# Sim-to-Lab-to-Real: Safe Reinforcement Learning with Shielding and Generalization Guarantees

Kai-Chieh Hsu<sup>a,1</sup>, Allen Z. Ren<sup>b,1</sup>, Duy P. Nguyen<sup>a</sup>, Anirudha Majumdar<sup>b,2</sup>, Jaime F. Fisac<sup>a,2</sup>

<sup>a</sup>Department of Electrical and Computer Engineering, Princeton University, United States <sup>b</sup>Department of Mechanical and Aerospace Engineering, Princeton University, United States

#### Abstract

Safety is a critical component of autonomous systems and remains a challenge for learning-based policies to be utilized in the real world. In particular, policies learned using reinforcement learning often fail to generalize to novel environments due to unsafe behavior. In this paper, we propose Sim-to-Lab-to-Real to bridge the reality gap with a probabilistically guaranteed safety-aware policy distribution. To improve safety, we apply a dual policy setup where a performance policy is trained using the cumulative task reward and a backup (safety) policy is trained by solving the Safety Bellman Equation based on Hamilton-Jacobi (HJ) reachability analysis. In Sim-to-Lab transfer, we apply a supervisory control scheme to shield unsafe actions during exploration; in Lab-to-Real transfer, we leverage the Probably Approximately Correct (PAC)-Bayes framework to provide lower bounds on the expected performance and safety of policies in unseen environments. Additionally, inheriting from the HJ reachability analysis, the bound accounts for the expectation over the worst-case safety in each environment. We empirically study the proposed framework for ego-vision navigation in two types of indoor environments with varying degrees of photorealism. We also demonstrate strong generalization performance through hardware experiments in real indoor spaces with a quadrupedal robot. See https://sites.google.com/princeton.edu/sim-to-lab-to-real for supplementary material.

Keywords: Reinforcement Learning, Sim-to-Real Transfer, Safety Analysis, Generalization

# 1. Introduction

Reinforcement Learning (RL) techniques have been increasingly popular in training autonomous robots to perform complex tasks such as traversing uneven outdoor terrains [1] and navigating through cluttered indoor environments [2]. Through interactions with environments and feedback in the form of a reward signal, robots learn to reach target locations relying on onboard sensing (e.g., RGB-D cameras). In order to achieve good empirical generalization performance in different environments, the robot needs to be trained in multiple environments and collect experiences continuously. Due to tight hardware constraints and high sample complexities of RL techniques, in most cases, the training is performed solely in simulated environments.

However, the robots' performance often degrades sharply when they are deployed in the real world, where there can be substantial changes in environments such as different lighting conditions and noise in robot actuation. This performance drop opens the need for research on Sim-to-Real transfer. The typical approach is to simulate a large number of environments with randomized properties and train the policy to work well across environments, with the expectation that real environments at deployment time will be well captured by the rich distribution of training variations. This technique, namely domain randomization, has helped bridge the Sim-to-Real gap substantially [3–5]. In the field of visual navigation, conditions such as camera poses, scene layout, and wall textures can be randomized. However, previous Sim-to-Real techniques

Email addresses: kaichieh@princeton.edu (Kai-Chieh Hsu), allen.ren@princeton.edu (Allen Z. Ren)

<sup>&</sup>lt;sup>1</sup>Equal contributions in alphabetical order

<sup>&</sup>lt;sup>2</sup>Equal contributions in advising

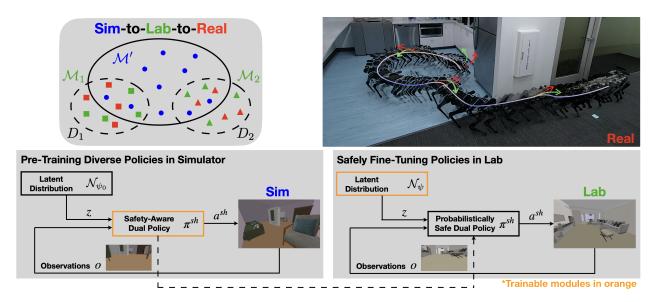


Figure 1: Overview of the Sim-to-Lab-to-Real framework. Top Left: During Sim stage, we train robot policies in a wide variety of environments and conditions M' (Blue circles). Then the same policies from Sim can be fine-tuned in different, more specific settings  $M_{1,2,...}$  (Green triangles/rectangles) during Lab training, which are in the same distribution  $D_{1,2,...}$  of Real environments (Red triangles/rectangles). For example, we may first train using environments of randomized furniture configurations in Sim, and then fine-tune policies in realistic room layouts [6] before deploying in Real indoor spaces. Bottom: In Sim stage, Sim-to-Lab-to-Real trains a safety-aware dual policy conditioned on latent variable sampled from a distribution, and then safely fine-tunes the latent distribution in Lab stage to adapt to a specific environment distribution. Top right: Sample trajectory of a quadrupedal robot running trained policy in a real kitchen environment. The backup policy (Green arrow) overrides the performance policy (Red arrow) when the safety critic value (colored trajectory) exceeds some threshold, steering the robot away from obstacles.

do not explicitly address safety of the robots. Usually, it is worth compromising the performance (e.g., success rate and time needed for reaching the target) to allow better safety of the system (e.g., avoiding dangerous collisions with humans or furniture). While safety violations are inconsequential in simulation, robots trained without safety considerations will tend to exhibit similar unsafe behavior once deployed in real environments. Another drawback of these techniques is that they do not provide any guarantees on robots' performance or safety when they are deployed in different real environments. A "certificate" of robots' generalization performance and safety is necessary before they are deployed in safety-critical environments (e.g., households with children).

In this work, we explore an intermediate training stage between Sim and Real, which we call Lab, that aims to systematically bridge the Sim-to-Real gap by explicitly enforcing hard safety constraints on the robot and certifying the performance and safety before deployment in Real. The proposed Sim-to-Lab-to-Real framework is motivated by the conventional engineering practice whereby, before deploying autonomous systems in the real world, designers usually test them in a realistic but controlled environment, such as a test track for autonomous cars or testing warehouse facilities for quadrupeds at Boston Dynamics [7]. This standard pipeline opens up an opportunity for autonomous systems to further improve performance and safety in the Lab stage. Our insight is that (1) in simulation, where environments can be easily randomized and data is easily collected, the robot can be trained in a wide range of environments and conditions; (2) after that, the robot needs to fine-tune in more specific environments before being deployed in similar ones in the real world; (3) this training stage can also certify the system before Real deployment, especially if the training can provide guarantees on its performance and safety in the real world. Through such extensive training and validation, we can deploy the system confidently in the real environments. Fig. 1 demonstrates the overview of the proposed Sim-to-Lab-to-Real framework.

Fine-tuning in the Lab stage differs from training in the Sim stage in that the Lab stage is more safety-critical. In other words, we want the autonomous system to safely explore in this stage to improve the

performance. In order to realize safe Sim-to-Lab transfer, we need to consider both (1) how safety is formulated and (2) how safety is ensured throughout training. Typical approaches in safe RL combine the safety objective with the performance objective, including adding large negative reward when violating the safety constraints or minimizing the worst-case performance using conditional value at risk (CVaR)

- formulation [8, 9]. However, these methods do not attempt to explicitly enforce hard safety constraints and face a fragile balance between performance and safety. Specifically, these methods require hand-tuning the weights of different components in the objective function, which makes them often fail to generalize well to unseen tasks or environments. Our approach instead builds upon a dual policy setup, where a performance policy optimizes task reward and a backup (safety) policy keeps the robot away from failure conditions.
  - We then apply a least-restrictive control law [10] (or shielding) with the safety state-action value function from the backup policy: the performance policy is only overridden by the backup policy when the safety state-action value function predicts that the proposed performance action would result in an inevitable safety violation in the future. The backup policy is pre-trained in the Sim stage and ready to ensure safe exploration once Lab training starts. Based on the safety policy training developed in [11, 12] using Hamilton-Jacobi (HJ) reachability-analysis, our backup agent can learn from near-failure with dense signals. Unlike previous work that uses binary safety failure indicators [13, 14], our training does not rely on experiencing safety violations, which enables the backup policy to be updated in safety-critical conditions. As we show in Sec. 7.2.1, the number of safety violations is reduced by 4%–77% compared to previous safe R L work in different settings.

We also would like to provide "certificates" on the performance and safety of the robot after Lab training. However, this can be very challenging since typically it is not possible to fully specify the environment distribution where the robot is deployed (e.g., range of wind velocities for drone navigation, or minimum distance between obstacles for home robot navigation). Traditional techniques that provide performance guarantees for control policies, such as from robust control [15, 16] and model-based reachability analysis [17, 18], typically assume an explicit description of such uncertainty affecting the system (e.g., bound on actuation noise) and/or the environment. To tackle the challenge, we apply the Probably Approximately Correct (PAC)-Bayes Control framework [19-21], which provides lower bounds on the expected performance and safety when testing learned policies in unseen environments, while (1) not assuming explicit knowledge of the environment and (2) suited for systems with high-dimensional observations like vision. The framework also naturally fits our setup, as the two training stages of PAC-Bayes Control, prior and posterior, can be assigned to the Sim and Lab stages. As required by the PAC-Bayes setup, we train a distribution of policies by conditioning the performance (and backup) policy on latent variables sampled from a distribution. After training a prior policy distribution in Sim stage, we fine-tune the distribution in Lab, obtaining a posterior policy distribution and its associated generalization guarantee. While previous work in PAC-Bayes Control does not consider explicit policy architecture for safety, we now combine it with HJ reachability analysis and improve the generalization bounds for performance and safety by 40% (Sec. 7.2.2).

#### 1.1. Statement of Contributions

The primary contribution of this work is to propose Sim-to-Lab-Real, a framework that combines HJ reachability analysis and the PAC-Bayes Control framework to improve safety of robots during training and real-world deployment, and provide generalization guarantees on robots' performance and safety in real environments. Additionally, we make the following contributions:

- We propose an algorithm for concurrently training the performance policy that optimizes task reward and a backup policy that follows the Safety Bellman Equation (5) in Sec. 5. We introduce annealing parameters that allow gradual learning of performance and safety in the Sim stage. We also demonstrate that HJ-reachability R L can learn the safety state-action value function end-to-end from images and enable safe exploration with a shielding scheme.
- We propose a modification of off-policy actor-critic algorithms that incorporates the policy distribution regularization from PAC-Bayes Control in Sec. 6. By constraining the KL divergence between the prior and posterior policy distribution in expectation with batch samples and a weighting coeficient, we optimize the generalization bound in the Lab stage efficiently. With a shielding-based policy architecture, we are able to significantly improve the bound compared to previous PAC-Bayes Control works.

• We demonstrate the ability of our framework to reduce safety violations during training and improve empirical performance and safety, as well as generalization guarantees, compared to other safe learning techniques and previous work in PAC-Bayes Control in Sec. 7. We set up ego-vision navigation tasks in two types of environments including one with realistic indoor room layout and visuals. We also validate our approach and generalization guarantees with a quadruped robot navigating in real indoor environments (Sec. 7.2.3).

#### 2. Related Work

95

Safe Exploration. Ensuring safety during training has long been a problem in the reinforcement learning community. On one hand, the RL agent usually needs to experience failure in order to learn to be safe. On the other hand, being too conservative hinders exploring the state/action space suficiently. Constrained MDP (CMDP) is a frequently used framework in safe exploration to satisfy constraints by changing the optimization objective to include some forms of risk [22]. CMDP faces two main challenges: how to incorporate the safety constraints in RL algorithms and how to eficiently solve the constrained optimization problem. Chow et al. [9] use Lagrangian methods to transform the constrained optimization into an unconstrained one over the primal variable (policy) and the dual variable (penalty coeficient). A recent line of works building on reachability analysis argues that optimizing the sum of rewards and penalties is not an accurate encoding of safety [12]. Instead, they encode the safety specification of dynamical systems by finding the optimal safety value function, which is a solution to a Hamilton-Jacobi-Bellman/Isaacs variational inequality [23, 24]. With this safety value function, they apply a least-restrictive control law to shield any performance-oriented policy by overriding with an optimal safe action only when the agent is at states with critically low safety values [10, 25]. Cheng et al. [26] propose a similar shielding framework, utilizing the related control barrier function (CBF) concept. If the system dynamics are control-afine and a CBF is available, a smooth safety override can be computed eficiently by solving a quadratic program with a linear CBF constraint. However, these methods all assume that the dynamics and the environment are at least approximately known. Moreover, they also require that the reachability value function or CBF is available before the learning starts, which is

To mitigate the curse of dimensionality and address generalization to novel environments, we build upon reachability R L [11, 12], which finds an approximate safety value function. Recent methods in [13, 14, 27, 28] address the safety problem by similar learning-based methods and shielding schemes as proposed in this work. However, the major differences lie in how the safety state-action value function, or safety critic, is trained and where the backup actions come from. Dalal et al. [27] assume the safety of systems can be ensured by adjusting the action in a single time step (no long-term effect). Thus, they learn a linear safety-signal model and formulate a quadratic program to find the closest control to the reference control such that the safety constraints are satisfied. Srinivasan et al. [13] and Thananjeyan et al. [14] learn the safety state-action value function from only sparse and binary safety labels. Srinivasan et al. [13] use this function to filter out the unsafe actions from the performance policy and resample actions until the backup agent deems the proposed actions safe, while Thananjeyan et al. [14] let the backup agent directly take over. The concurrent work [28] uses the same reachability RL to learn the backup agent. Our method is distinct in that (1) we propose the two-stage training to further reduce the safety violations in training and (2) we train the reachability RL end-to-end from images without pre-training the visual encoder. We compare our reachability critic with risk critic [13, 14] as detailed in Sec. 7.2.1. Our method reduces the number of safety violations by up to 77% in Lab training and 38% in testing.

non-trivial for high-dimensional dynamics and/or unknown environments.

