

## 439 A Environment Details

440 Here we provide additional details for our simulation environments. An unabridged version of the  
441 description from Section 5 is as follows:

- 442 • **Push (2D)**: Push a round puck using the tool such that it stops at the specified goal loca-  
443 tion. The goal space is a subset of 2D final puck locations  $\mathcal{G} \subset \mathbb{R}^2$ , and the control action  
444 space  $\mathcal{A}_C \in \mathbb{R}^2$  specifies the  $x$  and  $y$  tool velocities.
- 445 • **Catch balls (2D)**: Use the tool to catch three balls that fall from the sky. The agent’s  
446 goal is to catch all three balls, which start from varying locations on the  $x$ - $y$  plane. We use  
447 a 1-dimensional control action space that specifies the  $x$  velocity of the tool at each step.
- 448 • **Scoop (2D)**: Use the tool to scoop balls out of a reservoir containing 40 total balls. Here  
449 we specify goals of scooping  $x \in \{1, 2, \dots, 7\}$  balls. The control action space  $\mathcal{A}_C \in \mathbb{R}^3$   
450 specifies the velocity of the rigid tool in  $x$  and  $y$  directions, along with its angular velocity.
- 451 • **Fetch cube (3D)**: Use the tool to retrieve an object randomly positioned beneath a ver-  
452 tical overhang. This task is additionally challenging because the position of the robot end  
453 effector is restricted to a rectangular region in the  $x$ - $y$  plane of dimensions  $0.8\text{m} \times 0.2\text{m}$  to  
454 avoid collision with the overhang. The tool is a three-link chain where each link is a box  
455 parameterized by its width, length, and height. The design space also includes the relative  
456 angle between two connected links, with a total of  $n = 11$  parameters. The control action  
457 space  $\mathcal{A}_C \in \mathbb{R}^3$  represents a change in end-effector position.
- 458 • **Lift cup (3D)**: Use the tool to lift a cup of randomized geometry from the ground into  
459 the air. This task requires careful design of the tool to match cup geometry. The tool is a  
460 four-link fork with two prongs parameterized by the separation, tilt angle, width, length,  
461 and height of the prongs. The same parameters are applied to both prongs to maintain  
462 symmetry. The handle dimensions are fixed. The design space has  $n = 5$  parameters.  
463  $\mathcal{A}_C \in \mathbb{R}^3$  represents a change in end-effector position.
- 464 • **Scoop (3D)**: A 3D analog of the 2D scoop task. This task has the same goal space as the  
465 2D scoop task, but the tool in 3D is a six-link scoop composed of a rectangular bottom  
466 plate parameterized by its width and length, and four rectangular side plates attached to  
467 each side of the bottom plate. Each side plate is parameterized by its height and relative  
468 angle to the bottom plate. A handle with fixed dimensions is attached to one of the side  
469 plates. There are  $n = 10$  total design parameters.  $\mathcal{A}_C \in \mathbb{R}^6$  represents a change in end-  
470 effector pose.

## 471 B Experimental Details

### 472 B.1 Training Hyperparameters & Architecture Details for Our Framework

473 In Tables 5, 6, 7, 8, 9, and 10, we provide detailed hyperparameters for our framework for each  
474 environment. Unless otherwise specified, we use the neural network architectures for the design  
475 policy, control policy, and value function from [7].

### 476 B.2 Training Hyperparameters & Architecture Details for Baselines

477 For the CMA-ES baseline, we perform hyperparameter sweeps for a fair comparison with our frame-  
478 work. For the CMA-RL baseline, we use the same set of best performing hyperparameters for the  
479 outer CMA-ES loop. The tested hyperparameter configurations for each baseline are listed in Ta-  
480 ble 2. Except model architecture differences, we use the same optimization hyperparameters for  
481 Ours, Ours(shared arch.), and HWasP-minimal.

