
Cautious Adaptation For Reinforcement Learning in Safety-Critical Settings

Jesse Zhang¹ Brian Cheung¹ Chelsea Finn² Sergey Levine¹ Dinesh Jayaraman³

Abstract

Reinforcement learning (RL) in real-world safety-critical target settings like urban driving is hazardous, imperiling the RL agent, other agents, and the environment. To overcome this difficulty, we propose a “safety-critical adaptation” task setting: an agent first trains in non-safety-critical “source” environments such as in a simulator, before it adapts to the target environment where failures carry heavy costs. We propose a solution approach, CARL, that builds on the intuition that prior experience in diverse environments equips an agent to estimate risk, which in turn enables relative safety through risk-averse, cautious adaptation. CARL first employs model-based RL to train a probabilistic model to capture uncertainty about transition dynamics and catastrophic states across varied source environments. Then, when exploring a new safety-critical environment with unknown dynamics, the CARL agent plans to avoid actions that could lead to catastrophic states. In experiments on car driving, cartpole balancing, half-cheetah locomotion, and robotic object manipulation, CARL successfully acquires cautious exploration behaviors, yielding higher rewards with fewer failures than strong RL adaptation baselines. Website at <https://sites.google.com/berkeley.edu/carl>.

1. Introduction

An experienced human driving a rental car for the first time is initially very aware of her lack of familiarity with the car. How sensitive is it to acceleration and braking? How does it respond to steering? How wide is the vehicle and what is its turning radius? She drives mindfully, at low speeds and making wide turns, all the while observing the car’s

¹UC Berkeley, CA, USA ²Stanford, CA, USA ³University of Pennsylvania, PA, USA. Correspondence to: Jesse Zhang <jessezhang@berkeley.edu>.



Figure 1. The Safety-Critical Adaptation (SCA) task framework. In an initial “pretraining” phase, an agent is trained in various sandbox source environments without safety concerns, such as in a simulator. In the second “adaptation” phase, it aims to adapt to a safety-critical target environment, producing high rewards quickly and safely.

responses and adapting to it. Within minutes, once she is familiar with the car, she begins to drive more fluently and efficiently. Humans draw upon their prior experiences to perform this kind of safe, quick adaptation to unfamiliar situations all the time, such as when playing with a new tennis racquet, or walking on a new slippery surface.

Such problems are critical to address in autonomous systems: such as when a self-driving car must learn to drive in a new country, or when a planetary rover might have to learn to explore a harsh new environment. Missteps in real-world situations can cause real damage to robots and their environments. An important bottleneck in applying today’s standard machine learning approaches to control in these real-world situations is that they are trained without any notion of safe behavior under uncertainty. Recent works have attempted to address this by proposing methods for safe exploration during reinforcement learning — in other words, how might an agent avoid risky actions during training time? This still requires that the robot acquire its notions of uncertainty and risks at the same time as it is learning to perform tasks in the new environment, which is difficult and hazardous.

Could we instead rely on transferring notions of uncertainty

and risk acquired from prior experience in other related domains, such as in simulated environments, where safety may not be as much of a concern? In other words, we could make the safe learning problem easier through knowledge transfer, relaxing the problem to safe *adaptation*, like the human driver above. Rather than making our planetary rover learn from scratch on a new planet, we would ask: how might it draw on its experience in varied terrains on Earth to perform cautious adaptation to the new planet? This “safety-critical adaptation” task framework, providing a route to real-world reinforcement learning, is the first contribution of this work.

Next, we propose a solution approach based on model-based reinforcement learning, called Cautious Adaptation in RL (CARL). CARL works by first pretraining a probabilistic model of transition dynamics and catastrophe on a population of training domains with varied, unknown dynamics. Through this experience over many environments, the model learns to estimate the epistemic uncertainty (model uncertainty) of unknown environment dynamics, thus permitting estimation of a distribution of outcomes for any action executed by the agent. When introduced into a new target environment, CARL uses this estimated distribution of outcomes to select *cautious* actions that have a low chance of leading to catastrophic outcomes. Much like the human driver in the opening example above, all the information collected during this cautious phase of exploration is fed back into the model to adapt it to the new environment, leading to increasingly well-tuned, confident predictions.

In our experiments, we evaluate CARL in four safety-critical adaptation settings: driving a wider car, balancing a longer cartpole, controlling a half-cheetah robot with one disabled joint, and manipulating larger balls. CARL consistently achieves higher rewards with fewer catastrophic failures than meta-learning and other RL baselines within few adaptation episodes in each setting. Further, CARL successfully acquires interpretable risk-averse exploration behaviors appropriate to each setting. We view this safety-critical adaptation framework and the results of CARL as promising initial steps towards making reinforcement learning possible in safety-critical real-world settings like robotics.

2. Related Work

Cautious or risk-averse learning has close connections to learning robust control policies, as well as the uncertainty estimation derived from Bayesian reinforcement learning (Ghavamzadeh et al., 2015; Strens, 2000). Rather than conventionally maximizing a reward function, accounting for risk usually involves allocating more attention to ‘worst-case’ outcomes in an environment.

Risk-Aversion in RL. Incorporating safety requires man-

aging risks and reducing the impact of unforeseen negative outcomes. Risk management is extensively studied in quantitative finance. In portfolio optimization, a commonly used quantity that measures the expected return considering the worse α -% of cases is Conditional Value at Risk (CVaR) (Rockafellar et al., 2000). Rather than optimizing the expected reward, risk averse policies optimize the lower α -quartile of the distribution of rewards. To explicitly address safety, the RL community has adopted measures like CVaR into the objective (Morimura et al., 2010; Borkar & Jain, 2010; Chow & Ghavamzadeh, 2014; Tamar et al., 2015; Chow et al., 2015) to create risk-averse policies. Rajeswaran et al. (2016) propose learning robust policies by sampling from the α -quartile of an ensemble of models. While the model ensemble is trained on a given source distribution, the policy is only trained on the lower α -quartile rewards from trajectories sampled from this ensemble. This leads to more conservative and robust policies.

