000	1	Appendix
001		
002	1.1	Overview
003		
004	The	Appendix contains the following content.
005		
006		• Failure faxonomy (Appendix 1.2): more thorough definition and figure to discussions
007		
800		• FailGen Data Generation Pipeline (Appendix 1.3): more discussion of FailGen imple-
009		mentation with example configurations files.
010		• AHA Datasets (Appendix 1.4): more details on the instruction-tuning dataset and evaluation
012		datasets.
013		• Additional Experimental Results (Appendix 1.5): more details and experiments with instruction finetuning.
014 015		• Downstream Robotic Application: VLM Reward Generation (Appendix 1.6): more
016		poncy ronouts, generated reward function examples, and prompts.
017 018		• Downstream Robotic Application: VLM Task-plan Generation (Appendix 1.7): more policy rollouts, generated task-plan examples, and prompts.
019		• Downstream Robotic Application: VLM Sub-task Verification (Appendix 1.8): more
020		policy rollouts.
021		
022	1.2	Failure Taxonomy
023		
024	We	conducted an in-depth study of recent real-world, diverse robot datasets (such as Open-X (Padalkar
025	et al	., 2023), DROID (Knazatský et al., 2024), and EGO4D (Grauman et al., 2022)) and the policies
026	whi	sh can be categorized into seven types: incomplete grasp inadequate grip retention misaligned
027	keyf	rame, incorrect rotation, missing rotation, wrong action sequence, and wrong target object.
028	Ince	Amplete Crean (No. Crean) Failures No. Crean is an object contribution that account when
029 030	the g	gripper reaches the desired grasp pose but fails to close before proceeding to the next keyframe.
031	Inac	lequate Grip Retention (Slip) Failure: Slip is an object-centric failure that occurs after the
032	obje	ct has been successfully grasped. As the gripper moves the object toward the next task-specific
033	keyf	rame, the grip weakens, causing the object to slip from the gripper. For generating the AHA
034	data	set for training and evaluation, we configured a 5-timestep activation for the Slip failure mode,
035	trigg	gering the object to drop from the gripper.
036	Mis	aligned keyframe (Translation) Failure: This action-centric failure occurs when the gripper
037	mov	es toward a task keyframe, but a translation offset along the X, Y, or Z axis causes the task to fail.
038	FOr In th	the AHA training and evaluation dataset, we introduced a translation offset of $[-0.5, 0.5]$ meters.
039	or Z	axis from the original waypoint. The translation coordinate system is depicted in Figure 3 (Left)
040	T	
042	Inco	prect Rotation (Rotation) Failure: Rotation is an action-centric failure that occurs when
043	vaw	or pitch leading to task failure. For the AHA dataset, we set a rotation offset of [-3, 14, 3, 14] in
044	radi	ans along roll, vaw, or pitch. The rotation coordinate system is depicted in Figure 3 (Right).
045	ъ <i>л</i> !-	
046	han	sing Kotation (No_Rotation) Failure: No_Rotation is an action-centric failure that
047	rota	tion (roll, vaw, or pitch) for the sub-task, resulting in task failure.
048	11 7	$\mathbf{x} = \mathbf{x} + \mathbf{y} + $
049	that	mg Action Sequence (wrong_action) ranure: wrong_action is an action-centric failure
050	corr	ect one. For example, in the task put, cube, in drawer, the robot moves the cube toward
051	the	drawer before opening it, leading to task failure.
052	Wm	ang Target Object (Wrong, object) Failure: Wrong, object is an object centric failure
053	that	occurs when the robot acts on the wrong target object, not matching the language instruction. For



Figure 1: Failure mode reference coordinate systems. (Left) Translation coordinate system, and (Right) rotation coordinate system.

example, in the task pick_the_red_cup, the gripper picks up the green cup, causing failure. We perform a sweep through all manipulable objects, swapping them with the target object in the scene.

1.3 FAILGEN DATA GENERATION PIPELINE

We developed FailGen, an environment wrapper that can be easily integrated into any simulator. It leverages pre-defined or hand-crafted robot demonstrations for imitation learning, where each 083 trajectory is represented as a waypoint-based system. Two consecutive waypoints form a sub-084 task, with each sub-task linked to a predefined set of language descriptions. FailGen allows for 085 modifications to environment parameters, such as gripper end-effector translation, rotation, and open/close state. By altering these parameters, we systematically generate failures at every waypoint. 087 However, for the 79 tasks collected from RLBench, we do not initially know which sub-task will fail 088 due to specific failure modes. To address this, we perform a systematic sweep, using RLBench's builtin success conditions to explore all possible combinations. This generates a configuration of failures for each task, which we then use to procedurally generate all failure training data. Additionally, we 091 manually annotate each sub-task with natural language instructions describing the task, and pair this with failure mode explanations to serve as language input for instruction-tuning. Example of the 092 configuration files are depicted at Figure 5. 093

095 1.4 Aha Dataset

Using FailGen, we curated two datasets from RLBench (James et al., 2020). The first is the training dataset, AHA dataset (train), which is used for instruction-tuning AHA alongside the co-train dataset. The second is the testing dataset, AHA dataset (test), used for evaluation. AHA dataset (train) contains approximately 49k image-query pairs of failures derived from 79 tasks, while AHA dataset (test) consists of around 11k image-query pairs from 10 hold-out tasks.