A different line of works use learning-based methods to capture the residual error between the nominal model and the real dynamics, which results from model mismatch and/or uncertainties. Then, they combine learned models with model-based RL or model predictive control (MPC) to allow safe exploration. Berkenkamp et al. [29] use Gaussian process (GP) to estimate the performance of control parameters and they only deploy parameters that are predicted to be higher than a predefined threshold. On the other hand, Koller et al. [30] use GP to estimate the residual error and then utilize this model to over-approximate the forward-reachable set (FRS). They formulate a terminal constraint in MPC to only deploy policies whose

FRS reaches a known control-invariant set (under some safety controllers). Liu et al. [31] learn the unknown residual with regression and quantify the residual error bound by formulating a covariance shift problem. Our method is distinct in that we use a model-free approach since we only assume we have high-dimensional measurements like RGB images, which cannot be easily handled with model-based methods. Secondly, we do not assume having access to a safe set a priori.

Generalization Theory and Guarantees. In supervised learning, generalization theory provides a principled guarantee on the true expected loss on new samples drawn from the underlying (but unknown) data distribution, after training a model using a finite number of samples. Foundational frameworks include Vapnik-Chervonenkis (VC) theory [32] and Rademacher complexity [33]; however the resulting bounds are generally extremely loose for neural networks. More recent approaches based on PAC-Bayes generalization theory [34] have provided non-vacuous bounds for neural networks in supervised learning [35, 36]. Majumdar et al [19] apply the PAC-Bayes framework in policy learning settings and provide generalization guarantees for control policies in unseen environments. Follow-up work has provided strong guarantees in different robotics settings including for learning neural network policies for vision-based control [21, 37–40]. Unlike supervised learning settings where picking a PAC-Bayes prior can be dificult, previous work in control settings has encoded different domain knowledge in the prior, such as diverse trajectories from human demonstrations [37].

Our work also encodes diverse navigation strategies into the prior through maximum entropy learning [41]. In addition, previous work has not adopted safety-related policy architectures nor considered safety during training. Combining PAC-Bayes theory with reachability safety analysis, we are able to provide stronger guarantees on performance and safety.

Safe Visual Navigation in Unseen Environments. Robot navigation has witnessed a long history of research [42], and many of the approaches have focused on explicit mapping of the environment combined with longhorizon planning in order to reach a goal location [43, 44]. Some recent works apply a map-less approach [2, 45] or builds a map-like belief of the world [46] instead. They often take an end-to-end learning approach and start to tackle generalization to previously unseen environments. Similar to them, we train from pixels to actions, and use RGB images as the policy input without any depth information or mapping of the environment. Furthermore, we place more emphasis on the safety of the robot; we aim to train the robot avoiding any collision with obstacles and reaching some target location without the need of explicit mapping (e.g., initial and target locations can be in the same living room). There has been work that explicitly aims to improve safety of the navigating agent. A popular approach is to detect any novel environment or location (often using a neural network) and resort to conservative actions when novelty is detected [47, 48]. A slightly different approach is to estimate the uncertainty of the policy output and act cautiously when the policy is uncertain where to go [49, 50]. However, these work learn the notion of novelty and uncertainty purely from data, often in the form of binary signals, which can be sample ineficient and not generalizable to unseen domains. Closer to our work, there has been a line of work in applying Hamilton-Jacobi reachability analysis in visual navigation. Bajcsy et al [51] solves for the reachability set at each step but relies on a map generated using onboard camera. Li et al [52] proposes supervising the visual policy using expert data generated by solving a reachability problem. As detailed in the following section, our work also leverages reachability analysis but does not build a map of the environment nor relies on ofline data generated by a different (expert) agent.

Adaptive Sim-to-Real Transfer. Directly applying policies trained in simulation to real environments can lead to bad performance and safety, and there has been work that adaptively bridges the Sim-to-Real gap. One line of work addresses the mismatch in robot and environment dynamics by explicitly searching for simulation parameters (e.g., mass, friction coeficient) that result in trajectories matching the real rollouts [53–55]. Mehta et al. [56] propose active domain randomization that looks for simulation parameters that leads to different trajectories than reference ones, and those parameters are deemed important to train upon. A different approach [57] searches for simulation parameters by directly optimizing task reward in real environments without matching the dynamics. A work closer to ours is Multi-Fidelity R L by Cutler

et al. [58], in which lower-fidelity environment (i.e., simulation) determines exploration heuristics for higher-fidelity environment (i.e., real world), and higher-fidelity environment learns model parameters for lower-fidelity environment. In a similar spirit, we learn safety-aware policy in lower-fidelity simulation for safer exploration in the Lab stage, where the policy distribution is fine-tuned. An important distinction of our work from previous ones is that we jointly address the Sim-to-Real gap in robot perception, environment configuration and dynamics. In addition, we provide probabilistic guarantees on the performance and safety of policies being deployed in real environments.

#### 95 3. Problem Formulation and Preliminaries

We consider a robot with discrete-time dynamics given by

$$s_{t+1} = f_E(s_t, a_t),$$
 (1)

with state s 2 S 2 R  $n_s$ , control input a 2 A 2 R  $n_s$ , and environment E 2 E that the robot interacts with (e.g., an indoor space with furniture including initial and goal locations of the robot). Below we introduce the different conditions of the environments considered in the three stages. See Figure. 1a for visualization.

Environment - Sim. In the Sim stage, we assume there is a set of training environments  $M^{'} \supseteq E$  (e.g., synthetic indoor spaces with randomized arrangement of furniture),  $M^{'} := \{E_1, E_2, \cdots, E_{N^{'}}\}$ . There is no assumption on how  $M^{'}$  is distributed in E.

Environment - Lab. In the Lab stage, we are concerned with more specific conditions, and there can be different distributions of environments  $D_1, D_2, \ldots$  (e.g., ofice or home spaces, dimensions of the obstacles), with which the policies trained in Sim can be fine-tuned. We assume no explicit knowledge of each distribution  $D_i$ ; instead, we assume there is a set of  $N_i$  training environments drawn i.i.d. from  $D_i$  available for the robot to train in; we denote these training datasets by  $M_i := \{E_1, E_2, \cdots, E_N \}$   $\mathbb{D}_{i_1}$ . With a slight abuse of notations and for convenience, we consider a single target condition when introducing the rest of formulation and the approach, and denote the concerned distribution D, the training set M, and the number of training environments N.

Environment - Real. In the Real stage, we assume the robot is deployed in environments from the same distribution D but unseen during the Lab stage.

Next we introduce the rest of problem settings including the robot sensor, the policy, and robot's task involving the reward function and the failure set. These settings hold the same for all three stages, except for the failure set which we do not require knowledge of at Real deployment.

Sensor. In all environments, we assume the robot has a sensor (e.g., RGB camera) that provides an observation  $o = h_E(s)$  using a sensor mapping  $h : S \times E \rightarrow O$ .

Task and Policy. Suppose the robot's task can be defined by a reward function, and let  $R_E(\pi)$  denote the cumulative reward gained over a (finite) time horizon by a deterministic policy  $\pi:O\to A$  when deployed in an environment E. We assume the policy  $\pi$  belongs to a space  $\Pi$  of policies. We also allow policies that map histories of observations to actions by augmenting the observation space to keep track of observation sequences. We assume  $R_E(\pi)$  [0, 1], but make no further assumptions such as continuity or smoothness. We use  $\xi_E^{s,\pi}:[0,T]\times E\to S$  to denote the trajectory rollout from state s using policy  $\pi$  in the environment E up to a time horizon T.

#### 3.1. Goal

230

Our goal is to learn policies that provably generalize to unseen environments at the Real stage. This is very challenging since we do not have explicit knowledge of the underlying distribution D. We employ a slightly more general formulation where a distribution P over policies  $\pi \ 2 \ \Pi$  instead of a single policy is used. In addition to maximizing the policy reward, we want to minimize the number of safety violations, i.e., the number of times that the robot enters failure sets. Our goal can then be formalized by the following optimization problem, which we would like to lower bound as the guarantee:

$$R^{2} := \sup_{P \supseteq P} R_{D}(P), \text{ where } R_{D}(P) := \mathop{\mathbb{E}}_{E \supseteq D} \mathop{\mathbb{E}}_{\pi \boxtimes P} R_{E}(\pi) , \qquad (2)$$

$$R_{E}(\pi) := \mathop{\mathbb{E}}_{E}(\pi) \mathbf{1} \ \ \exists t \ \ \exists \ [0,T], \xi^{s,\pi}(t) \ \not \exists \ F_{E} , \qquad (3)$$

$$R_{E}(\pi) := \overline{R}_{E}(\pi) \mathbf{1} \quad \exists t \, \exists \, [0,T], \, \xi^{s,\pi}(t) \not \exists \, F_{E}, \qquad (3)$$

where  $\overline{R_E}(\pi)$   $\mathbb{P}[0,1]$  denotes the task reward that does not penalize safety violation, P denotes the space of probability distributions on the policy space  $\Pi$ , and  $1\{\cdot\}$  is the indicator function. Here the task reward can be either dense (e.g., normalized cumulative reward) or sparse (e.g., reaching the target or not).

#### 3.2. Generalization Bounds

Recently, PAC-Bayes generalization bounds have been applied to policy learning settings in order to provide formal generalization guarantees in unseen environments. We briefly introduce this framework here, as it will be used in our overall approach presented in Section 4. First it requires training a prior policy distribution Po, which we do in the Sim stage with the set of environments M'. Then in the Lab stage, we fine-tune Po with environments M to obtain the posterior distribution P. Now, define the empirical reward of P as the average expected reward across training environments in M:

$$R_{M}(P) := \frac{1}{N} \sum_{E \supseteq M}^{X} \sum_{\pi \supseteq P}^{h} R_{E}(\pi) . \tag{4}$$

The following theorem can then be used to lower bound the true expected reward  $R_D(P)$ .

Theorem 1 (PAC-Bayes Bound for Control Policies; adapted from [19]). Let Po 2 P be a prior distribution. Then, for any P  $\mathbb{P}$  P, and any  $\delta \mathbb{P}$  (0, 1), with probability at least  $1-\delta$  over sampled environments M  $\mathbb{P}$  D  $\mathbb{P}$ , the following inequality holds:

$$R_D(P) \ge R_{PAC}(P, P_0) := R_M(P) - \frac{D}{C(P, P_0)}, \text{ where } C(P, P_0) := \frac{KL(P \boxtimes P_0) + \log(\frac{2^{\sqrt{N}}}{\delta})}{2N},$$

and KL(P | | Q) stands for Kullback-Leibler (KL) divergence between probability distribution P and Q.

Maximizing the lower bound  $R_{PAC}$  can be viewed as maximizing the empirical reward  $R_M(P)$  along with a regularizer C that prevents overfitting by penalizing the deviation of the posterior P from the prior  $P_0$ . By fine-tuning  $P_0$  to P and maximizing the bound in the Lab stage, we can provide a generalization guarantee for trained policies in unseen environments in the Real stage.

Remark 1. In exchange for assuming almost nothing about the environment distribution D and providing statistical guarantees that hold in arbitrarily high confidence  $(1 - \delta)$  instead of only in expectation over sampled environments (e.g., conformal prediction [59]), the PAC-Bayes framework requires at least a few hundred Lab environments (N ≥ 100) to achieve reasonably tight generalization bounds. This requires substantial resources for training the policies in the Lab stage. In this work we use simulated environments for Lab training, but we envision that training in real environments is well scalable for industry practitioners with extensive training resources. Please refer to Sec. 8 for more discussion.

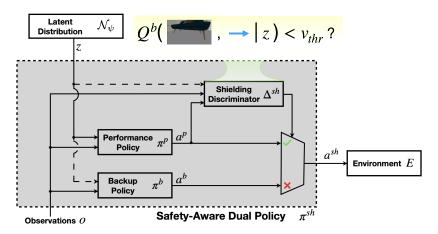


Figure 2: Architecture of the safety-aware policy distribution: we consider a dual policy setup where the performance policy  $\pi_p$  (and backup policy  $\pi_b$ , optionally) is conditioned on latent variables sampled from a distribution encoding diverse behavior. The safety state-action value function  $Q_b(o_t, a_t)$  from the backup policy is used as the shielding discriminator  $\Delta^{sh}$ , which determines whether the proposed action by the performance policy,  $a_p$ , is safe. The action from the backup policy,  $a_b$ , overrides only if necessary.