Model-Based Safety Approaches. Prior work has utilized dynamics models to define notions of safety. One approach is to calculate safety constraints and then utilize these models, typically Gaussian Processes, to satisfy them (Fisac et al., 2017; Sadigh & Kapoor, 2016; Berkenkamp et al., 2017; Ostafew et al., 2016; Hakobyan et al., 2019; Hanssen & Foss, 2015; Hewing et al., 2019). Another approach is to incorporate the models’ transition predictions into the cost function for model-free policies (Berkenkamp et al., 2017). Other methods have used these models to perform robust model-predictive control (Aswani et al., 2013). Generally, these works make strict assumptions, e.g. linear dynamics, quadratic cost functions, or additive state perturbations.

There are also model-based RL approaches that aim to reduce other notions of risk. Yu et al. (2019) estimate future states as additional inputs to a policy maximizing financial portfolio returns and Lecarpentier & Rachelson (2019) propose risk-averse tree search for non-stationary MDPs.

Epistemic Uncertainty in RL. Prior work has shown that model-based RL models can be quickly adapted during test time by meta-learning for fast adapting parameters during training (Nagabandi et al., 2018; Sæmundsson et al., 2018). These fast-adapting parameters offer greater flexibility in adapting to new circumstances which an agent may encounter at test time. Nagabandi et al. (2018) show that real robots can quickly adapt to miscalibrations during evaluation through this fast adaptation acquired through meta-learning. Such approaches are complementary to our approach, as they provide a means to explicitly train for fast adaptation to disturbances in the environment, while they do not account for any notion of safety.

Henaff et al. (2019) propose using the uncertainty of a model to regularize policy learning. The policy is encouraged to create trajectories which are distributionally close to tra-

jectories observed in training data. In other words, the policy attempts to keep the trajectories within its training domain. In our work, our policy is encouraged to behave cautiously in unfamiliar environments rather than to explicitly remain in familiar ones. Kahn et al. (2017) train a collision prediction model in a model-based framework to favor safe, low-velocity collisions for robot navigation. Rather than learning to estimate uncertainties from scratch within a safety-critical environment, our approach exploits pretraining to make learning safe from the first instant in the target environment.

Domain randomization. Our safety-critical adaptation setting closely follows the domain randomization (Sadeghi & Levine, 2017; Peng et al., 2018; Tobin et al., 2017) framework. While both attempt to pretrain policies in a mix of source environments and transfer them to some target environment, CARL focuses on safe adaptation to the target environment — to accomplish this, it follows an explicitly cautious action policy at adaptation time, different from the policy used in the pretraining environments.

3. Background: PETS

Before discussing CARL, we first lay out some preliminaries. CARL can be built upon any model-based RL approach with ensembled dynamics models and MPC planning. For this work, we build upon PETS (Chua et al., 2018), a recently proposed approach for model-based reinforcement learning. We describe the main features of the PETS framework below:

Probabilistic dynamics model. PETS trains an ensemble of probabilistic dynamics models within its environment. Each model in the ensemble is a probabilistic neural network that outputs a distribution over the next state s' conditioned on the current state s and action a . The data for training these models comes from trajectories executed by following the action selection scheme described below.

Action selection. This action selection scheme is sampling-based model-predictive control (MPC): an evolutionary search method that finds action sequences with the highest predicted reward. The reward of an action sequence in turn is computed by propagating action outcomes autoregressively through the learned probabilistic models.

Reward computation. Specifically, starting from a state s_0 , for each sampled action sequence $A = [a_1, \dots, a_H]$, where H is the planning horizon, the dynamics model f first predicts a distribution over s_1 after executing a_0 . A particle propagation method samples Monte Carlo samples from this distribution. For each sample, the dynamics model then predicts the state distribution for s_2 , conditioned on executing a_1 , and the process repeats. This recursive particle propagation results in a large number N of particles

$\{\hat{s}_H^i\}_{i=1}^N$ after H steps. These N particles represent samples from the distribution of possible states after executing A . Each such particle $i \in [1, N]$ is now assigned a predicted reward r^i , which is a function of its full state trajectory starting from s_0 . Finally, the mean of those predicted rewards is considered the score of the action sequence:

$$R(A) = \sum_i r^i / N. \quad (1)$$

We call this the action score. Then, the winning action sequence $A^* = \arg \max_A R(A)$ with the highest action score is selected, the first action in that sequence is executed, and the whole process repeats starting from the resulting new state s_1 .

4. Safety-Critical Adaptation with CARL

We first define the “safety-critical adaptation” (SCA) task setting for enabling RL in safety-critical settings, as motivated in Sec 1. SCA involves two phases:

- **Pretraining:** An agent trains in various “source” sandbox environments where safety is not a concern. In general, catastrophic states, even if encountered, do not cause an RL episode to terminate here.
- **Adaptation:** The agent must adapt to the safety-critical target environment, which generally lie outside the distribution of training environments, such as in a sim-to-real transfer setting. Catastrophic states such as car collisions, when encountered, lead to immediate RL episode termination.

We evaluate SCA approaches for their ability to achieve *high task rewards, quickly, and while avoiding catastrophe*.

To meet these requirements, we propose CARL, Cautious Adaptation in RL, which builds upon PETS (Sec 3). Although we use PETS in this work, we note CARL can be applied on most model-based RL approaches with sampling-based planners. We demonstrate this by applying CARL to another model-based algorithm, as described in Sec 5. We describe its behavior during the two SCA phases below.