101 102

103

094

074

075 076 077

078

079 080

081

1.5 ADDITIONAL EXPERIMENTAL RESULTS

We conducted additional experiments to better understand and visualize AHA's predictions. We
 trained two versions of the AHA model with 13B parameters, using different language models for
 fine-tuning: Llama-2-13B and Vicuna-1.5-13B. The results showed less than a 2.5% performance
 difference between the two models, indicating that our fine-tuning data is effective regardless of the
 base language model. These results are presented in Table 3. Additionally, we visualized the output

108		data:
109		# Where to save the demos save_path: /home/data
110	<pre>1 save_path: /home/\${oc.env:USER}/data/failgen_data</pre>	<pre># The size of the images to save waypoints: [0, 1, 2, 3] failures:</pre>
111	2 obs mode: rgb	- type: grasp
112	3 render mode: sensors	enabled: False waypoints: [1]
112	4 shader: default	- type: translation v
113	5 sim backend: auto	name: trans_y enabled: False
114	6 image size: [256 256]	waypoints: [1,2,3] range: [-0.5, 0.5]
115	7 stages [0, 1, 2, 2]	- type: rotation x
116	7 Stayes: [0, 1, 2, 5]	name: rot_x enabled: False
117	8 failures:	waypoints: [0]
118	9 – type: grasp	
110	10 enabled: false	name: bad_seq
119	11 stages: [2]	waypoints: [2,3]
120	12 - type: trans_x	sub-tasks:
121	13 enabled: false	- task_no: 0 enabled: False
122	14 stages: [0. 1. 3]	type: dummy targets: [ball]
123	15 noise: 0.1	processes: [waypoint0, waypoint1] task_description: [
194	16 - type trans y	"grasp onto the clock knob", "pick on the clock knob"
105	17 enabled: false	- task_no: 1
120		enabled: False type: dummy
126	18 Stages: [0, 1, 3]	targets: [ball] processes: [waypoint1, waypoint2]
127	19 noise: 0.1	task_description: ["rotate the knob",
128	20 - type: trans_z	"turn the knob"]
129	21 enabled: false	- task_no: 2 enabled: False
130	22 stages: [0, 1, 3]	type: dumny targets: [ball]
100	23 noise: 0.1	processes: [waypoint2, waypoint3] task_description: [
131	24	"let go", "release the gripper"
132)



133

152

153 154



Figure 3: Data distribution of AHA dataset for both train and test.

predictions from various baselines compared to our model and evaluated performance across all three datasets, with the results shown in Figure 4.

AHA model performance uncertainty estimation. To evaluate the relationship between uncertainty estimation and model performance, we conducted additional experiments across various evaluation datasets. Specifically, we compared the sentence token prediction probabilities of AHA-13B with those of LLaVA v1.5-13B. AHA-13B exhibited significantly higher average prediction probabilities, reflecting its superior accuracy. These findings underscore the positive impact of fine-tuning with the AHA failure dataset on model confidence and performance as depicted in Table 2.



Figure 4: Examples of different failure modes. Row 1: No_grasp and Rotation_x. Row 2: Rotation_y and Rotation_z. Row 3: Slip and Wrong_sequence. Row 4: Translation_x and Translation_y. Row 5: Translation_z and Wrong_object.

Effects of Viewpoints on Evaluation. We evaluated the reasoning capabilities of the AHA model on the ManiSkill-Failure dataset across three different viewpoint configurations (one, two, and three viewpoints). Interestingly, we observed a slight performance advantage when using single-viewpoint images. We attribute this to the resolution limitations of the LLaVA-1.5 visual encoder (256x256), where single-viewpoint inputs offer clearer and more focused visual information for failure reasoning, as summarized in Table 1.

Model: AHA-13B (Viewpoints)	Binary Success	ROUGE-L	LLM Fuzzy Match	Cosine Similarity
One viewpoint	1.000	0.673	0.587	0.712
Two viewpoints	1.000	0.615	0.587	0.671
Three viewpoints	1.000	0.600	0.633	0.681

Table 1: Performance comparison across different numbers of viewpoints for AHA-13B

Dataset	AHA-13B (Output Probabilities / Cosine Similarity)	LLaVA-13B-v1.5 (Output Probabilities / Cosine Similarity)
AHA Dataset (Test)	0.670 / 0.583	0.066 / 0.208
Maniskill Fail	0.457 / 0.681	0.024 / 0.208
RoboFail	0.292 / 0.471	0.000 / 0.203

Table 2: Performance against model prediction sentence probabilities likelihood evaluated across datasets.