### 482 B.3 Generalization to Unseen Goals Experiment Details

483 For **Fetch cube**, the rectangular region of initial poses is defined by  $x \in [-0.395, 0.395]$  and  
484  $y \in [0.4, 0.7]$ . The cutout region corresponds to two disconnected rectangular patches contained in

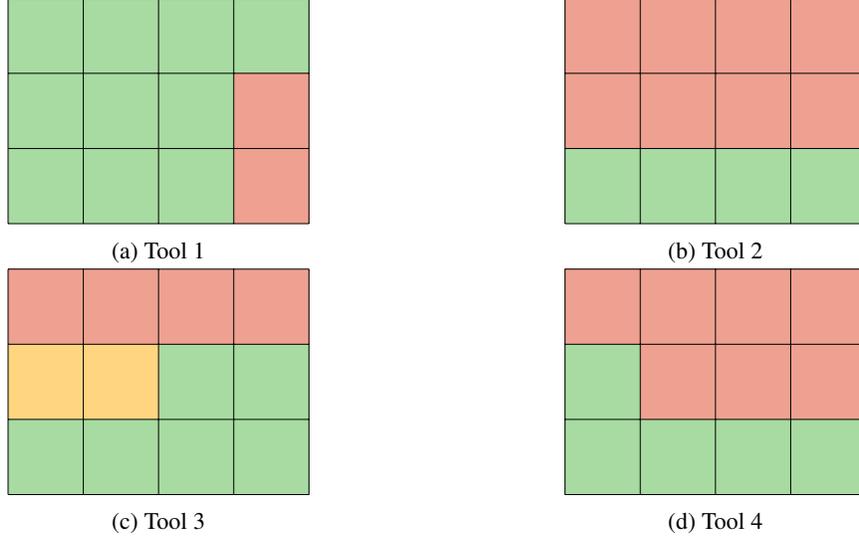


Figure 10: Heatmaps of success and failures for trials where fixed tools are used for a range of initializations. The grid here directly maps to the set of 2D initial cube locations tested for each tool, from the point of view of the robot. Here **green** indicates a success, **red** indicates failure, and **orange** indicates a failure where the final cube position is within 5cm of success.

Method	Hyperparameters	Values
CMA-ES	Population Size	10, <b>24</b> , 100, 1000
	Initial Stdev	<b>0.1</b> , 1.0, 10.0
	Center Learning Rate	0.01, 0.1, <b>1.0</b>
	Covariance Learning Rate	0.01, 0.1, <b>1.0</b>
	Rank $\mu$ Learning Rate	0.01, 0.1, <b>1.0</b>
	Rank One Learning Rate	0.01, 0.1, <b>1.0</b>
CMA-RL	Poicy Net	(256, 256, 256, 256, 256)
	Value Net	(256, 256, 256, 256, 256)
	Learning Rate	3e-4
	Batch Size	(50000, 20000(3D scoop))
	Minibatch Size	2000
Ours(shared arch.)	Poicy Net	(256, 256, 256, 256, 256)
	Value Net	(256, 256, 256, 256, 256)
	Learning Rate	1e-4
	Batch Size	(50000, 20000(3D scoop))
	Minibatch Size	2000

Table 2: We tune over these values for hyperparameters of baseline methods. Bolded values indicate the best performing settings for CMA-ES, which we use in our comparisons.

485 the training region defined by  $x_1 \in [-0.350, -0.045]$ ,  $y_1 \in [0.434, 0.666]$  and  $x_2 \in [0.045, 0.350]$ ,  
 486  $y_2 \in [0.434, 0.666]$  respectively.

487 **Zero-shot performance.** We train six policies using our framework where the cutout region re-  
 488 moves a fraction of the total training area equal to 0.1, 0.2, 0.4, 0.6, 0.8, and 0.9 respectively.

489 **Fine-tuning performance.** For the fine-tuning experiment, we specifically select  
 490 four initializations that we find our policies do not complete successfully zero-shot:  
 491  $\{(-0.167, 0.367), (-0.129, 0.357), (0.430, 0.493), (0.415, 0.610)\}$ .

492 **B.4 Trading off Design and Control Complexity Experiment Details**

493 We train four agents independently on `catch balls`, setting the value of  $\alpha$ , the tradeoff reward  
 494 parameter defined in Equation 1, to 0, 0.3, 0.7, and 1.0 respectively.

495 **B.5 Real Robot Experiment Details**

496 For our real-world experiments, we use a Franka Emika Panda arm. We control the robot using an  
 497 impedance controller from the Polymetis [38] library.

498 Tools are 3D printed using polylactic acid (PLA) on  
 499 commercially available Ender 3 and Ender 5 printers  
 500 with nozzle diameter 0.4mm. We print using a layer  
 501 height of 0.3mm and 10% infill. We perform slicing  
 502 using the Ultimaker CURA software.

