

# TRUSTDeHANDS: A MASSIVELY PARALLEL BENCHMARK FOR SAFE DEXTEROUS MANIPULATION

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Safe Reinforcement Learning (Safe RL) aims to maximize expected total rewards meanwhile avoiding violating safety constraints. Although a plethora of safety-constrained environments have been developed to evaluate Safe RL methods, most of them focus on navigation tasks, which are rather simple and have non-trivial gap with real-world applications. For robotics studies, dexterous manipulation is becoming ubiquitous; however, the idea of safe dexterous manipulations are rarely studied in robotics applications. In this paper, we propose **TrustDeHands**, a massively parallel benchmark for Safe RL studies on safe dexterous manipulation tasks. TrustDeHands is built within the Isaac Gym, a GPU-level parallel simulator that enables highly efficient RL training process. To stay close to real world settings, TrustDeHands offers multi-modal visual inputs, including RGB, RGB-D and point cloud, and supports a variety of arms and dexterous hands from different brands. Moreover, TrustDeHands provides a solid implementation of eight popular safe policy optimization algorithms; this facilitates trustworthy validation for Safe RL methods outside navigation tasks. TrustDeHands include a myriad of challenging tasks that require safety awareness (e.g., Jegna). Results on these tasks show that Safe RL methods can achieve better performance than classical RL algorithms, indicating the effectiveness of Safe RL in safe robot manipulation tasks. To our best knowledge, TrustDeHands is the first benchmark targeting at safe dexterous manipulation. We expect this benchmark to consistently serve as a reliable evaluation suite for future Safe RL developments and further promote the integration between the lines of research of Safe RL and dexterous manipulation. The code and demonstration can be found at <https://sites.google.com/view/trustdehands/>.

## 1 INTRODUCTION

Reinforcement Learning (RL) is a powerful way to solve sequence decision problems and has achieved superhuman performance in games (Silver et al., 2016; 2017; Vinyals et al., 2019; OpenAI, 2018), robotic (Andrychowicz et al., 2020; Chen et al., 2022b), and financial (Hambly et al., 2021). Before the RL deploy in the real world, researchers are tasked with proving their trustworthiness to maximize the benefits of AI systems while minimizing their risks (Xu et al., 2022). The fundamental principle of RL is that an agent tries to maximize the cumulative returns by trial and error, but the agents may play dangerous or harmful behaviors during the learning process. Thus, it is important to consider safe exploration that is known as safe RL. Most existing RL simulators usually ignore safety learning issue, where failure is acceptable and even desirable to learn from bad outcomes. But in the real world, such exploration can produce undesired miseries.

In recent years, Robot manipulation (Billard & Kragic, 2019) is an important direction for the application of RL, which covers many reserach (Kim et al., 2021; Okamura et al., 2000; Kumar et al., 2016). Among them, dexterous multi-fingered manipulation is the most challenging task, which puts forward higher requirements for control (Bircher et al., 2017; Rahman et al., 2016). To deal with dynamic environment and policy generalization, OpenAI et al. (2019); Chen et al. (2022a) have studied dexterous manipulation tasks to achieve some significant results. Moreover, trustworthiness of the manipulation (Xu et al., 2022) is an important issue needed to be considered. If a robot wants to manipulate in the real world, ensuring safe and trustworthy is the highest priority. But in previous researches, there is no benchmark for safe manipulation. So we propose TrustDeHands, which uses safe policy learning to learn dexterous manipulation, hoping to fill this research gap.

Additionally, many existing Safe RL methods is unavailable, which result in researchers suffer from incorrect implementations, unfair comparisons and misleading conclusions. Safe policy learning is critical in real-world RL applications, where dangerous decisions are undesirable. For example, a robot agent should avoid taking actions that irreversibly damage its hardware (Ray et al., 2019a; Dulac-Arnold et al., 2019). Due to its importance, the community has been actively researching safe policy learning (e.g., (Alshiekh et al., 2018; Stooke et al., 2020b; Gu et al., 2021; Yuan et al., 2021; Gronauer, 2022; Yang et al., 2022; Liu et al., 2022)). However, most of the existing work mainly focuses on algorithm design. Among these works, either the authors did not publish the source code (e.g. P3O (Zhang et al., 2022)), or the algorithms were implemented using different frameworks (e.g. PCPO (Yang et al., 2020b) in Theano (Al-Rfou et al., 2016), CPPO-PID (Stooke et al., 2020b) in PyTorch), with divergent approaches (FOCOPS (Zhang et al., 2020b) does not parallelize sample collection while others do), and on separate tasks (FOCOPS is tested solely on MuJoCo-Velocity (Todorov et al., 2012) and CPPO-PID solely on Safety-Gym (Ray et al., 2019a)). While there exists safety-starter-agents (Ray et al., 2019a) as a publicly available collection of algorithms, it was implemented using TensorFlow1, required old hardware and system, lack recent updates, and was no longer maintained. As a result, the Safe RL community has experienced serious difficulty in reproducing the experimental results, comparing algorithms fairly, and deriving correct insights. An open-source, standardized algorithm implementation for algorithm verification and empirical study is desperately needed.

To facilitate the consideration we mentioned above, we developed a bimanual dexterous manipulation environment: **TrustDeHands**, with a unified re-implementation of Safe RL algorithms. We highlight three particularly desirable features of TrustDeHands:

- **For Safe RL researchers.** We provide a series of complex and challenging safe dexterous manipulation tasks. The design of these tasks stems from the need for safety robot manipulation in our daily life (e.g., sweeping the floor without touching other furniture). In these environments, we have done exhaustive experiments with the implemented algorithms and contributed the results, our observations, and analysis for the reference of the community.
- **For robotic researchers.** We are the first collection of tasks focused on safe dexterous manipulations. In addition to safety research, we also provide a variety of features, including 1) multi-modal information as the policy input (e.g., contact force, RGB image, RGB-D image, point cloud...). 2) customizable dexterous hands and a robotic arm drive to the dexterous hand. These features provide a comprehensive platform for robotic research.
- **Unified, highly-optimized, and extensible Safe RL algorithms.** We re-implement widely used Safe RL algorithms, which support TrustDeHands and all popular environments in a single well-designed algorithms framework. We have done maximum abstraction and encapsulation, deriving a similar model structure and update paradigm, thus enabling code reuse, ensuring a clean code style, and making it extremely extensible.

## 2 RELATED WORK

**Safe RL Environments** Simulator plays a critical role to the training for RL since it is very expensive to collect data in the real world. Safety-gym (Ray et al., 2019b) introduces a robot that has to navigate through a cluttered environment to achieve a task, which is a suite of complex continuous control environments for Safe RL. Safe-control-gym (Yuan et al., 2021) introduces cart-pole, 1D, and 2D quadrotor dynamic systems to achieve control tasks like stabilization or trajectory tracking, which allows us for constraint specification and disturbance injection onto a robot’s inputs, states, and inertial properties through a portable configuration system. AI Safety Gridworlds (Leike et al., 2017) proposes an environment for evaluating various safe properties of intelligent agents, including safe interruptibility, avoiding side effects, safe exploration, distributional shift, etc. MuJoCo-Velocity, originally proposed in (Zhang et al., 2020a), consists of a series of safety tasks like constrained velocity based on MuJoCo environment (Todorov et al., 2012). However, there still lacks a safe environment for safe robot manipulation, which the difficulty lies in requiring safe high-dimensional continuous space control and dealing with the dynamic environment. So we introduce TrustDeHands, which aims to apply Safe RL to dexterous manipulation, providing a more challenging environment for evaluating Safe RL algorithms.

**Safe RL Algorithms** Since we formulate safe RL under CMDPs (Altman, 1999), in this section we mainly review algorithms w.r.t. CMDPs. For more discussions about safe RL algorithms, please re-

fer to the recent survey (Xu et al., 2022; Gu et al., 2022). With the rise of deep RL, CMDPs are also moving to more high-dimensional continuous control problems. CPO (Achiam et al., 2017b) proposes the general-purpose policy search algorithm for Safe RL with guarantees for near-constraint satisfaction at each iteration. PCPO (Yang et al., 2020a) utilizes a different two-step approach (i.e. first finds the policy with the maximum return, then projects this policy back into the safety region in terms of the minimum KL divergence.). FOCOPS (Zhang et al., 2020a) has adopted a similar idea by directly solving the constrained policy optimization problem via the primal-dual approach (Boyd et al., 2004) then projecting the solution back into the parametric policy space. Traditional robot control also considers the safety problem. Chow et al. (2018; 2019) presents a method via constructing Lyapunov function to guarantees the constraint satisfaction during training. Stooke et al. (2020a) combines PID control with Lagrangian methods which dampens cost oscillations resulting in reduced constraint violations. It is still lacking a unified and efficient framework to cover these algorithms. Therefore, we provide PyTorch-version re-implementations of widely used safe policy optimization algorithms, hoping to facilitate experimental validation in Safe RL research.

**Dexterous Manipulation** Manipulation is one of the essential research topics in robotics, researchers have long tried to establish a stable theory of manipulation (Billard & Kragic, 2019). However, traditional methods mostly rely on various assumptions, such as knowing the environmental dynamics model or having no uncertainty in the process. In recent years, learning-based approaches have been successful in this regard, coping with uncertainty in perception and even generalizing to unseen objects (Bohg et al., 2013). There are many learning-based benchmarks for robotic manipulation in recent years (Yu et al., 2020; James et al., 2020; Zhu et al., 2020), but none of them use dexterous hands or consider safe constraints. Dexterous multi-finger hands provide intrinsic dexterity for better manipulation in unstructured scenes and contact-rich situations, but additionally bring the challenges of high-dimensional control and complex contact models (Bircher et al., 2017; Rahman et al., 2016). Previous research methods have mostly focused on trajectory optimization or model prediction, which highly relied on accurate dynamics models (Kim et al., 2021; Okamura et al., 2000; Kumar et al., 2016). For example, Williams et al. (2015) performs in-hand manipulation of a cube using a trajectory optimization technique known as Model Predictive Path Integral (MPPI). Charlesworth & Montana (2021) extended the MPPI method to allow objects to be thrown and catch between two hands. OpenAI et al. (2019) solved a Rubik’s cube using model-free RL and domain randomization techniques. Chen et al. (2022a) proposed an in-hand manipulation system to learn how to manipulate a large number of objects of different shapes, and even generalize to unseen objects. Qin et al. (2022; 2021) studied dexterous manipulation learning from human demonstration. Chen et al. (2022c) studied the bimanual dexterous manipulation to solve cooperative manipulation and skill generalization problem. While most of them focus on unconstrained dexterous manipulation, how to do dexterous manipulation safely is an unstudied topic. In this paper, we provide a massively parallel benchmark for safe dexterous manipulation, hoping to facilitate research on how to manipulate safely.

### 3 THE SAFETY LEARNING ENVIRONMENT

TrustDeHands is consist of two parts: the safety learning environment and the safe policy optimization algorithms. In this section, we present the high-level design of a safety learning environment.

#### 3.1 SYSTEM DESIGN AND DATASETS

TrustDeHands is a collection of challenging dexterous manipulation tasks, underpinned by Isaac Gym (Makovychuk et al., 2021) and capable of high parallelism on the GPU. In TrustDeHands, all tasks require two dexterous hands to manipulate one or more objects. We design a series of tasks that require policies to perform safe dexterous manipulation, including throwing, grasping, jerking, pulling, etc. At the same time, each task provides the customizability of dexterous hands and objects to support a diverse task.

The construction of the dataset includes the configuration of dexterous hands and objects. The core goal of our dataset is to generate a wide variety of scenarios for learning constrained dexterous manipulation. We collected a variety of dexterous multi-finger hands as manipulators, including most of the dexterous hands currently used in robotics. In addition to manipulators, objects also play a crucial role in building datasets. Our manipulation objects are mainly from the YCB (Calli et al., 2017) and SAPIEN (Xiang et al., 2020) datasets. Both datasets contain many objects used in everyday life.

### 3.2 TASKS REPRESENTATION

TrustDeHands contains 10+ tasks focused on dexterous manipulation. Each task contains two dexterous hands and one or more manipulated objects, such as balls, blocks, etc., with the ultimate goal is to manipulate objects placed at the task-specified locations while to make the agent satisfy. The default dexterous hand used by our framework is the Shadow Hand (ShadowRobot, 2005), more details are provided in Appendix A. The agent performs each task according to its observation, action representation, and its reward and cost function definition. We provide more underlying technical details about the tasks in Appendix B.

**Constrained Markov Decision Processes (CMDPs)** Constrained Markov Decision Processes (CMDPs) is defined as  $(\mathcal{S}, \mathcal{A}, \mathbb{P}, r, \rho_0, \gamma, \mathcal{C})$ , where  $\mathcal{S}$  is the state space,  $\mathcal{A}$  is the action space,  $\mathbb{P} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$  is the transition probability function,  $r : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$  is the reward function,  $\rho_0(\cdot) \in \mathcal{P}(\mathcal{S})$  is the initial state distribution ( $\mathcal{P}(X)$  denotes the set of probability distributions over a set  $X$ ),  $\gamma \in [0, 1]$  is the discount factor, and  $\mathcal{C} = \{(c, b)\}$  is the constraint set, where  $c : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$  is a cost function, and  $b$  is the cost threshold.