Table 3: Ablation on Different Base LLMs for Fine-Tuning. We fine-tuned AHA-13B using both LLaMA-2-13B and Vicuna-1.5-13B as base LLM models. The quantitative results show that the average performance difference between the two models is less than 2.5%, indicating that our failure formulation and the AHA dataset are effective regardless of the base model selection.

	AHA dataset (Test)			ManiSkill-Fail			RoboFail					
Models	$ROUGE_L \uparrow$	Cos Sim ↑	BinSucc(%) ↑	Fuzzy Match↑	$ROUGE_L \uparrow$	Cos Sim ↑	BinSucc(%)↑	Fuzzy Match↑	$ROUGE_L \uparrow$	Cos Sim ↑	BinSucc(%) ↑	Fuzzy Match ↑
AHA-13B (Llama-2)	0.446	0.583	0.702	0.768	0.600	0.681	1.000	0.633	0.280	0.471	0.643	0.465
AHA-13B (Vicuna-1.5)	0.458	0.591	0.709	0.695	0.574	0.657	1.000	0.851	0.290	0.468	0.661	0.605

229 230 231

232

240

254 255

256

220

221

222

224

225

226 227 228

1.6 VLM REWARD GENERATION

In this section, we present reward functions generated by GPT-40 and AHA for comparison, as shown
in Figure 5. Additionally, we demonstrate RL policy rollouts improved through AHA 's failure
feedback across all five tasks along with all the final dense reward function modified by AHA shown
in Figure 6 and 7. For all tasks, except **put_sphere_on_holder** (trained with PPO for 10M steps),
PPO was trained for 25M steps prior to reflection and evaluation.

Simulation task Details We describe each of the 4 tasks in detail, along with their Maniskill variations
 and success condition.

241 1.6.1 PICKUP YCB 242

243 Filename: pick_single_ycb.py

Task: Pick up the single YCB object and lift it up to target height.

Success Metric: The object position is within goal_thresh (default 0.025) euclidean distance of the goal position.

- 248 249 1.6.2 PUSH T
- 250 Filename: push_T.py
 251
- **Task:** Push the T into the T shaped area.

Success Metric: The 3D T block covers at least 90

1.6.3 PLACE SPHERE

257 Filename: place_sphere_v1.py

Task: Pick up the sphere and place it into the bin.

Success Metric: the sphere is on top of the bin. That is, the sphere's xy-distance to the bin goes near
 0, and its z-distance to the bin goes near the sphere radius + the bottom bin block's side length the object is static. That is, its linear and angular velocities are bounded with a small value the gripper is not grasping the object.

264 1.6.4 STACK CUBE

266 Filename: stack_cube_v1.py

Task: Pick up the red cube and stack it onto the green cube.

Success Metric: the red cube is on top of the green cube (to within half of the cube size), the red cube is static, the red cube is not being grasped by the robot (robot must let go of the cube).

Table 4: Examples of detection and reasoning over failures by various VLMs. We sampled one
 evaluation data point from each of the three evaluation datasets and compared the reasoning predic tions generated by different state-of-the-art VLMs alongside our AHA-13B. Our model consistently
 provides the most semantically similar reasoning to the ground-truth answers for the given queries.

274				
275	Evaluation Datasets	AHA dataset (Test)	Maniskill-Fail	RoboFail (Liu et al., 2023)
276		50 58 81 0 0		
277		50 58 81 0 0		
278		10 10 0	1 1 TO TO	1 2 2 2
279		50 58 81 0 0	1 2 3 4 5	
280				
281	Question	At the current sub-task, the	At the current sub-task, the robot	At the current sub-task, the robot
282		robot is picking up the hockey stick. First, determine whether it	is picking up the red cube. First, determine whether it has suc-	is placing the apple into the bowl. First, determine whether it
203		has succeeded by choosing from	ceeded by choosing from ["yes", "po"] If not explain why it	has succeeded by choosing from
285		it failed?	failed?	it failed?
286	LLaVA 1.5-13B	The image displays a time-lapse	Based on the image, it depicts a	For the pick-up task, i would say
287		sequence of a robotic arm grab- bing a hockey stick.	different sub-task of a robot grasp- ing a red block.	the robot arm has succeeded in picking.
288	GPT4o	No, the robot arm fails to hold the	Yes, the robot arm has success-	No, the robot gripper failed to pick
289		hockey stick securely.	fully picked up the red block.	up the bowl in the second frame.
290	Gemini 1.5 Flash	Yes.	Yes, Yes, Yes.	Frame 1, no the gripper is not
291				
292	AHA-13B (Ours)	No, the robot slip the object out of its gripper.	No, the robot gripper fails to close the gripper.	No, this is not the right action se- quence for the task.
294	Ground-truth Answer	No, the robot slip the object out of its gripper.	No, the robot gripper fails to close.	No, this is not the right action se- quence for the task.