4.1. Pretraining: Modeling Epistemic Uncertainties

Consider that each environment is specified by a variable z , which controls the state transition dynamics as well as state safety, i.e., whether a given state is safe or unsafe in that environment. In general, the target environment ID z is unseen among pretraining environments.

In PETS (Chua et al., 2018), the probabilistic ensemble is trained to represent uncertainty within a single environment. In CARL, we would instead like to train the ensemble to

Algorithm 1 Pretraining

```

1: Initialize probabilistic ensemble dynamics model  $f$  and state
  safety model  $g$ .
2: Collect data  $\mathcal{D}$  by executing a random controller in one random
  training environment for one trial.
3: for environment ID  $z \sim$  training environments do
4:   Train the models  $f$  and  $g$  to predict state transitions and
  state safety respectively, on  $\mathcal{D}$ 
5:   for  $t = 0$  to task horizon do
6:     for evolutionary search stage= $1,2,\dots$  do
7:       for sampled action sequence  $A$  do
8:         Run state propagation to produce  $N$  particles
9:         Evaluate  $A$  as  $R(A) = \sum_i r_i/N$ 
10:        end for
11:        Refine search to find  $A^* = \arg \max R(A)$ 
12:       end for
13:       Set  $a_t$  to first action of  $A^*$ , and execute  $a_t$ .
14:       Record the state transition and state safety label ( $w_{t+1}$ )
  as a tuple  $(s_t, a_t, s_{t+1}, w_{t+1})$  in  $\mathcal{D}$ 
15:     end for
16:   end for

```

capture the uncertainty associated with being dropped into a new environment, with unknown environment ID z .

To do this, during the SCA pretraining phase, CARL trains a single PETS agent across all pretraining environments, with varying, unknown values of z . Specifically, at the beginning of each pretraining episode, we randomly sample one of the pretraining z 's from a uniform distribution. Since z is unavailable to the agent, the dynamics model must capture epistemic uncertainty about state transitions and state safety due to unknown z . Algorithm 1 shows pseudocode.

4.2. Adaptation: Risk-Averse Action Selection

After this pretraining, CARL exploits the uncertainty captured in the ensemble to perform risk-averse action selection during adaptation. We provide two versions of this scheme, representing aversion to two different notions of risk.

Case 1: Low Reward Risk-Aversion. In this case, CARL minimizes the risk of producing low rewards over an episode. To do this, it replaces the action score $R(A)$ of equation 1 in PETS with a newly defined ‘‘generalized action score’’ $R_\gamma(A)$, in which the ‘‘caution parameter’’ $\gamma \in [0, 100]$ controls the degree of caution exercised in evaluating action sequences in the new environment. $R_\gamma(A)$ is defined as:

$$R_\gamma(A) = \sum_{i:r^i \leq v_{100-\gamma}(r)} r^i/N, \quad (2)$$

where $v_k(r)$ denotes the value of the k^{th} percentile of predicted rewards $\{r^j\}_{j=1}^N$ among the N particles after particle propagation.

Unpacking this definition, R_γ measures the mean score of the bottom $100 - \gamma$ percentile of the predicted outcomes

from the PETS model. When $\gamma = 50$, for instance, it measures the mean of the worst 50 percentile of predicted rewards. This is a pessimistic evaluation of the prospects of the action sequence A — it only pays attention to the worst performing particles in the distribution. At caution $\gamma = 0$, R_γ exactly matches the definition of R in equation 1: it measures the mean predicted reward of all the particles produced after particle propagation through the model. In our experiments, we heuristically set $\gamma = 50$.

To implement low reward risk-aversion, CARL simply selects actions that optimize this generalized action score $A_\gamma^* = \arg \max_A R_\gamma(A)$ at adaptation time in the target environment.

CARL with low reward risk aversion does not require anything more than the reward function to be available at each step during pretraining. Yet, in many cases, it might automatically identify and avoid catastrophic states: for the task of driving a car to a goal, driving into a wall is catastrophic since it is highly predictive of low reward.

Case 2: Catastrophic State Risk-Aversion. On the other hand, in many settings, a notion of state safety that is decoupled from the task reward is appropriate. For instance, states where the wheels of a car are on the curb may be deemed catastrophic, even if a strategy involving such states might yield high task rewards for steering around a corner.

Aversion to such catastrophic states is only possible if the pretraining environments also provide state safety annotations $w(s)$ for each state s encountered by the agent. Given such annotations, during pretraining, CARL trains an ensemble state safety predictor $g(s)$ to predict $w(s)$ alongside the ensemble dynamics models f from PETS.

Then, during adaptation, rather than relying on reward-based pessimism to obtain risk-averse behavior, we directly evaluate risk by predicting the probability of encountering catastrophic states and penalize it. Specifically, rather than directly maximizing the action score $R(A)$ of Eq 1 during planning, we include a new penalty proportional to the estimated probability of encountering catastrophic states, as predicted by g .

$$R_\lambda(A) = R(A) - \lambda g(A), \quad (3)$$

where $g(A)$ is predicted probability of catastrophic failure, computed as the ensemble mean of catastrophe probabilities, summed over the predicted state trajectory. Actions are chosen to maximize R_λ . See Appendix section A for a description of how $g(A)$ is computed.