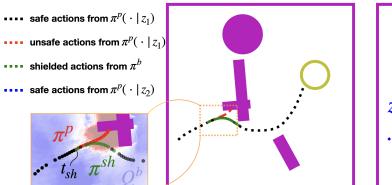
#### 4. Method Overview

Our proposed Sim-to-Lab-to-Real framework bridges the reality gap with probabilistic guarantees by learning a safety-aware policy distribution Fig. 2 shows the architecture of the safety-aware dual policy. It explicitly handles safety through the use of a shielding discriminator, which monitors the candidate actions from the performance policy and overrides them with backup actions only when deemed necessary. We condition the performance policy  $\pi_p$  (and the backup policy  $\pi_b$  and the shielding discriminator  $\Delta^{sh}$ ) on a latent variable z sampled from some distribution, encoding different trajectories to follow (and different shielding strategies to take). With these tools, we divide the Sim-to-Real gap into two components, i.e., Sim-to-Lab and Lab-to-Real, which we tackle by a two-stage training pipeline as shown in Fig. 1. We show how to jointly train a dual policy conditioning on a latent distribution in Sec. 5. The details of Lab training and derivations of generalization guarantees are provided in Sec. 6.

For training, we use a proxy reward function  $r_E:S\times A\times E\to R$ , such as dense reward in distance to target, as a single-step surrogate for the task reward  $\overline{R_E}(\pi)$ . Additionally, for every interaction with the environment, the robot receives a safety cost  $g_E(s)$  (e.g., distance to nearest obstacle). We train both performance and backup policies with modifications of the off-policy Soft Actor-Critic (SAC) algorithm [60]. We denote the neural network weights of the actor and the critic  $\theta$  and w. We use superscripts  $(\cdot)^p$  and  $(\cdot)^b$  to denote critics, actors, and actions from the performance or backup agent. In order to parameterize the policy distribution, we condition the performance (and the backup) policy on a latent variable  $z \ \mathbb{R}^{n_z}$ . We assume the latent variable is sampled from a multivariate Gaussian distribution with diagonal covariance as  $z \ \mathbb{R}^n$   $(\mu, \Sigma)$ , where  $\mu \ \mathbb{R}^n$  is the mean and  $\Sigma \ \mathbb{R}^n$   $\mathbb{R}^n$  is the diagonal covariance matrix. For notational convenience, we denote  $\sigma \ \mathbb{R}^n$  the element-wise square-root of the diagonal of  $\Sigma$ , and define  $\psi = (\mu, \sigma)$ ,  $N_{\psi} := N\left(\mu, \text{diag}(\sigma^2)\right)$ . This parameterization enables our framework to quantify the difference between the policy distribution after Sim training and Lab training, by which we can use PAC-Bayes Control to give probabilistic guarantees.

#### 5. Pre-Training a Diverse Dual Policy in Simulation

The goal of the first training stage is to train the dual policy jointly with the fixed latent distribution in simulation, where training is not safety-critical (safety violations are not restricted). In this training stage, we use the environment dataset M that contains environments that are not necessarily similar to those from the target environment distribution D. Similar to domain randomization techniques, we use environments



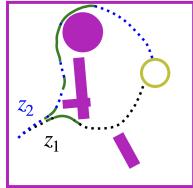


Figure 3: Rollout trajectories of the safety-aware policy distribution: the latent variables sampled from the distribution induce a diverse exploration motives and value-based shielding manages to override the unsafe actions. Red dashed line shows the unshielded actions; Black/Blue dotted lines show the safe actions by the performance policy; Green lines show the backup actions overriding unsafe actions. The inset shows safety values  $Q(o, \pi^b(o))$  with the observation o taken when the heading angle fixed to the one at time instant  $t_{ch}$ .

and conditions with randomized properties, such as random arrangement of furniture in indoor space and random camera tilting angle on the robot.

In the following subsections, we first review how to learn a backup policy by reachability R L optimizing for the worst-case safety. Then, we propose a shielding scheme with physical meaning to override unsafe candidate actions proposed by the performance policy. Additionally, we incorporate information-theoretic objectives to induce diversity into the learned policy distribution, which helps with fine-tuning the policy distribution and achieve stronger generalization guarantees in the next training stage. Finally, we show how to jointly train two agents, performance and backup, to realize all the above-mentioned goals.

# 5.1. Safety through Reachability Reinforcement Learning

Failures are usually catastrophic in safety-critical settings; thus worst-case safety, instead of an average safety over the trajectory, should be considered. For training the backup policy, we incorporate tools from reachability reinforcement learning [11, 12] and optimize the discounted safety Bellman equation (DSBE) as below,

$$Q^{b}(o_{t},a_{t}):=(1-\gamma)g_{E}(s_{t})+\gamma\max_{g_{E}(s_{t})}^{n}\max_{a_{t+1}\boxtimes A}^{n}Q^{b}o_{t+1},a_{t+1}^{n}, \tag{5}$$

where  $o_t = h_E(s_t)$  and  $\gamma$  is the discount factor. This discount factor represents how much attention the RL agent places on future outcomes: if  $\gamma$  is small, the RL agent only cares about "imminent danger", and as  $\gamma \rightarrow 1$ , one recovers the infinite-horizon safety state-action value function. In the training, we initialize  $\gamma = 0.8$  and gradually anneal  $\gamma$  towards 1 during the process.

The safety state-action value function in (5) captures the maximum cost  $g_E$  along the trajectory starting from  $s_t$  with action  $a_t$  assuming that the safest control input is applied at every instant thereafter. Thus,  $\min_{a_t \, \Box A} \, Q(o_t, a_t) > 0$  indicates that the robot is predicted to inevitably violate safety in the future if  $a_t$  is taken. By utilizing this (annealed) DSBE, we have an exact encoding of the property we want our system to satisfy. The DSBE allows the backup agent to learn the safety state-action value function from near-failure executions, which significantly reduces failure events during training. Additionally, the DSBE enables the backup agent to update using a dense learning signal, which is suitable for the joint training of performance and backup agents. To our knowledge, this work demonstrates the first instance of reachability RL in fully end-to-end training with extremely high-dimensional inputs (RGB images), without the need for pre-training a vision encoder as in [28].

#### 5.2. Shielding

We leverage a least-restrictive control law, i.e., shielding, to reduce the number of safety violations in both training and deployment. Suppose we have two policies: performance-pursuing policy  $\pi^p$  and safety-pursuing (backup) policy  $\pi^b$ . Before we apply a candidate action from the performance-pursuing policy, we use a shielding discriminator  $\Delta^{sh}$  to check if it is safe. We replace the proposed action with the action from the backup policy if and only if that candidate action is deemed to result in safety violations in the future. The shielding criterion is summarized in (6). This ensures minimum intervention by the backup policy while the performance policy guides the robot towards the target [10, 61].

$$\pi^{sh}(0) =$$
 $\frac{\pi^{p}(0), \quad \Delta^{sh}(0, \pi^{p}, Q^{b}) = 1}{\pi^{b}(0), \quad \text{otherwise}}$ 
(6)

The safety value function learned by DSBE represents the maximum cost along the trajectory in the future if following the learned policy. If we define the safety margin function  $g_E(s)$  to be the closest distance to the obstacles, then  $Q^b(o,a)$  represents the closest distance of the robot to the obstacles in the future. Based on this, we propose a value-based shielding with the threshold having a physical interpretation, i.e., a margin from the failure. Once the robot receives the current observation o and uses performance policy to generate action  $\pi^p(o)$ , the backup policy overrides the action if and only if  $Q^b(o,a^p) > v_{thr}$ . In other words, the shielding discriminator is defined as below

$$\Delta^{sh}(0, \pi^p, Q^b) := 1Q^b(0, \pi^p(0)) \le v_{thr} . \tag{7}$$

Fig. 3 shows an example of shielding that prevents applying unsafe actions from the performance policy (replace the red dotted lines with green dotted lines in the inset). We compare the safety state-action value function based on DSBE with ones by sparse safety indicators [13, 14] in Sec. 7.2.1 and Fig. 6; our approach affords much better safety during training and deployment.

#### 5.3. Diversity through Maximization of Latent-based Mutual Information

During Sim training, we also maximize the diversity of robot behavior encoded by the latent distribution, which has shown in different work [37, 41] to result in better performance after fine-tuning the distribution, which we perform in the Lab stage. Each latent variable is sampled from the distribution. As the policy is conditioned on the latent variable, it should lead to different trajectories around obstacles and towards the target (Fig. 3). With a single policy instead of a distribution, it is prone to overfit to some set of environments and fails to adapt in new environments (Fig. 11).

In order to distinguish the resulting trajectories using latent variables, we maximize mutual information between observations of trajectories  $\xi_o$  and latent variables z, which can be lowered bounded by sum of mutual information between each observation and the latent variable  $I(\xi_o;z) \geq \prod_{t=1}^T I(o_t;z)$  [62] (observation-marginal MI). We can further lower bound  $I(O;Z) \geq E_{z \boxtimes N_{\psi_o},o \boxtimes p(\cdot \mid \pi^{sh},z)}[\log q_{\phi}(z \mid o)] - E_{z \boxtimes N_{\psi_o}}[\log p(z)]$ , where the posterior  $p(z \mid o)$  is approximated by a learned discriminator  $q_{\phi}(z \mid o)$ , parameterized by a neural network with weights  $\phi$  [41]. Intuitively, in order to make trajectories recognizable by the discriminator, the trajectories need to be diverse. Similar to [41], before updating the policies after sampling a batch of experiences, we augment the proxy reward by a weighted mutual information reward with coeficient  $\beta$ :

$$r_{aug}(s_t, a_t, o_t, z) = r(s_t, a_t) + \beta \log q_{\phi}(z|o_t) - \log p(z).$$
 (8)

This encourages the value function to assign higher reward to regions more recognizable by the discriminator. Concurrently, we train the discriminator by maximizing log  $q_{\phi}(z|o)$  with Stochastic Gradient Descent (SGD),

As shown in Fig. 2, we can additionally condition backup policies with the latent distribution; the robot may avoid obstacles in different directions, and such skills might be beneficial when there is a distributional shift of obstacle placement and geometry in the Lab stage.

The backup agent can also depend on a latent variable. Since the backup agent can intervene at any state and condition on any latent, we instead optimize the conditional mutual information between action and latent given the current observation I(A; Z|O) (observation-conditional MI). We modify the backup policy training objective (5) as below

$$h \qquad i \\ \theta^{b^{\otimes}} = \arg\min_{\theta} L(\theta) := E_{o,z} E_{a^{\otimes} g(\cdot|o,z)} Q(o,a;z) - vI(A;Z|O), \tag{10}$$

where the Q-function is now conditioned on a latent variable and  $\nu$  is the coeficient balancing the safety cost and the diversity. Through derivations in Appendix A, we modify the SAC formulation and the backup actor is updated as,

$$\theta \leftarrow \theta - \mathbb{Z}_{\theta} \ E_{(o,z)\mathbb{Z}B,a\mathbb{Z}} \ \frac{\theta}{\pi(\cdot \mid o,z)} Q(o,a;z) - \nu \log \pi_{\theta}(a \mid o,z) + \nu \log \frac{1}{n_{s}} \ \frac{X_{h_{s}}}{n_{s}} \ \pi_{\theta}(a \mid o,z_{i}), \tag{11}$$

where B is the replay buffer. Intuitively, for specific action a given current observation o, we want it to have high probability for policy conditioned on a specific latent variable z and low probability for other latent variables  $\{z_i\}$  sampled from the distribution. Note that when the backup agent is also conditioned on latent variable z, the shielding discriminator in (7) now becomes  $Q^b(o,\pi^p(o),z) \leq v_{thr}$ . While similar formulations have been explored in previous work [41, 63, 64] to achieve diverse trajectories/skills in RL, to our best knowledge, we are the first to consider a continuous distribution of latent variables instead of a discrete one. We find this brings difficulty in training, exacerbated by using robot observations instead of true states (e.g., ground-truth locations of the robot); nonetheless, we show effectiveness of such diversity-induced training in Sec. 7.2.4.

# 5.4. Joint Training of Performance and Backup Policies.

355

Now we are ready to perform joint training of the dual policy. In the Sim stage, we fix the latent distribution to be a zero-mean Gaussian distribution with diagonal covariance  $N_{\psi_0}$ , where  $\psi_0$  =  $(0,\sigma_0)$ . For each episode during training, we sample a latent variable z  $\mathbb{Z}$   $N_{\psi_0}$  and condition the performance policy (and the backup policy) on it for the whole episode. The training procedure is illustrated in Algorithm 1.

Since we train both policies with modifications of the off-policy SAC algorithm, we can use transitions from actions proposed by either backup policy or performance policy. The transitions from both policies are stored in a shared replay buffer and are sampled at random to update the parameters of actors and critics for both performance and backup agents. At every step during training, the robot needs to select a policy to follow. We introduce a parameter  $\rho$ , the probability that the robot chooses an action proposed by the backup policy. We initialize  $\rho$  to 1, meaning that at the beginning, all actions are sampled from the backup policy. Our intuition is that the backup policy needs to be trained well before shielding mechanism is introduced in the training. We gradually anneal  $\rho$  to 0. Additionally, to realize a safe Sim-to-Lab transfer, we want the performance agent to be aware of the backup agent. Thus, we also apply shielding during Sim training. However, since the backup actor and critic may not be able to shield successfully in the beginning, we introduce a parameter  $\epsilon$ , which is the probability that the shielding is activated at this time step. This parameter can be viewed as how much we trust the backup policy and to what extent we want it to shield the performance policy. We typically initialize  $\epsilon$  to 0 and anneal it to 1 gradually. The influence of  $\rho$  and  $\epsilon$  are further analyzed in Sec. 7.2.4.

The details of updating the policies and the discriminator are shown in the Algorithm 2. Notice that while we train the backup policy  $\pi^b$  using the executed action  $a_t^{sh}$ , the performance policy  $\pi^p$  is trained using the originally proposed action  $a_t$  ("action re-labeling"). This ensures that the performance agent learns to associate its proposed action with the transition outcome, and avoids keeping proposing unsafe actions.

