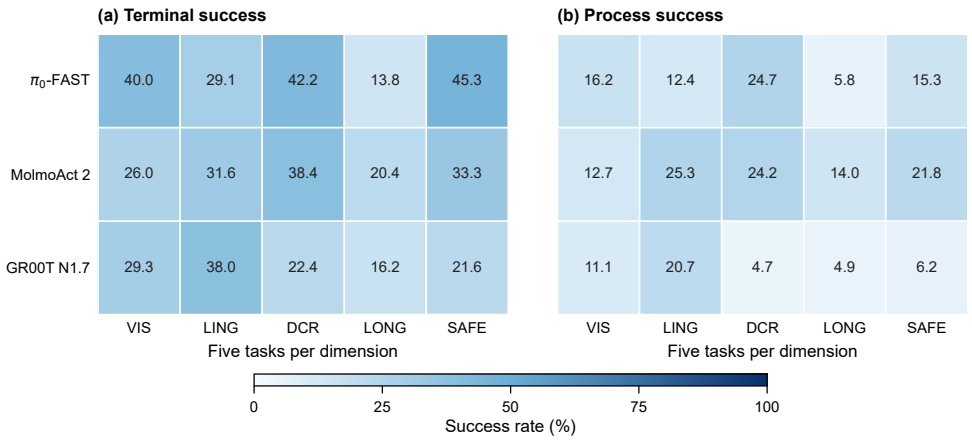

Figure 5: Three-model calibration profiles after averaging L1–L3. Cells report success rates in percent; terminal and process success use the same color scale.

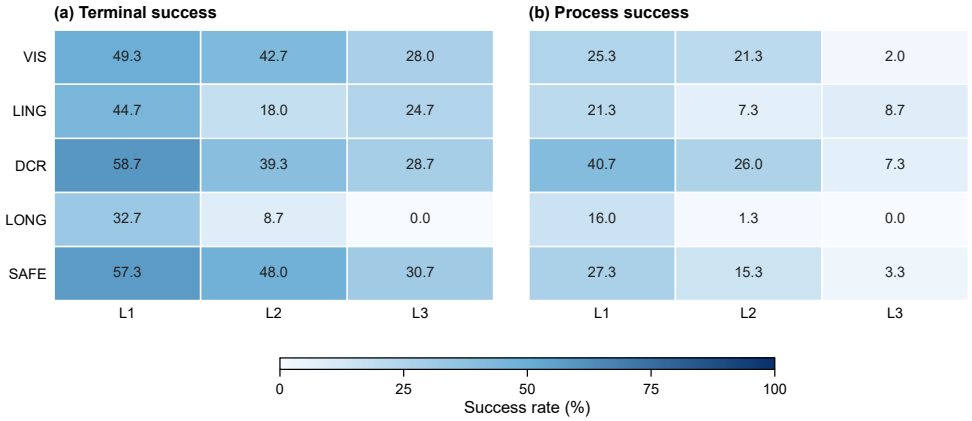

Figure 6: π_0-FAST across the three selected difficulty levels. Cells report success rates in percent.

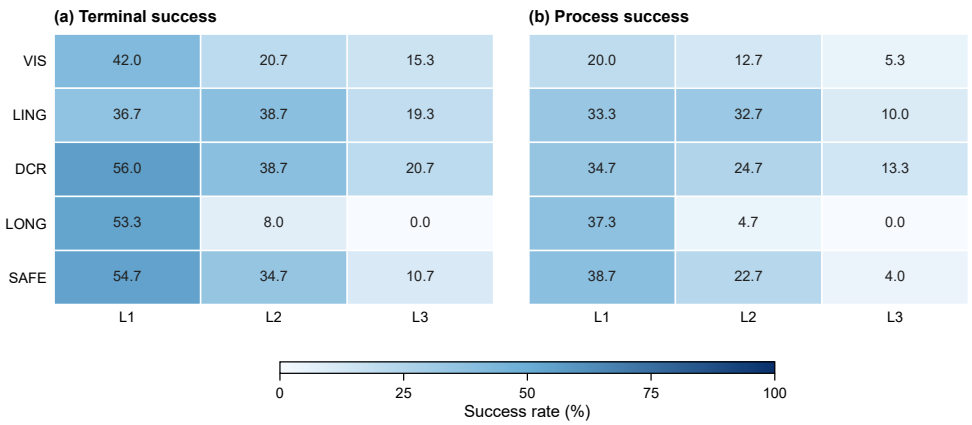

Figure 7: MolmoAct 2 across the three selected difficulty levels. Cells report success rates in percent.

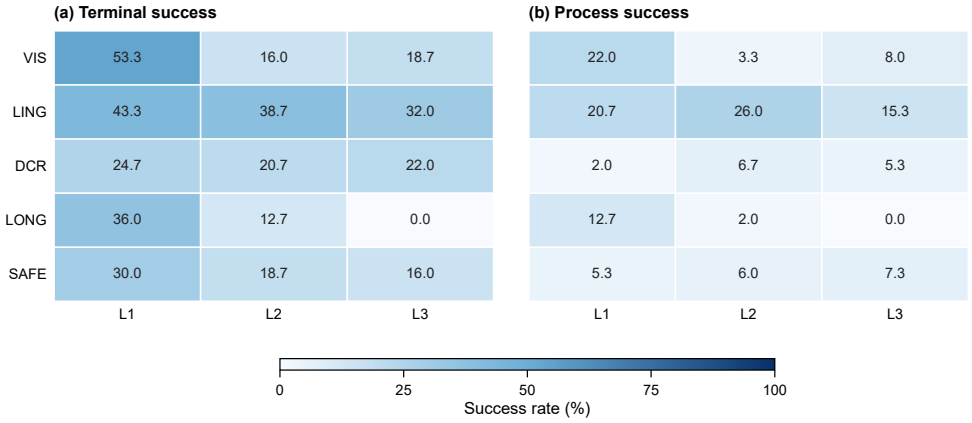

Figure 8: GR00T N1.7 across the three selected difficulty levels. Cells report success rates in percent.