4.3. Adaptation: Model Finetuning

As it gathers experience in the target environment using its risk-averse policy, CARL also improves its dynamics model

Algorithm 2 Adaptation

```

1: Inputs: Pretraining dataset  $\mathcal{D}$ 
2: for target environment adaptation episode=1,2,... do
3:   for  $t = 0$  to task horizon do
4:     for evolutionary search stage=1,2,... do
5:       for sampled action sequence  $A$  do
6:         Run state propagation
7:         Evaluate  $A$  with generalized score (Eq 2 or 3)
8:       end for
9:       Refine search to find  $A^* = \arg \max R_\gamma(A)$ 
10:    end for
11:    Execute first action of  $A^*$ 
12:    Record outcome in  $\mathcal{D}$ 
13:    Finetune the probabilistic ensemble model  $f$  on  $\mathcal{D}$ 
14:  end for
15: end for

```

over time to finetune it to the new environment. Since dynamics models do not need manually specified reward functions during training, the ensemble model can continue to be trained in the same way as during pretraining.

As the model improves over time, the distribution of predicted outcomes becomes more and more narrow over time. The adaptation procedure in Algorithm 2 sums up cautious action selection and model finetuning.

Empirically, we find it useful to stabilize adaptation by maintaining data from the pretraining episodes in the training set and adding target environment data as it is acquired. CARL computes all model updates during adaptation on this combined dataset. We use probabilistic neural network ensembles for the dynamics model (Chua et al., 2018), and training proceeds through stochastic gradient descent.

5. Experiments

We evaluate CARL for safety-critical adaptation (SCA) in four separate settings: cartpole, half-cheetah, simulated car driving, and object manipulation with a robotic hand. Recall that in SCA, the aim is to adapt effectively and quickly to a safety-critical target environment with few catastrophic events, after having trained on some sandbox source environments in the pretraining phase.

5.1. Experimental Setup

First, we describe our four safety-critical adaptation settings in detail:

- **CartPole balancing:** We modify the standard OpenAI Gym cartpole task so that the cartpole is a different length in each environment (see top left pane of Fig 2). We randomly sample this cartpole length from the range [0.4, 0.8] during the pretraining phase. We define catastrophic states as those where the pole falls below the cart or the cart falls off the edge of the rail. We test on one environ-




Figure 2. Images of our test environments: CartPole, Half-Cheetah Duckietown car driving, and Baoding ball manipulation.

ment from the pretraining source environment distribution (“in-distribution”) with cartpole length 0.6, and five out-of-distribution environments, with lengths 0.3, 1.0, 1.5, 2.0, 2.5. As discussed in Sec 4, we expect that cautious behavior is most pertinent when the target environment is significantly different from the source pretraining environments. This out-of-distribution setting also reproduces what happens when a car driving policy must be transferred sim-to-real, or when a rover is trained on Earth terrains but must work on a new planet.

- **Half-cheetah locomotion:** To test SCA in the Gym half-cheetah setting, we randomly disable a joint in each training environment. This joint is marked in red in visualizations, see the top right pane of Fig 2. During adaptation, a held-out joint is disabled. We define states where the head is in contact with the floor as catastrophic.
- **Car driving:** Our driving environment is based on Duckietown (Chevalier-Boisvert et al., 2018) The task, illustrated in the bottom left pane of Fig 2, is to make a right turn around a corner to reach a fixed goal. The agent observes its current x, y coordinates, current velocity, and steering angle. The reward at each time step is the negative Manhattan distance from the goal, with a completion bonus of 100 if it successfully reaches the goal. The agent controls the steering angle and velocity. We do not permit the car to drive off the road, and define catastrophic states as those where the car touches the road boundaries. Cars of different widths have different catastrophic states. Tight turns close to the corner are risky, but can yield high rewards if executed correctly. During pretraining, the car width is sampled uniformly from the range [0.05, 0.1] before each training episode. We evaluate adaptation to widths 0.075, 0.1, 0.125, 0.15, 0.175, and 0.20.
- **Hand manipulation:** The robotic manipulation environment, originally presented in PDDM (Nagabandi et al.,

2019), involves controlling a 24-DoF hand for a Baoding ball rotation task, where the objective is to rotate two Baoding balls in the palm for a fixed time horizon without dropping either ball (see Fig 2, bottom right). A large negative reward is obtained upon dropping either ball. More details are presented in (Nagabandi et al., 2019), but we modify this environment by sampling the ball size, shared by both balls, from the range $[0.026, 0.027]$ for each pretraining episode. We evaluate adaptation to ball sizes 0.0255, 0.0265, 0.0275, 0.0280, 0.0285, and 0.029. PETS cannot solve this task, so we build CARL upon PDDM (Nagabandi et al., 2019) instead. PDDM differs from PETS in various ways, such as utilizing deterministic dynamics models and a stronger MPC optimization algorithm.

Baselines. Our first baseline is metalearning with MAML (Finn et al., 2017) applied to a PPO-trained RL agent (Schulman et al., 2017) (PPO-MAML) in both cartpole and half-cheetah. Metalearning is performed in the SCA pretraining phase. In our experiments, PPO-MAML failed to train stably on Duckietown car driving and Baoding ball manipulation, so we omit those results. To allow PPO-MAML to train stably, we allowed it 1500x the number of pretraining episodes as CARL.

Our second baseline is robust adversarial reinforcement learning (RARL) (Pinto et al., 2017). Since we found that this approach also required more data to train well, we compare against RARL at 2x and 20x the number of pretraining episodes as CARL. RARL works by training a model-free RL agent jointly with an adversary that perturbs the actions emitted by the agent. We train the adversary to apply 2D destabilizing forces on the pendulum in cartpole and 2D forces on the torso and feet of the cheetah in half-cheetah in exactly the same way as was originally done by the authors. We train the adversary to perturb the motor torques in Duckietown. It was unable to learn to manipulate the balls in the Baoding task, so we omit RARL from that comparison.