We use  $\pi : \mathcal{S} \rightarrow \mathcal{P}(\mathcal{A})$  to denote a stationary policy, and use  $\Pi$  to denote the set of all stationary policies. Let  $\tau = \{s_t, a_t, r_{t+1}, c_{t+1}\}_{t \geq 0} \sim \pi$  be a trajectory generated by  $\pi$ , where  $s_0 \sim \rho_0(\cdot)$ ,  $a_t \sim \pi(\cdot|s_t)$ ,  $s_{t+1} \sim \mathbb{P}(\cdot|s_t, a_t)$ ,  $r_{t+1} = r(s_{t+1}|s_t, a_t)$ , and  $c_{t+1} = c(s_{t+1}|s_t, a_t)$ . The *state value function* of  $\pi$  is defined as  $V_\pi(s) = \mathbb{E}_\pi[\sum_{t=0}^{\infty} \gamma^t r_{t+1} | s_0 = s]$ . The goal of reinforcement learning is to maximize the *expected total reward*, defined as  $J(\pi) = \mathbb{E}_{s \sim \rho_0(\cdot)}[V_\pi(s)]$ .

We define the *cost return function* as  $J^c(\pi) = \mathbb{E}_{s \sim \rho_0(\cdot)}[\sum_{t=0}^{\infty} \gamma^t c_{t+1} | s_0 = s]$ , and the feasible policy set  $\Pi_{\mathcal{C}}$  as  $\Pi_{\mathcal{C}} = \{\pi | \pi \in \Pi, J^c(\pi) \leq b, \forall (c, b) \in \mathcal{C}\}$ . The goal of Safe RL is to learn the optimal policy  $\pi_*$  such that

$$\pi_* = \arg \max_{\pi \in \Pi_{\mathcal{C}}} J(\pi). \quad (1)$$

**Observation** Here we briefly describe the observation space of the tasks, more details can be seen in Appendix B.1. The observation of all tasks consisted of three parts: state information of the left and right Shadow Hand, and information about the task specification. In each task, the state information of the left and right Shadow Hand is the same, each Shadow Hand contains 24 minimum drive units (which contains four underdriven fingertip units) and its state consists of the following information:

- $\mathcal{D}_p, \mathcal{D}_v, \mathcal{D}_f \in \mathbb{R}^{24}$ , corresponds to all joint *DoF* (Degree of Freedom) of angle, velocity, and force with drive units, respectively.
- $\mathcal{P}_w, \mathcal{R}_w \in \mathbb{R}^3$  represents the position and rotation of the base of the hand.
- $\mathcal{FT}_i = [FT_{\text{pose}}, FT_{v_l}, FT_{v_a}, FT_f, FT_t] \in \mathbb{R}^{19}$ , corresponds to the pose, linear velocity, angular velocity, force magnitude, and torque of each fingertip, respectively.
- $\mathcal{A} \in \mathbb{R}^{20/26}$ , indicates the action executed by the hand in the previous step, which is consistent with the action space.

With above definitions, the state information of one Shadow Hand can be represented as  $\mathcal{H}and = \{\mathcal{D}_p, \mathcal{D}_v, \mathcal{D}_f, \mathcal{P}_w, \mathcal{R}_w, \{\mathcal{FT}_i\}_{i=1}^5, \mathcal{A}\}$ . We character the observation of each task by the following information:

$$\{\mathcal{H}and_{\text{left}}, \mathcal{H}and_{\text{right}}, \mathcal{G}_{\text{task}}\}, \quad (2)$$

where  $\mathcal{G}_{\text{task}}$  represents some observation information specific to different tasks.

**Action** The dual Shadow Hands have more than 40 dimensions of action space, where each Shadow Hand has five fingers with a total of 24 degrees of freedom, the thumb has 5 joints and 5 degrees of freedom, and all other fingers have 3 degrees of freedom and 4 joints (where the joint at the end of each finger is uncontrollable). Therefore, the action space of each hand is 20 dimensions. If the base of the hand is not fixed, there are six dimensions to represent the translation and rotation of the hand base. For the lower and upper limit of the joint angle, see Table 2. In each step, we use the absolute value of each joint angle as the target, and use the PD controller to make it move.

**Reward** We designed some auxiliary rewards to help RL agents learn more consistently, and each task contains a task-specific bonus. In general, our reward design is goal-based and follows the same set of logic. For object-catching tasks, our reward is simply related to the difference between






Figure 1: Six representative tasks of TrustDeHands, including Safe Finger, Hand Over Wall, Hand Up Wall, Pick Bottles, Jenga, Clean House. The diagram below shows the Jenga task, where two robotic arms need to work with each other. The left hand needs to push the blocks hidden in the middle to the right, while the right hand needs to extract the blocks.

the pose of the object and the target. For other tasks that require the hand to hold the object, our reward generally consists of three parts: the distance from the left hand to a left-hand grip point on the object, the distance from the right hand to a right-hand grip point on the object, and the distance from the object to the object’s target.

**Cost** Each task contains different constraints (e.g., the ball needs to be thrown to a specified height or a specified angle to prevent damage to other items; the robots need to clean the floor without hitting other furniture). The specific constraint design depends on the safety requirements of each task.

### 3.3 EVALUATION SUITE

TrustDeHands has a collection of 10+ different tasks. These tasks form an evaluation suite for benchmarking the performance of Safe RL algorithms. In this section, we describe six of these representative tasks (see Figure 1), and the remaining tasks are described in detail in Appendix B.

**Safe Finger** In this task, one hand needs to throw the ball to the other hand, and both hands except the same horizontal plane. We design two different constraints, including constraints on the minimum drive units, and constraints on the finger joints.

**Hand Up Wall** In this task, one hand needs to throw the ball at a certain height to the other hand. The constraint of this task is concerned with the magnitude of the force of the minimum drive unit.

**Hand Over Wall** In this task, one hand needs to throw the ball at a certain angle to the other hand. It is more concerned with the constraints in a certain behavioral paradigm and needs to take into account the collaboration of the whole hand-driven unit.

**Pick Bottles** There are five bottles in a tight row, and the dual hands need to pick up two of the bottles smoothly without touching the others.

**Jenga** The dual hands need to collaborate to extract the specified blocks from an unstable stacking structure and avoid breaking the rest of the blocks apart.

**Clean House** There is a broom, trash, and dustpan in this environment. We need to use both hands to manipulate the broom to sweep the trash into the dustpan. There will also be a chair as an obstacle on the way.




Figure 2: **Left:** the original scene in the simulation; **Right:** the RGB view of the scene; the RGB-Depth view of the scene; the point cloud view of the scene.

### 3.4 VISUAL INFORMATION

It is very difficult to obtain the state information of the robot in the real world. One way to solve this problem is to use the vision sensor as the input to train the policy. Therefore, we provide multiple modalities of visual information as input, including RGB, RGB-D, and point cloud, see Figure 2. It is generated using the camera in the Isaac Gym, and the pose and toward of the camera can be customized by the user to obtain the desired visual information. We also propose a point cloud parallel acceleration function to adapt Isaac Gym and provide an example of using it to train Hand Over task, see Appendix.

### 3.5 CUSTOMIZABLE DIVERSIFORM MANIPULATORS AND ADAPTATION CHALLENGE

There are more type of dexterous hands than the shadow hand like allegro hand, trifinger, et al, and supporting other dexterous hands helps to advance research and community development. Therefore, in addition to the Shadow Hand, we also provide five kinds of other dexterous multi-finger hands in TrustDeHands. In addition, using a robotic arm drive at the base of the dexterous hand not only matches the real-world setting but also an important step for sim to real transfer. Becasue it is very difficult to match the real dynamics of the flying hand, the TrustDeHands provide a way to reduce the reality gap by adjusting the dynamics and physics parameters of the arm, which simplifies the deployment process from simulation to the real world applications.

Moreover, we offer a variety of arms and a variety of dexterous hand combinations, which has many benefits. For example, researchers can choose the hand they want according to their own conditions, which brings wider applicability to our benchmark. At the same time, we can use different arms and different hands to study the adaptability and generalization ability of policies, which challenges multi-task learning and meta learning research in the future. An schematic of this feature is shown in Figure 3.

## 4 EXPERIMENTS

### 4.1 SAFE RL ALGORITHMS IMPLEMENTATION

Based on their original papers or public code base, we re-implement eight algorithms (CPO (Achiam et al., 2017a), PCPO (Yang et al., 2020b), FOCOPS (Zhang et al., 2020b), P3O (Zhang et al., 2022), PPO-Lag (Ray et al., 2019a), TRPO-Lag (Ray et al., 2019a), CPPO-PID (Stooke et al., 2020b), and IPO (Liu et al., 2020)), covering major safe policy optimization algorithms. A brief introduction to each algorithm is given in Appendix E.

We abstract similar structure of the safe policy optimization algorithms, and modularize the code into interaction with environments, parallel sample collection, buffer storage and computation, algorithm core update, and auxiliary functionalities such as visualization and logger. Maximum abstraction and encapsulation take place at the implementation of algorithms core, where each algorithm inherits directly from its

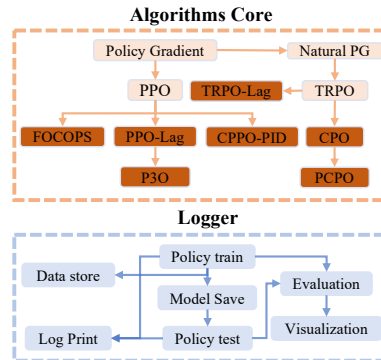


Figure 4: An overview of algorithms core design and logger.




Figure 3: Using different dexterous hands and robot arms in TrustDeHands provides diversity. The left image is the dexterous hands with a robot arm driver doing the Jenga task, the demo can be found in <https://sites.google.com/view/trustdehands/>. where the left hand is Kuka connected with ShadowHand, and the right hand is X-ar connected with Allegro. Among them, we support five kinds of the dexterous hands shown in the upper right corner, and eight kinds of robot arms shown in the lower right corner, which can be both customized by the user.

base algorithm, thus only unique features have to be implemented and all other code can be reused. An overview of the core of algorithms and logger is shown in Figure 4.

For algorithms implementation, it is critical to ensure its correctness and reliability. To achieve this goal, we examine the implementation of our algorithms carefully. To test the performance of our implementation, we run the eight algorithms on 30 tasks (for a complete list of the tasks, please refer to Appendix F.1) contained in the four environment suites and present our experimental results for the reference of the community in Appendix F.2.

#### 4.2 EVALUATION PROTOCOL

**Metrics** We define the following metrics to depict the safety performance of an agent in different tasks. (1) the average return of trajectories,  $J^r(\theta)$ ; (2) the average cumulative cost of trajectories,  $J^c(\theta)$ . In Safe RL domain, for any two agents, the superiority of the agents is determined by the following priority comparisons. On the one hand, the agent that satisfies the constraint will definitely outperform the unconstrained one. On the other hand, two agents that satisfy the constraint are determined by comparing the magnitude of their cumulative returns.

**Algorithms** For the UnSafe RL algorithm we uniquely use PPO (Schulman et al., 2017), where the reward function contains no information about the auxiliary costs. For Safe RL algorithms, we evaluate the performance of PPO-Lag (Ray et al., 2019b), FOCOPS, P3O, and PCPO algorithms on TrustDeHands, and the remaining Safe RL algorithms we implemented are in our anonymous Github repository.

#### 4.3 RESULTS

We mainly conduct three experiments and analyze the results in this section: 1) The performance of PPO, PPO-Lag, FOCOPS, P3O, CPPO-PID, and algorithms on six representative tasks 2) The performance of eight Safe RL algorithms on the Safe Finger task 3) The performance of point cloud RL on the Hand Over Wall task.

For 1), We evaluate the performance of PPO, PPO-Lag, FOCOPS, CPPO-PID, and P3O algorithms on six tasks, and we implemented the rest of the Safe RL algorithms in our anonymous Github repository. The performance of each algorithm is shown in Figure 5. It can be observed that PPO-Lag can achieve high performance within the range allowed by the cost, and is the best performing algorithm here. Comparing the performance of PPO and PPO-Lag, it can be found that PPO-Lag can perform similarly to PPO in Jenga, and Safe Finger tasks, but the cost is constrained to a lower range, which indicates that the model has learned how to safely manipulate. A remarkable result is that in the Jenga and Pick Bottle task, the performance of PPO-Lag is far superior to PPO. This

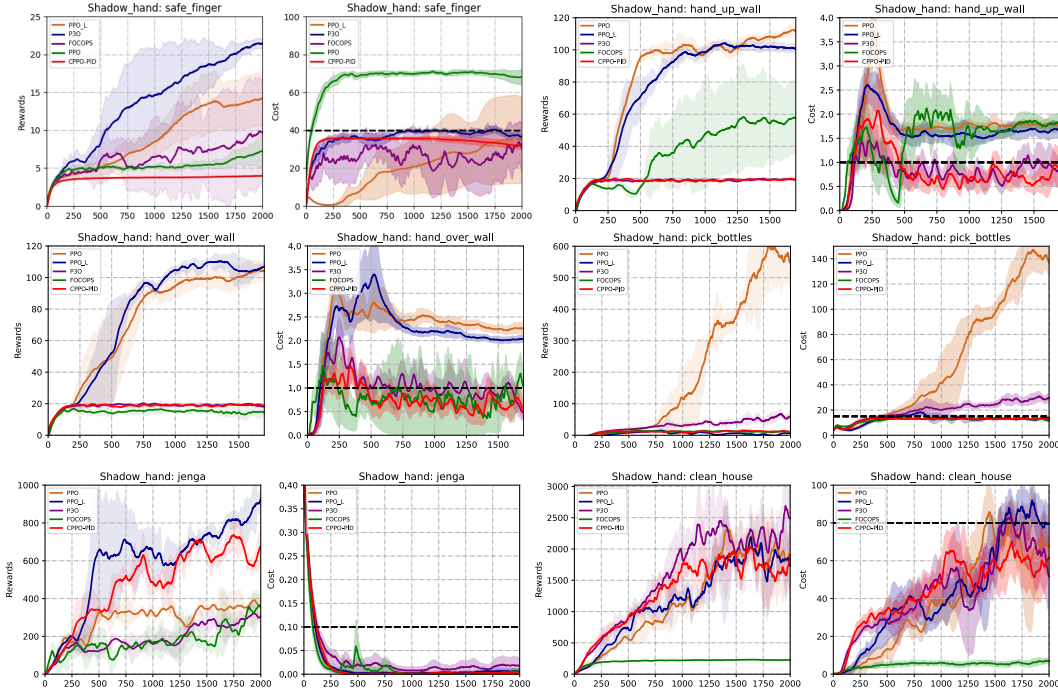


Figure 5: Learning curves for all six tasks. The shaded region represents the standard deviation of the score over 3 trials. Curves are smoothed uniformly for visual clarity. All algorithms interact with environments in 100M steps and the number of parallel simulations is 2048. is because in these environments, learning safe manipulation is benefits. For example, in Jenga, when the policy learns to not mess up other blocks, the target objects are also easier to be removed, resulting in a higher reward. It fully illustrates the advantage of Safe RL in learning better policies on manipulation tasks. However, on most of the tasks, FOCOPS and P3O are basically unable to achieve the performance of PPO-Lag, or even can not complete the task. Therefore, the performance of the current Safe RL algorithm on manipulation still has a lot of room for exploration.