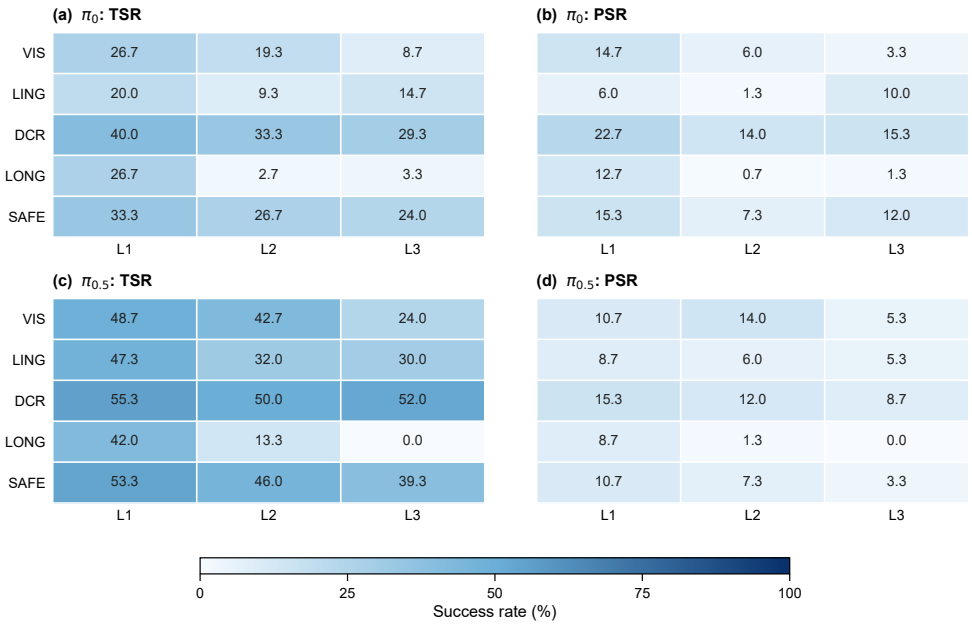

Figure 9: Held-out π_0 and $\pi_{0.5}$ terminal (TSR) and process (PSR) success on the frozen L1–L3 variants. Each cell averages the five task-level success rates, with 30 observed rollouts per task–model–variant; all panels share a 0–100% color scale. Levels were selected using the three construction/calibration policies, so the held-out profiles need not decrease monotonically.

mutation step a difficulty scale; L1–L3 are frozen calibration ranks, not a guarantee of identical ordering for unseen policies.

H.2 Complete results for five model

Table 18 to 22 show detailed statics for five model across 215 variants.

Table 18: Complete depth-wise results for **π_0-FAST** on the 215-artifact FORGE-Bench. **PSR** is process success rate; **T–P Gap** is TSR − PSR in percentage points; **Act.** is the median number of action steps to process success in Isaac Sim. D1–D10 denote mutation depth. The depth-free LING activation artifacts are reported separately as LING setup. Colors mark the best value across the five models at the same dimension and depth. Missing or incomplete model–artifact cells follow the frozen 30-episode zero-imputation policy.

Dim.	Metric	π_0-FAST									
		D1	D2	D3	D4	D5	D6	D7	D8	D9	D10
Root (L0)	PSR ↑	29.3									
	T–P Gap ↓	14.0									
	Act. ↓	150.5									
VIS	PSR ↑	**28.0**	22.0	22.0	**20.0**	**22.0**	**22.5**	13.3	**7.5**	5.8	4.2
	T–P Gap ↓	20.0	11.3	19.3	26.7	24.0	42.5	35.8	34.2	38.3	36.7
	Act. ↓	156.0	166.0	166.0	**136.0**	**152.0**	**147.0**	129.0	112.0	127.0	**92.0**
LING	PSR ↑	28.0	**26.0**	16.0	19.3	7.3	18.7	24.0	**22.7**	17.3	8.7
	T–P Gap ↓	24.0	23.3	16.0	13.3	10.7	17.3	26.0	20.0	20.0	16.0
	Act. ↓	**143.0**	151.0	**146.0**	152.0	**138.0**	**140.5**	**127.0**	144.5	139.0	127.0
DCR	PSR ↑	**27.3**	29.3	**38.7**	16.0	20.0	28.3	30.8	32.2	35.6	13.3
	T–P Gap ↓	16.0	19.3	14.7	28.7	28.0	20.8	15.8	20.0	24.4	42.2
	Act. ↓	168.0	**147.0**	142.5	**128.5**	131.0	160.5	170.0	149.0	155.0	161.5
LONG	PSR ↑	22.0	2.0	0.0	0.0	0.0	0.0	0.0	0.0	–	–
	T–P Gap ↓	18.0	9.3	0.0	0.0	0.0	0.0	0.0	0.0	–	–
	Act. ↓	146.0	561.0	–	–	–	–	–	–	–	–
SAFE	PSR ↑	**24.0**	16.7	**15.3**	16.7	12.0	23.3	15.0	23.3	22.2	15.6
	T–P Gap ↓	19.3	26.0	31.3	33.3	38.7	31.7	29.2	46.7	42.2	47.8
	Act. ↓	143.0	157.0	**157.0**	**152.0**	130.0	127.0	145.0	**137.0**	146.0	**137.5**

H.3 Mutation trajectories and failure modes

The trajectory plots use the matched three-model all-variant snapshot with 215 trajectory artifacts. It contains observed episodes for π_0-FAST, MolmoAct 2, and GR00T, respectively, and includes the five LING activation setups. Figures 10 and 11 show that search step is not uniformly monotonic across tasks or dimensions, while Figure 13 shows the accompanying failure composition. The non-monotonicity is expected because the sequence may change operators and because LING and LONG do not define cumulative physical ladders.