Third, we compare against a simple model-based adaptation approach, MB + Finetune. For this, we pretrain PETS on the same set of pretraining environments as CARL and then adapt it in the target environment by finetuning the dynamics model, while continuing to select reward-optimal actions as in standard PETS. In the Baoding ball task, we replace PETS with PDDM (Nagabandi et al., 2019) as PETS fails during pretraining.

We also compare two variants of CARL corresponding to the two notions of risk-aversion described in Sec 4.2. In all results, we call the low-reward risk-aversion variant as CARL (Reward), and the catastrophic state risk-aversion variant simply as CARL (State). CARL (State) and CARL (Reward) are built upon PDDM in the Baoding manipulation task, and upon PETS for all other environ-

ments. See Appendix section B for more experiment details.

Performance Metrics. SCA approaches aim to adapt safely to the target environment, as quickly as possible, and achieve as high rewards as possible. To evaluate this, for each method, we report the “average maximum reward” over adaptation time t , which is the average over 10 random seeds of the maximum over t adaptation episodes of the reward obtained in the target environment. In CartPole, Duckietown, and Half-Cheetah, $t = 10$ adaptation episodes, where each adaptation episode corresponds to one environment rollout. In Baoding, $t = 15$ adaptation episodes where each episode contains 35 environment rollouts. (Baoding policies are pre-trained on 100 episodes of 35 environment rollouts each). Finally, to measure the safety of the adaptation process, we also report the cumulative number of catastrophes suffered by each method during adaptation, which more directly measures the extent to which different methods avoid catastrophic failures and perform safe adaptation. We keep the number of pretraining episodes fixed for all methods unless otherwise mentioned. Because SCA assumes sandbox source environments and safety-critical target environments, catastrophic events are allowed to occur at detriment to the agent’s reward during the pretraining phase for our first three environments, but during adaptation these events will terminate the episode. In the Baoding environment, we also terminate episodes at pretraining time when the ball drops to avoid the discontinuous restoration of the balls to the hand.

5.2. Results

We now analyze the performance of CARL and the baselines, qualitatively and quantitatively, in all four safety-critical adaptation settings.

Epistemic Uncertainty Visualization. The first question we ask is: *how well does the CARL ensemble model epistemic uncertainties about state and catastrophe predictions over the course of pretraining?* To visualize this, we plot the model’s predictions for the same action sequences throughout the course of training for the Duckietown environment in Fig 3. On top of these trajectories predicted by the model, we overlay curves corresponding to ground truth trajectories of cars of various widths, executing the same action sequence. As pretraining proceeds in Fig 3, the predicted trajectory distribution fits these ground truth trajectories more and more closely, indicating that the model captures the epistemic uncertainty due to varying car widths. Furthermore, the catastrophe predictor learns the shape of the state space correctly, assigning increasing catastrophe probabilities to states as they get closer to the true catastrophe states at the road boundaries.




Figure 3. Trajectory and catastrophe prediction generated from a CARL model for a fixed action sequence, at various stages of pretraining, progressing from left to right. Blue-yellow curves in the background are predicted trajectories for a fixed action sequence, and the thick curves in the foreground correspond to ground truth trajectories of different width cars when executing that action sequence. Over pretraining, the predicted trajectories converge well to the distribution of ground truth car behaviors. Predicted states are colored by the catastrophe probabilities (see color bar).

Adaptation Rewards and Catastrophes. Given that CARL can capture epistemic uncertainty about state transitions and catastrophic states, we now present our main quantitative results, seeking to answer, *how effective is CARL at reaching high task rewards with minimal catastrophic events during training?* Fig 4 demonstrates the average maximum reward after all adaptation episodes and the average total number of catastrophes for each method in CartPole, Duckietown, and Baoding.

Across CartPole and Duckietown, CARL (State) performs slightly worse in-distribution, but attains relatively higher rewards as the testing environment moves further away from the training environments where caution becomes important. In Duckietown, CARL (State) generally avoids collisions better than all other methods as the car width is increased. This leads to the highest rewards relative to other methods at the largest car width, 0.2. Even in CartPole, in which nearly all methods suffered catastrophe during every adaptation step at pole heights of 1.5 and greater, CARL (State) prevents catastrophe for longer, leading to higher rewards. Finally, in Baoding, CARL (State) achieves higher rewards across 5 out of 6 target ball sizes than MB + Finetune, and drops the ball less than MB + Finetune on all target ball sizes. CARL (Reward) performs the worst among the three methods here, matching performance with MB + Finetune in-distribution but obtaining lower rewards out-of-distribution and dropping the ball more across all ball sizes. However, CARL (State)’s ability to perform well when applied to PDDM instead of PETS demonstrates the versatility of our catastrophic state risk aversion approach.

MB + Finetune achieves high reward with few catastrophes on in-distribution environments, however its reward drops off drastically at out-of-distribution CartPole lengths. Its rewards also decrease while collisions increase in Duckietown as the car width increases. RARL: 2x Itr generally performs poorly in CartPole and Duckietown. RARL: 20x Itr’s performance is competitive with both CARL variants, however it is still outperformed at the environments farthest away from the training distribution despite being

pretrained on 20x the number of trajectories. PPO-MAML is omitted in our Duckietown results, and both PPO-MAML and RARL are omitted in our Baoding results, because they did not train stably. On CartPole, PPO-MAML performs poorly, achieving low reward and reaching catastrophic states nearly every episode across in-distribution and out-of-distribution environments, indicating that it may require many more adaptation episodes.