After the joint training, we obtain the trained dual policies  $\pi^p$  and  $\pi^b$ , and the latent distribution  $N_{\psi_0}$  that encodes diverse solutions in the environments. We now fix the weights of the two policies, and consider the latent variable z also part of their parameterization. This gives rise to the space of policies  $\Pi := \{\pi_z^p, \pi_z^b: O \to A \mid z \ \mathbb{R}^{n_z}\}$ ; hence, the latent distribution  $N_{\psi_0}$  can be equivalently viewed as a

# Algorithm 1 Joint training in simulator

```
Require: M', \pi^p, \pi^b, q_{\varphi}, N_{\psi_0} := N(0, \sigma I), \rho = 1, \epsilon = 0, \gamma = \gamma_{init}
 1: Sample E 2 M' and z 2 N_{\psi} reset environment
                                                                                                          Same latent for whole episode
 2: for t ← 1 to num_prior_step do
          With probability \rho, sample action a_t \ \mathbb{D} \pi^b(\cdot | o_t, z); else sample a_t \ \mathbb{D} \pi^p(\cdot | o_t, z)
 3:
          With probability \epsilon, apply shielding a_t^{sh} = \pi^{sh}(\pi^b, Q^b, o_t, a_t, z)
 4:
 5:
          Step environment r_t, o_t, s_{t+1} = f_E(s_t, a_t^{sh})
          Save (o_{t+1}, o_t, a_t, a_t^{sh}, z, r_t) to replay buffer
 6:
          Update \pi^p, \pi^b, q_{\varphi}
 7:
                                                                                                                              See Algorithm 2
          Anneal \rho \rightarrow 0, \epsilon \rightarrow 1, \gamma \rightarrow 1
 8:
          if timeout or failure then
 9:
               Sample E ② M′ and z ② N<sub>ψ η</sub> reset environment
10:
11:
12: end for
13: return \pi^p, \pi^b, N_{\psi_0}
```

# Algorithm 2 Updating the performance policy, backup policy, and the discriminator

```
 for t ← 1 to num_policy_update do

2:
       Sample batch \{(o_{t+1}, o_t, a_t, z, r_t)\} from the replay buffer
                                                                                           2 Action re-labeling
       Augment rt with mutual information reward (8)
3:
       Update \pi^p to maximize r_{aug} with SAC
4:
                                                                                    Observation-marginal MI
       Sample batch \{(o_{t+1}, o_t, a_t^{sh}, z, r_t)\} from the replay buffer
5:
       Update \pi^b to minimize g_E(s) with modified SAC by (5) and (11)
                                                                                 Observation-conditional MI
7: end for
8: for t ← 1 to num discrminator update do
       Sample batch \{(o_t, z)\} from the replay buffer
       Update q_{\varphi} with SGD (9)
10:
                                                                                    Observation-marginal MI
11: end for
```

distribution on the space  $\Pi$  of policies. In the next section, we will consider  $N_{\psi_0}$  as a prior distribution  $P_0$  on  $\pi$  and "fine-tune" it by searching for a posterior distribution  $P = N_{\psi}$ , which comes with the generalization guarantee from PAC-Bayes Control.

# 6. Safely Fine-Tuning Policies in Lab

In the second training stage, we consider more safety-critical training environments such as test tracks for autonomous cars or indoor lab space, where the conditions can be more realistic and closer to real environments. After pre-training the performance and backup policies with shielding, the robot can safely explore and fine-tune the prior policy distribution  $P_0$  in a new set of environments M sampled from the unknown distribution D. Leveraging the PAC-Bayes Control framework, we can provide "certificates" of generalization for the resulting posterior policy distribution P.

The PAC-Bayes generalization bound  $R_{PAC}$  associated with P from Eq. (1) consists of two parts: (1)  $R_M(P)$ , the empirical reward of P as the average expected reward across training environments in M (4), which can be optimized using SAC algorithm; (2) a regularizer  $C(P, P_0)$  that penalizes the posterior P for deviating significantly from the prior  $P_0$ ,

$$C(P, P_0) := \frac{KL(P \boxtimes P_0) + \log(\frac{2}{\delta} \frac{N}{\delta})}{2N}.$$
 (12)

Note that the only term in  $C(P, P_0)$  that involves P is the KL divergence term between P and  $P_0$ . To minimize  $C(P, P_0)$ , we modify the SAC objective to include minimization of the KL divergence term. Also,

we consider stochasticity of the policy from the latent distribution instead of the policy network; this leads to removing the policy entropy regularization in SAC and adding a weighted KL divergence term to the actor loss:

h

i

where  $\alpha$   $\square$  R is a weighting coeficient to be tuned. In practice, we find the gradient of the K L divergence term heavily dominates the noisy gradient of actor and critic, and thus we approximate the K L divergence with an expectation on the posterior:

$$\max_{P} E_{o,z} E_{a \otimes g(\cdot \mid o,z)} Q^{p}(o,a) - \alpha \log \frac{P(\dot{z})}{P_{o}(z)}$$
(14)

Below we show the algorithm for this stage of training. To avoid safety violations, we always apply value-based shielding to the proposed action, and continue to apply action-relabeling when updating P.

# Algorithm 3 Safely fine-tuning the policy distribution

```
Require: M, \pi^p, \pi^b, P = P_0
 1: Sample E 2 M and z 2 P, reset environment
 2: for t ← 1 to num_posterior_step do
 3:
        Sample a_t \ \mathbb{Z} \pi^p(\cdot | o_t, z)
        Apply value-based shielding a_t^{sh} = \pi^{sh}(\pi^b, Q^b, o_t, a_t, z)
 4:
        Step environment r_t, o_t, s_{t+1} \supseteq P(\cdot | s_t, a_t^{sh})
 5:
        Save (o_{t+1}, a_t, z, r_t) to replay buffer
                                                                                                       2 Action re-labeling
 6:
        Update P using SAC with weighted regularization (14)
 7:
 8:
        if timeout or failure then
            Sample E 2 M, z 2 P, reset environment
 9:
10:
        end if
11: end for
12: return P
```

# 6.1. Computing the Generalization Bound.

After training, we can calculate the generalization bound using the optimized posterior P . First, note that the empirical reward R  $_{\rm M}$  (P) involves an expectation over the posterior and thus cannot be computed in closed form. Instead, it can be estimated by sampling a large number of policies  $z_1,...,z_L$  from P: R  $_{\rm M}$  (P):=  $_{\rm N}$  L  $_{\rm L}$  P  $_{\rm EBM}$  P  $_{\rm i=1}$  R  $_{\rm E}$  ( $\pi^{\rm p}_{z_i}^{\rm p}$ ), and the error due to finite sampling can be bounded using a sample convergence bound R  $_{\rm M}$  [65]. The final bound R  $_{\rm bound}$  (P)  $\leq$  R  $_{\rm D}$  (P) is obtained from R  $_{\rm M}$  and C (P, P  $_{\rm O}$ ) by a slight tightening of C  $_{\rm PAC}$  from Theorem 1 using the KL-inverse function [19]. Please refer to Appendix A2 in [37] for detailed derivations. Overall, our approach provides generalization guarantees in novel environments from the distribution D: as policies are randomly sampled from the posterior P and applied in test environments, the expected success rate over all test environments is guaranteed to be at least R  $_{\rm bound}$  (P) (with probability  $_{\rm L}$  O over the sampling of training environments;  $_{\rm L}$  = 0.01 for all experiments in Sec. 7). Through reachability shielding during training and generalization guarantees for the resulting policies, we bridge the Lab-to-Real gap with a probabilistically guaranteed safety-aware policy distribution.

# 7. Experiments

Through extensive simulation and hardware experiments, we aim to answer the following questions: does our proposed Sim-to-Lab-to-Real achieve (1) lower safety violations during Lab training compared to other safe learning methods, (2) stronger generalization guarantees on performance and safety compared to previous work in PAC-Bayes Control, and (3) better empirical performance and safety during deployment

compared to all baselines? We also evaluate (a) the relative importance of the Sim stage and Lab stage, (b) how the value threshold in shielding affects safety and eficiency, (c) how the regularization weight in (14) affects generalization guarantees and empirical performance after training, (d) how the two annealing parameters during Sim training,  $\epsilon$  and  $\rho$ , affect training performance, and (e) how diversity components, mutual information maximization during Sim training and latent dimension, affect performance after Lab training.

#### 7.1. Experiment Setup

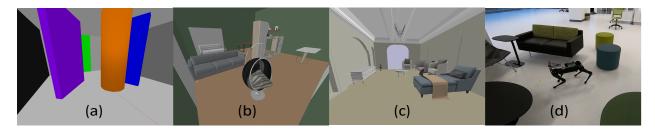


Figure 4: Samples of environments used in experiments: (a) Sim training in Vanilla-Env; (b) Sim training in Advanced-Env; (c) Advanced-Realistic training; (d) physical robot deployment.

#### 7.1.1. Environments

425

We evaluate the proposed methods by performing ego-vision navigation task in two types of environment. The first type (Vanilla-Env) consists of undecorated rooms of  $2m \times 2m$  with randomly placed cylindrical and rectangular obstacles of different dimensions and poses, and the robot needs to bypass them and the walls to reach a green door (a smaller circular region in front it) (Fig. 4a). A virtual camera is simulated with 120 degree field of view both vertically and horizontally, outputting RGB images of  $48 \times 48$  pixels. We treat the robot as a point mass when checking collision.

The second type of environments (Advanced-Env) uses realistic furniture models from the 3D-FRONT dataset [6] (Fig. 4b); the robot needs to safely reach some target location (a smaller circular region) using given distance and heading angle towards the target. A virtual camera is simulated at the front of the robot with 72 degree field of view vertically and 128 degree horizontally (matching the ZED 2 camera used in Real deployment), outputting RGB images of 90×160 pixels. When checking collisions, we approximate the robot as a circular shape of radius 25cm, roughly the same as the quadrupedal robot in Real deployment.

For both types of environments, the control loop runs at 10Hz and the maximum number of steps is 200. The robot is commanded with forward velocity ([0.5,1] m/s for performance policy and [0.2,0.5] m/s for backup policy) and angular velocity ([-1,1] rad/s for both policies). We use dense proxy reward that is proportional to the percentage of distance traveled between initial location and goal, and the safety signal is calculated as the minimum distance to obstacles and walls. Additionally, we assume the robot is given  $\ell_E(s)$  and  $\Delta_E(s)$ , distance and relative bearing to the goal.

For Sim training, we randomize obstacle and furniture configurations to cover possible scenarios as much as possible. We also randomize camera poses (tilt and roll angles) in Advanced-Env to account for possible noise in real experiments. Sim training uses 100 environments in Vanilla-Env and 500 environments in Advanced-Env. After Sim training, we can fine-tune the policies in different types of Lab environments. For Vanilla-Env, we consider:

- Vanilla-Normal: shares the same environment parameters as ones in the Sim stage.
- Vanilla-Dynamics: increases the lower bound of forward and angular velocity (more aggressive maneuvers).
- Vanilla-Task: adds an additional condition on success that the the robot needs to enter the target region with a yaw angle within a small range instead of  $2\pi$  (no restriction) in the Sim stage. The robot

may pass through the target region and re-enter it with the required yaw orientation. The robot knows the lower bound and the upper bound of the required yaw range.

and for Advanced-Env, we consider:

- Advanced-Dense: assigns a higher density of furniture in the rooms, resulting in smaller clearances between them.
- Advanced-Realistic: uses realistic room layouts (Fig. 4c) and associated furniture configurations from the 3D-FRONT dataset, which are similar to real environments. We perform Lab-to-Real transfer with policies trained in this Lab (Fig. 4d). More details about the dataset and room layouts can be found in Appendix C.

### 7.1.2. Policy

430

435

460

465

We parameterize the performance and backup agents with neural networks consisting of convolutional layers and then fully connected layers. The actor and critic of each agent share the same convolutional layers. In Vanilla-Env, a single RGB image is fed to the convolution layers, and the latent variable is appended to the output of the last convolutional layer before fully connected layers. In Advanced-Env, we stack 4 previous RGB images while skipping 3 frames between two images to encode the past trajectory of the robot. Then, the stacked images are concatenated with the first 10 dimensions of the latent variable by repeating each dimension to the image size. Rest of the dimensions is appended to the output of the last convolutional layer. In addition to the image observation, the actors and critics also receive two auxiliary signals  $\ell_E(s)$  and  $\Delta_E(s)$ , which are also appended to the output of the last convolutional layer. The details of neural network architecture and training can be found in Appendix B.

#### 7.1.3. Baselines

We compare our methods to five prior RL algorithms that neglect safety violations (Base and PAC\_Base [19]) or address safety by reward shaping (RP and PAC\_RP) or use a separate safety agent (SQRL [13] and Recovery RL [14]). Sim-to-Lab-to-Real varies from SQRL and Recovery RL in that the latter trains the safety critic by the sparse safety indicators as below,

$$Q^{b}(o_{t},a_{t}) := I_{E}(s_{t}) + \gamma 1 - I_{E}(s_{t}) \min_{\substack{t+1 \ a}} Q^{b} o_{t+1}, a_{t+1},$$

where  $I_E(s_t) = 1\{g_E(s_t) > 0\}$  is the indicator function of the safety violations. In Sim-to-Lab-to-Real, the safety state-action values represent the robot's closest distance to the obstacles in the future, while in Recovery RL and SQRL, the values represent the probability that the robot will hit the obstacle (but the probability strongly depends on the discount factor used). The major distinction between Sim-to-Lab-to-Real and PAC-Bayes control is that the latter does not handle the safety explicitly but instead hopes to use diverse policies and fine-tuning to prevent unsafe maneuver. We give a brief description of these methods below and summarize the similarities and differences in Table 1.