1.6.5 OPEN DRAWER

Filename: open_cabinet_drawer_v1.py

Task: Pull open the drawer.

Success Metric: The drawer is open at least 90% of the way, and the angular/linear velocities of the drawer link are small.



Figure 5: (Left) Example of improved dense reward function using GPT-40 for reflection. (Right) Example of improved dense reward function using AHA for reflection



Figure 6: RL policy roll-outs via improved with AHA. Row 1: pickup_YCB. Row 2: push_T. Row 3: Place_sphere. Row 4: stack_cube. Row 5: open_drawer

1.7 VLM TASK-PLAN GENERATION

In this section, we present the policy rollouts improved by AHA in Figure 8, along with the modified task plans in Figure 9.

Simulation task Details We describe each of the 3 tasks in detail, along with their PyBullet variations and success condition.

1.7.1 PUT BANANA CENTRE

367 Filename: ours_raven_ycb_pick.py
368

Task: Pick up the banana and place it onto the centre of the table.

Success Metric: The success condition on the final location of the banana with respect to the table area.
 372

373 1.7.2 STACK BANANA 374

355

356 357 358

359 360

361

362

363

364 365

366

375 Filename: ours_ycb_banana_spam_stack.py

- **Task:** Pick up the banana and place it onto the spam can.
 - Success Metric: The position of the banana should be on the spam can, and rest stably.