In Half-Cheetah, our experiments evaluate only one out-of-distribution setting: adapting to running with a disabled front foot. We report the results after 10 adaptation episodes here (also see Fig 5): CARL (State) performs the best by obtaining rewards around 1400 and reaching the least catastrophic states across all methods. MB + Finetune and RARL: 20x Itr reach rewards nearing 1200, however RARL: 20x encounters catastrophe every iteration. CARL (Reward) reaches rewards of 750, while RARL: 2x Itr and PPO-MAML perform poorly. Most of these results match CartPole and Duckietown, with the exception of the much poorer performance of CARL (Reward). We see that in Half-Cheetah, reward-based risk aversion is insufficient to completely prevent reaching low-value states such as the head colliding with the ground, and state-based aversion is required.

Adaptation speed. Next, we ask, *how does cautious action selection help with adaptation speed?* We plot results over adaptation time in each environment for both the average maximum reward and the running total catastrophe count to show adaptation speed in the target environments. Fig 5 shows the average maximum reward over adaptation time for all methods in the farthest out-of-distribution environments of CartPole, Duckietown, and Baoding. It also demonstrates the reward over time on the one target environment of Half-Cheetah. In CartPole (top row), risk-aversion allows both CARL methods to immediately start with higher returns on the first iteration. As adaptation progresses their lead over other methods continues and it results in both achieving max returns above 70. All other methods are unable to significantly improve within 10 adaptation steps.

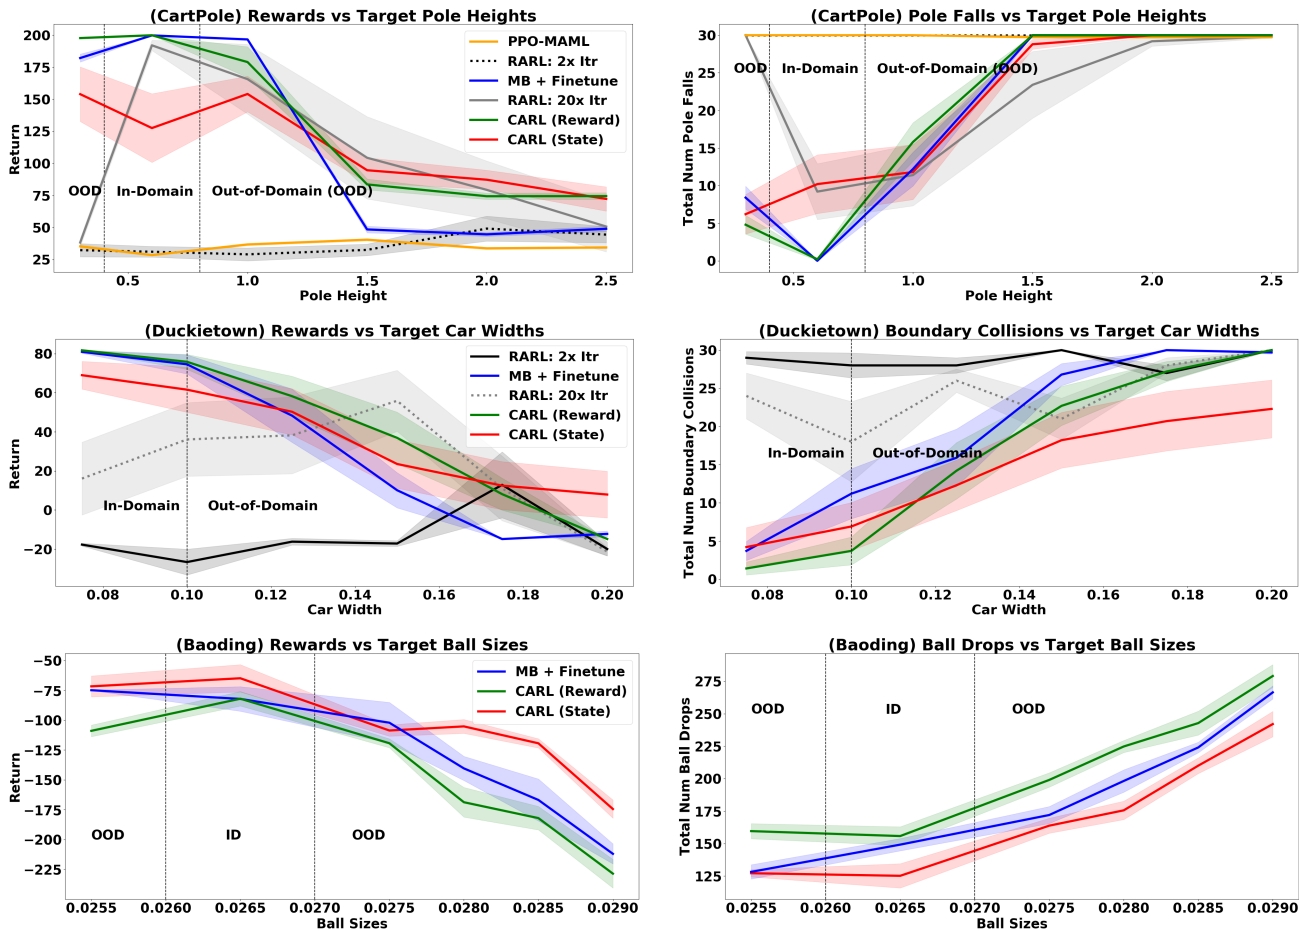


Figure 4. Evaluation of the final average maximum reward and average total catastrophic events at various target environments in the CartPole, Duckietown, and Baoding environments.

Even though all methods experience catastrophe at nearly every iteration, the CARL approaches both obtain higher returns by delaying catastrophe.

For the Half-Cheetah task (second row), RARL: 20x Itr, MB + Finetune, and CARL (State) obtain similar rewards upon starting. However, halfway through adaptation, CARL (State) achieves and maintains the best performance afterwards. CARL (Reward) starts off with much lower rewards, but eventually improves significantly. MB + Finetune and RARL: 20x Itr are also able to improve, however RARL: 2x Itr and PPO-MAML fail to adapt to the target environment.