Held-out trajectories and cross-model transfer. Figure 12 adds the complete observed search-step trajectories of the two held-out policies, π_0 and $\pi_{0.5}$, over the same 215-artifact candidate universe. As in the source-model plots, search step records mutation order rather than a calibrated difficulty rank.

For a matched transfer comparison, Table 23 tests a root-relative criterion: among candidates for which a 20-point PSR drop is possible, 85/93 for π_0 and 73/95 for $\pi_{0.5}$ reach that threshold. The frozen L3 variants also retain diagnostic yield for the held-out policies (8/10 and 6/10 eligible variants, respectively). These results support transfer of *diagnostic value*, not transfer of a single difficulty ordering. The rank associations are moderate, and the trajectories in Figure 12 include model-specific rebounds and divergent responses to the same mutation. A variant can thus expose one policy's failure case while leaving another little changed or even improving it, as the LIBERO-04 color-change case study below illustrates. Search step records the adaptive mutation sequence, not a scalar mutation depth or a universal difficulty axis; the calibrated L1–L3 labels likewise need not induce identical PSR curves on previously unseen policies.

Table 19: Complete depth-wise results for **MolmoAct 2** on the 215-artifact FORGE-Bench. **PSR** is process success rate; **T–P Gap** is TSR − PSR in percentage points; **Act.** is the median number of action steps to process success in Isaac Sim. D1–D10 denote mutation depth. The depth-free LING activation artifacts are reported separately as LING setup. Colors mark the best value across the five models at the same dimension and depth. Missing or incomplete model–artifact cells follow the frozen 30-episode zero-imputation policy.

Dim.	Metric	MolmoAct 2									
		D1	D2	D3	D4	D5	D6	D7	D8	D9	D10
Root (L0)	PSR ↑					30.0					
	T–P Gap ↓					10.0					
	Act. ↓					147.5					
VIS	PSR ↑	25.3	29.3	26.0	17.3	12.0	22.5	20.0	5.0	8.3	23.3
	T–P Gap ↓	20.7	6.7	2.0	5.3	3.3	7.5	10.8	14.2	32.5	26.7
	Act. ↓	156.0	140.0	140.0	210.5	297.5	151.0	243.0	193.5	174.0	172.0
LING	PSR ↑	34.0	24.7	34.0	13.3	32.7	32.0	28.0	22.7	32.0	10.0
	T–P Gap ↓	9.3	10.7	4.0	7.3	6.0	9.3	7.3	11.3	4.0	9.3
	Act. ↓	168.0	167.0	168.5	170.5	183.0	171.0	167.0	175.0	156.0	163.0
DCR	PSR ↑	27.3	34.0	32.0	32.0	27.3	34.2	39.2	42.2	52.2	20.0
	T–P Gap ↓	6.0	10.7	13.3	18.7	28.0	10.8	9.2	17.8	13.3	15.6
	Act. ↓	178.0	158.0	176.0	163.0	154.0	166.0	243.0	164.0	168.0	179.5
LONG	PSR ↑	22.7	20.7	0.0	0.0	0.0	0.0	6.7	0.0	–	–
	T–P Gap ↓	8.7	4.0	0.0	8.3	0.0	0.0	10.0	0.0	–	–
	Act. ↓	181.0	353.0	–	–	–	–	386.0	–	–	–
SAFE	PSR ↑	10.0	28.0	15.3	19.3	20.7	25.0	25.0	40.0	52.2	36.7
	T–P Gap ↓	22.0	13.3	23.3	15.3	8.0	11.7	25.8	26.7	18.9	22.2
	Act. ↓	155.0	156.0	164.0	166.0	184.0	193.5	180.0	187.5	183.0	184.0

Table 20: Complete depth-wise results for **GR00T N1.7** on the 215-artifact FORGE-Bench. **PSR** is process success rate; **T–P Gap** is TSR − PSR in percentage points; **Act.** is the median number of action steps to process success in Isaac Sim. D1–D10 denote mutation depth. The depth-free LING activation artifacts are reported separately as LING setup. Colors mark the best value across the five models at the same dimension and depth. Missing or incomplete model–artifact cells follow the frozen 30-episode zero-imputation policy.