378	"Push T to shaped T area"	"Stack the red cub	e onto green cube"	"Put sphere into sphere holder"			
379	<pre>def compute_dense_reward(self, obs: Any, action: torch.Tensor, info: Dict):</pre>	y, action: torch.Tensor, info: Dict): def compute dense reward(self, obs: Any, & Distance from the TCP (rool center top to cubeA dist = torch.linaig.nor act will cube the section and the torch.linaig.		<pre>def compute_dense_reward(self, obs: Any, action: torch.Tensor, info: Dict):</pre>			
380	# Calculate the distance from the TCP to the T-shaped peg top_to_obj_dist = torch.linalg.norm(self.tee.pose.p = self.agent.tcp.pose.p, axisel) reaction reward = 1 = torch tanh(5 t top to obj_dist)	<pre>self.cubeA.pose.p - self.agent.tcp.p) reaching_reward = 1 - torch.tanh(5 * reward = reaching_reward</pre>	<pre>>cose.p, axis=1 * tcp_to_cubeA_dist)</pre>	<pre># Compute the distance from the TCP (gripper) to the object (sphere) top_to_obj_dist = torch.linalg.norm(self.obj.pose.p - self.agent.tcp.pose.p, szis=1) reaching reward = 1 - torch_tanh(5 * tor to obt dist)</pre>			
381	reward = reaching_reward	<pre># Reward for grasping cubeA is cubeA grasped = info["is cubeA gr</pre>	casped"]	reward = reaching_reward # Add grasping_reward if the object is being grasped			
382	# Calculate the pseudo-rendered intersection area for the T-shaped peg in the goal region intersection_area = self.pseudo_render_intersection()	reward += is_cubeA_grasped # Distance from cubeA to cubeB along	the x and y directions (ignore z)	<pre>is_grasped = info["is_obj_grasped"] reward += is_grasped</pre>			
383	intersection_reward = intersection_area / seif.intersection_thresh reward += intersection_reward	<pre>cubes_to_cubes_dist_xy = torch_lins; self.cube8.pose.p[, :2] - self.cu) stacking reward xy = 1 - torch_tanh;</pre>	<pre>(g.norm(ubeA.pose.p[, :2], axis=1 (5 * cubeA to cubeB dist xv)</pre>	<pre># Compute the distance from the object (sphere) to the bin obj_tc_bin_dist = torch.linlg.norm(self.bin.pose.p - self.obj.pose.p, axis=1) placing reward = 1 - torch.tah/(5 * obj to bin dist)</pre>			
207	# Calculate the distance from the TCP to the end-effector goal position tcp_to_goal_dist = torch.linalg.norm(self.eg_goal_pos.pose.p - self.agent.tcp.pose.p, axis=1)	reward += stacking_reward_xy * is_co # Penalty for misalignment in the z	direction	reward += placing_reward + is_grasped # Add an additional z-axis penalty to ensure precise placement into the bin			
304	end_sfreetor_goal_ceward = 1 - torch.tanh(5 * top_to_goal_dist) reward *= end_sffector_goal_reward * (intersection_area >= self.intersection_thresh)	<pre>z_offset_penalty = torch.abs(self.co self.cubeA.pose.p[, 2]) reward -= z_offset_penalty * is_cube</pre>	mes.pose.p[, 2] = m_grasped	<pre>z_offset_penalty = torch.abs(self.bin.pose.p[, 2] - self.obj.pose.p[, 2]) reward -= 5 * torch.tanh(5 * z_offset_penalty) * is_grasped</pre>			
385	<pre># Adding a penalty term for misalignment along the x direction tcp_to_obj_x_dist = torch.abs(seif.tee.pose.p(:, 0) -</pre>	<pre># Reward for keeping cubeA static is_cubeA_static = info["is_cubeA_sta static_reward = 1 - torch.tanh(5 *</pre>	stic"]	<pre># Add reward if the object is on the bin and the robot is static reward += info["is_obj_static"] * info["is_obj_on_bin"]</pre>			
386	<pre>self.agent.tcp.pose.p(;, 0)) alignment_penality = -0.5 * torch.tanh(5 * tcp_to_obj_x_dist) reward += alignment_penality</pre>	<pre>torch.linalg.norm(self.agent.robot.get_q reward += static_reward * info["is_c f Final success reward</pre>	<pre>typel()[, :-2], axis=1)) typeA_on_cubeB"]</pre>	# Set reward to maximum if the task is successfully completed reward[info["success"]] = 13 return reward			
387	<pre># Reward for achieving the task success condition reward[info["success"]] = 3.0</pre>	<pre>reward[info["success"]] = 8 return reward</pre>					
388	"Distant the		"Distant the sec				
389	def compute dense reward(self, obser Any,	, action: torch.Tensor, info: Dict):	def compute_dense_reward(self, obs: Any, act # Compute the distance between the TCP	STATC DOTTIE :: ion: torch.Tenser, infe: Dict): tool center woith! and the handle			
390	<pre>top_to_obj_dist = torch.limitg.m self.obj.pose.p - self.agent</pre>	orm(.tcp.pose.p, axis=1	<pre>top_to_handle_dist = torch.linalg.norm(</pre>	.tcp.pose.p. axis=1			
391	reaching reward = 1 - torch.tanh reward = reaching_reward # Beaard for graming the object	(5 * tcp_to_obj_dist)	reward = reaching_reward # Compute the reward for grasping the he				
392	i_grasped = info('i_grasped') reward += is_grasped		reward += is_grasped # Compute the distance the drawer has by	an pulse out			
393	# Introducing a penalty for miss: missed grasp penalty = torch.when reward += missed_grasp_penalty	ing the grasp re(is_grasped == 0, -0.5, 0.0)	<pre>self.handle_link.joint.limits(, 1) open_reward = 1 - torch.tanh(5 * (self.s))</pre>	<pre>- self.head_link.joint.linte[, 0] - self.head_link.joint.linte[, 0] in_open_frac - open_frac)]</pre>			
394	# Reward for grasping with proper rotation alignment) rotation_alignment = 1 - torch.tr	r rotation (introducing a reward for anh(5 * torch.abs(self.agent.tcp.pose.q -	# Encourage a progressive motion of open drawer's open_frac is increasing	ing the drawer: add positive reward only if			
395	<pre>self.obj.pose.qj.sum() reward += rotation_alignment * i f introducing a penalty for misa.</pre>	s_grasped lignment during grasp	drawer_change = open_frac - getattr(sell progressive_reward = drawer_change.where reward += progressive_reward + 10 # Mul selfprev_open_frac = open_frac	, "grew open [fee", open_Eced) (derwer_champs >0, torch: Lamoor(0.0)) tiplier adjusted to emphasize its importance			
396	misalignment_penalty = torch.whe reward += misalignment_penalty =	re(rotation_alignment < 0.5, -0.5, 0.0) is_grasped	<pre># Compute the static reward to encourage link is_static = { Tarch !</pre>	the agent to keep the drawer open			
397	<pre># Distance from object to goal objto_goal_dist = torch.linalg. </pre>	norm(.obj.pose.p, axis=1	<pre>torch.linaig.norm(self.handle_link) & (torch.linaig.norm(self.handle_link. static_reward = 1 - torch.tanh(5 * (1 - reward += static_reward * info["open_enc</pre>	mguas_vescoscy, akie] <= 1 linas_velocity, akie] <= 0.1) link_is_tatic.float()) ggh [*]]			
398	place reward = 1 - torch tanks reward += place reward * ianks # Reward for 1157the the object	* obj_to_goal_dist) ped	# Penalty for y-direction offset movement y_offset_penalty = torch.abs(self.agent.	t after grasping top.pose.p[1, 1] - infe["handle_link_pos"][;,			
399	<pre>if it is the set of the set</pre>	1.0, 0.0 Id height for lifting	<pre>1)) reward -= y_offset_penalty * is_grasped</pre>				
400	<pre>reward += lift_reward * is_grasp # Final success reward reward(info("success")] = 6</pre>	ed	<pre># Add a large reward if the task is succ reward(info["success"]) = 5.0 return reward</pre>	esfully completed			
401	return reward						
402	Figure 7: Ex	amples of modif	ied reward fund	tion via AHA			
402							
403							
405	1.7.3 STACKS CUBES						
406	F ¹						
407	Fliename: ours_raven_bow	1_stack.py					
408 409	Task: Pick up the green cube and it on top of the green.	l place into the g	reen bowl, and th	en take the yellow cube and stack			
410	Success Metric: When the yello	w cube is stably	stack on top of tl	ne green in the green bowl.			
412	1.8 VLM SUB-TASK VERIFIC	CATION					
414	In this section, we leverage Ma	nipulate-An	ything (Duan	et al., 2024) as the main policy			
415	framework, integrating it with A	HA. AHA functi	ions as a sub-tas	k verifier VLM, playing a crucial			
415	role in ensuring task success whe	n using Manipu	late-Anythi	ng. Examples of the roll-outs are			
/117	shown in Figure 10.	- 1	-	-			
417	- Simulation task Datails We does	ribe each of the 1	tacks in datail at	ong with their DI Ranch variations			
410	and success condition	1100 cach of the 4	usis in uctail, al	ong with their KEDenell variations			
419	and success condition.						
420							
421	1.8.1 PUI BLOCK						
422 423	Filename: put_block.py						
424	Task: Pick up the green block ar	nd place it on the	red mat.				
425 426	Success Metric: The success con	ndition on the rec	d mat detects the	target green block.			
427 428	1.8.2 PICKUP CUP						
429	Filename: pickup_cup.py						
430 431	Task: Pick up the red cup.						
	Success Metric: Lift up the red	cup above the pro	e-defined location	n.			