CARL (State) performs much better than the other methods in Duckietown (third row). Catastrophic state risk-aversion allows it to achieve higher returns from the start, and continue to adapt on the test environment. While all other methods collide with the border nearly every iteration, CARL (State) collides significantly less, and even re-

duces the collision rate throughout the course of adaptation. Interestingly, RARL: 2x Itr is also able to increase its rewards during adaptation, whereas every other method fails to significantly improve its performance. We note that in the driving task, positive rewards are received only upon reaching the goal. Thus CARL (State) is the only method able to successfully make the turn in this extremely out-of-distribution task.

Finally, in the Baoding task, CARL (State) is able to achieve substantially higher reward from the first adaptation episode, and it continues to adapt and achieve higher reward than both other methods throughout the course of adaptation. This is mirrored in the catastrophe plot, in which CARL (State) consistently drops the ball the least number of times during adaptation. Here, CARL (Reward) and MB + Finetune perform similarly.

Cautious Behavior Finally, given the benefits of risk-aversion, we ask, *what qualitative behaviors does risk aver-*

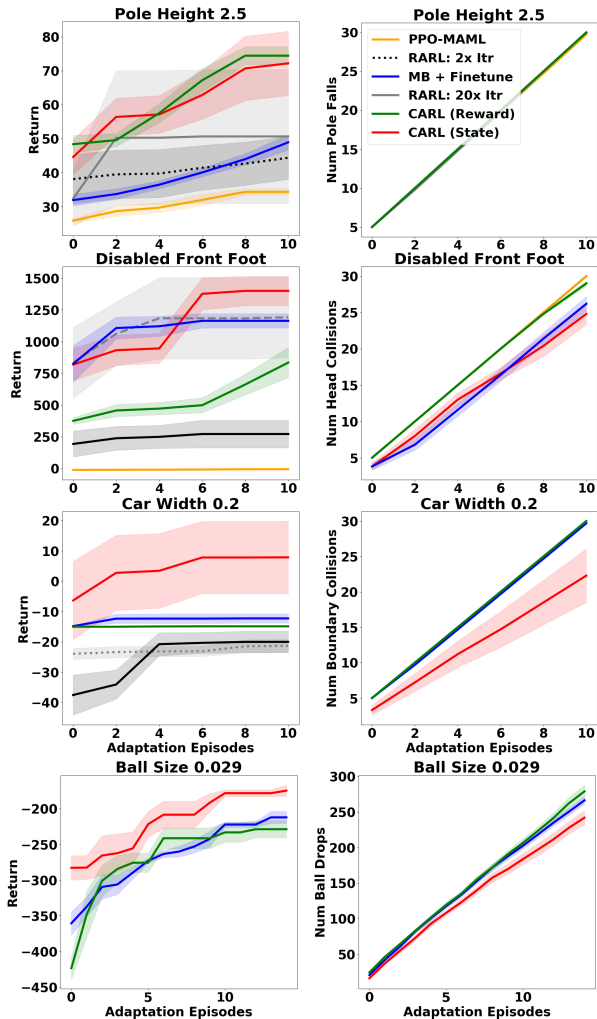


Figure 5. Average maximum reward and average total catastrophes over time during adaptation. *Top*: CartPole at the largest pole height, 2.5. *Second*: Half-Cheetah with the front foot disabled. *Third*: Duckietown at the largest car width, 0.2. *Bottom*: Baoding at the largest ball size, 0.029.

...ion produce? To answer this, Fig 6 presents state trajectories produced using planning with the same dynamics model, with and without the risk aversion of CARL (State). We note that for all visualizations, our dynamics models’ weights are exactly the same and we perform a frame-by-frame comparison starting from the same timestep in the rollouts; the only difference is the risk-aversion. In three environments, we see that CARL (State) produces distinct, interpretably risk-averse behaviors. In CartPole, CARL (State) is more willing to pay action costs in order to keep the pendulum upright. This behavior likely comes from how a lower pole can lead to states with higher catastrophe probability. When viewing rollouts, we notice that this behavior results in the cart using much more of the rail to do so, trading off the risk of the pole falling for the




Figure 6. A comparison of standard vs. risk-averse behavior, under the same out-of-distribution conditions and with the same dynamics model weights in three settings. In each setting, the top row contains frames from an episode of standard model-based planning (arranged left-to-right) while the bottom row shows CARL’s catastrophic state risk-averse planning. The specific environments are: *(Top)* CartPole with pole length 1.0. *Middle*: Duckietown with car width 0.15. *Bottom*: Half-Cheetah with the front foot disabled.

risk of falling off the rail. In Duckietown, the risk-averse agent drives more slowly and takes a wider turn, resulting in it being able to fully complete the turn while the standard agent clips the corner. We even notice risk-averse behaviors such as making an extremely wide turn, then backing up and completing the turn to reach the goal. The Half-Cheetah comparison highlights a key behavior that allows CARL (State) to stay upright for longer. When the disabled front foot is too far back, it is still able to recover from the precarious situation by angling the front leg joints differently.

6. Discussion

To address the difficulty of applying reinforcement learning in safety-critical situations such as robotics, we propose a general safety-critical adaptation task setting that transfers knowledge gained in less safety-critical source environments to enable safe learning in the target environment. Further, we have proposed CARL, a new method for model-based reinforcement learning approaches that works by implementing risk-averse adaptation strategies within the safety-critical adaptation framework. Our experiments demonstrate that CARL consistently outperforms strong adaptation baselines including metalearning and robust RL on four control tasks, including a difficult object manipulation task where CARL is applied to a new model-based RL algorithm. CARL achieves higher rewards, more quickly, while suffering fewer catastrophic events during learning.