For 2), On the other hand, we tested the entire eight algorithms we implemented on the Safe Finger task, which is shown in Figure 1. We the specific details can be found in Appendix C.

It can be seen that CPO has almost no work, and the PPO algorithm will cause a high cost. This may be due to various approximation methods in CPO, which puts an emergency on Safe RL towards complex manipulation areas.

Table 1: Performance on TrustDeHands.

	HandOver.Finger		HandOver.Joint		HandOver.Underarm.Finger		HandOver.Underarm.Joint	
Performance	Reward	Cost	Reward	Cost	Reward	Cost	Reward	Cost
CPO	$4.57 \pm 0.02$	$32.41 \pm 0.01$	$6.18 \pm 0.01$	$29.54 \pm 0.02$	$21.63 \pm 0.01$	$34.72 \pm 0.01$	$2.68 \pm 0.01$	$30.3 \pm 0.02$
TRPO.L	$3.62 \pm 0.01$	$36.02 \pm 0.01$	$3.83 \pm 0.01$	$10.77 \pm 0.02$	$4.55 \pm 0.03$	$15.6 \pm 0.01$	$3.37 \pm 0.02$	$33.13 \pm 0.01$
PPO.L	$17.4 \pm 0.04$	$27.69 \pm 0.01$	$22.44 \pm 0.04$	$71.15 \pm 0.02$	$25.5 \pm 0.01$	$65.56 \pm 0.02$	$24.11 \pm 0.03$	$64.39 \pm 0.05$
P3O	$21.93 \pm 0.07$	$40.54 \pm 0.02$	$21.9 \pm 0.03$	$31.34 \pm 0.02$	$22.72 \pm 0.0$	$47.3 \pm 0.02$	$20.09 \pm 0.05$	$55.84 \pm 0.09$
PCPO	$3.08 \pm 0.02$	$70.15 \pm 0.01$	$3.08 \pm 0.12$	$70.15 \pm 0.01$	$0.3 \pm 0.01$	$3.22 \pm 0.02$	$0.26 \pm 0.01$	$5.0 \pm 0.04$
FOCOPS	$13.89 \pm 0.04$	$39.68 \pm 0.01$	$19.79 \pm 0.37$	$33.45 \pm 0.15$	$18.95 \pm 0.02$	$38.57 \pm 0.01$	$4.1 \pm 0.01$	$32.54 \pm 0.02$
CPPO-PID	$3.63 \pm 0.01$	$29.41 \pm 0.02$	$5.21 \pm 0.03$	$28.54 \pm 0.04$	$0.31 \pm 0.1$	$3.28 \pm 0.2$	$3.81 \pm 0.07$	$32.6 \pm 0.023$
IPO	$3.12 \pm 0.03$	$69.23 \pm 0.02$	$3.01 \pm 0.21$	$69.21 \pm 0.03$	$0.27 \pm 0.02$	$4.18 \pm 0.01$	$0.32 \pm 0.02$	$5.24 \pm 0.05$

## 5 POTENTIAL RESEARCH PROBLEMS TO STUDY USING TRUSTDEHANDS

TrustDeHands provides ample opportunities to study trustworthy manipulation of dexterous hands based on Safe RL. We found that the primal-dual based approach (Boyd et al., 2004) result in great volatility in the update of lagrange multipliers. A potential research direction is to consider combining feedback control methods in control systems, such as PID (Ziegler et al., 1942; Ang et al., 2005), ADRC (Han, 2009), etc., to mitigate the instability and volatility of Lagrange multipliers in the learning process. Therefore, it would be interesting to combine control methods of complex systems with Safe RL methods to solve complex manipulation problems of dexterous hands.

Sim to real is an important research direction about transferring the simulation result to the real robot. Around this theme, our benchmark includes many of the robot arms and dexterous hands, which were accepted by many research labs. It is convenient for different researchers to choose their own arms and hands for training in the simulation. Meanwhile, the tasks in our benchmark, such as picking bottles<sup>1</sup>, Jenga<sup>2</sup>, etc., are meaningful in the real world but also needed to ensure safety if transferring the trained policy from simulation to the real world. So our benchmark can also be used to study how to perform sim to real more safely from the perspective of Safe RL.

Training policy with a state-based observation space is difficult for sim to real transfer because such inputs are not available in the real world. So it also makes sense to study the more readily available policy inputs in the real world, such as point clouds. Our environment supports a multimodal input such as visual and forces information, which can support research in this direction. We hope that our benchmark can serve as a tool to study the sim to real transfer of dexterous hands.

Finally, generalization is an important direction to explore, which is a potential strength of RL. TrustDeHands supports self-customization, enabling switching and linking different hands and arms to evaluate the generality of different algorithms. Users can use TrustDeHands as a platform for modification or secondary development to design richer and more challenging target tasks, and we hope that this work will contribute to the flourishing of the RL community.

## 6 CONCLUSION AND FUTURE WORK

In this work, we presented TrustDeHands, which is the first benchmark focused on safe dexterous manipulation. We standardize the safe policy optimization methods for solving CMDPs and introduce a unified, highly-optimized, extensible, and comprehensive algorithms re-implementation. We checked the correctness of our algorithms on the existing Safe RL benchmark and tested it on TrustDeHands. The results show that the Safe RL algorithm can better solve the safety problem in dexterous manipulation. For example, using Safe RL can grab the target bottle without touching other bottles, and avoid collision with obstacles when sweeping the floor. However, it is difficult for unconstrained RL algorithms to have this guarantee. These situations are very important for robots in real-world environments because RL-based methods tend to lead to unpredictable behaviors that are prone to danger and damage to robots.

Additionally, we support two features regarding visual policy input and various arms and dexterous hands. Some sort of visual input is becoming increasingly common in real-world RL-trained robots, and so benchmarks for this setting are important. Diverse arms and hands increase the applicability of our benchmarks and allow us to study policy generalization between different robots.

We believe that TrustDeHands can significantly accelerate the progress of future research on safe manipulation, facilitate the integration of reinforcement learning with robotic control, and will make a singular contribution to the reinforcement learning community.

## REFERENCES

- Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained policy optimization. In *International conference on machine learning*, pp. 22–31. PMLR, 2017a.
- Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained policy optimization. In *International conference on machine learning*, pp. 22–31. PMLR, 2017b.
- Rami Al-Rfou, Guillaume Alain, Amjad Almahairi, Christof Angermueller, Dzmitry Bahdanau, Nicolas Ballas, Frédéric Bastien, Justin Bayer, Anatoly Belikov, Alexander Belopolsky, et al. Theano: A python framework for fast computation of mathematical expressions. *arXiv e-prints*, pp. arXiv–1605, 2016.
- Mohammed Alshiekh, Roderick Bloem, Rüdiger Ehlers, Bettina Könighofer, Scott Niekum, and Ufuk Topcu. Safe reinforcement learning via shielding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- Eitan Altman. *Constrained Markov decision processes*. CRC Press, 1999.

<sup>1</sup><https://www.youtube.com/watch?v=hvibZrLxYyQ>

<sup>2</sup>More details in <https://en.wikipedia.org/wiki/Jenga>



- OpenAI: Marcin Andrychowicz, Bowen Baker, Maciek Chociej, Rafal Jozefowicz, Bob McGrew, Jakub Pachocki, Arthur Petron, Matthias Plappert, Glenn Powell, Alex Ray, et al. Learning dexterous in-hand manipulation. *The International Journal of Robotics Research*, 39(1):3–20, 2020.
- Kiam Heong Ang, Gregory Chong, and Yun Li. Pid control system analysis, design, and technology. *IEEE transactions on control systems technology*, 13(4):559–576, 2005.
- Aude Billard and Danica Kragic. Trends and challenges in robot manipulation. *Science*, 364(6446): eaat8414, 2019.
- Walter G Bircher, Aaron M Dollar, and Nicolas Rojas. A two-fingered robot gripper with large object reorientation range. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3453–3460. IEEE, 2017.
- Jeannette Bohg, Antonio Morales, Tamim Asfour, and Danica Kragic. Data-driven grasp synthesis—a survey. *IEEE Transactions on robotics*, 30(2):289–309, 2013.
- Stephen Boyd, Stephen P Boyd, and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- Berk Calli, Arjun Singh, James Bruce, Aaron Walsman, Kurt Konolige, Siddhartha Srinivasa, Pieter Abbeel, and Aaron M Dollar. Yale-cmu-berkeley dataset for robotic manipulation research. *The International Journal of Robotics Research*, 36(3):261–268, 2017.
- Henry J Charlesworth and Giovanni Montana. Solving challenging dexterous manipulation tasks with trajectory optimisation and reinforcement learning. In *International Conference on Machine Learning*, pp. 1496–1506. PMLR, 2021.
- Tao Chen, Jie Xu, and Pulkit Agrawal. A system for general in-hand object re-orientation. In *Conference on Robot Learning*, pp. 297–307. PMLR, 2022a.
- Yuanpei Chen, Yaodong Yang, Tianhao Wu, Shengjie Wang, Xidong Feng, Jiechuang Jiang, Stephen Marcus McAleer, Hao Dong, Zongqing Lu, and Song-Chun Zhu. Towards human-level bimanual dexterous manipulation with reinforcement learning. *arXiv preprint arXiv:2206.08686*, 2022b.
- Yuanpei Chen, Yaodong Yang, Tianhao Wu, Shengjie Wang, Xidong Feng, Jiechuang Jiang, Stephen Marcus McAleer, Hao Dong, Zongqing Lu, and Song-Chun Zhu. Towards human-level bimanual dexterous manipulation with reinforcement learning. *CoRR*, abs/2206.08686, 2022c.
- Yinlam Chow, Ofir Nachum, Edgar Duenez-Guzman, and Mohammad Ghavamzadeh. A lyapunov-based approach to safe reinforcement learning. *Advances in neural information processing systems*, 31, 2018.
- Yinlam Chow, Ofir Nachum, Aleksandra Faust, Edgar Duenez-Guzman, and Mohammad Ghavamzadeh. Lyapunov-based safe policy optimization for continuous control. *arXiv preprint arXiv:1901.10031*, 2019.
- Gabriel Dulac-Arnold, Daniel Mankowitz, and Todd Hester. Challenges of real-world reinforcement learning. *arXiv preprint arXiv:1904.12901*, 2019.
- Sven Gronauer. Bullet-safety-gym: Aframework for constrained reinforcement learning. 2022.
- Shangding Gu, Jakub Grudzien Kuba, Muning Wen, Ruiqing Chen, Ziyang Wang, Zheng Tian, Jun Wang, Alois Knoll, and Yaodong Yang. Multi-agent constrained policy optimisation. *arXiv preprint arXiv:2110.02793*, 2021.
- Shangding Gu, Long Yang, Yali Du, Guang Chen, Florian Walter, Jun Wang, Yaodong Yang, and Alois Knoll. A review of safe reinforcement learning: Methods, theory and applications. *arXiv preprint arXiv:2205.10330*, 2022.
- Ben Hambly, Renyuan Xu, and Huining Yang. Recent advances in reinforcement learning in finance. *arXiv preprint arXiv:2112.04553*, 2021.