Dim.	Metric	GR00T N1.7									
		D1	D2	D3	D4	D5	D6	D7	D8	D9	D10
Root (L0)	PSR ↑					6.7					
	T–P Gap ↓					20.0					
	Act. ↓					180.0					
VIS	PSR ↑	16.0	7.3	12.7	6.7	5.3	3.3	7.5	5.0	3.3	20.0
	T–P Gap ↓	30.7	21.3	16.0	14.7	16.0	12.5	16.7	13.3	14.2	17.5
	Act. ↓	318.0	260.5	191.0	187.0	164.5	153.5	173.0	181.5	163.0	145.0
LING	PSR ↑	24.7	20.7	18.0	20.0	26.0	8.7	19.3	15.3	9.3	15.3
	T–P Gap ↓	17.3	16.0	21.3	17.3	12.7	12.0	22.7	30.0	16.7	16.7
	Act. ↓	245.0	230.0	245.0	283.0	207.0	311.0	280.0	290.0	204.0	280.5
DCR	PSR ↑	5.3	6.0	8.7	14.0	4.0	14.2	11.7	18.9	17.8	12.2
	T–P Gap ↓	20.0	14.7	15.3	8.0	22.0	22.5	21.7	20.0	21.1	20.0
	Act. ↓	190.0	211.0	186.0	242.5	137.0	198.0	195.5	198.0	188.0	194.0
LONG	PSR ↑	12.0	2.0	2.5	1.7	20.0	0.0	0.0	0.0	–	–
	T–P Gap ↓	15.3	9.3	9.2	11.7	13.3	4.4	3.3	33.3	–	–
	Act. ↓	242.0	472.0	523.0	159.5	94.0	–	–	–	–	–
SAFE	PSR ↑	4.0	8.7	7.3	9.3	5.3	0.8	1.7	3.3	4.4	2.2
	T–P Gap ↓	14.0	17.3	30.0	19.3	22.0	23.3	13.3	20.0	14.4	30.0
	Act. ↓	159.5	226.0	238.0	240.0	171.0	–	184.5	161.0	166.5	190.5

H.4 Operator relationships

Figures 14 and 15 group parent–child changes by the operator recorded for the child. Success impact is the change in TSR or PSR from the matched parent batch; failure impact is the corresponding

Table 21: Complete depth-wise results for π_0 on the 215-artifact FORGE-Bench. **PSR** is process success rate; **T–P Gap** is TSR−PSR in percentage points; **Act.** is the median number of action steps to process success in Isaac Sim. D1–D10 denote mutation depth. The depth-free LING activation artifacts are reported separately as LING setup. Colors mark the best value across the five models at the same dimension and depth. Missing or incomplete model–artifact cells follow the frozen 30-episode zero-imputation policy.

Dim.	Metric	π_0									
		D1	D2	D3	D4	D5	D6	D7	D8	D9	D10
Root (L0)	PSR ↑	26.7									
	T–P Gap ↓	14.0									
	Act. ↓	159.0									
VIS	PSR ↑	14.7	4.0	14.7	2.0	5.3	2.5	0.0	0.8	0.8	0.0
	T–P Gap ↓	**12.0**	9.3	10.0	12.0	7.3	10.0	**6.7**	**12.5**	**9.2**	**7.5**
	Act. ↓	145.0	356.0	256.0	239.0	229.0	230.0	–	196.0	–	–
LING	PSR ↑	1.3	2.0	1.3	0.7	1.3	0.7	2.0	2.0	2.7	10.0
	T–P Gap ↓	**8.0**	**4.7**	**3.3**	**7.3**	8.0	**2.7**	**4.7**	**6.0**	6.0	**4.7**
	Act. ↓	160.0	259.5	284.0	488.0	172.5	215.0	273.0	302.0	214.0	117.0
DCR	PSR ↑	14.7	13.3	20.0	16.0	18.7	6.7	5.8	3.3	3.3	11.1
	T–P Gap ↓	17.3	**8.0**	**12.7**	10.7	**14.7**	**9.2**	15.8	**10.0**	**6.7**	**10.0**
	Act. ↓	**146.0**	268.0	177.0	254.5	148.0	273.0	210.0	188.0	485.0	289.0
LONG	PSR ↑	13.3	0.7	**6.7**	0.0	1.1	**4.4**	0.0	**3.3**	–	–
	T–P Gap ↓	15.3	**2.0**	14.2	0.0	**0.0**	**4.4**	0.0	**10.0**	–	–
	Act. ↓	162.0	505.0	245.0	–	526.0	**485.0**	–	**373.0**	–	–
SAFE	PSR ↑	4.0	16.7	2.0	2.0	8.7	0.0	5.0	5.6	10.0	6.7
	T–P Gap ↓	16.0	**8.7**	**13.3**	19.3	**8.0**	15.0	14.2	**5.6**	20.0	**6.7**
	Act. ↓	171.0	181.0	163.0	271.0	147.0	–	148.0	320.0	144.0	234.0

Table 22: Complete depth-wise results for $\pi_{0.5}$ on the 215-artifact FORGE-Bench. **PSR** is process success rate; **T–P Gap** is TSR−PSR in percentage points; **Act.** is the median number of action steps to process success in Isaac Sim. D1–D10 denote mutation depth. The depth-free LING activation artifacts are reported separately as LING setup. Colors mark the best value across the five models at the same dimension and depth. Missing or incomplete model–artifact cells follow the frozen 30-episode zero-imputation policy.