- Sim-to-Lab-to-Real (ours): trains a distribution over dual policies conditioned on latent variables with guarantees on generalization to novel environments. We present two variants: PAC\_Shield\_Perf, whose performance policy is conditioned on latent variables, and PAC\_Shield\_Both, whose both performance and backup policies are conditioned on latent variables (Fig. 2).
- Shield (ours): trains a dual policy without conditioning on latent variables, thus no distribution over policies nor generalization guarantees.
- PAC-Bayes Control [19]: trains a distribution over policies conditioned on latent variables that optimizes for either only task reward (PAC\_Base) or reward with penalty (PAC\_RP).
- Base: trains a single policy that optimizes the task reward only.

- Reward Penalty (RP): trains a single policy but augments the task reward with penalty on safety violations,  $\hat{r}_E(s, a) = r_E(s, a) \lambda 1\{g_E(s) > 0\}$ .
- Safety Critic for R L (SQRL) [13]: trains a dual policy. The backup critic optimizes the Lagrange relaxation of CMDP,  $\hat{J}(\pi) = J(\pi) + \nu E_{all} \pi [(\nu_{thr} Q^b(o,a)]$ , with a rejection sampling method that re-samples action if  $Q^b(o,a) > \nu_{thr}$ .
- Recovery RL [14]: trains a dual policy. The backup critic is trained in the same method as SQRL, but the backup action is from the backup actor instead of being re-sampled from performance policy.

Methods	Dual Policy	Safety Treatment	Generalization Guarantees
Sim-to-Lab-to-Real (ours)	<b>√</b>	√(Reachability safety critic)	<b>√</b>
Shield (ours)	✓	√(Reachability safety critic)	X
PAC+Base [19]	X	X	$\checkmark$
PAC+Reward Penalty [19]	X	√(Reward with safety penalty)	$\checkmark$
Base	X	X	X
Reward Penalty Safety Critic for R L [13]	X	√(Reward with safety penalty)	X
Recovery R L [14]	X	√(Risk safety critic)	X
Necovery NL [14]	$\checkmark$	√(Risk safety critic)	X

Table 1: Major distinctions among Sim-to-Lab-to-Real and baseline methods.

#### 7.2. Results

480

485

490

495

We compare all the methods by (1) safety violations in Lab training and (2) success and safety at deployment (Figure 5). We calculate the ratio of number of safety violations to the number of episodes collected during training. For deployment, we show the percentage of failed trials (solid bars in Figure 5) and unfinished trials (hatched bars). We summarize the main findings below:

- 1. Across Lab training, our proposed Sim-to-Lab-to-Real (PAC\_Shield\_Perf and PAC\_Shield\_Both) achieves the fewest safety violations. This demonstrates the eficacy of the reachability safety state-action value function for shielding. Compared to the risk-based safety critics in SQRL [13] and Recovery RL [14], our safety critics can learn from near-failure and with dense cost signals, as discussed in 7.2.1. Adding penalty in the reward function does not reduce safety violations significantly.
- 2. In testing environments, Sim-to-Lab-to-Real achieves the lowest unsuccessful fraction of trajectories (solid bars plus hatched bars). This indicates that training a diverse and safe policy distribution achieves better generalization performance to novel environments. Sim-to-Lab-to-Real also achieves the fewest safety violations (solid bars) at test time. This suggests that explicitly enforcing hard safety constraints improves the safety not only in training but also in testing. In Sec. 7.2.2 we show stronger generalization guarantees (for both performance and safety) compared to PAC-Bayes baselines.
- 3. In Sec. 7.2.3 we show Sim-to-Lab-to-Real achieves the best performance and safety among baselines when the policies are deployed on a quadrupedal robot navigating through real indoor environments. The empirical performance and safety also validate the theoretical generalization guarantees from PAC-Bayes Control.
- 4. In Sec. 7.2.4 we show that high diversity of trajectories from the latent distribution results in better generalization at test time. Without diversity maximization in Sim training, the resulting trajectories can concentrate close to a single one and hinder downstream fine-tuning in Lab. However, we also find that in Advanced-Env, PAC\_Base and PAC\_RP (distribution over policies) perform worse than Base and RP (single policy). We find that high diversity without shielding may hinder training progress due to frequent safety violations interfering with strategy exploration.

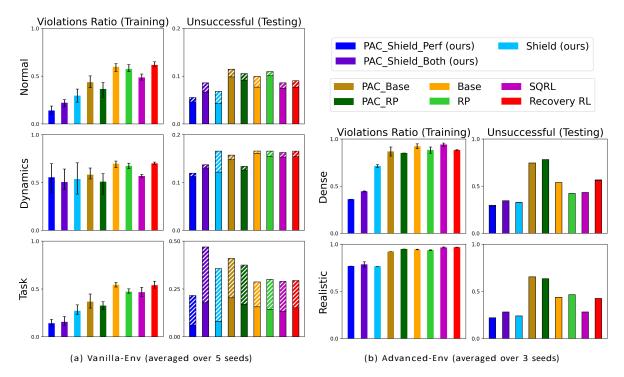


Figure 5: Comparison of safety violations during Lab training and unsuccessful trials at test time: Sim-to-Lab-to-Real (PAC\_Shield\_Perf and PAC\_Shield\_Both) has the lowest safety violations in both training and deployment. First, it showcases the benefits of using a shielding scheme in contrast with Base, RP and vanilla PAC-Bayes. Second, reachability safety critic enables safer exploration and safety satisfaction in deployment as compared to SQRL and recovery RL. Additionally, Sim-to-Lab-to-Real has lower unsuccessful ratio in deployment than Shield, which shows a diverse but safe policy distribution not only provides a generalization bound but also improves the empirical performance to novel environments.

- 5. We find that adding latent distribution to the backup policy introduces dificulty during Sim training, and leading to similar, if not worse, performance and safety at test time. We suspect that PAC\_Shield\_Both would take more samples to converge well in training and requires more careful tuning of hyperparameters. Following discussions focus on results of PAC\_Shield\_Perf, in which only the performance policy is conditioned on the latent variable.
- 6. Compared to other Labs, violation ratios in Advanced-Realistic tend to be higher, although our methods still reduce safety violations by 20-25%. Also, there are few unfinished trials at test time (the robot neither reaches the target nor collides with obstacles). Given the tight spacing in realistic indoor environment (Fig. 6b), the non-trivial dimensions of the quadruped robot, and the complex visuals, the backup policy can fail to ensure safety in some environments.

# 7.2.1. Reachability vs. Risk-Based Safety Critic

500

505

510

Sim-to-Lab-to-Real and previous safe RL methods differ in the metric used to quantify safety and train the backup agent. By utilizing reachability RL, we have an exact encoding of the property we want our system to satisfy, i.e., the distance should be no closer to obstacles than a specific threshold. In contrast, SQRL and Recovery RL define safety by the risk of colliding with obstacles in the future and use binary safety indicators. We argue that risk-based threshold can easily overfit to specific scenarios since the probability heavily depends on the discount factor used. In addition, reachability objective allows the backup agent to learn from near failure, while the risk critic in SQRL and Recovery RL needs to learn from complete failures. Fig. 6 shows 2D slices of the safety state-action values in both environment settings. Reachability critics provide thicker unsafe regions next to obstacles, while risk-based critics fail to recognize many unsafe regions or consider unsafe only when very close to obstacles. Among different Lab setups, compared to the

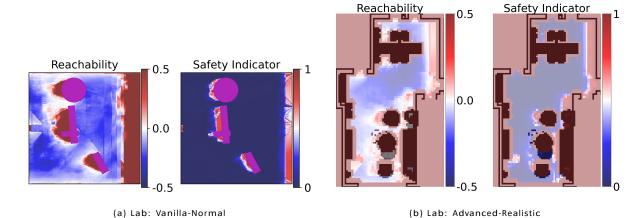


Figure 6: 2 D slices of safety state-action value functions when the robot is facing to the right: we train the safety critic using RL modified from Hamilton-Jacobi reachability analysis ("Reachability"), while SQRL and recovery RL train it with sparse binary indicators ("Safety Indicator"). Reachability safety critic better captures the unsafe region - there is a more gradual change in safety value near the obstacles (from blue to red color, lower to higher value), indicating that the robot is getting closer to the obstacles. In contrast, risk-based critic shows a more binary separation between safe and unsafe regions, leaving the robot little room and time to steer away from obstacles. The unsafe regions are also thinner than those learned with reachability. Thus, our reachability-based approach achieves fewer safety violations during both training and deployment.

baselines, our method reduces safety violations by 77%, 4%, 76%, 62%, and 23% in training and 38%, 26%, 54%, 34%, and 28% in deployment.

Sensitivity analysis: value threshold. Through experiments, we find the value threshold used in shielding essential to performance and safety. We first investigate how the threshold using during training affects the final results among the three Lab settings in Vanilla-Env, which are shown in Fig. 8e. v<sub>thr</sub> = 0 natu-rally results in more safety violations during training compared to  $v_{thr} = -0.05$  and  $v_{thr} = -0.10$ . Policies trained with v<sub>thr</sub> = 0 also performs the worst at test time, which indicates that less shielding during training makes the robot learn unsafe or aggressive maneuver. Next we evaluate how the value threshold affects robot trajectories at test time. Fig. 7 shows the trajectories using different thresholds in the two settings. Small threshold leads to robot passing very closely next to obstacles, while a bigger threshold leads to more conservative behavior. We also would like to highlight the challenges of learning safe policies in Advanced-Env. As shown in the figure, with  $v_{thr} = -0.15$  the robot avoids the first obstacle, and then the backup policy steers the robot away from the target, potentially deeming the clearance next to the target not suficient. However, this brings the robot near the wall, and due to imperfect training of the backup actor, the robot fails to escape. With tight spacing and large dimensions of the robot in Advanced-Env, we find the backup agent more dificult to train, and the final test performance and safety can be sensitive to the shielding threshold. In Advanced-Realistic, average test success rate with  $v_{thr}$  = -0.05, -0.1, -0.15 are 0.678, 0.786, and 0.762 respectively. Future work could look into adapting the threshold after short experiences in different environments online.

#### 7.2.2. Generalization Guarantees

In this subsection, we evaluate the PAC-Bayes generalization guarantees obtained after Lab training, and the effect of adding reachability shielding in the policy architecture to the bounds. Table 2 shows the bounds and test results on safety (not colliding with obstacles) and success (safely reaching the goal) among Lab training. The true expected success and safety are tested with environments that are similar to the Lab training environments (of the same distribution) but unseen before. In all settings, the true expected success and safety are higher than the bound in all settings, which validates the guarantees derived using PAC-Bayes Control. Furthermore, we compare the bound trained using PAC\_Shield\_Perf with previous PAC-Bayes Control method (PAC Base) in the Vanilla-Env and Advanced-Realistic. With shielding, the generalization

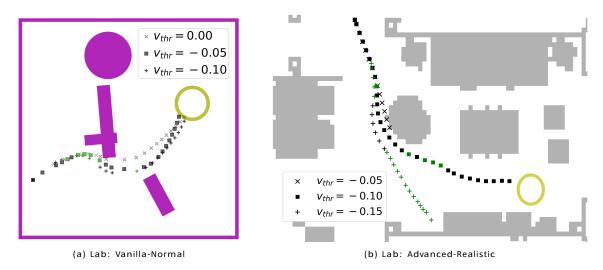


Figure 7: Rollout trajectories using different value threshold for shielding: higher threshold (more negative) results in more conservative maneuver, i.e., keeping farther away from obstacles (purple in (a) and grey in (b)). In Advanced-Env, the complex visuals and tight spacing cause challenges in learning the backup agent. We tend to find a relatively conservative threshold ( $v_{thr} = -0.10$ ) works well in practice, and too high threshold can prevent the robot from reaching the goal and accidentally steer it towards tight space.

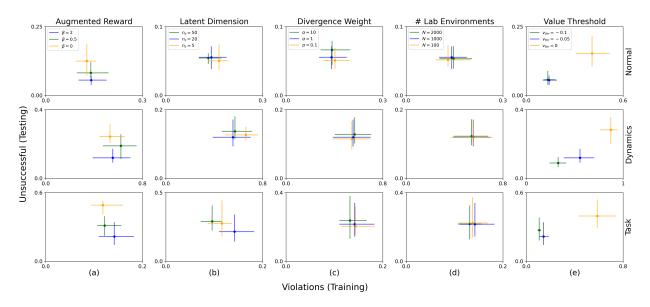


Figure 8: Sensitivity analyses: we study the influence of different hyper-parameters to Sim-to-Lab-to-Real. The results are averaged over 5 seeds in Vanilla-Env. If not specified, the hyper-parameters default to  $\beta$  = 2,  $n_z$  = 20,  $\alpha$  = 1, N = 1000, and  $v_{thr}$  = -0.05, as shown in blue. Results suggest the augmented reward in Sim training and the value threshold in shielding are the two most important hyper-parameters.

bound improves in all settings. In the dificult setting of Advanced-Realistic, the bound improves from 0.366 to 0.786 for task completion and from 0.367 to 0.794 for safety satisfaction. Thus, explicitly enforcing hard safety constraints not only improves empirical outcomes but also provides stronger certification to policies in novel environments. In Sec. 7.2.3 we also demonstrate empirical results of physical robot experiments validating the guarantees.