503 We roll out each policy for 100 environment steps or  
 504 until a success is detected.

505 For the `fetch cube` task, we measure the success  
 506 based on whether the center of mass of the 5cm cube  
 507 is closer than 0.5m from the base of the robot. Please  
 508 see Table 3 for per-tool details. The tool images are  
 509 shown in Figure 9, from left to right: Tools 4, 2, 3, 1  
 510 respectively.



Figure 11: Real world setup. We use a Franka Panda arm and five RealSense D435 cameras for tracking.

Tool	Initial cube position (x, y, z)
Tool 1	(-0.110, -0.803, 0.025)
Tool 2	(-0.339, -0.588, 0.025)
Tool 3	(0.155, -0.731, 0.025)
Tool 4	(0.211, -0.633, 0.025)

Table 3: Initial cube positions corresponding to tools fabricated in real experiments.

511 For the `lift cup` task, we measure success based on whether the cup has been lifted higher than  
 512 0.4m off of the plane of the workspace. Please see Table 4 for per-tool details. The tool images are  
 513 shown in Figure 9, from left to right: Tools 1, 2, 3, 4.

Tool	Cup geometry parameters (length/width, height)
Tool 1	(0.3, 0.6)
Tool 2	(0.3, 0.9)
Tool 3	(0.5, 0.8)
Tool 4	(0.9, 0.6)

Table 4: Cup geometry parameters corresponding to tools fabricated in real experiments. Note that the length and width parameters share a single value.

514 We also present detailed results for the `fetch cube` experiments using tools generated for a specific  
 515 initial position for a range of initializations. Recall that we test the policies on a  $3 \times 4$  grid of initial  
 516 positions that span a range of  $12\text{cm} \times 85.6\text{cm}$ , for a total of 12 trials per tool. We plot the successes  
 517 and failures for each tool according to geometric position in Figure 10. We can see that the control  
 518 policy is able to use each tool to solve the task for several initializations, but each tool is specialized  
 519 for particular regions.

<b>Hyperparameter</b>	<b>Value</b>
Tool Position Init.	(20, 10)
Control Steps Per Action	1
Max Episode Steps	150
Slack Reward	-0.001
Tool Length Ratio	(-0.5, 0.5)
Tool Length Init.	(2.0, 2.0, 2.0)
Tool Angle Init.	(0.0, 0.0, 0.0)
Tool Angle Ratio	(-1.0, 1.0)
Tool Angle Scale	90.0
Control GNN	(64, 64, 64)
Control Index MLP	(128, 128)
Design GNN	(64, 64, 64)
Design Index MLP	(128, 128)
Control Log Std.	-1.0
Design Log Std.	-2.3
Fix Design & Control Std.	True
Policy Learning Rate	2e-5
Entropy $\beta$	0.01
Value Learning Rate	1e-4
KL Divergence Threshold	0.005
Batch Size	50000
Minibatch Size	2000
PPO Steps Per Batch	10

Table 5: Hyperparameters used for our framework on the push task.

<b>Hyperparameter</b>	<b>Value</b>
Tool Position Init.	(20, 10)
Control Steps Per Action	1
Max Episode Steps	150
Slack Reward	-0.001
Tool Length Ratio	(-0.5, 2.0)
Tool Length Init.	(2.0, 1.0, 1.0)
Tool Angle Init.	(0.0, 0.0, 0.0)
Tool Angle Ratio	(-1.0, 1.0)
Tool Angle Scale	60.0
Control GNN	(64, 64, 64)
Control Index MLP	(128, 128)
Design GNN	(64, 64, 64)
Design Index MLP	(128, 128)
Control Log Std.	0.0
Design Log Std.	0.0
Fix Design & Control Std.	True
Policy Learning Rate	2e-5
Entropy $\beta$	0.01
Value Learning Rate	1e-4
KL Divergence Threshold	0.002
Batch Size	50000
Minibatch Size	2000
PPO Steps Per Batch	10

Table 6: Hyperparameters used for our framework on the catch balls task.