- Jingqing Han. From pid to active disturbance rejection control. *IEEE transactions on Industrial Electronics*, 56(3):900–906, 2009.
- Stephen James, Zicong Ma, David Rovick Arrojo, and Andrew J. Davison. Rlbench: The robot learning benchmark & learning environment. *IEEE Robotics Autom. Lett.*, 5(2):3019–3026, 2020.
- Michael A Johnson and Mohammad H Moradi. *PID control*. Springer, 2005.
- Uikeyum Kim, Dawoon Jung, Heeyoen Jeong, Jongwoo Park, Hyun-Mok Jung, Joono Cheong, Hyouk Ryeol Choi, Hyunmin Do, and Chanhun Park. Integrated linkage-driven dexterous anthropomorphic robotic hand. *Nature communications*, 12(1):1–13, 2021.
- Vikash Kumar, Emanuel Todorov, and Sergey Levine. Optimal control with learned local models: Application to dexterous manipulation. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 378–383. IEEE, 2016.
- Jan Leike, Miljan Martic, Victoria Krakovna, Pedro A Ortega, Tom Everitt, Andrew Lefrancq, Laurent Orseau, and Shane Legg. Ai safety gridworlds. *arXiv preprint arXiv:1711.09883*, 2017.
- Yongshuai Liu, Jiaxin Ding, and Xin Liu. Ipo: Interior-point policy optimization under constraints. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 4940–4947, 2020.
- Zuxin Liu, Zijian Guo, Zhepeng Cen, Huan Zhang, Jie Tan, Bo Li, and Ding Zhao. On the robustness of safe reinforcement learning under observational perturbations. *arXiv preprint arXiv:2205.14691*, 2022.
- Viktor Makoviychuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey, Miles Macklin, David Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, and Gavriel State. Isaac gym: High performance gpu-based physics simulation for robot learning, 2021.
- Arkadi Nemirovski. Interior point polynomial time methods in convex programming. *Lecture notes*, 42(16):3215–3224, 2004.
- Allison M Okamura, Niels Smaby, and Mark R Cutkosky. An overview of dexterous manipulation. In *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No. 00CH37065)*, volume 1, pp. 255–262. IEEE, 2000.
- OpenAI. Openai five. <https://blog.openai.com/openai-five/>, 2018.
- OpenAI, Ilge Akkaya, Marcin Andrychowicz, Maciek Chociej, Mateusz Litwin, Bob McGrew, Arthur Petron, Alex Paino, Matthias Plappert, Glenn Powell, Raphael Ribas, Jonas Schneider, Nikolas Tezak, Jerry Tworek, Peter Welinder, Lilian Weng, Qiming Yuan, Wojciech Zaremba, and Lei Zhang. Solving rubik’s cube with a robot hand. *CoRR*, abs/1910.07113, 2019.
- Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 652–660, 2017.
- Yuzhe Qin, Yueh-Hua Wu, Shaowei Liu, Hanwen Jiang, Ruihan Yang, Yang Fu, and Xiaolong Wang. Dexmv: Imitation learning for dexterous manipulation from human videos. *CoRR*, abs/2108.05877, 2021.
- Yuzhe Qin, Hao Su, and Xiaolong Wang. From one hand to multiple hands: Imitation learning for dexterous manipulation from single-camera teleoperation. *CoRR*, abs/2204.12490, 2022.
- Nahian Rahman, Luca Carbonari, Mariapaola D’Imperio, Carlo Canali, Darwin G Caldwell, and Ferdinando Cannella. A dexterous gripper for in-hand manipulation. In *2016 IEEE International Conference on Advanced Intelligent Mechatronics (AIM)*, pp. 377–382. IEEE, 2016.
- Alex Ray, Joshua Achiam, and Dario Amodei. Benchmarking safe exploration in deep reinforcement learning. *arXiv preprint arXiv:1910.01708*, 7:1, 2019a.

- Alex Ray, Joshua Achiam, and Dario Amodei. Benchmarking safe exploration in deep reinforcement learning. *arXiv preprint arXiv:1910.01708*, 7:1, 2019b.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- ShadowRobot. Shadowrobot dexterous hand. <https://www.shadowrobot.com/dexterous-hand-series/>, 2005.
- David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.
- Adam Stooke, Joshua Achiam, and Pieter Abbeel. Responsive safety in reinforcement learning by pid lagrangian methods. In *International Conference on Machine Learning*, pp. 9133–9143. PMLR, 2020a.
- Adam Stooke, Joshua Achiam, and Pieter Abbeel. Responsive safety in reinforcement learning by pid lagrangian methods. In *International Conference on Machine Learning*, pp. 9133–9143. PMLR, 2020b.
- Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ international conference on intelligent robots and systems*, pp. 5026–5033. IEEE, 2012.
- Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.
- Grady Williams, Andrew Aldrich, and Evangelos A. Theodorou. Model predictive path integral control using covariance variable importance sampling. *CoRR*, abs/1509.01149, 2015.
- Fanbo Xiang, Yuzhe Qin, Kaichun Mo, Yikuan Xia, Hao Zhu, Fangchen Liu, Minghua Liu, Hanxiao Jiang, Yifu Yuan, He Wang, et al. Sapien: A simulated part-based interactive environment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11097–11107, 2020.
- Mengdi Xu, Zuxin Liu, Peide Huang, Wenhao Ding, Zhepeng Cen, Bo Li, and Ding Zhao. Trustworthy reinforcement learning against intrinsic vulnerabilities: Robustness, safety, and generalizability. *arXiv preprint arXiv:2209.08025*, 2022.
- Long Yang, Jiaming Ji, Juntao Dai, Yu Zhang, Pengfei Li, and Gang Pan. Cup: A conservative update policy algorithm for safe reinforcement learning. *arXiv preprint arXiv:2202.07565*, 2022.
- Tsung-Yen Yang, Justinian Rosca, Karthik Narasimhan, and Peter J Ramadge. Projection-based constrained policy optimization. *arXiv preprint arXiv:2010.03152*, 2020a.
- Tsung-Yen Yang, Justinian Rosca, Karthik Narasimhan, and Peter J Ramadge. Projection-based constrained policy optimization. *arXiv preprint arXiv:2010.03152*, 2020b.
- Tianhe Yu, Deirdre Quillen, Zhanpeng He, Ryan Julian, Karol Hausman, Chelsea Finn, and Sergey Levine. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In *Conference on Robot Learning*, pp. 1094–1100. PMLR, 2020.
- Zhacong Yuan, Adam W Hall, Siqi Zhou, Lukas Brunke, Melissa Greeff, Jacopo Panerati, and Angela P Schoellig. safe-control-gym: a unified benchmark suite for safe learning-based control and reinforcement learning. *arXiv preprint arXiv:2109.06325*, 2021.
- Linrui Zhang, Li Shen, Long Yang, Shixiang Chen, Bo Yuan, Xueqian Wang, Dacheng Tao, et al. Penalized proximal policy optimization for safe reinforcement learning. *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 2022.

- Yiming Zhang, Quan Vuong, and Keith Ross. First order constrained optimization in policy space. *Advances in Neural Information Processing Systems*, 33:15338–15349, 2020a.
- Yiming Zhang, Quan Vuong, and Keith Ross. First order constrained optimization in policy space. *Advances in Neural Information Processing Systems*, 33:15338–15349, 2020b.
- Yuke Zhu, Josiah Wong, Ajay Mandlekar, and Roberto Martín-Martín. robosuite: A modular simulation framework and benchmark for robot learning. *CoRR*, abs/2009.12293, 2020.
- John G Ziegler, Nathaniel B Nichols, et al. Optimum settings for automatic controllers. *trans. ASME*, 64(11), 1942.

## A MORE DETAILS ABOUT SHADOW HAND

The Shadow Dexterous Hand (ShadowRobot, 2005) is an example of a robotic hand designed for human-level dexterity; it has five fingers with a total of 24 degrees of freedom. The hand has been commercially available since 2005; however it still has not seen widespread adoption, which can be attributed to the daunting difficulty of controlling systems of such complexity (Andrychowicz et al., 2020).

The limits of each joint in Shadow hand are as Table 2. The thumb has 5 joints and 5 degrees of freedom, while all other fingers have 3 degrees of freedom and 4 joints. It should be noted that the joints at the end of each finger are not controllable. The distal joints of the fingers are coupled like that of human fingers, making the angle of the middle joint always bigger or equal to the angle of the distal joint. This allows the middle phalange is curved, while the distal phalange is straight. There is an extra joint (LF5) at the end of the little finger to allow the little finger to rotate in the direction of the thumb. There are two joints at the wrist, which guarantees that the entire hand can rotate 360 degrees.




Figure 6: Illustration of the joints on a dexterous robotic hand.

Stiffness, damping, friction, and armature are also important physical parameters in robotics. For each Shadow hand’s joint, we show our DoF properties in Table 3. This part can be adjusted in the Isaac Gym simulator.

Table 2: Finger range of motion.

Joints	Corresponds to the number of ??	Min	Max
Finger Distal (FF1,MF1,RF1,LF1)	15, 11, 7, 3	0°	90°
Finger Middle (FF2,MF2,RF2,LF2)	16, 12, 8, 4	0°	90°
Finger Base Abduction (FF3,MF3,RF3,LF3)	17, 13, 9, 5	-15°	90°
Finger Base Lateral (FF4,MF4,RF4,LF4)	18, 14, 10, 6	-20°	20°
Little Finger Rotation(LF5)	19	0°	45°
Thumb Distal (TH1)	20	-15°	90°
Thumb Middle (TH2)	21	-30°	30°
Thumb Base Abduction (TH3)	22	-12°	12°
Thumb Base Lateral (TH4)	23	0°	70°
Thumb Base Rotation (TH5)	24	-60°	60°
Hand Wrist Abduction (WR1)	1	-40°	28°
Hand Wrist Lateral (WR2)	2	-28°	8°

Table 3: DoF properties of Shadow Hand.

Joints	Stiffness	Damping	Friction	Armature
WR1	100	4.78	0	0
WR2	100	2.17	0	0
FF2	100	3.4e+38	0	0
FF3	100	0.9	0	0
FF4	100	0.725	0	0
MF2	100	3.4e+38	0	0
MF3	100	0.9	0	0
MF4	100	0.725	0	0
RF2	100	3.4e+38	0	0
RF3	100	0.9	0	0
RF4	100	0.725	0	0
LF2	100	3.4e+38	0	0
LF3	100	0.9	0	0
LF4	100	0.725	0	0
TH2	100	3.4e+38	0	0
TH3	100	0.99	0	0
TH4	100	0.99	0	0
TH5	100	0.81	0	0

## B TASK SPECIFICATIONS

### B.1 BASIC STATE SPACE AND ACTION SPACE

The state space dimension of each environment is up to 400 dimensions in total, and the action space dimension is up to 40 dimensions. All environments are goal-based, and each epoch will randomly reset the object’s starting pose and target pose to improve generalization. We only use the shadow hand and object state information as observation at present. The observation of all tasks is composed of three parts: the state information of the left and right hands, and the information of objects and



target. The state information of the left and right hands were the same for each task, including hand joint and finger positions, velocity, and force information. The state information of the object and goal are different for each task, which we will describe in the following. Table 4 shows the specific information of the left-hand and right-hand state.

Table 4: Observation space of bimanual shadow hands.

Index	Description
0 - 23	right shadow hand dof position
24 - 47	right shadow hand dof velocity
48 - 71	right shadow hand dof force
72 - 136	right shadow hand fingertip pose, linear velocity, angle velocity (5 x 13)
137 - 166	right shadow hand fingertip force, torque (5 x 6)
167 - 169	right shadow hand base position
170 - 172	right shadow hand base rotation
173 - 198	right shadow hand actions
199 - 222	left shadow hand dof position
223 - 246	left shadow hand dof velocity
247 - 270	left shadow hand dof force
271 - 335	left shadow hand fingertip pose, linear velocity, angle velocity (5 x 13)
336 - 365	left shadow hand fingertip force, torque (5 x 6)
366 - 368	left shadow hand base position
369 - 371	left shadow hand base rotation
372 - 397	left shadow hand actions

## B.2 SAFE FINGER

This environment contains two dexterous hands. At the beginning of each episode, a ball falls randomly around the right hand, and the two hands have to collaborate to place the ball to a given position. Since the target is out of the reach of the right hand, and the right hand cannot pass the ball to the left hand directly, a possible solution is that the right hand grabs the ball, throws it to the left hand; the left hand catches the ball, and puts it to the target. Note that the base of the hand is fixed.

**Observations** The 398-dimensional observational space for Hand Over task is shown in Table 5. It should be noted that since the base of the dual hands in this task is fixed, the observation of the dual hands is compared to the Table 4 of reduced 24 dimensions.

Table 5: Observation space of Safe Finger.

Index	Description
0 - 373	dual hands observation shown in Table 4
374 - 380	object pose
381 - 383	object linear velocity
384 - 386	object angle velocity
387 - 393	goal pose
394 - 397	goal rot - object rot

**Actions** The 40-dimensional action space for one hand in Safe Finger task is shown in Table 6.

Table 6: Action space of Safe Finger.

Index	Description
0 - 19	right shadow hand actuated joint
20 - 39	left shadow hand actuated joint

**Reward** For timestep  $t$ , let  $x_{b,t}$  be the position of the ball and  $x_{g,t}$  be the position of the goal. We use  $d_{p,t}$  to denote the positional distance between the ball and the goal  $d_{p,t} = \|x_{b,t} - x_{g,t}\|_2$ . Let  $d_{a,t}$  denote the angular distance between the object and the goal, and the rotational difference is  $d_{r,t} = 2 \arcsin \min\{|d_{a,t}|, 1.0\}$ . The reward is defined as follows,

$$r_t = \exp\{-0.2(\alpha d_{p,t} + d_{r,t})\}, \quad (3)$$

where  $\alpha$  is a constant balances positional and rotational rewards.

**Cost** In these tasks, we constrain the freedom of joints ②, ③ and ④ of forefinger (please refer to figure 6 (b)). Without the constraint, joints ② and ③ have freedom of  $[0^\circ, 90^\circ]$  and joint ④ of  $[-20^\circ, 20^\circ]$ . The safety tasks restrict joints ②, ③, and ④ within  $[22.5^\circ, 67.5^\circ]$ ,  $[22.5^\circ, 67.5^\circ]$ , and  $[-10^\circ, 10^\circ]$  respectively. Let  $\text{ang\_2}$ ,  $\text{ang\_3}$ ,  $\text{ang\_4}$  be the angles of joints ②, ③, ④, and the cost is defined as:

$$c_t = \mathbb{I}(\text{ang\_2} \notin [22.5^\circ, 67.5^\circ], \text{ or } \text{ang\_3} \notin [22.5^\circ, 67.5^\circ], \text{ or } \text{ang\_4} \notin [-10^\circ, 10^\circ]). \quad (4)$$

### B.3 HAND UP WALL

Similarly, this environment is similar to the Hand Over Wall, the difference is that the wall in this environment only retains the lower half, so the ball needs to be thrown high to prevent it from hitting the wall, requiring different motion skill.