Table 2: Results of PAC-Bayes guarantees and test success and safety: to compute the bound, each environment has 1000 policies sampled from the latent distribution and tested. The results in the first two rows are based on PAC\_Shield\_Perf.

	Advanced-Realistic				
Method	PAC_Shield_Perf	PAC_Base	S Q R L		
# Lab Environments	1000	1000	1000		
Success Bound	0.701	0.297	-		
True Expected Success	0.786	0.366	0.712		
Real Robot Success	0.767	0.433	0.667		
Safety Bound	0.708	0.304	-		
True Expected Safety	0.794	0.367	0.713		
Real Robot Safety	0.867	0.433	0.667		

	Vanilla-Normal					Va	ınilla-Dy	namics		
Method	PAC_Shield_Perf			PAC_Base	PAC_Shield_Perf			PAC_Base		
Divergence Weight	1	1	1	10	1	1	1	1	10	1
# Lab Environments	100	1000	2000	1000	1000	100	1000	2000	1000	1000
Success Bound	0.778	0.876	0.900	0.896	0.735	0.692	0.820	0.839	0.828	0.778
True Expected Success	0.948	0.945	0.947	0.934	0.886	0.881	0.880	0.878	0.872	0.843
Safety Bound	0.793	0.911	0.917	0.913	0.816	0.717	0.835	0.851	0.837	0.815
True Expected Safety	0.954	0.954	0.954	0.953	0.902	0.888	0.887	0.887	0.883	0.852

	Vanilla-Task			Advanced-Dense							
Method	P	A <u>C</u> Shi	ield_Per	f	PAC_Base		PAC	_Shield_	_Perf		PAC_Base
Divergence Weight # Lab Environments	1	1	1	10	1	2	2	2	1	5	2
	100	1000	2000	1000	1000	100	500	1000	500	500	1000
Success Bound	0.578	0.757	0.792	0.777	0.468	0.402	0.578	0.623	0.512	0.557	0.254
True Expected Success	0.847	0.851	0.844	0.853	0.590	0.577	0.663	0.703	0.621	0.644	0.327
Safety Bound	0.769	0.884	0.899	0.887	0.663	0.412	0.579	0.630	0.518	0.564	0.259
True Expected Safety	0.939	0.939	0.940	0.938	0.796	0.583	0.671	0.709	0.629	0.652	0.332

Sensitivity analysis: weight of policy distribution regularization ( $\alpha$ ). When optimizing the generalization bound (14), we place a weighting coeficient  $\alpha$  to balance gradients of the training reward and of the estimated KL divergence between the prior and posterior policy distribution,  $P_0$  and P. Here we study the effect of using different values of  $\alpha$  in the generalization bound and test performance. Fig. 8c shows that too strong regularization ( $\alpha$  = 10) prevents the Lab training from tuning the prior distribution suficiently, resulting in worse testing performance after training. The effect of different  $\alpha$  is more prominent in Advanced-Dense training. With same 500 training environments,  $\alpha$  = 2 achieves 0.578 on success bound while 0.512 for  $\alpha$  = 1 and 0.557 for  $\alpha$  = 5.

Sensitivity analysis: number of Lab environments (N). Thm. 1 indicates the PAC-Bayes bound depends on the number of environments used in the Lab training. Fig. 8d demonstrates that in Vanilla-Env, N does not have a significant effect on training safety violations and test performance. We suspect that training in Vanilla-Env does not require a large number of environments for generalization. In the more dificult Advanced-Dense, with the same  $\alpha=2$ , higher N = 1000 achieves the best test success (0.703) and safety (0.709) compared to smaller N = 100 and N = 500 (Table. 2), which demonstrates that a higher number of Lab environments help fine-tuning the policies achieving strong generalization in complex environments.

#### 7.2.3. Physical Experiments

To demonstrate empirical performance and safety of trained policies in real environments (Lab-to-Real transfer) and verify the generalization guarantees, we evaluate the policies in real indoor environments in

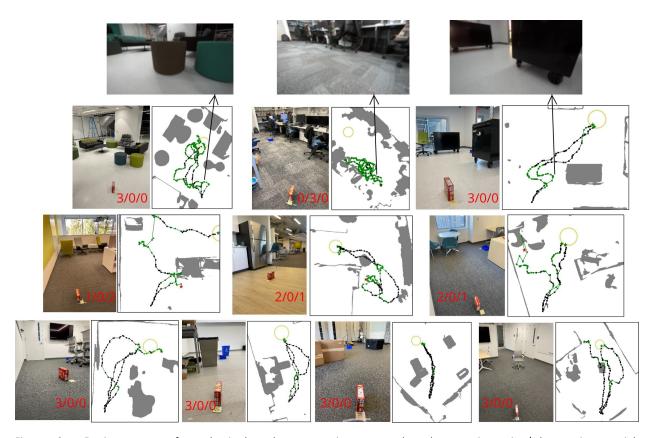


Figure 9: Environments for physical robot experiments and robot trajectories/observations with PAC\_Shield\_Perf: we run the policy three times in each environment by sampling different latent variables from the posterior distribution. The three numbers in images indicates success/unfinished/failure split. Green dots indicates shielding in effect. Green star indicates success in reaching the target. Red star indicates colliding with obstacles. We scan the environment using an iPad Pro tablet before experiments to generate the 2D map (which the robot does not have access to). The robot trajectory is obtained using localization algorithm of the onboard camera, and is inaccurate at places (intersecting obstacles; not exactly reaching the target but the robot deems so, which we consider success).

the Engineering Quadrangle building at Princeton University. We deploy a Ghost Spirit quadrupedal robot equipped with a ZED 2 camera at the front (Fig. 4d), matching the same dynamics and observation model used in Advanced-Realistic Lab. For the distance and relative bearing to the goal, before each trial the robot is given the ground-truth measurement at the initial location, and then it uses the localization algorithm native to the stereo camera to update the measurement at each step.

We pick ten different locations with furniture configurations and dificulty similar to those in Advanced-Realistic Lab. Based on test results after Lab training, we run policies trained with PAC\_Shield\_Perf (best performance overall), PAC\_Base (PAC-Bayes baseline with low generalization guarantees), and SQRL (best overall among other baselines). Each policy is evaluated at one environment 3 times (30 trials total). The results are shown in Table. 2. Our policy is able to achieve the best performance (0.767) and safety (0.867), validating the theoretical guarantees from PAC-Bayes Control. The upper-right of Fig. 1 shows a trajectory when running policies trained with PAC\_Shield in a kitchen environment where the backup policy and the shielding discriminator help the robot avoid hitting the obstacles and reach the target successfully.

Fig. 9 shows the 10 real environments and robots' trajectories when running policies trained with PAC\_Shield\_Perf. Green dots indicate shielding in effect, which is activated often near obstacles. The first and third images on top of the figure show the robot's view when shielding successfully guides robot away from the sofa stool and the cabinet. In the second environment, the backup policy keeps shielding the robot away from center of the room with value threshold  $v_{thr} = -0.10$ , and all three trials ended as

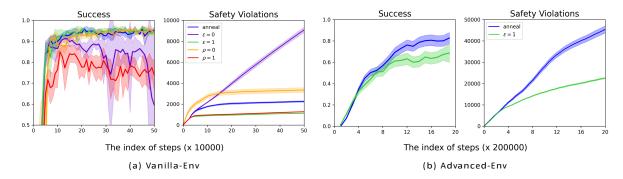


Figure 10: Effect of  $\rho$  and  $\epsilon$  scheduling in Sim training: annealing  $\rho$  and  $\epsilon$  helps balance between safety violations and task completion. If not specified, for Vanilla-Env,  $\rho$  initializes at 1 and decays by 0.5 every 25000 steps, and  $\epsilon$  initializes at 0 with 1 –  $\epsilon$  decaying by 0.5 every 50000 steps. For Advanced-Env,  $\rho$  initializes at 0.5 and decays by 0.5 every 500000 steps, and  $\epsilon$  initializes at 0 with 1 –  $\epsilon$  decaying by 0.5 every 200000 steps. The results are over 5 random seeds for Vanilla-Env and 3 random seeds for Advanced-Env.

unfinished. This is possibly due to the cluttered scene of desks at the top half of the observation. We also test with small value threshold  $v_{thr} = -0.05$  during experiments, and the robot is able to reach the target without shielding always activated. This highlights the need for adapting the shielding value threshold online in future work.

#### 7.2.4. Other Studies

Ablation Study: importance of two-stage training. We evaluate the significance of Lab training by testing the prior policy distribution (without fine-tuning in Lab) in Vanilla-Env. Without Lab training, the unsuccessful ratio in deployment increases by 16%, 8% and 14%. This suggests that Lab training is essential to policies adapting to real dynamics and new environment distributions. Additionally, we test the importance of Sim training with Shield (no policy distribution). Without Sim training, the safety violations in Lab training increases by 60%, 11% and 65%. This demonstrates that Sim training enables the backup agent to monitor and override unsafe behavior from the beginning of Lab training.

Sensitivity analysis: the probability of sampling actions from the backup policy (p) and the probability of activating shielding  $(\epsilon)$ . One of the main contributions of our work is the effective joint training of both performance and back agents (realized in Sim training). The two parameters,  $\rho$  and  $\epsilon$ , directly affect the exploration in Sim training. With high  $\rho$  or high  $\epsilon$ , the RL agent basically only explores conservatively within a small safe region. However, in the beginning of the training, we should allow the RL agent to collect diverse state-action pairs. On the other hand, we also gradually anneal  $\rho \to 0$  and  $\epsilon \to 1$ since we want the performance policy to be aware of the backup policy. In other words, the performance policy is effectively in shielded environments towards end of Sim training. Fig. 10 shows the Sim training progress under different p and  $\epsilon$  scheduling. With constant  $\rho = 0$  or  $\epsilon = 0$ , the number of safety violations is much higher than that with both parameters annealing. Even worse,  $\epsilon = 0$  results in the number of safety violations increase at constant speed and the training success fluctuates significantly. On the other hand, with  $\rho = 1$  or  $\epsilon = 1$ , the number of safety violations is only half as that with both parameters annealing. However, this is at the expense of exploration and leads to worse success rate in deployment. In Vanilla-Env ρ = 1 leads to very poor training success. Although in Vanilla-Env  $\epsilon$  = 1 does not have significant effect on training success, in the Advanced-Env, insuficient exploration hinders training progress. Also note that Sim training is not safety-critical and we do not aim to reduce safety violations then.

Sensitivity analysis: diversity-induced Sim training. We argue that training a diverse and safe policy distribution helps improve safety and performance in novel environments. There are two hyper-parameters in our algorithm affecting the diversity, i.e., augmented reward coeficient  $\beta$  and latent dimension  $n_z$ . Fig.8a

# High Diversity ( $\beta = 2$ )

# Without Diversity ( $\beta = 0$ )

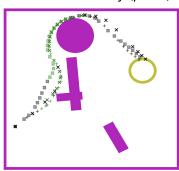


Figure 11: High augmented reward coeficient induces a diverse policy distribution: the diversity is essential to finetuning the latent distribution in Lab training and to good generalization to novel environments. Black markers indicate actions from the performance policy being executed, and green markers are for actions from the backup policy.

and Fig.8b show the violation ratio in Lab training and unsuccessful ratio in testing under different  $(\beta, n_z)$  choices. We find that training without augmented reward  $(\beta = 0)$  results in the lowest violation ratio; however, the unsuccessful ratio in testing is the highest. In fact, we observe that with  $\beta = 0$ , rollout trajectories conditioned on different latent variables almost converge to a single trajectory as shown in Fig. 11. This reflects why safety is better satisfied but at the expense of generalization. On the other hand, when the coefficient is sufficiently large  $(\beta = 2)$ , the policy distribution becomes diverse and generalizes well to unseen testing environments. Note that even with high diversity, safety can still be well ensured with shielding. For the second source of diversity, our proposed Sim-to-Lab-to-Real is robust to different latent dimension.

### 8. Conclusion

In this work, we propose the Sim-to-Lab-to-Real framework that combines Hamilton-Jacobi reachability analysis and PAC-Bayes generalization guarantees to bridge the Sim-to-Real reality gap with a probabilistically guaranteed safety-aware policy distribution. Joint training of a performance and a backup policy in Sim training (1st stage) enables a safety-aware exploration during Lab training (2nd stage). By optimizing the generalization bounds in Lab training, our approach is able to probabilistically certify robot performance and safety before deployment. We demonstrate significant reduction in safety violations in training and stronger performance and safety during test time. Results from experiments with a quadrupedal robot in real indoor space validate the theoretical guarantees.

# 8.1. Discussion: Environment distribution.

As elaborated in Sec. 3, the generalization guarantees obtained through our framework assumes no distribution shift between Lab and Real in terms of environments. To bridge the discrepancy, we model the real environments by using (1) photorealistic dataset of indoor room layouts and furniture models and (2) dynamics from system identification of the real robot and camera poses. Additionally, we note that previous works in PAC-Bayes Control [19, 21, 37] have consistently shown real deployment validating the bounds. Even under a slight of shift in distribution, we believe that a certificate of performance and safety is useful and provides confidence for deploying the system.