<b>Hyperparameter</b>	<b>Value</b>
Tool Position Init.	(15, 10)
Control Steps Per Action	5
Max Episode Steps	30
Slack Reward	-0.001
Tool Length Ratio	(-0.7, 0.2)
Tool Length Init.	(6.0, 3.0, 3.0)
Tool Angle Init.	(0.0, 0.0, 0.0)
Tool Angle Ratio	(-0.1, 0.7)
Tool Angle Scale	90.0
Control GNN	(64, 64, 64)
Control Index MLP	(128, 128)
Design GNN	(64, 64, 64)
Design Index MLP	(128, 128)
Control Log Std.	0.0
Design Log Std.	0.0
Fix Design & Control Std.	True
Policy Learning Rate	2e-5
Entropy $\beta$	0.01
Value Learning Rate	3e-4
KL Divergence Threshold	0.1
Batch Size	50000
Minibatch Size	2000
PPO Steps Per Batch	10

Table 7: Hyperparameters used for our framework on the `scoop` task.

<b>Hyperparameter</b>	<b>Value</b>
Tool Position Init.	(0.0, 0.5, 0.02)
Control Steps Per Action	10
Max Episode Steps	100
Slack Reward	-0.001
Success Reward	10.0
Box Dimensions Min	(0.005, 0.05, 0.005)
Box Dimensions Max	(0.015, 0.1, 0.02)
Tool Angle Min	(-90.0, -90.0, -90.0)
Tool Angle Max	(90.0, 90.0, 90.0)
Control Action Min	(-1.0, -1.0, -1.0, -0.2, -0.2, -0.2)
Control Action Max	(1.0, 1.0, 1.0, 0.2, 0.2, 0.2)
Control Action Scale	0.1
Control GNN	(128, 128, 128)
Control Index MLP	(128, 128)
Design GNN	(128, 128, 128)
Design Index MLP	(128, 128)
Control Log Std.	0.0
Design Log Std.	0.0
Fix Design & Control Std.	False
Policy Learning Rate	1e-4
Entropy $\beta$	0.0
Value Learning Rate	3e-4
KL Divergence Threshold	0.5
Batch Size	50000
Minibatch Size	2000
PPO Steps Per Batch	10

Table 8: Hyperparameters used for our framework on the `fetch_cube` task.

<b>Hyperparameter</b>	<b>Value</b>
Tool Position Init.	(0.0, 1.2, 0.05)
Control Steps Per Action	10
Max Episode Steps	150
Slack Reward	-0.001
Success Reward	10.0
Box Dimensions Min	(0.005, 0.02, 0.01)
Box Dimensions Max	(0.01, 0.1, 0.03)
Tool Angle Min	(-30.0, -30.0, -30.0)
Tool Angle Max	(30.0, 30.0, 30.0)
Control Action Min	(-1.0, -1.0, -1.0, -1.57, -1.57, -1.57)
Control Action Max	(1.0, 1.0, 1.0, 1.57, 1.57, 1.57)
Control Action Scale	0.1
Control GNN	(128, 128, 128)
Control Index MLP	(128, 128)
Design GNN	(128, 128, 128)
Design Index MLP	(128, 128)
Control Log Std.	0.0
Design Log Std.	-1.0
Fix Design & Control Std.	True
Policy Learning Rate	2e-5
Entropy $\beta$	0.01
Value Learning Rate	3e-4
KL Divergence Threshold	0.5
Batch Size	50000
Minibatch Size	2000
PPO Steps Per Batch	5

Table 9: Hyperparameters used for our framework on the `lift` cup task.

<b>Hyperparameter</b>	<b>Value</b>
Tool Position Init.	(0.0, 0.05, 0.1)
Control Steps Per Action	10
Max Episode Steps	100
Slack Reward	-0.001
Success Reward	10.0
Box Dimensions Min	(0.04, 0.005, 0.02)
Box Dimensions Max	(0.08, 0.005, 0.05)
Tool Angle Min	(-15.0, -15.0, -15.0)
Tool Angle Max	(15.0, 15.0, 15.0)
Control Action Min	(-1.0, -1.0, -1.0, -1.57, -1.57, -1.57)
Control Action Max	(1.0, 1.0, 1.0, 1.57, 1.57, 1.57)
Control Action Scale	0.05
Control GNN	(128, 128, 128)
Control Index MLP	(128, 128)
Design GNN	(128, 128, 128)
Design Index MLP	(128, 128)
Control Log Std.	0.0
Design Log Std.	0.0
Fix Design & Control Std.	False
Policy Learning Rate	1e-4
Entropy $\beta$	0.01
Value Learning Rate	3e-4
KL Divergence Threshold	0.5
Batch Size	20000
Minibatch Size	2000
PPO Steps Per Batch	5

Table 10: Hyperparameters used for our framework on the 3D scoop task.