Figure 8: TAMP policy roll-outs via improved with AHA. Row 1: put_banana_centre. Row 2: stack_banana. Row 3: stack_cubes

458 1.8.3 SORT MUSTARD

454

455

456 457

470 471

472

- 459
 460
 Filename: sort_mustard.py
- **Task:** Pick up the yellow mustard bottle, and place it into the red container.
- 462 **Success Metric**: The yellow mustard bottle inside red container.
- 464 1.8.4 PICK & LIFT
- **466 Filename:** pick_and_lift.py
- **Task:** Pick up the red cube.
- 469 **Success Metric**: The red cube is lifted up.
 - 1.9 LIMITATIONS AND OPPORTUNITIES
- While AHA is quite capable, it is not without limitations. In the following sections, we discuss some of these limitations and potential risks for real-world deployment.
- 475 Does not cover all failure cases. While AHA is instruction-tuned to address the seven most common 476 failure modes in robotic tabletop manipulation, and provides fine-grained failure identification down 477 to axis-level translation or orientation offsets of the gripper, it is inevitable that certain edge cases remain uncovered. These edge cases fall outside the distribution of failure modes considered in our 478 approach and may occur in real-world scenarios. To address this limitation, future work could focus 479 on distilling policies already trained with action-based "artifacts" to generate naturally occurring 480 failures. These could then be labeled using human annotations or a hybrid approach combining AHA 481 and a GPT-4o-based labeling system. Such efforts could enhance the robustness and applicability of 482 failure detection across a broader range of scenarios. 483
- Improving policies directly with AHA Currently, we have demonstrated AHA's effectiveness
 in enhancing existing robotics systems that utilize off-the-shelf VLMs/LLMs to automate various
 pipeline components. AHA has shown significant impact in improving code generation, task-plan