Dim.	Metric	$\pi_{0.5}$									
		D1	D2	D3	D4	D5	D6	D7	D8	D9	D10
Root (L0)	PSR ↑	15.3									
	T–P Gap ↓	42.7									
	Act. ↓	**146.0**									
VIS	PSR ↑	8.0	0.7	5.3	0.0	10.7	0.8	2.5	4.2	0.8	3.3
	T–P Gap ↓	32.0	16.0	14.7	13.3	12.0	19.2	15.0	25.8	12.5	22.5
	Act. ↓	**132.5**	143.0	**106.0**	–	185.5	179.0	**119.0**	**104.0**	**83.0**	159.5
LING	PSR ↑	2.0	1.3	0.0	1.3	6.0	1.3	1.3	0.7	1.3	5.3
	T–P Gap ↓	8.7	12.0	13.3	10.0	26.0	9.3	10.7	11.3	7.3	24.7
	Act. ↓	158.0	**150.0**	–	**129.0**	143.0	141.5	138.5	**125.0**	**120.5**	**116.0**
DCR	PSR ↑	8.7	1.3	14.0	4.0	4.0	4.2	0.8	8.9	2.2	3.3
	T–P Gap ↓	30.0	14.7	29.3	31.3	28.7	15.0	30.0	12.2	**6.7**	20.0
	Act. ↓	150.0	179.0	**133.0**	166.0	**107.0**	**126.0**	**108.0**	**147.0**	**144.5**	**144.0**
LONG	PSR ↑	9.3	1.3	0.8	0.0	1.1	0.0	0.0	0.0	–	–
	T–P Gap ↓	26.7	17.3	17.5	**7.5**	3.3	8.9	0.0	26.7	–	–
	Act. ↓	**94.5**	**208.0**	**208.0**	–	145.0	–	–	–	–	–
SAFE	PSR ↑	2.7	4.0	2.0	1.3	5.3	2.5	5.8	0.0	6.7	0.0
	T–P Gap ↓	22.7	24.7	18.0	22.7	25.3	25.8	32.5	7.8	25.6	21.1
	Act. ↓	**120.5**	**130.0**	182.0	188.5	**125.0**	**81.0**	**118.0**	–	**92.0**	–

change in per-episode canonical failure incidence. Shared operators are interpreted within their dimension-bound axes. These plots are descriptive: operators were selected from feedback rather than randomized, so their measured deltas must not be read as causal treatment effects.

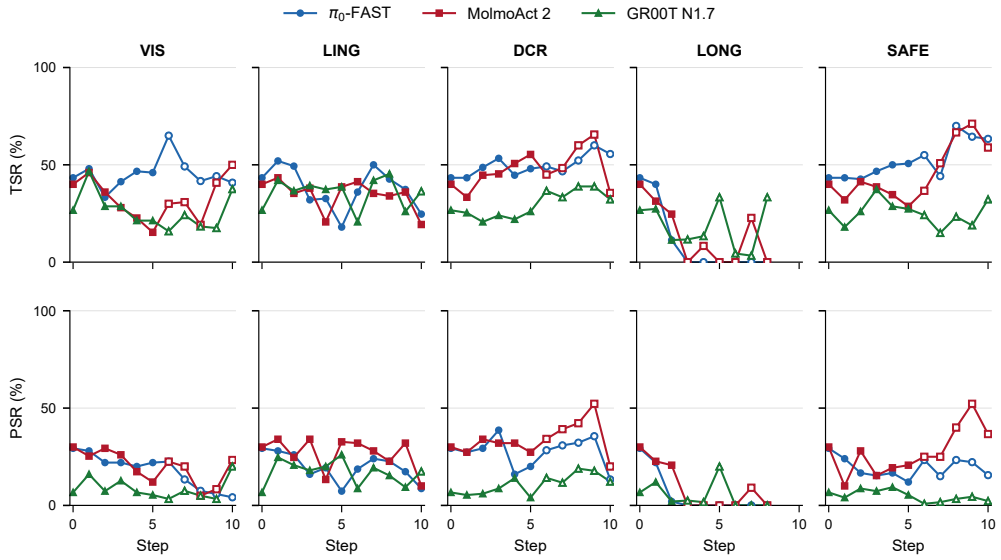

Figure 10: Terminal and process success across matched search steps for the three calibration models. Coverage at later steps varies by dimension and is shown in the source visualization.

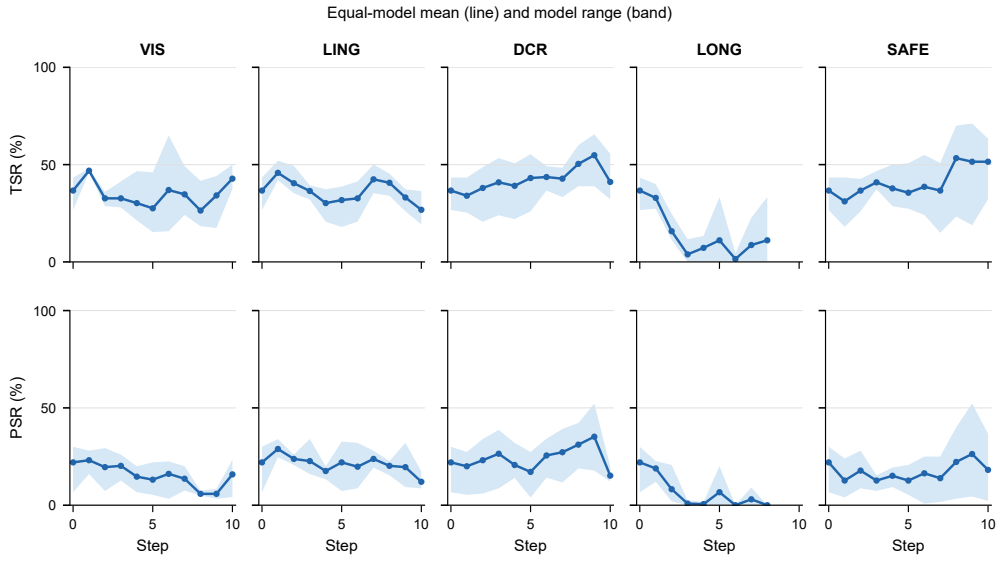

Figure 11: Average three-model terminal and process success over matched tasks at each mutation step. These curves describe search trajectories, not the post-calibration L1–L3 ordering.