# 8.2. Discussion: Large-scale Lab training.

We acknowledge that one limitation of our framework is that, in exchange for assuming close to nothing about the environment distribution and providing statistical guarantees that hold in arbitrarily high confidence instead of in expectation only (e.g., conformal prediction [59]), we require at least a few hundred

environments for "Lab" training to achieve tight PAC-Bayes generalization guarantees (e.g., < 10% difference between empirical performance and theoretical guarantee), which means performing "Lab" training with real conditions can be dificult for us researchers in university labs with limited hardware, computation, and human resources. In this work, we resort to performing "Lab" training in realistic simulated environments.

Nonetheless, we envision that our framework is well suited for industry practitioners who have access to either extensive training facilities (e.g. Google's robot "farms" [66], Boston Dynamics' testing warehouse [7]), large-scale distributed systems (e.g. Amazon's warehouses [67]), or vast amounts of "Lab-like" data collection (e.g. Cruise and Waymo's thousands—millions of test driver miles [68]). For these practical and often safety-critical applications, our framework can improve safety during training and provide generalization guarantees for performance and safety at deployment. For university labs achieving similar scales of data collection and training, it would be promising to explore (1) crowdsourcing robots training across labs [69] and (2) mechanisms for automatically resetting the robot [70] and randomizing the environments.

On the theoretical front, first it would be worth identifying the most representative environments for training (e.g., using coresets [71]). PAC-Bayes guarantee holds as long as the policies are "evaluated" in the training environments M and the training reward R  $_{\rm M}$  is evaluated. We could potentially obtain similar tight generalization guarantees by training on a much smaller set of environments compared to M used in this work. Second, recent growing interest in PAC-Bayes bound [72] and other types of generalization guarantees [73] could lead to tighter and also more sample-eficient bounds for certifying the generalization performance and safety.

# 675 Acknowledgement

660

Allen Z. Ren and Anirudha Majumdar were supported by the Toyota Research Institute (TRI), the NSF CARER award [2044149], the Ofice of Naval Research [N00014-21-1-2803], and the School of Engineering and Applied Science at Princeton University through the generosity of William Addy '82. This article solely reflects the opinions and conclusions of its authors and not ONR, NSF, TRI or any other Toyota entity. We would like to thank Zixu Zhang for his valuable advice on the setup of the physical experiments.

### References

685

- [1] A. Kumar, Z. Fu, D. Pathak, J. Malik, RMA: Rapid Motor Adaptation for Legged Robots, in: Proceedings of Robotics: Science and Systems (RSS), Virtual, 2021. doi:10.15607/RSS.2021.XVII.011.
- [2] Y. Zhu, R. Mottaghi, E. Kolve, J. J. Lim, A. Gupta, L. Fei-Fei, A. Farhadi, Target-driven visual navigation in indoor scenes using deep reinforcement learning, in: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), 2017, pp. 3357–3364. doi:10.1109/ICRA.2017.7989381.
- [3] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, P. Abbeel, Domain randomization for transferring deep neural networks from simulation to the real world, in: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2017, pp. 23–30. doi:10.1109/IROS.2017.8202133.
- [4] F. Muratore, F. Ramos, G. Turk, W. Yu, M. Gienger, J. Peters, Robot learning from randomized simulations: A review, 2021. arXiv:2111.00956.
- [5] F. Sadeghi, S. Levine, Cad2rl: Real single-image flight without a single real image, in: Proceedings of Robotics: Science and Systems (RSS), Cambridge, Massachusetts, 2017. doi:10.15607/RSS.2017.XIII.034.
- [6] H. Fu, B. Cai, L. Gao, L.-X. Zhang, J. Wang, C. Li, Q. Zeng, C. Sun, R. Jia, B. Zhao, H. Zhang, 3D-FRONT: 3D Furnished Rooms With layOuts and semaNTics, in: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), 2021, pp. 10933–10942.

[7] Boston-Dynamics, Inside the Lab: Robotics After Hours, https://www.youtube.com/watch?v= Jq0GknnKvXM, 2022.

700

710

- [8] Y. Chow, M. Ghavamzadeh, Algorithms for cvar optimization in mdps, in: Proceedings of Advances in Neural Information Processing Systems (NeurIPS), Montreal, Quebec, Canada, 2014, pp. 3509–3517.
- [9] Y. Chow, M. Ghavamzadeh, L. Janson, M. Pavone, Risk-constrained reinforcement learning with percentile risk criteria, Journal of Machine Learning Research (JMLR) 18 (2017) 6070–6120.
- [10] J. F. Fisac, A. K. Akametalu, M. N. Zeilinger, S. Kaynama, J. Gillula, C. J. Tomlin, A general safety framework for learning-based control in uncertain robotic systems, IEEE Transactions on Automatic Control (TAC) 64 (2019) 2737–2752.
  - [11] J. F. Fisac, N. F. Lugovoy, V. Rubies-Royo, S. Ghosh, C. J. Tomlin, Bridging hamilton-jacobi safety analysis and reinforcement learning, in: Proceedings of the International Conference on Robotics and Automation (ICRA), 2019, pp. 8550–8556. doi:10.1109/ICRA.2019.8794107.
  - [12] K.-C. Hsu, V. Rubies-Royo, C. J. Tomlin, J. F. Fisac, Safety and liveness guarantees through reach-avoid reinforcement learning, in: Proceedings of Robotics: Science and Systems, Virtual, 2021. doi:10.15607/ RSS.2021.XVII.077.
  - [13] K. Srinivasan, B. Eysenbach, S. Ha, J. Tan, C. Finn, Learning to be safe: Deep rl with a safety critic, 2020. arXiv:2010.14603.
  - [14] B. Thananjeyan, A. Balakrishna, S. Nair, M. Luo, K. Srinivasan, M. Hwang, J. E. Gonzalez, J. Ibarz, C. Finn, K. Goldberg, Recovery RL: Safe reinforcement learning with learned recovery zones, IEEE Robotics and Automation Letters (RAL) 6 (2021) 4915–4922.
- [15] K. Zhou, J. C. Doyle, Essentials of robust control, volume 104, Prentice hall Upper Saddle River, NJ, 1998.
  - [16] S. Xu, T. Chen, Robust h-infinity control for uncertain stochastic systems with state delay, IEEE Transactions on Automatic Control (TAC) 47 (2002) 2089–2094.
  - [17] A. Majumdar, R. Tedrake, Funnel libraries for real-time robust feedback motion planning, The International Journal of Robotics Research (IJRR) 36 (2017) 947–982.
- [18] S. Singh, A. Majumdar, J.-J. Slotine, M. Pavone, Robust online motion planning via contraction theory and convex optimization, in: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), 2017, pp. 5883–5890. doi:10.1109/ICRA.2017.7989693.
  - [19] A. Majumdar, A. Farid, A. Sonar, PAC-Bayes Control: Learning policies that provably generalize to novel environments, The International Journal of Robotics Research (IJRR) 40 (2021) 574–593.
- [20] A. Farid, S. Veer, A. Majumdar, Task-driven out-of-distribution detection with statistical guarantees for robot learning, in: Proceedings of the Conference on Robot Learning (CoRL), 2021.
  - [21] S. Veer, A. Majumdar, Probably approximately correct vision-based planning using motion primitives, in: Proceedings of the 2020 Conference on Robot Learning (CoRL), volume 155 of Proceedings of Machine Learning Research, PMLR, 2021, pp. 1001–1014.
- [22] J. García, F. Fernández, A comprehensive survey on safe reinforcement learning, Journal of Machine Learning Research (JMLR) 16 (2015) 1437–1480.
  - [23] S. Bansal, M. Chen, S. Herbert, C. J. Tomlin, Hamilton-jacobi reachability: A brief overview and recent advances, in: Proceedings of the IEEE 56th Annual Conference on Decision and Control (CDC), 2017, pp. 2242–2253. doi:10.1109/CDC.2017.8263977.

- [24] J. F. Fisac, M. Chen, C. J. Tomlin, S. S. Sastry, Reach-Avoid Problems with Time-Varying Dynamics, Targets and Constraints, in: Proceedings of the 18th International Conference on Hybrid Systems: Computation and Control, HSCC '15, New York, NY, USA, 2015, p. 11–20. doi:10.1145/2728606.2728612.
  - [25] K. Leung, E. Schmerling, M. Zhang, M. Chen, J. Talbot, J. C. Gerdes, M. Pavone, On infusing reachability-based safety assurance within planning frameworks for human–robot vehicle interactions, The International Journal of Robotics Research 39 (2020) 1326–1345.
  - [26] R. Cheng, G. Orosz, R. M. Murray, J. W. Burdick, End-to-end safe reinforcement learning through barrier functions for safety-critical continuous control tasks, in: Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence, AAAI'19/IAAI'19/EAAI'19, AAAI Press, 2019. doi:10.1609/aaai.v33i01.33013387.
- [27] G. Dalal, K. Dvijotham, M. Vecerik, T. Hester, C. Paduraru, Y. Tassa, Safe exploration in continuous action spaces, 2018. arXiv:1801.08757.
  - [28] B. Chen, J. Francis, J. Oh, E. Nyberg, S. L. Herbert, Safe autonomous racing via approximate reachability on ego-vision, 2021. arXiv:2110.07699.
- [29] F. Berkenkamp, A. P. Schoellig, A. Krause, Safe controller optimization for quadrotors with gaussian processes, in: Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), 2016, pp. 491–496. doi:10.1109/ICRA.2016.7487170.
  - [30] T. Koller, F. Berkenkamp, M. Turchetta, A. Krause, Learning-based model predictive control for safe exploration, in: Proceedings of the IEEE Conference on Decision and Control (CDC), 2018, pp. 6059–6066. doi:10.1109/CDC.2018.8619572.
- [31] A. Liu, G. Shi, S.-J. Chung, A. Anandkumar, Y. Yue, Robust regression for safe exploration in control, in: Proceedings of the 2nd Conference on Learning for Dynamics and Control, volume 120 of Proceedings of Machine Learning Research, PMLR, 2020, pp. 608–619. URL: https://proceedings.mlr.press/v120/liu20a.html.
- [32] V. N. Vapnik, A. Y. Chervonenkis, On the uniform convergence of relative frequencies of events to their probabilities, in: Measures of complexity, Springer, 2015, pp. 11–30.
  - [33] O. Bousquet, S. Boucheron, G. Lugosi, Introduction to statistical learning theory, in: Summer school on machine learning, Springer, 2003, pp. 169–207.
  - [34] D. A. McAllester, Some pac-bayesian theorems, Machine Learning 37 (1999) 355-363.
- [35] G. K. Dziugaite, D. M. Roy, Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data, in: Proceedings of the Thirty-Third Conference on Uncertainty in Artificial Intelligence (UAI), Sydney, Australia, August 11-15, 2017.
  - [36] M. Pérez-Ortiz, O. Rivasplata, J. Shawe-Taylor, C. Szepesvári, Tighter risk certificates for neural networks, Journal of Machine Learning Research (JMLR) 22 (2021).
- [37] A. Z. Ren, S. Veer, A. Majumdar, Generalization guarantees for imitation learning, in: Proceedings of the 2020 Conference on Robot Learning (CoRL), volume 155 of Proceedings of Machine Learning Research, PMLR, 2021, pp. 1426–1442.
  - [38] A. E. Gurgen, A. Majumdar, S. Veer, Learning provably robust motion planners using funnel libraries, arXiv preprint arXiv:2111.08733 (2021).
- [39] A. Agarwal, S. Veer, A. Z. Ren, A. Majumdar, Stronger generalization guarantees for robot learning by combining generative models and real-world data, arXiv preprint arXiv:2111.08761 (2021).

- [40] A. Farid, D. Snyder, A. Z. Ren, A. Majumdar, Failure prediction with statistical guarantees for vision-based robot control, in: Proceedings of the Robotics: Science and Systems (RSS), 2022.
- [41] B. Eysenbach, A. Gupta, J. Ibarz, S. Levine, Diversity is all you need: Learning skills without a reward function, in: Proceedings of the International Conference on Learning Representations (ICLR), 2019.
- [42] F. Bonin-Font, A. Ortiz, G. Oliver, Visual navigation for mobile robots: A survey, Journal of Intelligent and Robotic Systems 53 (2008) 263–296.
  - [43] R. Sim, J. J. Little, Autonomous vision-based exploration and mapping using hybrid maps and Rao-Blackwellised particle filters, in: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2006, pp. 2082–2089. doi:10.1109/IROS.2006.282485.
- [44] S. Thrun, A. Bücken, Integrating grid-based and topological maps for mobile robot navigation, in: Proceedings of the AAAI Conference on Artificial Intelligence, 1996, pp. 944–951.
  - [45] S. Bansal, V. Tolani, S. Gupta, J. Malik, C. Tomlin, Combining optimal control and learning for visual navigation in novel environments, in: Proceedings of the 2020 Conference on Robot Learning (CoRL), volume 100 of Proceedings of Machine Learning Research, PMLR, 2020, pp. 420–429.
- [46] S. Gupta, J. Davidson, S. Levine, R. Sukthankar, J. Malik, Cognitive mapping and planning for visual navigation, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017, pp. 2616–2625.
  - [47] C. Richter, N. Roy, Safe visual navigation via deep learning and novelty detection, in: Proceedings of Robotics: Science and Systems (RSS), Cambridge, Massachusetts, 2017. doi:10.15607/RSS.2017.XIII.064.