"Stack banana	onto spam can"	"Stack two cubes into a bowl"		
Thank you for the clarification. Given that the robot gripper mo we need to account for this offset when placing the banana on to Let's assume the offset is a known constant, 'GRIPPER_OFFSET_Z'. Here's the updated 'gen_plan' and 'gen_domain' functions:	ves to the desired position with an offset along the z direction, p of the potted meat can. We'll use this offset to adjust the placement pose accordingly.	These you for pointing that out. Let's correct the plan to account for the z-offset when placing the second cube. ""Python Gef gen_land(initial:RavemBelier, x_offset:float, y_offset:float, z_offset:float): plan =[]		
``_python ORIPPER_OFFSET_Z = 0.02 ∦ Example offset value, adjust as neces	sary	# Pick up the green block green, block = initial.objects["mojpet_6"] plum_energical data of debug argue block and a pair block		
def gen_plan(initial: RavenBelief, grasp: RavenGrasp): plan = []		plan approximation could gate / green_stock.powr.powr.y/ # Place the green block into the green book mean howitelts_howing for the for		
<pre># Pick up the banana plan.append(Action("pick", ["object_1", grasp])) </pre>		grven_st = green_bowl.pojectgrouper x, y, z = green_bowl.point plan.append(Aclon('place', [x + x_offset, y + y_offset, z + z_offset]))		
<pre>a det the pose of the ported way that the spam_pose = initial.cojects["doject_0"].pose # Calculate the placement pose for the banana on top of the</pre>	potted meat can	<pre># Pick up the yellow block yellow_block = initial.objects["object_4"] plan.approAction("pick", yellow_block.pose.point))</pre>		
ummanic_pister_pose = movembose(x=spam_pose.x, y=spam_pose.y, z=spam_pose.z + BLOCK_SIZE + ORIPPER_OFFSET_Z, # Adjust	for the gripper offset	# Stack the yellow block onto the green block green_block_new_pose = [x + x.offset, y + y.offset, z + z.offset + BLOCK_SIZE]		
pitch=spam_pose.pitch, yaw=spam_pose.yaw)		<pre>refrum blau refrum blau r</pre>		
<pre># Place the banana on top of the potted meat can plan.append(Action(*place*, ["object_1", grasp, banana_place</pre>	_pose]))	<pre>def gen_domain(initial:RavenBellef): return {</pre>		
return plan def gen_domain(initial: RavenBelief):		"y_offSet": ContinuousSampler(-0.02, 0.02), "z_offSet": ContinuousSampler(0.01, 0.03), 		
"grasp": GraspSampler()		<pre>imail_pulanition i. **Pick up the green block**: The plan starts by picking up the green block ('object_6'). 2. **Place the green block into the green bowl*: The green block is then placed into the green bowl ('object_5'). The offsets</pre>		
### Explanation		1 2019st, yourset, yourset, attom for some walability in the exact platement will be used. 3. "Bick the yellow block onto the green block": Enalty, the yellow block is placed on top of the green block, which is now in the book. The high is adjusted by adding be BOGK-31E to be z-coordinate of the green block's memorying the block.		
 "Picking the Banana": The plan starts by picking up the ban 2. "Calculating Placement Pose": The pose of the potted meat c pose for the banana is calculated by adding the height of the ba (CapTODE DFOFF) to the 1 second part of the part meat and 	ana ('object_1') using a grasp sampled from the 'GraspSampler'. an ('object_0') is retrieved from the initial state. The placement nana (approximated as 'BLOK, SIZE') and the gripper offset once The roll pick and with one heat the come to expert preserve	Heiden stock as alsocked on texts.		
alignment. 3. "*Placing the Banana**: The plan then places the banana at th	e calculated pose on top of the potted meat can.	 The 'x_offset', y_offset', and 'z_offset' parameters allow for slipht variations in the placement of the green block within the bool, ensuring it fits regulations of minor positional differences? The yellow block is always placed directly on top of the green block, ensuring a stack is formed. The weather for the offset is chosen to be commit accounts are block as an interd correctly but large another to green the start is an interd correctly but large another to green the start is an interd correctly but large another to green the start is an interd correctly but large another to green the start is an interd correctly but large another to green the start is an interd correctly but large another to green the start is an interd correctly but large another to green the start is an interd correctly but large another to green the start is an interd correctly but large another to green the start is an interd correctly but large another to green the start is an interd correct but the start is another to green the start is an interd correctly but large another to green the start is an interd correctly but large another to green the start is an interd correctly but large another to green the start is an interd correct by the large another to green the start is an interd correct but the large another to green the start is an interd correct by the large another to green the start is an interd correct by the large another to green the start is an interd correct by the large another to green the start is an interd correct by the large another to green the start is an interd correct by the large another to green the start is an interd correct by the start is another to green the start is an interd correct by the large another to green the start is an interd correct by the start is another to green the start is an interd correct by the start is an interd		
www.Why This Will Always Satisfy the Goal	ensuring that a valid grass is used for micking up the bagana	The ordinal lifty.		
 "Placement Calculation": The placement pose is calculated ba gripper offset, ensuring that the banana is placed directly on t "Generalization": The plan is general enough to handle diffe 	set on the current pose of the potted meat can and adjusted for the op of it. rent initial positions and orientations of the objects, as it	If the offsets are too large, the green block might not be placed correctly within the bowl, or the yellow block might not stack properly on the green block.		
This approach ensures that the banana will always be stacked on	top of the potted meat can, accounting for the gripper offset along	 This domain for the orivers is chosen to minimize these fixes, but if a failure occurs, the orisets can be arguined to smaller ranges. This plan should achieve the goal of placing the green block into the bowl and then stacking the yellow block onto the green block, there are an arguined to show the green block. 		
the 2 direction.		given une inicial state and une versien vonain.		
	"Move bahar	ha to centre.		
	We will add an additional parameter to the 'gen_plan' function to position.	ng the object we can adjust the posterior postation accordingly. pecify the y offset and use it to correct the placement		
	Were's the updated implementation:			
	<pre>def gen_plan(initial: RavenBelief, grasp: RavenGrasp, y_offset: fl plan = [] object_name = "object_0" # The banana is identified as "objec plan amend/Artico("pick" inhight page or available</pre>	oat): t_0*		
	<pre>a Adjust the y position by the offset when placing the object adjusted_y = TABLE_CENTER[1] + y_offset plan.append(Action("place", [object_name, grasp, RavenPose(x=T. extension)</pre>	ABLE_CENTER[0], y=adjusted_y, z=0.02)]))		
	def gen_domain(initial: RavenBelief):			
	"grasp": GraspSampler(), "y_offset": ContinuousSampler(-0.05, 0.05) # Allowing for 	a small offset range		
	### Explanation:			
	 The function now takes an additional parameter 'y_offset' to When placing the object, the y position is adjusted by adding 	account for the slight offset along the y direction. the 'y_offset' to the table center's y coordinate.		
	 gen_domain Function: The function returns a dictionary with two entries: 'grasp': Sampled using the 'GraspSampler'. 'y offset': Sampled using a 'ContinuousSampler' with a range 	e of -0,85 to 9,05, allowing for a small adjustment in the v		
	direction. ann Why This Will Always Satisfy the Goal:			
	 The plan still explicitly picks up the banana and places it at to offset. The 'y_offset' parameter allows for correcting the placement pos. The table conter coordinates are fixed and known, ensuring the b 	he center of the table, but now with an adjustment for the y liken to account for the slight offset during grasping. mana is always nowed to the correct location with the necessary		
	adjustment. This approach ensures that the goal is achieved regardless of the	specific grasp sampled and the slight offset during grasping, as		
	long as the offset is within the specified range.			
_				
F	figure 9: Examples of mo	odified task-plan via AHA		
generation, and sub-task	success detection. Howe	ver, AHA is not yet capable of directly influencing		
training low level policie	es unrougn failure languag	ge reasoning. A potential next step would involve		
corrective actions tied t	o various failure modes	This would enable low-level policies to interpret		
failure reasoning in a co	unterfactual manner and	generate corrective actions directly		
iunure reasoning in a co	anternaetuur manner anu	generate concentre actions uncerty.		
References				
Lafe Duan Wester V	we William David N	De Ware King Ebeen: D' () Estat D		
Krishna Manipulata	an, wildert Pumacay, Yi	Ku wang, Kiana Ensani, Dieter Fox, and Ranjay		
nreprint arYiw 2406	18915 2024	-world robots using vision-ranguage models. <i>arxiv</i>		
ртертни игліv.2400.1	0713, 2027.			
Kristen Grauman And	rew Westhury Fugene F	Byrne Zachary Chavis Antonino Eurnari Pohit		
Girdhar Jackson Ham	burger. Hao Iiang Miao l	Lin. Xingyn Lin. et al. Ego4d. Around the world in		
3,000 hours of egocen	tric video. In Proceeding	s of the IEEE/CVF Conference on Computer Vision		
and Pattern Recogniti	<i>ion</i> , pp. 18995–19012, 20	122.		