I Oracle Validation, Diagnostic Cases, and Limitations

I.1 Human validation

Human reviewers annotated 396 rollouts, yielding 792 terminal/process binary labels. The aggregate terminal accuracy is 96.2%, process accuracy is 95.7%, label-level micro accuracy is 96.0%, and both labels match the reviewer on 94.7% of rollouts.

The validation set does not contain GR00T, π_0, $\pi_{0.5}$, or LIBERO-23, and its model/task composition should therefore be stated whenever the aggregate number is used.

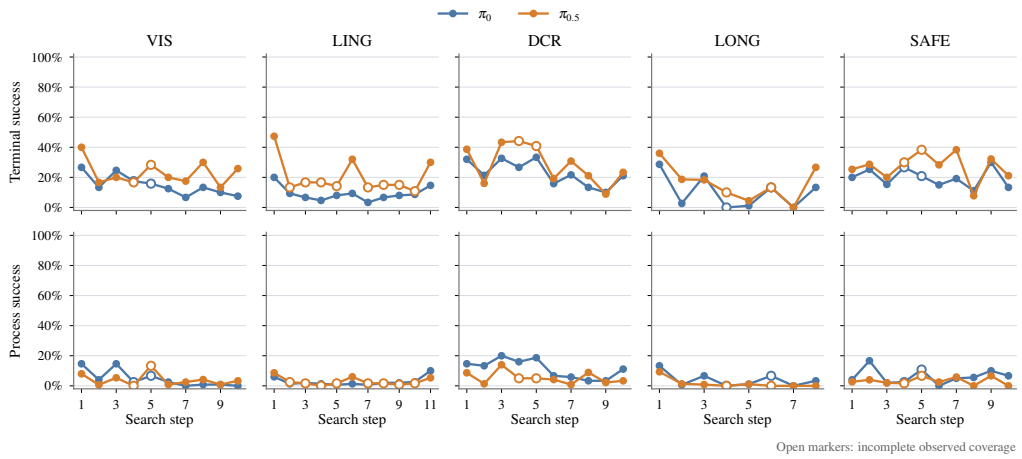

Figure 12: Terminal and process success of π_0 and $\pi_{0.5}$ across all observed candidate mutation steps in the five dimensions. Each point pools the observed rollouts from the canonical task variants present at that step; the number of tasks varies with step. Open markers indicate that one or more canonical model–artifact cells at that step have incomplete observed coverage. Unobserved cells are excluded rather than assigned zero success.

Table 23: Root-relative held-out diagnostic yield on the four fully observed LIBERO tasks. A variant is eligible for the 20-percentage-point (pp) drop test only if that policy's root PSR is at least 20%. L3 denotes the frozen hard-level variants selected using the three calibration models. Eligible roots differ by policy, so the rows are not a model ranking.

Held-out policy	Candidates	$\geq$20 pp drop	L3 candidates	L3 $\geq$20 pp drop
π_0	93	85/93 (91.4%)	10	8/10 (80.0%)
$\pi_{0.5}$	95	73/95 (76.8%)	10	6/10 (60.0%)

I.2 Diagnostic cases on frozen variants

The selected cells illustrate several behaviors that averages obscure. LIBERO-04/VIS L1 is a clean high-performance case for π_0-FAST (100.0% TSR and PSR). At LIBERO-04/SAFE L3, the same model reaches 70.0% TSR but 0.0% PSR, showing how terminal completion can conceal pervasive process violations. LIBERO-23/VIS L3 illustrates model disagreement: GR00T obtains 83.3/40.0 TSR/PSR, compared with 63.3/6.7 for π_0-FAST and 20.0/10.0 for MolmoAct 2. On the same LIBERO-23/SAFE L3 variant, the two held-out policies both reach 56.7% TSR, yet π_0 reaches 43.3% PSR and $\pi_{0.5}$ only 6.7% PSR. Thus, identical terminal success can hide policy-specific process failures. Several LING rows are also non-monotonic under the fixed semantic protocol, confirming that mutation index and linguistic difficulty are not interchangeable.

Case study: model-specific response to a controlled color mutation. Figure 16 follows one parent–child transition in the complete mutation trajectory for LIBERO-04, whose instruction is to pick the banana and place it on the plate. From VIS step 9 to step 10, `change_material` changes one non-target orange from orange to red while preserving the task instruction, object geometry, and scene layout. Each model is evaluated for 30 rollouts on both variants.

Table 24: FORGE Oracle agreement with human review. False success means that the Oracle accepts an outcome the reviewer rejects; false failure is the opposite.

Label	n	Acc.	Both succ.	Both fail	False succ.	False fail
Terminal	396	96.2	204	177	8	7
Process	396	95.7	131	248	4	13