**Observations** The 398-dimensional observational space for Hand Up Wall task is shown in Table 7. It should be noted that since the base of the dual hands in this task is fixed, the observation of the dual hands is compared to the Table 4 of reduced 24 dimensions.

Table 7: Observation space of Hand Up Wall.

Index	Description
0 - 373	dual hands observation shown in Table 4
374 - 380	object pose
381 - 383	object linear velocity
384 - 386	object angle velocity
387 - 393	goal pose
394 - 397	goal rot - object rot

**Actions** The 40-dimensional action space for one hand in Hand Up Wall task is shown in Table 8.

Table 8: Action space of Hand Up Wall.

Index	Description
0 - 19	right shadow hand actuated joint
20 - 39	left shadow hand actuated joint

**Reward** For timestep  $t$ , let  $x_{b,t}$  be the position of the ball and  $x_{g,t}$  be the position of the goal. We use  $d_{p,t}$  to denote the positional distance between the ball and the goal  $d_{p,t} = \|x_{b,t} - x_{g,t}\|_2$ . Let  $d_{a,t}$  denote the angular distance between the object and the goal, and the rotational difference is  $d_{r,t} = 2 \arcsin \min\{|d_{a,t}|, 1.0\}$ . The reward is defined as follows,

$$r_t = \exp\{-0.2(\alpha d_{p,t} + d_{r,t})\}, \quad (5)$$

where  $\alpha$  is a constant balances positional and rotational rewards.

**Cost** The ball is thrown from the right hand to the left hand, and the curve of the ball throwing out is not consistent, and we set a wall with height in the middle of the two hands. This requires more delicate hand manipulation, and when the right hand does not throw the ball with the proper force or angle, it will be difficult to throw the ball over the wall. If the ball hits the wall, the cost is 1, otherwise it is 0. The size of the wall and the hole can be customized by the user.

### B.4 HAND OVER WALL

This environment is similar to Safe Finger, except that it has a wall between each hand and a hole in the middle of the wall. We need to learn policy to keep the ball from hitting the wall during the toss.

**Observations** The 398-dimensional observational space for Hand Over Wall task is shown in Table 9. It should be noted that since the base of the dual hands in this task is fixed, the observation of the dual hands is compared to the Table 4 of reduced 24 dimensions.

**Actions** The 40-dimensional action space for one hand in Hand Over Wall task is shown in Table 10.

Table 9: Observation space of Hand Over Wall.

Index	Description
0 - 373	dual hands observation shown in Table 4
374 - 380	object pose
381 - 383	object linear velocity
384 - 386	object angle velocity
387 - 393	goal pose
394 - 397	goal rot - object rot

Table 10: Action space of Hand Over Wall.

Index	Description
0 - 19	right shadow hand actuated joint
20 - 39	left shadow hand actuated joint

**Reward** For timestep  $t$ , let  $x_{b,t}$  be the position of the ball and  $x_{g,t}$  be the position of the goal. We use  $d_{p,t}$  to denote the positional distance between the ball and the goal  $d_{p,t} = \|x_{b,t} - x_{g,t}\|_2$ . Let  $d_{a,t}$  denote the angular distance between the object and the goal, and the rotational difference is  $d_{r,t} = 2 \arcsin \min\{|d_{a,t}|, 1.0\}$ . The reward is defined as follows,

$$r_t = \exp\{-0.2(\alpha d_{p,t} + d_{r,t})\}, \quad (6)$$

where  $\alpha$  is a constant balances positional and rotational rewards.

**Cost** This constraint is more demanding than Wall Down, where we require the ball thrown to fit through a specified narrow hole. If the ball hits the wall, the cost is 1, otherwise it is 0. The size of the wall and the hole can be customized by the user.

## B.5 PICK BOTTLES

This environment contains two hands, a table and five bottles. The five bottles were placed in a row on the table horizontally with very little space between them. We need to pick up two bottles with two dexterous hands, and not touch the bottle around it to cause possible damage.

**Observations** The 400-dimensional observational space for Pick Bottles task is shown in Table 11. It should be noted that since the base of the dual hands in this task is fixed, the observation of the dual hands is compared to the Table 4 of reduced 24 dimensions.

Table 11: Observation space of Pick Bottles.

Index	Description
0 - 397	dual hands observation shown in Table 4
398 - 404	left bottle pose
405 - 407	left bottle linear velocity
408 - 410	left bottle angle velocity
411 - 417	right bottle pose
418 - 420	right bottle linear velocity
421 - 423	right bottle angle velocity

**Actions** The 52-dimensional action space for one hand in Pick Bottles task is shown in Table 12.

**Reward** The reward consists of three parts: the distance from the left hand to the left bottle cap, the distance from the right hand to the right bottle cap, and the height of the two bottles that need to be picked. The height of the two bottles that need to be picked is given by  $d_{height}$ . The position difference between the left hand to the left bottle cap  $d_{left}$  is given by  $d_{left} = \|x_{lhand} - x_{lbcap}\|_2$ . The position difference between the right hand to the right bottle cap  $d_{right}$  is given by  $d_{right} = \|x_{rhand} - x_{rbcap}\|_2$ . The reward is given by this specific formula:

$$r = d_{height} * 20 - d_{left} - d_{right} \quad (7)$$

Table 12: Action space of Pick Bottles.

Index	Description
0 - 19	right Shadow Hand actuated joint
20 - 22	right Shadow Hand base translation
23 - 25	right Shadow Hand base rotation
26 - 45	left Shadow Hand actuated joint
46 - 48	left Shadow Hand base translation
49 - 51	left Shadow Hand base rotation

**Cost** The constraint of this environment is that we can't touch other bottles when we pick the bottle. When our hand, or the bottle we picked, touches other bottles, the cost is set to 1, otherwise it is 0.

### B.6 JENGA

Jenga is a fitness game which is very suitable for Safe RL algorithm evaluation. Players take turns removing one block at a time from a tower made up of many blocks. In this environment, we need to remove the one we want from the 16 blocks without knocking over the others.

**Jenga** The 411-dimensional observational space for Jenga task is shown in Table 13. It should be noted that since the base of the dual hands in this task is fixed, the observation of the dual hands is compared to the Table 4 of reduced 24 dimensions.

Table 13: Observation space of Jenga.

Index	Description
0 - 397	dual hands observation shown in Table 4
398 - 404	object pose
405 - 407	object linear velocity
408 - 410	object angle velocity

**Actions** The 52-dimensional action space for one hand in Jenga task is shown in Table 14.

Table 14: Action space of Jenga.

Index	Description
0 - 19	right Shadow Hand actuated joint
20 - 22	right Shadow Hand base translation
23 - 25	right Shadow Hand base rotation
26 - 45	left Shadow Hand actuated joint
46 - 48	left Shadow Hand base translation
49 - 51	left Shadow Hand base rotation

**Reward** For timestep  $t$ , let  $x_{b,t}$  as the position of the left middle finger,  $x_{g,t}$  as the position of the left end of the object, and  $d_{p,t} = \|x_{b,t} - x_{g,t}\|_2$ . Define  $d_{y,t}$  as the y-axis direction of the position of the object center, the reward is defined as follows:

$$r_t = 30 * (d_{y,t} + 0.6) - d_{p,t} \quad (8)$$

**Cost** The constraint of this environment is that we can not touch other blocks in the Jenga. The cost is 1 if all blocks move more than 0.01 cm, and 0 otherwise.

### B.7 CLEAN HOUSE

This environment is in a scene we usually clean at home. We need to control the broom with both hands to sweep the trash from the ground into the dustpan without touching other furniture (e.g. chairs).

**Observations** The 431-dimensional observational space for Clean House task is shown in Table 15. It should be noted that since the base of the dual hands in this task is fixed, the observation of the dual hands is compared to the Table 4 of reduced 24 dimensions.

Table 15: Observation space of Clean House.

Index	Description
0 - 397	dual hands observation shown in Table 4
398 - 404	object pose
405 - 407	object linear velocity
408 - 410	object angle velocity
411 - 417	goal pose
418 - 421	goal rot - object rot
422 - 424	the bottom of broom position
425 - 427	the left handle position of the broom
428 - 430	the right handle position of the broom

**Actions** The 52-dimensional action space for one hand in Clean House task is shown in Table 16.

Table 16: Action space of Clean House.

Index	Description
0 - 19	right Shadow Hand actuated joint
20 - 22	right Shadow Hand base translation
23 - 25	right Shadow Hand base rotation
26 - 45	left Shadow Hand actuated joint
46 - 48	left Shadow Hand base translation
49 - 51	left Shadow Hand base rotation

**Reward** The reward consists of four parts: the distance from the left hand to the left handle position of the broom, the distance from the right hand to the right handle position of the broom, the object (trash) position to the bottom of broom position, and the distance from the object to the target (dustpan) point. The distance from the object to the target point is given by  $d_{target}$ . The position difference from the left hand to the left handle position of the broom is given by  $d_{left}$ . The position difference from the right hand to the right handle position of the broom is given by  $d_{right}$ . The object position to the bottom of broom position is given by  $d_{bottom}$ . The reward is given by this specific formula:

$$r = 50 - d_{target} * 10 - 5 * d_{left} - 5 * d_{right} \quad (9)$$

**Cost** The constraint of this environment is that we can not damage other furniture when we sweep the floor. So there is a chair in the path of the trash and the dustpan. The cost is 1 when the broom touches the chair and make it move, and 0 otherwise.

## C FOUR ENVIRONMENTS IN SAFE FINGER.

All environments are comes from Safe Finger. The difference between Safe Finger and Safe Joint is whether it is a joint constrained or a finger constrained, which is described as follows:

**Safety Joint.** In these tasks, we constrain the freedom of joint ④ of forefinger (please refer to Figure 6 (a) and (f)). Without the constraint, joint ④ has freedom of  $[-20^\circ, 20^\circ]$ . The safety tasks restrict joint ④ within  $[-10^\circ, 10^\circ]$ . Let  $\text{ang\_4}$  be the angle of joint ④, and the cost is defined as:

$$c_t = \mathbb{I}(\text{ang\_4} \notin [-10^\circ, 10^\circ]). \quad (10)$$

**Safety Finger.** In these tasks, we constrain the freedom of joints ②, ③ and ④ of forefinger (please refer to Figure 6 (b) and (f)). Without the constraint, joints ② and ③ have freedom of  $[0^\circ, 90^\circ]$  and joint ④ of  $[-20^\circ, 20^\circ]$ . The safety tasks restrict joints ②, ③, and ④ within  $[22.5^\circ, 67.5^\circ]$ ,  $[22.5^\circ, 67.5^\circ]$ , and  $[-10^\circ, 10^\circ]$  respectively. Let  $\text{ang\_2}$ ,  $\text{ang\_3}$ ,  $\text{ang\_4}$  be the angles of joints ②, ③,



④, and the cost is defined as:

$$c_t = \mathbb{I}(\text{ang\_2} \notin [22.5^\circ, 67.5^\circ], \text{ or } \text{ang\_3} \notin [22.5^\circ, 67.5^\circ], \text{ or } \text{ang\_4} \notin [-10^\circ, 10^\circ]). \quad (11)$$

Hand over stands for the situation that two Shadow Hands with palms facing up, opposite each other, and an object that needs to be passed in Safe Finger, and it stands the object that needs to be thrown from the vertical hand to the palm-up hand in Safe Finger. Specific information can refer to Appendix B.2.

## D POINT CLOUD

We replace the object state information with point clouds in the case of 128 parallel environments. The point cloud is captured by the depth camera and downsampled to 2048 points. The features are extracted using PointNet (Qi et al., 2017) to a 128-dimensional vector and concatenate with other observations. It can be seen that under the same episode and the same number of environments, the performance of point cloud input is not as good as full state input, but it can also achieve some performance. But also using an RTX 3090 GPU, the point cloud RL has only 200+ fps, and the full state can reach 30000+. In fact, we can only open up to 128 environments when using point clouds. This was a problem with Isaac Gym’s poor parallel support for cameras. We further refined the method to enhance the parallelization of the point cloud extraction in order to close this gap. When compared to Isaac Gym’s original code, the speedup is 1.46 times, going from 232 fps to 339 fps.

## E DETAILS OF BENCHMARK ALGORITHMS

In this section, we review the key steps of typical Safe RL algorithms implemented in this benchmark, which include CPO (Achiam et al., 2017a), PCPO (Yang et al., 2020b), FOCOPS (Zhang et al., 2020b), P3O (Zhang et al., 2022), PPO-Lag (Ray et al., 2019a), TRPO-Lag (Ray et al., 2019a), CPPO-PID (Stooke et al., 2020b), and IPO (Liu et al., 2020). We implemented all of these algorithms and check the correctness but only evaluated some of them in TrustDeHands. Firstly, we will give a brief introduction to these algorithms below and give the hyperparameters of the algorithms we used in our evaluation. Then we have verified our re-implementations in other Safe RL benchmarks.