Figure 10: Examples of zero-shot data generator trajectories with AHA as sub-tasks verifier. Row 1: pickup_cube, pickup_cup. Row 2: put_block, sort_mustard

Stephen James, Zicong Ma, David Rovick Arrojo, and Andrew J Davison. Rlbench: The robot learning benchmark & learning environment. *IEEE Robotics and Automation Letters*, 5(2):3019–3026, 2020.

Alexander Khazatsky, Karl Pertsch, Suraj Nair, Ashwin Balakrishna, Sudeep Dasari, Siddharth Karamcheti, Soroush Nasiriany, Mohan Kumar Srirama, Lawrence Yunliang Chen, Kirsty Ellis, et al. Droid: A large-scale in-the-wild robot manipulation dataset. *arXiv preprint arXiv:2403.12945*, 2024.

Zeyi Liu, Arpit Bahety, and Shuran Song. Reflect: Summarizing robot experiences for failure explanation and correction. *arXiv preprint arXiv:2306.15724*, 2023.

Abhishek Padalkar, Acorn Pooley, Ajinkya Jain, Alex Bewley, Alex Herzog, Alex Irpan, Alexander Khazatsky, Anant Rai, Anikait Singh, Anthony Brohan, et al. Open x-embodiment: Robotic learning datasets and rt-x models. *arXiv preprint arXiv:2310.08864*, 2023.