References

- Aswani, A., Gonzalez, H., Sastry, S. S., and Tomlin, C. Provably safe and robust learning-based model predictive control. *Automatica*, 49(5):1216–1226, 2013.
- Berkenkamp, F., Turchetta, M., Schoellig, A. P., and Krause, A. Safe model-based reinforcement learning with stability guarantees, 2017.
- Borkar, V. and Jain, R. Risk-constrained markov decision processes. In *49th IEEE Conference on Decision and Control (CDC)*, pp. 2664–2669. IEEE, 2010.
- Chevalier-Boisvert, M., Golemo, F., Cao, Y., Mehta, B., and Paull, L. Duckietown environments for openai gym. <https://github.com/duckietown/gym-duckietown>, 2018.
- Chow, Y. and Ghavamzadeh, M. Algorithms for cvar optimization in mdps. In *Advances in neural information processing systems*, pp. 3509–3517, 2014.
- Chow, Y., Tamar, A., Mannor, S., and Pavone, M. Risk-sensitive and robust decision-making: a cvar optimization approach. In *Advances in Neural Information Processing Systems*, pp. 1522–1530, 2015.
- Chua, K., Calandra, R., McAllister, R., and Levine, S. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Advances in Neural Information Processing Systems*, pp. 4754–4765, 2018.
- Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks, 2017.
- Fisac, J. F., Akametalu, A. K., Zeilinger, M. N., Kaynama, S., Gillula, J., and Tomlin, C. J. A general safety framework for learning-based control in uncertain robotic systems, 2017.
- Ghavamzadeh, M., Mannor, S., Pineau, J., Tamar, A., et al. Bayesian reinforcement learning: A survey. *Foundations and Trends® in Machine Learning*, 8(5-6):359–483, 2015.
- Hakobyan, A., Kim, G. C., and Yang, I. Risk-aware motion planning and control using cvar-constrained optimization. *IEEE Robotics and Automation Letters*, 4(4):3924–3931, 2019.
- Hanssen, K. G. and Foss, B. Scenario-based mpc with gradual relaxation of output constraints. In *2015 54th IEEE Conference on Decision and Control (CDC)*, pp. 6530–6534, 2015.
- Henaff, M., Canziani, A., and LeCun, Y. Model-predictive policy learning with uncertainty regularization for driving in dense traffic. *arXiv preprint arXiv:1901.02705*, 2019.
- Hewing, L., Kabzan, J., and Zeilinger, M. N. Cautious model predictive control using gaussian process regression. *IEEE Transactions on Control Systems Technology*, pp. 1–8, 2019. ISSN 2374-0159. doi: 10.1109/tcst.2019.2949757. URL <http://dx.doi.org/10.1109/TCST.2019.2949757>.
- Kahn, G., Villafior, A., Pong, V., Abbeel, P., and Levine, S. Uncertainty-aware reinforcement learning for collision avoidance. *arXiv preprint arXiv:1702.01182*, 2017.
- Lecarpentier, E. and Rachelson, E. Non-stationary markov decision processes, a worst-case approach using model-based reinforcement learning, extended version, 2019.
- Morimura, T., Sugiyama, M., Kashima, H., Hachiya, H., and Tanaka, T. Nonparametric return distribution approximation for reinforcement learning. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pp. 799–806, 2010.
- Nagabandi, A., Clavera, I., Liu, S., Fearing, R. S., Abbeel, P., Levine, S., and Finn, C. Learning to adapt in dynamic, real-world environments through meta-reinforcement learning. *arXiv preprint arXiv:1803.11347*, 2018.
- Nagabandi, A., Konoglie, K., Levine, S., and Kumar, V. Deep dynamics models for learning dexterous manipulation, 2019.
- Ostafew, C. J., Schoellig, A. P., and Barfoot, T. D. Robust constrained learning-based nmpc enabling reliable mobile robot path tracking. *The International Journal of Robotics Research*, 35(13):1547–1563, 2016.
- Peng, X. B., Andrychowicz, M., Zaremba, W., and Abbeel, P. Sim-to-real transfer of robotic control with dynamics randomization. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1–8, May 2018. doi: 10.1109/ICRA.2018.8460528.
- Pinto, L., Davidson, J., Sukthankar, R., and Gupta, A. Robust adversarial reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning—Volume 70*, pp. 2817–2826. JMLR. org, 2017.
- Rajeswaran, A., Ghotra, S., Ravindran, B., and Levine, S. Epopt: Learning robust neural network policies using model ensembles. *arXiv preprint arXiv:1610.01283*, 2016.
- Rockafellar, R. T., Uryasev, S., et al. Optimization of conditional value-at-risk. *Journal of risk*, 2:21–42, 2000.
- Sadeghi, F. and Levine, S. Cad2rl: Real single-image flight without a single real image. In *RSS*, 2017.
- Sadigh, D. and Kapoor, A. Safe control under uncertainty with probabilistic signal temporal logic. 2016.

- Sæmundsson, S., Hofmann, K., and Deisenroth, M. P. Meta reinforcement learning with latent variable gaussian processes. *arXiv preprint arXiv:1803.07551*, 2018.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms, 2017.
- Strens, M. A bayesian framework for reinforcement learning. In *ICML*, volume 2000, pp. 943–950, 2000.
- Tamar, A., Glassner, Y., and Mannor, S. Optimizing the cvar via sampling. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
- Tobin, J., Fong, R., Ray, A., Schneider, J., Zaremba, W., and Abbeel, P. Domain randomization for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 23–30. IEEE, 2017.
- Yu, P., Lee, J. S., Kulyatin, I., Shi, Z., and Dasgupta, S. Model-based deep reinforcement learning for dynamic portfolio optimization, 2019.