### E.1 CPO (ACHIAM ET AL., 2017A)

For a given policy  $\pi_{\theta_k}$ , CPO updates new policy  $\pi_{\theta_{k+1}}$  as follows:

$$\pi_{\theta_{k+1}} = \arg \max_{\pi_{\theta} \in \Pi_{\theta}} \mathbb{E}_{s \sim d_{\pi_{\theta_k}}^{\rho_0}(\cdot), a \sim \pi_{\theta}(\cdot|s)} [A_{\pi_{\theta_k}}(s, a)] \quad (12)$$

$$\text{s.t. } J^c(\pi_{\theta_k}) + \frac{1}{1-\gamma} \mathbb{E}_{s \sim d_{\pi_{\theta_k}}^{\rho_0}(\cdot), a \sim \pi_{\theta}(\cdot|s)} [A_{\pi_{\theta_k}}^c(s, a)] \leq b, \quad (13)$$

$$\bar{D}_{\text{KL}}(\pi_{\theta}, \pi_{\theta_k}) = \mathbb{E}_{s \sim d_{\pi_{\theta_k}}^{\rho_0}(\cdot)} [\text{KL}(\pi_{\theta}, \pi_{\theta_k})[s]] \leq \delta. \quad (14)$$

It is impractical to solve the problem (12) directly due to the computational cost. Achiam et al. (2017a) suggest to find some convex approximations to replace the term  $A_{\pi_{\theta_k}}(s, a)$  and  $\bar{D}_{\text{KL}}(\pi_{\theta}, \pi_{\theta_k})$  Eq.(12)-(14). Concretely, Achiam et al. (2017a) suggest to use first-order Taylor expansion of  $J(\pi_{\theta})$  to replace the objective (12) as follows,

$$\frac{1}{1-\gamma} \mathbb{E}_{s \sim d_{\pi_{\theta_k}}^{\rho_0}(\cdot), a \sim \pi_{\theta_k}(\cdot|s)} \left[ \frac{\pi_{\theta}(a|s)}{\pi_{\theta_k}(a|s)} A_{\pi_{\theta_k}}(s, a) \right] = J(\pi_{\theta}) - J(\pi_{\theta_k}) \approx (\theta - \theta_k)^{\top} \nabla_{\theta} J(\pi_{\theta}).$$

Similarly, Achiam et al. (2017a) use the following approximations to turn the constrained policy optimization (12)-(14) to be a convex problem,

$$\frac{1}{1-\gamma} \mathbb{E}_{s \sim d_{\pi_{\theta_k}}^{\rho_0}(\cdot), a \sim \pi_{\theta_k}(\cdot|s)} \left[ \frac{\pi_{\theta}(a|s)}{\pi_{\theta_k}(a|s)} A_{\pi_{\theta_k}}^c(s, a) \right] \approx (\theta - \theta_k)^{\top} \nabla_{\theta} J^c(\pi_{\theta}), \quad (15)$$

$$\bar{D}_{\text{KL}}(\pi_{\theta}, \pi_{\theta_k}) \approx (\theta - \theta_k)^{\top} \mathbf{H}(\theta - \theta_k), \quad (16)$$

where  $\mathbf{H}$  is Hessian matrix of  $\bar{D}_{\text{KL}}(\pi_{\theta}, \pi_{\theta_k})$ , i.e.,

$$\mathbf{H}[i, j] =: \frac{\partial^2}{\partial \theta_i \partial \theta_j} \mathbb{E}_{s \sim d_{\pi_{\theta_k}}^{\rho_0}(\cdot)} [\text{KL}(\pi_{\theta}, \pi_{\theta_k})[s]],$$

Eq.(16) is the second-order approximation of (14).

Let  $\lambda_*, \nu_*$  is the dual solution of the following problem

$$\lambda_*, \nu_* = \arg \max_{\lambda \geq 0, \nu \geq 0} \left\{ \frac{-1}{2\lambda} (\mathbf{g}^\top \mathbf{H}^{-1} \mathbf{g} - 2\nu r + s v^2) + \nu c - \frac{\lambda \delta}{2} \right\};$$

where  $\mathbf{g} = \nabla_{\boldsymbol{\theta}} \mathbb{E}_{s \sim d_{\pi_{\boldsymbol{\theta}_k}}^{\rho_0}(\cdot), a \sim \pi_{\boldsymbol{\theta}}(\cdot|s)} [A_{\pi_{\boldsymbol{\theta}_k}}(s, a)]$ ,  $\mathbf{a} = \nabla_{\boldsymbol{\theta}} \mathbb{E}_{s \sim d_{\pi_{\boldsymbol{\theta}_k}}^{\rho_0}(\cdot), a \sim \pi_{\boldsymbol{\theta}}(\cdot|s)} [A_{\pi_{\boldsymbol{\theta}_k}}^c(s, a)]$ ,  $r = \mathbf{g}^\top \mathbf{H} \mathbf{a}$ ,  $s = \mathbf{a}^\top \mathbf{H}^{-1} \mathbf{a}$ , and  $c = J^c(\pi_{\boldsymbol{\theta}_k}) - b$ .

Finally, CPO updates parameters according to conjugate gradient as follows: if approximation to CPO is feasible, then

$$\boldsymbol{\theta}_{k+1} = \boldsymbol{\theta}_k + \frac{1}{\lambda_*} \mathbf{H}^{-1} (\mathbf{g} - \nu_* \mathbf{a}),$$

else,

$$\boldsymbol{\theta}_{k+1} = \boldsymbol{\theta}_k - \sqrt{\frac{2\delta}{\mathbf{a}^\top \mathbf{H}^{-1} \mathbf{a}}} \mathbf{H}^{-1} \mathbf{a}.$$

## E.2 PCPO (YANG ET AL., 2020B)

Projection-Based Constrained Policy Optimization (PCPO) is an iterative method for optimizing policies in a two-step process: the first step performs a local reward improvement update, while the second step reconciles any constraint violation by projecting the policy back onto the constraint set.

### Reward Improvement.

$$\begin{aligned} \pi_{\boldsymbol{\theta}_{k+\frac{1}{2}}} &= \arg \max_{\pi_{\boldsymbol{\theta}} \in \Pi_{\boldsymbol{\theta}}} \mathbb{E}_{s \sim d_{\pi_{\boldsymbol{\theta}_k}}^{\rho_0}(\cdot), a \sim \pi_{\boldsymbol{\theta}}(\cdot|s)} [A_{\pi_{\boldsymbol{\theta}_k}}(s, a)], \\ \text{s.t. } \bar{D}_{\text{KL}}(\pi_{\boldsymbol{\theta}}, \pi_{\boldsymbol{\theta}_k}) &= \mathbb{E}_{s \sim d_{\pi_{\boldsymbol{\theta}_k}}^{\rho_0}(\cdot)} [\text{KL}(\pi_{\boldsymbol{\theta}}, \pi_{\boldsymbol{\theta}_k})[s]] \leq \delta; \end{aligned}$$

### Projection.

$$\begin{aligned} \pi_{\boldsymbol{\theta}_{k+1}} &= \arg \min_{\pi_{\boldsymbol{\theta}} \in \Pi_{\boldsymbol{\theta}}} D(\pi_{\boldsymbol{\theta}}, \pi_{\boldsymbol{\theta}_{k+\frac{1}{2}}}), \\ \text{s.t. } J^c(\pi_{\boldsymbol{\theta}_k}) &+ \frac{1}{1-\gamma} \mathbb{E}_{s \sim d_{\pi_{\boldsymbol{\theta}_k}}^{\rho_0}(\cdot), a \sim \pi_{\boldsymbol{\theta}}(\cdot|s)} [A_{\pi_{\boldsymbol{\theta}_k}}^c(s, a)] \leq b. \end{aligned}$$

Then, Yang et al. (2020b) follows CPO (Achiam et al., 2017a) uses convex approximation to original problem, and calculate the update rule as follows,

$$\boldsymbol{\theta}_{k+1} = \boldsymbol{\theta}_k - \sqrt{\frac{2\delta}{\mathbf{g}^\top \mathbf{H}^{-1} \mathbf{g}}} \mathbf{H}^{-1} \mathbf{g} - \max \left( 0, \frac{\sqrt{\frac{2\delta}{\mathbf{g}^\top \mathbf{H}^{-1} \mathbf{g}}} \mathbf{a}^\top \mathbf{H}^{-1} \mathbf{g} + c}{\mathbf{a}^\top \mathbf{L}^{-1} \mathbf{a}} \right) \mathbf{L}^{-1} \mathbf{a},$$

where  $\mathbf{L} = \mathbf{I}$  if  $D$  is  $\ell_2$ -norm, and  $\mathbf{L} = \mathbf{H}$  if  $D$  is KL-divergence.

## E.3 FOCOPS (ZHANG ET AL., 2020B)

Zhang et al. (2020b) propose the First Order Constrained Optimization in Policy Space (FOCOPS) that is a two-step approach. We present it as follows.

### Step1: Finding the optimal update policy.

Firstly, for a given policy  $\pi_{\boldsymbol{\theta}_k}$ , FOCOPS finds an optimal update policy  $\pi^*$  by solving the optimization problem (12)-(14) in the non-parameterized policy space.

$$\pi^* = \arg \max_{\pi \in \Pi} \mathbb{E}_{s \sim d_{\pi_{\boldsymbol{\theta}_k}}^{\rho_0}(\cdot), a \sim \pi(\cdot|s)} [A_{\pi_{\boldsymbol{\theta}_k}}(s, a)] \quad (17)$$

$$\text{s.t. } J^c(\pi_{\boldsymbol{\theta}_k}) + \frac{1}{1-\gamma} \mathbb{E}_{s \sim d_{\pi_{\boldsymbol{\theta}_k}}^{\rho_0}(\cdot), a \sim \pi(\cdot|s)} [A_{\pi_{\boldsymbol{\theta}_k}}^c(s, a)] \leq b, \quad (18)$$

$$\bar{D}_{\text{KL}}(\pi_{\boldsymbol{\theta}}, \pi_{\boldsymbol{\theta}_k}) = \mathbb{E}_{s \sim d_{\pi_{\boldsymbol{\theta}_k}}^{\rho_0}(\cdot)} [\text{KL}(\pi, \pi_{\boldsymbol{\theta}_k})[s]] \leq \delta. \quad (19)$$

If  $\pi_{\boldsymbol{\theta}_k}$  is feasible, then the optimal policy for (17)-(19) takes the following form:

$$\pi^*(a|s) = \frac{\pi_{\boldsymbol{\theta}_k}(a|s)}{Z_{\lambda, \nu}(s)} \exp \left( \frac{1}{\lambda} (A_{\pi_{\boldsymbol{\theta}_k}}(s, a) - \nu A_{\pi_{\boldsymbol{\theta}_k}}^c(s, a)) \right), \quad (20)$$

where  $Z_{\lambda,\nu}(s)$  is the partition function which ensures (20) is a valid probability distribution,  $\lambda$  and  $\nu$  are solutions to the optimization problem:

$$\min_{\lambda, \nu \geq 0} \lambda \nu + \nu \tilde{b} + \lambda \mathbb{E}_{s \sim d_{\pi_{\theta_k}}^{\rho_0}(\cdot), a \sim \pi^*(\cdot|s)} [Z_{\lambda,\nu}(s)],$$

the term  $\tilde{b} = (1 - \gamma)(b - J^c(\pi_{\theta_k}))$ .

### Step 2: Projection.

Then, FOCOPS projects the policy found in the previous step back into the parameterized policy space  $\Pi_{\theta}$  by solving for the closest policy  $\pi_{\theta} \in \Pi_{\theta}$  to  $\pi^*$  in order to obtain  $\pi_{\theta_{k+1}}$ :

$$\pi_{\theta_{k+1}} = \arg \min_{\pi_{\theta} \in \Pi_{\theta}} \mathbb{E}_{s \sim d_{\pi_{\theta_k}}^{\rho_0}(\cdot)} [\text{KL}(\pi_{\theta}, \pi^*)[s]].$$

We usually apply stochastic gradient decent to obtain the solution of above  $\theta_{k+1}$ .

### E.4 PPO-LAG

The Lagrangian approach is a standard way to solve CMDP (1), which is also known as primal-dual policy optimization:

$$(\pi_{\star}, \lambda_{\star}) = \arg \min_{\lambda \geq 0} \max_{\pi \in \Pi_{\theta}} \{J(\pi) - \lambda(J^c(\pi) - b)\}. \quad (21)$$

TRPO-Lag and PPO-Lag combine the Lagrangian approach with TRPO and PPO. Concretely, PPO using the following clip term to replace  $J(\pi)$  in (21),

$$\mathcal{L}_{\text{clip}}^r(\pi_{\theta}) = \mathbb{E}_{s \sim d_{\pi_k}^{\rho_0}(\cdot), a \sim \pi_k(\cdot|s)} \left[ -\min \left\{ \frac{\pi_{\theta}(a|s)}{\pi_k(a|s)} A_{\pi_k}(s, a), \text{clip} \left( \frac{\pi_{\theta}(a|s)}{\pi_k(a|s)}, 1 - \epsilon, 1 + \epsilon \right) A_{\pi_k}(s, a) \right\} \right],$$

where  $\pi_k$  is short for  $\pi_{\theta_k}$ . With  $A_{\pi_k}(s, a)$  replacing  $A_{\pi_k}^c(s, a)$  respectively, and obtain  $\mathcal{L}_{\text{clip}}^c$  as follows,

$$\mathcal{L}_{\text{clip}}^c(\pi_{\theta}) = \mathbb{E}_{s \sim d_{\pi_k}^{\rho_0}(\cdot), a \sim \pi_k(\cdot|s)} \left[ -\min \left\{ \frac{\pi_{\theta}(a|s)}{\pi_k(a|s)} A_{\pi_k}(s, a), \text{clip} \left( \frac{\pi_{\theta}(a|s)}{\pi_k(a|s)}, 1 - \epsilon, 1 + \epsilon \right) A_{\pi_k}^c(s, a) \right\} \right].$$

Then, PPO-Lag updates the policy as follows,

$$(\pi_{k+1}, \lambda_{k+1}) = \arg \min_{\lambda \geq 0} \max_{\pi_{\theta} \in \Pi_{\theta}} \{ \mathcal{L}_{\text{clip}}^r(\pi_{\theta}) - \lambda (\mathcal{L}_{\text{clip}}^c(\pi_{\theta}) - b) \}. \quad (22)$$

All of the above terms can be estimated according to the policy  $\pi_k$ . Then PPO-Lag updates the policy according to first-order optimizer as follows

$$\theta_{k+1} = \theta_k + \eta \frac{\partial}{\partial \theta} \left( \mathcal{L}_{\text{clip}}^r(\pi_{\theta}) - \lambda (\mathcal{L}_{\text{clip}}^c(\pi_{\theta}) - b) \right) \Big|_{\theta=\theta_k, \lambda=\lambda_k}, \quad (23)$$

$$\lambda_{k+1} = \lambda_k + \eta (\mathcal{L}_{\text{clip}}^c(\pi_{\theta}) - b) \Big|_{\theta=\theta_k}, \quad (24)$$

where  $\eta > 0$  is step-size.