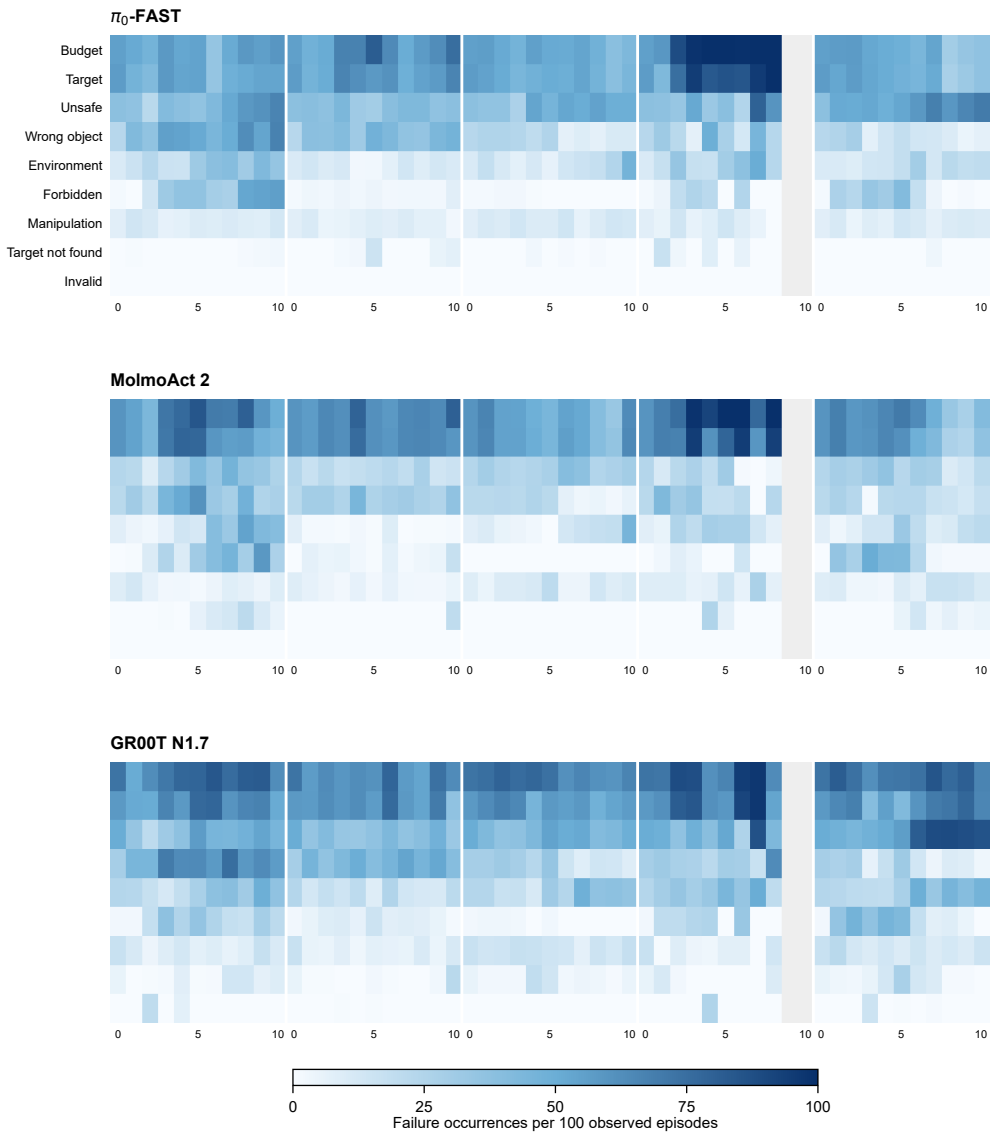

Figure 13: Canonicalized failure-mode incidence across search steps for the three construction/calibration models. Each mode is counted at most once per rollout, but different modes may co-occur.

The local appearance change produces sharply different responses across all five policies. MolmoAct 2 improves from 8/30 to 19/30 process-successful rollouts (26.7% to 63.3%, +36.7 points), while its TSR increases from 28/30 to 30/30. In contrast, π_0-FAST remains at 13/30 terminal successes and 0/30 process successes, and GR00T decreases slightly from 3/30 to 1/30 terminal successes while remaining at 0/30 process successes. The held-out π_0 stays at 1/30 terminal successes, while $\pi_{0.5}$ rises from 1/30 to 3/30; neither achieves process success at either step. The improvement in MolmoAct 2 PSR therefore does not generalize to the held-out policies on this pair. One plausible explanation is that the red appearance makes the non-target fruit more visually separable from the yellow banana for MolmoAct 2, but the observation alone does not identify a model's internal mechanism. More generally, the opposite response directions show that mutation difficulty is model-dependent rather than an intrinsic property of the edited scene. This example therefore motivates multi-model calibration and cautions against treating search step or operator identity as a universal difficulty ordering. Because the mutation was selected adaptively, the comparison is descriptive rather than a causal estimate of color sensitivity.