800

- [48] L. Wellhausen, R. Ranftl, M. Hutter, Safe robot navigation via multi-modal anomaly detection, IEEE Robotics and Automation Letters (RAL) 5 (2020) 1326–1333.
- [49] B. Lütjens, M. Everett, J. P. How, Safe reinforcement learning with model uncertainty estimates, in: Proceedings of the International Conference on Robotics and Automation (ICRA), IEEE, 2019, pp. 8662–8668.
- [50] G. Kahn, A. Villaflor, V. Pong, P. Abbeel, S. Levine, Uncertainty-aware reinforcement learning for collision avoidance, arXiv preprint arXiv:1702.01182 (2017).
- [51] A. Bajcsy, S. Bansal, E. Bronstein, V. Tolani, C. J. Tomlin, An eficient reachability-based framework for provably safe autonomous navigation in unknown environments, in: Proceedings of the IEEE 58th Conference on Decision and Control (CDC), IEEE, 2019, pp. 1758–1765.
- [52] A. Li, S. Bansal, G. Giovanis, V. Tolani, C. Tomlin, M. Chen, Generating robust supervision for learning-based visual navigation using hamilton-jacobi reachability, in: Proceedings of the 2nd Conference on Learning for Dynamics and Control, volume 120 of Proceedings of Machine Learning Research, PMLR, 2020, pp. 500–510.
- [53] F. Ramos, R. C. Possas, D. Fox, Bayessim: adaptive domain randomization via probabilistic inference for robotics simulators, arXiv preprint arXiv:1906.01728 (2019).
  - [54] Y. Chebotar, A. Handa, V. Makoviychuk, M. Macklin, J. Issac, N. Ratliff, D. Fox, Closing the sim-to-real loop: Adapting simulation randomization with real world experience, in: 2019 International Conference on Robotics and Automation (ICRA), 2019.
- [55] V. Lim, H. Huang, L. Y. Chen, J. Wang, J. Ichnowski, D. Seita, M. Laskey, K. Goldberg, Planar robot casting with real2sim2real self-supervised learning, arXiv preprint arXiv:2111.04814 (2021).

- [56] B. Mehta, M. Diaz, F. Golemo, C. J. Pal, L. Paull, Active domain randomization, in: Conference on Robot Learning, 2020, pp. 1162–1176.
- [57] F. Muratore, C. Eilers, M. Gienger, J. Peters, Data-eficient domain randomization with bayesian optimization, IEEE Robotics and Automation Letters 6 (2021) 911–918.
  - [58] M. Cutler, T. J. Walsh, J. P. How, Reinforcement learning with multi-fidelity simulators, in: Proceedings of the IEEE/RSJ International Conference on Robotics and Automation (ICRA), 2014.
  - [59] G. Shafer, V. Vovk, A tutorial on conformal prediction., Journal of Machine Learning Research 9 (2008).
- [60] T. Haarnoja, A. Zhou, P. Abbeel, S. Levine, Soft Actor-Critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor, in: Proceedings of the 35th International Conference on Machine Learning, volume 80 of Proceedings of Machine Learning Research, PMLR, 2018, pp. 1861– 1870.
- [61] M. Alshiekh, R. Bloem, R. Ehlers, B. Könighofer, S. Niekum, U. Topcu, Safe reinforcement learning via shielding, in: Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence and Thirtieth Innovative Applications of Artificial Intelligence Conference and Eighth AAAI Symposium on Educational Advances in Artificial Intelligence, AAAI Press, 2018.
  - [62] A. Jabri, K. Hsu, A. Gupta, B. Eysenbach, S. Levine, C. Finn, Unsupervised curricula for visual metareinforcement learning, in: Advances in Neural Information Processing Systems (NeurIPS), volume 32, 2019, pp. 10519–10530.

840

- [63] S. Kumar, A. Kumar, S. Levine, C. Finn, One solution is not all you need: Few-shot extrapolation via structured MaxEnt RL, in: Advances in Neural Information Processing Systems (NeurIPS), volume 33, Curran Associates, Inc., 2020, pp. 8198–8210.
- [64] A. Sharma, S. Gu, S. Levine, V. Kumar, K. Hausman, Dynamics-aware unsupervised discovery of skills, in: Proceedings of the International Conference on Learning Representations (ICLR), 2020.
- [65] J. Langford, R. Caruana, (Not) bounding the true error, in: Advances in Neural Information Processing Systems (NeurIPS), volume 14, MIT Press, 2002.
- [66] S. Levine, P. Pastor, A. Krizhevsky, J. Ibarz, D. Quillen, Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection, The International Journal of Robotics Research 37 (2018) 421–436.
- [67] Quartz, Amazon this company built one of the world's most efficient warehouses by embracing chaos, https://classic.qz.com/perfect-company-2/1172282/this-company-built-one-of-the-worlds-most-efficient-warehouses-by-embracing-chaos/, 2019.
- [68] FutureCar, A look at how waymo's self-driving test fleet safely traveled 2.7 million miles in san francisco last year, https://www.futurecar.com/5158/A-Look-at-How-Waymos-Self-Driving-Test-Fleet-Safely-Traveled-2-7-Million-Miles-in-San-Francisco-Last-Year, 2022.
  - [69] J. Ichnowski, K. Chen, K. Dharmarajan, S. Adebola, M. Danielczuk, V. Mayoral-Vilches, H. Zhan, D. Xu, R. Ghassemi, J. Kubiatowicz, et al., Fogros 2: An adaptive and extensible platform for cloud and fog robotics using ros 2, arXiv preprint arXiv:2205.09778 (2022).
- [70] B. Eysenbach, S. Gu, J. Ibarz, S. Levine, Leave no trace: Learning to reset for safe and autonomous reinforcement learning, in: Proceedings of the 6th International Conference on Learning Representations (ICLR), 2018. URL: https://openreview.net/forum?id=S1vuO-bCW.

- [71] Z. Borsos, M. Mutny, A. Krause, Coresets via bilevel optimization for continual learning and streaming, Advances in Neural Information Processing Systems 33 (2020) 14879–14890.
- [72] B. Guedj, A primer on pac-bayesian learning, arXiv preprint arXiv:1901.05353 (2019).
  - [73] S. Arora, R. Ge, B. Neyshabur, Y. Zhang, Stronger generalization bounds for deep nets via a compression approach, in: International Conference on Machine Learning, PMLR, 2018, pp. 254–263.

# Appendix A. Derivations for Inducing Diversity into Backup Policy Update

We add observation-conditional mutual information term to the loss function of backup policy.

We then approximate the expectation by the transitions sampled from the replay buffer as

$$L(\theta) \approx E_{(o,z)B,aB} \frac{\theta}{\pi}(\cdot|o,z) \quad Q(o,a;z) - \nu \log \pi_{\theta}(a|o,z) + \nu \log p(a|o) . \tag{A.2}$$

Finally, we approximate the marginal with the latent variables sampled from the distribution (empirical measure) as

$$L(\theta) \approx E_{(o,z)\text{BB},a\text{B}} \frac{\theta}{\pi(\cdot \mid o,z)} Q(o,a;z) - v \log \pi_{\theta}(a\mid o,z) + v \frac{1}{n_{s}} \frac{\chi^{n_{s}}}{n_{s}} \pi_{\theta}(a\mid o,z_{i}) . \tag{A.3}$$

Appendix B. Training Hyperparameters used in Experiments

We show the training hyperparameters used to generate the results in Fig. 5.

Table B.3: Hyperparameters for PAC\_Shield\_Perf in Sim training. Same neural network architecture is used for performance and backup policies.

	Environment Setting				
	Vanilla-Normal/Dynamics	Vanilla-Task	Advanced-Env		
# training steps	500000	1000000	4000000		
Replay buffer size	50000 (steps)	100000 (steps)	5000 (trajectories)		
Optimize frequency	2000	2000	20000		
# updater per optimize	1000	1000	1000		
Value shielding threshold	-0.05	-0.05	-0.05		
Latent Distribution					
Latent dimension (nz)	20	20	30		
Augmented reward coeficient (β)	2	2	2		
Prior standard deviation	2	2	2		
Optimization					
Optimizer	Adam	Adam	Adam		
Batch size (Performance)	128	128	128		
Discount factor (Performance)	0.99	0.99	0.99		
Learning rate (Performance)	0.0001	0.0001	0.0001		
Batch size (Backup)	128	128	128		
Discount factor (Backup)	$0.8 \to 0.999$	$0.8 \to 0.999$	$0.8 \to 0.99$		
Learning rate (Backup)	0.0001	0.0001	0.001		
N N Architecture					
Input channels	3	3	22 <sup>a</sup>		
CNN kernel size	[5,3,3]	[5,3,3]	[7,5,3]		
CNN stride	[2,2,2]	[2,2,2]	[4,3,2]		
CNN channel size	[8,16,32]	[8,16,32]	[16,32,64]		
MLP dimensions	[130+n <sub>z</sub> <sup>b</sup> ,128]	[132+n <sub>z</sub> <sup>b</sup> ,128]	[248+n <sub>z</sub> b ,256,256]		
Hardware Resource					
# CPU threads	8	8	16		
GPU	Nvidia V100 (16GB)	Nvidia V100 (16GB)	Nvidia A100 (40GE		
Runtime	8 hours	14 hours	12 hours		

<sup>&</sup>lt;sup>a</sup> We stack 4 previous RGB images while skipping 3 frames between two images and concatenate the stacked images with the first 10 elements of the latent variable (each element is repeated to match the same shape of a channel in an image).

# Appendix C. Environment Setup for Advanced-Env

In order to train the navigating agent in realistic environments before Real deployment, we use the 3D-FRONT (3D Furnished Rooms with layOuts and semaNTics) dataset [6] that offers a larger number of synthetic indoor scenes with professionally designed layouts and high-quality textured furniture. This is the richest dataset we find suitable to indoor navigation task, training with domain randomization and PAC-Bayes Control framework often requires more than 1000 environments.

For Sim training, we use 7m × 7m undecorated rooms as room layouts, and randomly placing 5 pieces of furniture from the dataset. We use 4 categories of furniture: Soft (2701 pieces available), Chair (1775 pieces available), Cabinet/Shelf/Desk (5725 pieces available), Table (1090 pieces available). We also randomly sample textures from the dataset to add to the walls and floor: for walls, we use categories Tile, Wallpaper, and Paint (911 images available in total), and for floor, we use Flooring, Stone, Wood, Marble, Solid Wood Flooring (466 images available in total). We set the minimum clearance between furniture, around the initial location, and around the goal to be 1m. The minimum distance between the initial location and the goal

<sup>&</sup>lt;sup>b</sup> The input of the first linear layer is composed of the output from the convolutional layers, latent variables and auxiliary signals, which is  $128 + n_z + 2$  in Vanilla-Normal/Dynamics,  $128 + n_z + 4$  in Vanilla-Task and  $256 + (n_z - 10) + 2$  in Advanced-Env.

Table B.4: Hyperparameters for PAC\_Shield\_Perf in Lab training.

	Environment Setting			
	Vanilla-Env	Advanced-Env		
# training steps	500000	3000000		
Replay buffer size	50000 (steps)	5000 (trajectories)		
Optimize frequency	2000	20000		
# updater per optimize	1000	1000		
Value shielding threshold	-0.05	-0.05		
The number of environments (N)	1000	1000		
Optimization				
Learning rate for latent mean	0.0001	0.0001		
Learning rate for latent std	0.0001	0.0001		
KL-divergence coeficient (α)	1	2		
Optimizer	Adam	Adam		
Batch size (Performance)	1024	128		
Discount factor (Performance)	0.99	0.99		
Learning rate (Performance)	0.0001	0.0001		
PAC-Bayes Bound				
The number of latent variables (L)	1000	1000		
Precision (δ)	0.01	0.01		
Hardware Resource				
# CPU threads	8	8		
GPU	Nvidia V100 (16GB)	Nvidia A100 (40GB)		
Runtime	6 hours	16 hours		

is 5m. Fig. C.12 shows samples of observations at the initial locations. For Advanced-Dense Lab where the furniture density is higher, we place 6 instead of 5 pieces of furniture, and the minimum clearance is 0.8m instead of 1m.

For Lab training, we instead use the professionally designed room layouts (with furniture configuration) from the dataset. The dataset contains 6813 different house layouts (each with multiple rooms). Since our focus is on obstacle avoidance with relatively short horizon, in each house, we try to sample initial and goal locations within one room. Unfortunately the dataset does not provide corresponding wall and floor textures in each layout, and we resort to random samples as in Vanilla-Env. Again we maintain a minimum clearance of 1m between furniture, around the initial and goal locations. To check the environment is solvable, we extract a 2D occupancy map for each room and run the Dijkstra algorithm. We also ensure there is at least one piece of furniture along the line connecting the initial and goal locations. We tend to find that many rooms are too crowded or the found path does not have enough clearance for the quadrupedal robot (about 0.5m wide). At the end, we are able to process about 2000 room environments, which are then split for training and testing. Fig. C.13 shows samples of observations at the initial locations.



Figure C.12: Samples of robot observations in Sim training of Advanced-Env: for better view here, the virtual camera is placed at a higher location than the robot.



Figure C.13: Samples of robot observations in Advanced-Realistic Lab: for better view here, the virtual camera is placed at a higher location than the robot.