### E.5 TRPO-LAG

TRPO-Lag shares a similar idea but it is adaptive to TRPO, where TRPO-Lag replaces  $J(\pi_{\theta})$  as follows,

$$J(\pi_{\theta}) \approx J(\pi_{\theta_k}) + (\theta - \theta_k)^{\top} \nabla_{\theta} J(\pi_{\theta}) =: \mathcal{L}^r(\pi_{\theta}). \quad (25)$$

Similarly,

$$J^c(\pi_{\theta}) \approx J^c(\pi_{\theta_k}) + (\theta - \theta_k)^{\top} \nabla_{\theta} J^c(\pi_{\theta}) =: \mathcal{L}^c(\pi_{\theta}), \quad (26)$$

and

$$\bar{D}_{\text{KL}}(\pi_{\theta}, \pi_{\theta_k}) \approx (\theta - \theta_k)^{\top} \mathbf{H}(\theta - \theta_k), \quad (27)$$

where  $\mathbf{H}$  is Hessian matrix of  $\bar{D}_{\text{KL}}(\pi_{\theta}, \pi_{\theta_k})$ , i.e.,

$$\mathbf{H}[i, j] =: \frac{\partial^2}{\partial \theta_i \partial \theta_j} \mathbb{E}_{s \sim d_{\pi_{\theta_k}}^{\rho_0}(\cdot)} [\text{KL}(\pi_{\theta}, \pi_{\theta_k})[s]].$$

Then, TRPO-Lag updates the policy as follows,

$$(\pi_{k+1}, \lambda_{k+1}) = \arg \min_{\lambda \geq 0} \max_{\pi_{\theta} \in \Pi_{\theta}} \{ \mathcal{L}^r(\pi_{\theta}) - \lambda (\mathcal{L}^c(\pi_{\theta}) - b) \} \Big|_{\theta=\theta_k, \lambda=\lambda_k}, \quad (28)$$

where the policy parameter  $\theta$  satisfies the following condition

$$(\theta - \theta_k)^{\top} \mathbf{H}(\theta - \theta_k) \leq \delta.$$

### E.6 P3O (ZHANG ET AL., 2022)

P3O solves the cumbersome constrained policy iteration via a single minimization of an equivalent unconstrained problem as follows,

$$\pi_{k+1} = \arg \min_{\pi \in \Pi_{\theta}} \left\{ \mathbb{E}_{s \sim d_{\pi_k}^{\rho_0}(\cdot), a \sim \pi_k(\cdot|s)} \left[ \frac{\pi(a|s)}{\pi_k(a|s)} A_{\pi_k}(s, a) \right] + \kappa B(\pi, b) \right\}, \quad (29)$$

where  $\kappa$  is a positive scalar, and the penalty term  $B(\pi, b)$  is defined as follows,

$$B(\pi, b) = \max \left\{ 0, \mathbb{E}_{s \sim d_{\pi_k}^{\rho_0}(\cdot), a \sim \pi_k(\cdot|s)} \left[ \frac{\pi(a|s)}{\pi_k(a|s)} A_{\pi_k}^c(s, a) \right] + (1 - \gamma)(J^c(\pi_k) - b) \right\}. \quad (30)$$

P3O utilizes a simple yet effective penalty approach to eliminate cost constraints and removes the trust-region constraint by the clipped surrogate objective.

For the practical implementation, P3O consider the following optimization objective:

$$\mathcal{L}_{\text{P3O}}(\theta) = \mathcal{L}_{\text{P3O}}^r(\theta) + \kappa \max \{0, \mathcal{L}_{\text{P3O}}^c(\theta)\}, \quad (31)$$

where

$$\begin{aligned} \mathcal{L}_{\text{P3O}}^r(\theta) &= \mathbb{E} \left[ -\min \left\{ \frac{\pi_{\theta}(a|s)}{\pi_k(a|s)} A_{\pi_k}(s, a), \text{clip} \left( \frac{\pi_{\theta}(a|s)}{\pi_{\theta_k}(a|s)}, 1 - \epsilon, 1 + \epsilon \right) A_{\pi_k}(s, a) \right\} \right], \\ \mathcal{L}_{\text{P3O}}^c(\theta) &= \mathbb{E} \left[ \max \left\{ \frac{\pi_{\theta}(a|s)}{\pi_k(a|s)} A_{\pi_k}^c(s, a), \text{clip} \left( \frac{\pi_{\theta}(a|s)}{\pi_{\theta_k}(a|s)}, 1 - \epsilon, 1 + \epsilon \right) A_{\pi_k}^c(s, a) \right\} + (1 - \gamma)(J^c(\pi_k) - b) \right], \end{aligned}$$

the notation  $\mathbb{E}[\cdot]$  is short for  $\mathbb{E}_{s \sim d_{\pi_k}^{\rho_0}(\cdot), a \sim \pi_k(\cdot|s)}[\cdot]$ . All of the term in (31) can be estimated according to the samples selected by  $\pi_k$ .

For each round, P3O chooses the parameter adaptively according to the following rule:

$$\kappa \leftarrow \min \{\rho\kappa, \kappa_{\max}\}, \quad (32)$$

where  $\rho > 1$  and  $\kappa_{\max}$  is a positive scalar.

Finally, P3O updates the policy parameter as follows,

$$\theta \leftarrow \theta - \eta \frac{\partial}{\partial \theta} \mathcal{L}_{\text{P3O}}(\theta). \quad (33)$$

### E.7 IPO (LIU ET AL., 2020)

IPO considers the objective with logarithmic barrier functions (Nemirovski, 2004) to learn the safe policy. Concretely, IPO considers the following way to update policy,

$$\pi_{k+1} = \arg \max_{\pi \in \Pi_{\theta}} \{ \mathcal{L}_{\text{clip}}^r(\pi) + \phi(\pi) \}, \quad (34)$$

where the clip objective is  $\mathcal{L}_{\text{clip}}^r(\pi)$ , and  $\phi(\pi)$  is the logarithm barrier function with respect to the CMDP problem,

$$\phi(\pi) = \frac{1}{m} \log \left( b - J^c(\pi) \right), \quad (35)$$

where  $m > 0$  is a hyper-parameter that needs to be tuned.

### E.8 CPPO-PID (STOOKE ET AL., 2020B)

CPPO-PID (Stooke et al., 2020b) also considers the primal-dual policy optimization method to solve the CMP problem,

$$(\pi_{\star}, \lambda_{\star}) = \arg \min_{\lambda \geq 0} \max_{\pi \in \Pi_{\theta}} \{ J(\pi) - \lambda(J^c(\pi) - b) \}. \quad (36)$$

The main difference between CPPO-PID and previous Lagrangian methods is that replacing the update rule (??), CPPO-PID considers PID control technique (Johnson & Moradi, 2005) to update Lagrange multiplier.

First, CPPO-PID updates the parameter  $\theta$  as follows,

$$\theta_{k+1} = \theta_k + \eta \frac{\partial}{\partial \theta} \left( \mathcal{L}_{\text{clip}}^r(\pi_{\theta}) - \lambda (\mathcal{L}_{\text{clip}}^c(\pi_{\theta}) - b) \right) \Big|_{\theta=\theta_k, \lambda=\lambda_k}, \quad (37)$$

where  $\eta > 0$  is step-size.

Then, CPPO-PID updates the parameter  $\lambda$  according to PID method (Algorithm 2 in (Stooke et al., 2020b)),

$$\lambda_{k+1} = \text{PID}(K_P, K_I, K_D, \pi_{\theta_k}, n), \quad (38)$$

where  $K_P, K_I, K_D$  are the parameter needed to be tuned, and  $n$  is the iteration number.

## F CORRECTNESS VERIFICATION

### F.1 TASK LIST

Below are four safety environments that provide interesting tasks for Safe RL. Among them, three are popular safety environments, MuJoCo-Velocity (Todorov et al., 2012), Safety-Gym (Ray et al., 2019a), Bullet-Safety-Gym (Gronauer, 2022). We test the 8 algorithms on 26 tasks in these safety environments to verified our re-implementations on TrustDeHands. The complete list of these tasks is shown in table 17.

Environment	Category	Task
MuJoCo	Velocity	1. Ant-Velocity
MuJoCo	Velocity	2. Hopper-Velocity
MuJoCo	Velocity	3. Swimmer-Velocity
MuJoCo	Velocity	4. Walk2d-Velocity
Safety-Gym	Goal	5. PointGoal1
Safety-Gym	Goal	6. CarGoal1
Safety-Gym	Button	7. CarGoal1
Safety-Gym	Button	8. CarButton1
Safety-Gym	Goal	9. PointGoal2
Safety-Gym	Goal	10. PointButton2
Bullet-Safety-Gym	Circle	11. BallCircle
Bullet-Safety-Gym	Circle	12. CarCircle
Bullet-Safety-Gym	Circle	13. DroneCircle
Bullet-Safety-Gym	Circle	14. AntCircle
Bullet-Safety-Gym	Gather	15. BallGather
Bullet-Safety-Gym	Gather	16. CarGather
Bullet-Safety-Gym	Gather	17. DroneGather
Bullet-Safety-Gym	Gather	18. AntGather
Bullet-Safety-Gym	Reach	19. BallReach
Bullet-Safety-Gym	Reach	20. CarReach
Bullet-Safety-Gym	Reach	21. DroneReach
Bullet-Safety-Gym	Reach	22. AntReach
Bullet-Safety-Gym	Run	23. BallRun
Bullet-Safety-Gym	Run	24. CarRun
Bullet-Safety-Gym	Run	25. DroneRun
Bullet-Safety-Gym	Run	26. AntRun

Table 17: The complete list of all 30 tasks we test our implementations on.

### F.2 TASK PERFORMANCE

we show the performance of our re-implementations on the tasks in 18, 19, 20, 21, 22, 23.



Table 18: Performance on MuJoCo-Velocity. We consider four MuJoCo environments where we attempt to train a robotic agent to move faster. We impose speed limits in our environments, which are calculated using 50% of the *undiscounted* speed attained by an unconstrained PPO agent after training for a million samples. In this table and the following tables, we specify the cost limit in the parenthesis after the task name.

	Swimmer-v3(205.6)		Hopper-v3(1047)		Walker2d-v3(2410)		Ant-v3(2147.5)	
	Reward	Constraint	Reward	Constraint	Reward	Constraint	Reward	Constraint
CPO	8.23 $\pm$ 2.23	25.82 $\pm$ 2.84	161.43 $\pm$ 3.99	83.65 $\pm$ 2.27	550.75 $\pm$ 368.6	82.66 $\pm$ 3.71	1030.34 $\pm$ 1.65	98.07 $\pm$ 5.29
TRPO-Lag	-2.8 $\pm$ 5.65	19.71 $\pm$ 3.14	911.61 $\pm$ 201.04	71.94 $\pm$ 12.39	1077.69 $\pm$ 1.46	79.82 $\pm$ 1.28	1032.12 $\pm$ 13.44	52.33 $\pm$ 21.15
PPO-Lag	-5.42 $\pm$ 5.41	24.66 $\pm$ 1.83	1075.08 $\pm$ 12.3	84.04 $\pm$ 11.31	1083.62 $\pm$ 1.57	99.23 $\pm$ 23.5	1008.73 $\pm$ 27.28	25.46 $\pm$ 10.73
P3O	29.95 $\pm$ 7.65	48.91 $\pm$ 11.34	1084.42 $\pm$ 9.19	87.12 $\pm$ 8.5	1084.93 $\pm$ 60.06	107.5 $\pm$ 37.56	954.65 $\pm$ 48.06	19.56 $\pm$ 9.12
PCPO	22.21 $\pm$ 5.53	49.45 $\pm$ 5.44	183.46 $\pm$ 7.71	88.1 $\pm$ 10.14	473.73 $\pm$ 184.77	115.22 $\pm$ 44.29	987.66 $\pm$ 8.32	63.2 $\pm$ 48.94
FOCOPS	65.01 $\pm$ 9.65	102.01 $\pm$ 22.35	1078.97 $\pm$ 17.74	92.07 $\pm$ 28.85	1165.08 $\pm$ 74.63	174.15 $\pm$ 69.06	1034.57 $\pm$ 12.13	50.72 $\pm$ 13.43
CPPO-PID	2.98 $\pm$ 3.02	26.1 $\pm$ 5.94	1045.29 $\pm$ 27.67	47.57 $\pm$ 26.24	1082.06 $\pm$ 13.2	84.66 $\pm$ 7.67	1049.61 $\pm$ 13.76	80.16 $\pm$ 5.69
IPO	127.26 $\pm$ 0.87	351.76 $\pm$ 4.75	1226.56 $\pm$ 116.8	220.62 $\pm$ 114.17	1485.74 $\pm$ 61.94	481.84 $\pm$ 63.2	1422.43 $\pm$ 414.35	415.49 $\pm$ 379.83

Table 19: Performance on Safety-Gym.