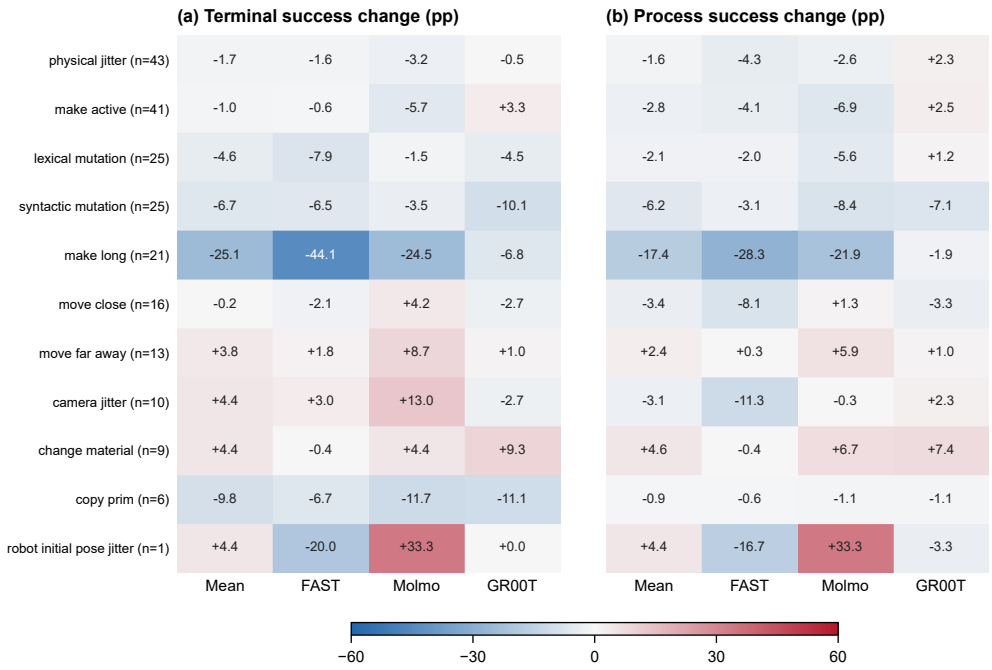

Figure 14: Observed parent–child terminal- and process-success changes grouped by mutation operator.

I.3 Limitations and Future Work

Limitations. Our current evaluation is conducted in simulation with a fixed-base Franka embodiment. Moreover, the reported operator–failure relationships are observational rather than causal because operators are chosen adaptively from execution feedback, and the resulting failure labels describe externally observed behavior rather than a policy's internal reasoning process.

Future work. Future versions of FORGE-Bench will broaden task families, simulators, robot embodiments, and real-world evaluation, while balancing task coverage across all five capability dimensions. Randomized operator interventions and larger held-out task suites would enable stronger causal analysis of mutation effects and more reliable estimates of benchmark generalization. Finally, extending the search to jointly compose capability dimensions may reveal failures that cannot be exposed by changing one dimension at a time.

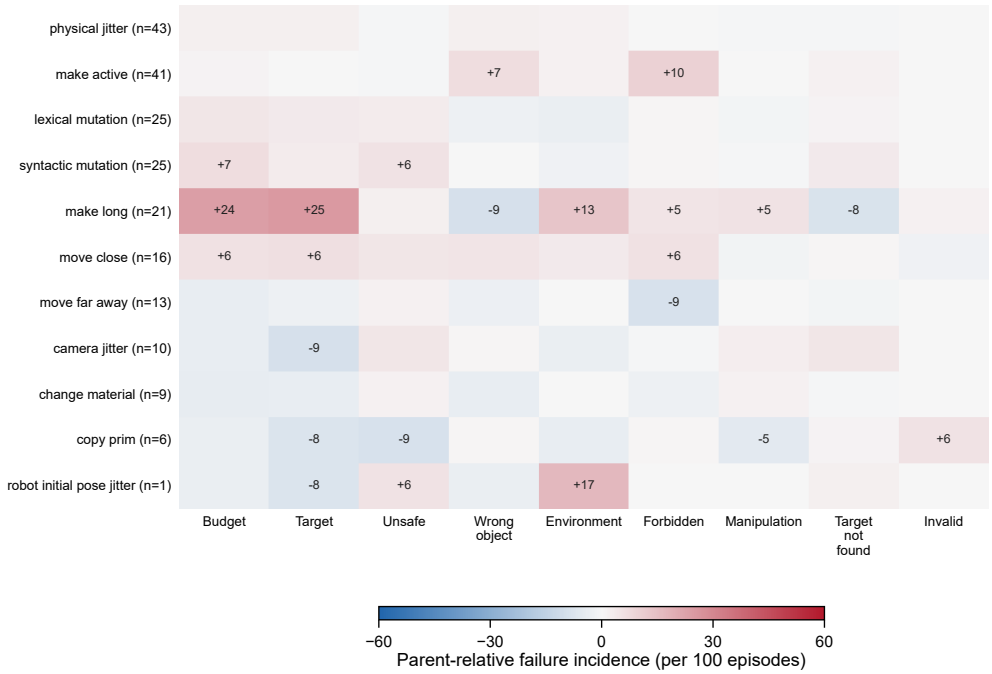

Figure 15: Observed parent–child changes in canonical failure-mode incidence grouped by mutation operator. Blue denotes fewer failures and red denotes more failures.

(a) Step 9: parent variant

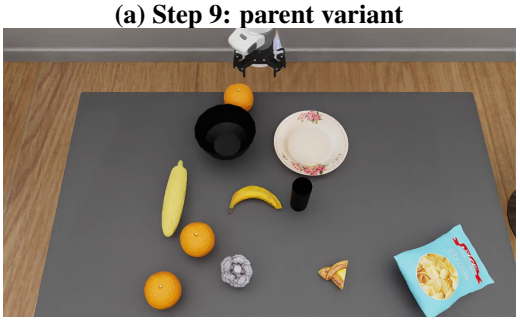

(b) Step 10: `change_material`



(c) Model response (30 rollouts per model and step)

	TSR (%)			PSR (%)		
Model	Step 9	Step 10	Δ	Step 9	Step 10	Δ
π_0-FAST	43.3	43.3	0.0	0.0	0.0	0.0
MolmoAct 2	93.3	100.0	+6.7	26.7	63.3	**+36.7**
GR00T N1.7	10.0	3.3	−6.7	0.0	0.0	0.0
π_0	3.3	3.3	0.0	0.0	0.0	0.0
$\pi_{0.5}$	3.3	10.0	+6.7	0.0	0.0	0.0

Figure 16: A controlled VIS mutation elicits divergent model responses on LIBERO-04. Panels (a) and (b) show the matched step-9 parent and step-10 child; the changed orange is located left of center below the yellow squash. Panel (c) reports terminal success rate (TSR), process success rate (PSR), and the step-10 minus step-9 change in percentage points.