	Safexp-PointGoal1-v0(25.0)		Safexp-PointButton1-v0(25.0)		Safexp-CarGoal1-v0(25.0)		Safexp-CarButton1-v0(25.0)	
	Reward	Constraint	Reward	Constraint	Reward	Constraint	Reward	Constraint
CPO	27.25 $\pm$ 0.11	44.34 $\pm$ 1.63	25.53 $\pm$ 1.77	100.74 $\pm$ 5.47	37.05 $\pm$ 0.23	52.92 $\pm$ 2.51	18.83 $\pm$ 1.82	160.06 $\pm$ 14.32
TRPO-Lag	12.82 $\pm$ 1.27	24.64 $\pm$ 2.24	2.98 $\pm$ 0.99	23.84 $\pm$ 3.4	24.38 $\pm$ 2.61	24.73 $\pm$ 1.93	0.45 $\pm$ 0.83	25.16 $\pm$ 2.14
PPO-Lag	15.7 $\pm$ 4.41	24.55 $\pm$ 6.72	4.49 $\pm$ 1.19	18.69 $\pm$ 3.96	19.02 $\pm$ 6.74	24.17 $\pm$ 9.84	0.93 $\pm$ 0.84	29.78 $\pm$ 12.71
P3O	13.83 $\pm$ 3.56	27.3 $\pm$ 5.96	2.28 $\pm$ 0.64	21.8 $\pm$ 6.1	18.58 $\pm$ 1.23	25.43 $\pm$ 3.31	0.2 $\pm$ 0.56	30.13 $\pm$ 10.03
PCPO	27.24 $\pm$ 0.23	53.03 $\pm$ 1.54	31.33 $\pm$ 0.49	131.04 $\pm$ 3.61	35.33 $\pm$ 1.32	56.84 $\pm$ 2.15	24.03 $\pm$ 2.21	274.06 $\pm$ 18.77
FOCOPS	23.0 $\pm$ 1.33	34.8 $\pm$ 5.1	4.79 $\pm$ 0.89	22.62 $\pm$ 6.3	18.42 $\pm$ 4.01	24.69 $\pm$ 6.79	1.19 $\pm$ 0.63	32.64 $\pm$ 13.28
CPPO-PID	2.84 $\pm$ 3.0	50.25 $\pm$ 32.53	0.04 $\pm$ 1.27	24.3 $\pm$ 21.7	1.82 $\pm$ 3.4	20.82 $\pm$ 16.04	1.01 $\pm$ 1.21	158.76 $\pm$ 83.54
IPO	25.59 $\pm$ 0.24	33.99 $\pm$ 1.29	7.83 $\pm$ 1.09	58.33 $\pm$ 2.7	27.38 $\pm$ 0.69	41.5 $\pm$ 2.07	3.66 $\pm$ 0.08	81.11 $\pm$ 2.55

Table 20: Performance on Bullet-Safety-Gym with agent Ant.

	SafetyAntRun-v0(25.0)		SafetyAntCircle-v0(25.0)		SafetyAntReach-v0(25.0)		SafetyAntGather-v0(25.0)	
	Reward	Constraint	Reward	Constraint	Reward	Constraint	Reward	Constraint
CPO	1787.14 $\pm$ 117.91	32.06 $\pm$ 5.34	633.87 $\pm$ 65.76	81.81 $\pm$ 9.12	36.34 $\pm$ 0.01	52.55 $\pm$ 0.03	0.74 $\pm$ 0.53	0.25 $\pm$ 0.06
TRPO-Lag	2274.23 $\pm$ 6.87	23.89 $\pm$ 0.38	494.43 $\pm$ 62.5	25.81 $\pm$ 3.66	22.26 $\pm$ 0.04	24.21 $\pm$ 0.02	0.74 $\pm$ 0.53	0.25 $\pm$ 0.06
PPO-Lag	2139.49 $\pm$ 33.15	10.96 $\pm$ 8.96	197.0 $\pm$ 55.92	19.27 $\pm$ 14.6	16.35 $\pm$ 0.12	20.78 $\pm$ 0.17	1.48 $\pm$ 0.46	0.21 $\pm$ 0.08
P3O	2161.45 $\pm$ 42.51	23.25 $\pm$ 0.39	341.08 $\pm$ 77.31	18.62 $\pm$ 6.14	15.99 $\pm$ 0.03	25.46 $\pm$ 0.05	2.5 $\pm$ 0.96	0.21 $\pm$ 0.08
PCPO	1519.33 $\pm$ 226.48	58.42 $\pm$ 51.41	353.08 $\pm$ 65.9	138.64 $\pm$ 63.84	34.33 $\pm$ 0.03	55.31 $\pm$ 0.04	3.95 $\pm$ 0.19	0.45 $\pm$ 0.1
FOCOPS	2261.37 $\pm$ 6.42	17.3 $\pm$ 5.44	218.78 $\pm$ 75.52	32.77 $\pm$ 17.84	16.38 $\pm$ 0.16	24.92 $\pm$ 0.05	10.12 $\pm$ 2.62	1.11 $\pm$ 0.18
CPPO-PID	2040.92 $\pm$ 374.26	15.01 $\pm$ 10.81	749.73 $\pm$ 104.48	19.28 $\pm$ 7.33	1.11 $\pm$ 0.05	24.99 $\pm$ 0.25	1.48 $\pm$ 0.46	0.21 $\pm$ 0.08
IPO	2050.19 $\pm$ 9.61	26.36 $\pm$ 1.27	755.53 $\pm$ 103.44	26.23 $\pm$ 1.67	26.25 $\pm$ 0.62	38.71 $\pm$ 0.34	4.91 $\pm$ 0.24	0.46 $\pm$ 0.04

Table 21: Performances on Bullet-Safety-Gym with agent Ball.

	SafetyBallRun-v0(25.0)		SafetyBallCircle-v0(25.0)		SafetyBallReach-v0(25.0)		SafetyBallGather-v0(25.0)	
	Reward	Constraint	Reward	Constraint	Reward	Constraint	Reward	Constraint
CPO	305.54 $\pm$ 262.19	28.19 $\pm$ 4.09	516.05 $\pm$ 10.63	25.02 $\pm$ 0.34	36.34 $\pm$ 0.01	52.55 $\pm$ 0.03	24.22 $\pm$ 0.48	0.58 $\pm$ 0.07
TRPO-Lag	296.85 $\pm$ 392.82	65.5 $\pm$ 60.53	596.27 $\pm$ 134.49	23.97 $\pm$ 0.51	22.26 $\pm$ 0.04	24.21 $\pm$ 0.02	24.22 $\pm$ 0.48	0.58 $\pm$ 0.07
PPO-Lag	574.19 $\pm$ 3.27	16.54 $\pm$ 2.3	428.04 $\pm$ 107.86	5.07 $\pm$ 6.6	16.35 $\pm$ 0.12	20.78 $\pm$ 0.17	19.63 $\pm$ 1.36	0.4 $\pm$ 0.04
P3O	574.43 $\pm$ 5.52	24.46 $\pm$ 1.28	596.43 $\pm$ 40.66	26.5 $\pm$ 1.7	15.99 $\pm$ 0.03	25.46 $\pm$ 0.05	21.82 $\pm$ 0.16	0.43 $\pm$ 0.07
PCPO	642.32 $\pm$ 139.1	131.88 $\pm$ 17.63	284.34 $\pm$ 139.48	64.04 $\pm$ 28.4	34.33 $\pm$ 0.03	55.31 $\pm$ 0.04	20.0 $\pm$ 4.47	0.97 $\pm$ 0.4
FOCOPS	576.68 $\pm$ 3.54	17.03 $\pm$ 1.74	535.54 $\pm$ 8.3	19.73 $\pm$ 5.06	16.38 $\pm$ 0.16	24.92 $\pm$ 0.05	21.63 $\pm$ 1.04	0.96 $\pm$ 0.21
CPPO-PID	558.42 $\pm$ 7.2	0.0 $\pm$ 0.0	766.7 $\pm$ 22.35	45.16 $\pm$ 9.39	1.11 $\pm$ 0.05	24.99 $\pm$ 0.25	19.63 $\pm$ 1.36	0.4 $\pm$ 0.04
IPO	489.38 $\pm$ 3.06	25.79 $\pm$ 0.29	790.14 $\pm$ 122.23	41.77 $\pm$ 15.78	26.25 $\pm$ 0.62	38.71 $\pm$ 0.34	24.05 $\pm$ 0.6	0.61 $\pm$ 0.03

Table 22: Performance on Bullet-Safety-Gym with agent Car.

	SafetyCarRun-v0(25.0)		SafetyCarCircle-v0(25.0)		SafetyCarReach-v0(25.0)		SafetyCarGather-v0(25.0)	
	Reward	Constraint	Reward	Constraint	Reward	Constraint	Reward	Constraint
CPO	660.59 $\pm$ 81.95	27.26 $\pm$ 0.85	516.05 $\pm$ 10.63	25.02 $\pm$ 0.34	8.68 $\pm$ 1.31	43.85 $\pm$ 1.83	23.1 $\pm$ 1.24	1.04 $\pm$ 0.1
TRPO-Lag	730.46 $\pm$ 2.82	25.45 $\pm$ 0.87	596.27 $\pm$ 134.49	23.97 $\pm$ 0.51	1.88 $\pm$ 0.23	23.68 $\pm$ 2.0	23.1 $\pm$ 1.24	1.04 $\pm$ 0.1
PPO-Lag	728.65 $\pm$ 2.35	21.05 $\pm$ 8.66	428.04 $\pm$ 107.86	5.07 $\pm$ 6.6	0.64 $\pm$ 0.16	35.23 $\pm$ 4.14	9.23 $\pm$ 2.39	0.4 $\pm$ 0.05
P3O	733.02 $\pm$ 0.1	25.78 $\pm$ 0.61	596.43 $\pm$ 40.66	26.5 $\pm$ 1.7	1.2 $\pm$ 0.21	25.48 $\pm$ 4.27	9.6 $\pm$ 2.26	0.4 $\pm$ 0.04
PCPO	402.96 $\pm$ 105.44	149.17 $\pm$ 138.0	284.34 $\pm$ 139.48	64.04 $\pm$ 28.4	6.22 $\pm$ 0.38	48.56 $\pm$ 4.14	20.53 $\pm$ 3.56	1.56 $\pm$ 0.38
FOCOPS	732.45 $\pm$ 1.13	12.99 $\pm$ 2.64	535.54 $\pm$ 8.3	19.73 $\pm$ 5.06	1.36 $\pm$ 1.23	31.13 $\pm$ 6.59	13.47 $\pm$ 4.42	0.91 $\pm$ 0.17
CPPO-PID	731.67 $\pm$ 4.24	42.12 $\pm$ 51.86	766.7 $\pm$ 22.35	45.16 $\pm$ 9.39	3.02 $\pm$ 0.32	27.74 $\pm$ 1.96	9.23 $\pm$ 2.39	0.4 $\pm$ 0.05
IPO	710.75 $\pm$ 0.55	58.16 $\pm$ 1.51	790.14 $\pm$ 122.23	41.77 $\pm$ 15.78	6.54 $\pm$ 0.16	27.79 $\pm$ 1.66	10.7 $\pm$ 2.83	0.71 $\pm$ 0.08

Table 23: Performance on Bullet-Safety-Gym with agent Drone.

	SafetyDroneRun-v0(25.0)		SafetyDroneCircle-v0(25.0)		SafetyDroneReach-v0(25.0)		SafetyCarGather-v0(25.0)	
	Reward	Constraint	Reward	Constraint	Reward	Constraint	Reward	Constraint
CPO	660.59 $\pm$ 81.95	27.26 $\pm$ 0.85	615.91 $\pm$ 99.46	51.44 $\pm$ 8.02	23.13 $\pm$ 5.68	20.08 $\pm$ 0.69	8.28 $\pm$ 0.24	1.08 $\pm$ 0.32
TRPO-Lag	730.46 $\pm$ 2.82	25.45 $\pm$ 0.87	809.99 $\pm$ 84.58	23.51 $\pm$ 1.5	20.22 $\pm$ 3.92	21.57 $\pm$ 1.42	7.32 $\pm$ 0.66	0.85 $\pm$ 0.12
PPO-Lag	728.65 $\pm$ 2.35	21.05 $\pm$ 8.66	336.91 $\pm$ 194.8	16.07 $\pm$ 1.97	6.51 $\pm$ 0.7	18.66 $\pm$ 1.01	5.02 $\pm$ 0.8	1.28 $\pm$ 0.08
P3O	733.02 $\pm$ 0.1	25.78 $\pm$ 0.61	302.52 $\pm$ 46.36	21.56 $\pm$ 0.79	4.09 $\pm$ 0.21	18.91 $\pm$ 0.36	5.59 $\pm$ 0.31	1.07 $\pm$ 0.22
PCPO	402.96 $\pm$ 105.44	149.17 $\pm$ 138.0	333.92 $\pm$ 89.46	80.82 $\pm$ 43.12	26.48 $\pm$ 3.21	26.44 $\pm$ 1.46	6.24 $\pm$ 1.13	0.99 $\pm$ 0.59
FOCOPS	732.45 $\pm$ 1.13	12.99 $\pm$ 2.64	312.17 $\pm$ 213.26	25.9 $\pm$ 3.78	15.45 $\pm$ 7.66	24.6 $\pm$ 0.46	4.83 $\pm$ 0.7	2.28 $\pm$ 1.97
CPPO-PID	731.67 $\pm$ 4.24	42.12 $\pm$ 51.86	697.05 $\pm$ 91.0	25.82 $\pm$ 4.45	15.2 $\pm$ 3.83	21.86 $\pm$ 0.9	5.02 $\pm$ 0.8	1.28 $\pm$ 0.08
IPO	710.75 $\pm$ 0.55	58.16 $\pm$ 1.51	871.97 $\pm$ 147.5	24.43 $\pm$ 1.2	18.71 $\pm$ 7.84	23.07 $\pm$ 1.16	6.84 $\pm$ 0.75	0.91 $\pm$ 0.05