

---

# Density Constrained Reinforcement Learning

---

Zengyi Qin<sup>1</sup> Yuxiao Chen<sup>2</sup> Chuchu Fan<sup>1</sup>

## Abstract

We study constrained reinforcement learning (CRL) from a novel perspective by setting constraints directly on *state density functions*, rather than the value functions considered by previous works. State density has a clear physical and mathematical interpretation, and is able to express a wide variety of constraints such as resource limits and safety requirements. Density constraints can also avoid the time-consuming process of designing and tuning cost functions required by value function-based constraints to encode system specifications. We leverage the duality between density functions and Q functions to develop an effective algorithm to solve the density constrained RL problem optimally and the constraints are guaranteed to be satisfied. We prove that the proposed algorithm converges to a near-optimal solution with a bounded error even when the policy update is imperfect. We use a set of comprehensive experiments to demonstrate the advantages of our approach over state-of-the-art CRL methods, with a wide range of density constrained tasks as well as standard CRL benchmarks such as Safety-Gym.

## 1. Introduction

Constrained reinforcement learning (CRL) (Altman, 1999; Achiam et al., 2017; Dalal et al., 2018; Paternain et al., 2019; Tessler et al., 2019; Yang et al., 2020; Zhang et al., 2020; Stooke et al., 2020) aims to find the optimal policy that maximizes the cumulative reward signal while respecting certain constraints such as safety requirements and resource limits. Existing CRL approaches typically involve constructing suitable cost functions and value functions to take into account the constraints. Then a crucial step is to choose appropriate parameters such as thresholds for the cost and value functions to encode the constraints. However, one

---

<sup>1</sup>Massachusetts Institute of Technology <sup>2</sup>California Institute of Technology. Correspondence to: Chuchu Fan <chuchu@mit.edu>.




Figure 1. Autonomous electric vehicle routing in Manhattan: Control the electric vehicles to reach the goals and always keep the remaining energy above a threshold. Vehicles can be charged at the charging stations. The *vehicle density* at each charging station is constrained by an upper bound due to resource limits. The roads and charging stations are from the real-world data (Blahoudek et al., 2020).

significant gap between the use of such methods and solving CRL problems is the correct construction of the cost and value functions, which is typically not solved systematically but relies on engineering intuitions (Paternain et al., 2019). Simple cost functions may not exhibit satisfactory performance, while sophisticated cost functions may not have clear physical meanings. When cost functions lack clear physical interpretations, it is difficult to formally guarantee the satisfaction of the performance specifications, even if the constraints on the cost functions are fulfilled. Moreover, different environments generally need different cost functions, which makes the tuning process time-consuming.

In this work, we study CRL from a novel perspective by imposing constraints on *state density functions*, which avoids the use of cost or value function-based constraints in previous CRL literature. Density is a measurement of state concentration in the state space, and is directly related to the state distribution. A variety of real-world constraints are naturally expressed as density constraints in the state space. For example, safety constraints is a special case of density

constraints where we require the entire density distribution of states being contained in safe regions. Resource limits can also be encoded by imposing maximum allowed density on certain states. For instance, consider the electric vehicle routing problem (Blahoudek et al., 2020) in Figure 1. In order to keep the remaining energy above a predefined threshold, vehicles can be charged at the charging stations. Due to resource limits, the *vehicle density* at each charging station is constrained by an upper bound. In our experiments, we will demonstrate that density constraints can be used in more general examples, which also include the standard value function-based constraints as special cases.

We cannot directly apply existing CRL optimization methods based on constraining cost or value functions to solve density constrained reinforcement learning (DCRL) problems, because the resulting solution cannot guarantee the satisfaction of required density constraints. Instead, we propose a novel general model-free algorithm (Algorithm 1) that exploits the duality between density functions and Q functions. Such algorithm enables us to incorporate state density constraints straightforwardly in CRL. Our proposed algorithm is applicable of handling both discrete and continuous state and action spaces, and can be flexibly combined with off-the-shelf RL methods to update the policy. We also prove that our algorithm can converge to a near-optimal solution with a bounded error even when the RL-based policy update is imperfect (Theorem 2).

Comprehensive experiments are conducted on various density constrained problems, which demonstrate the advantages of our approach over state-of-the-art CRL methods in terms of satisfying system specifications and improving the cumulative reward. A notable achievement is that although our method can handle much more general constraints than safety constraints, our DCRL algorithm still shows consistent improvement over existing CRL approaches on standard CRL benchmarks from MuJoCo and Safety-Gym (Ray et al., 2019) where constraints are primarily about safety.

**Our main contributions are:** 1) We are the first to introduce the DCRL problem with density constraints as a promising way to encode system specifications, which differs from the value function-based constraints in existing CRL literature. 2) We propose a general algorithm for solving DCRL problems by leveraging the duality between density functions and Q functions, and prove that our algorithm is guaranteed to converge to a near-optimal solution with a bounded error even when the policy update is imperfect. 3) We use an extensive set of experiments to examine the advantages of our method over leading CRL approaches, in a wide variety of density constrained tasks as well as standard CRL benchmarks.

## 2. Related Work

Constrained reinforcement learning (Garcia & Fernández, 2015) primarily focuses on two approaches: modifying the optimality criteria by combining a risk factor (Heger, 1994; Nilim & El Ghaoui, 2005; Howard & Matheson, 1972; Borkar, 2002; Basu et al., 2008; Sato et al., 2001; Dotan Di Castro & Mannor, 2012; Kadota et al., 2006; Lötjens et al., 2019) and incorporating extra knowledge to the exploration process (Moldovan & Abbeel, 2012; Abbeel et al., 2010; Tang et al., 2010; Geramifard et al., 2013; Clouse & Utgoff, 1992; Thomaz et al., 2006; Chow et al., 2018). Our method falls into the first category by imposing constraints and is closely related to constrained Markov decision processes (Altman, 1999) (CMDPs). CMDPs has been extensively studied in robotics (Gu et al., 2017; Pham et al., 2018), game theory (Altman & Shwartz, 2000), and communication and networks (Hou & Zhao, 2017; Bovopoulos & Lazar, 1992). Most previous works consider the constraints on value functions, cost functions and reward functions (Altman, 1999; Paternain et al., 2019; Altman & Shwartz, 2000; Dalal et al., 2018; Achiam et al., 2017; Ding et al., 2020). Instead, we directly impose constraints on the state density functions. Chen et al. (2019) and Chen et al. (2019) study the duality between density functions and value functions in CMDPs when the full model dynamics are known, while we consider the model-free reinforcement learning and take a further step to prove the duality of density functions to Q functions. In Geibel and Wysotzki (2005) density was studied as the probability of entering error states and thus has fundamentally different physical interpretations from us. In Dai et al. (2017) the duality was used to boost the actor-critic algorithm. The duality is also used in the policy evaluation community (Nachum et al., 2019; Nachum & Dai, 2020; Tang et al., 2019). The offline policy evaluation method proposed by Nachum et al. (2019) can also be used to estimate the state density in our paper, but their focus is policy evaluation rather than constrained RL. Therefore, we claim that this paper is the first work to consider density constraints and use the duality property to solve CRL.

## 3. Preliminaries

**Markov Decision Processes (MDP).** An MDP  $\mathcal{M}$  is a tuple  $\langle S, A, P, R, \gamma \rangle$ , where (1)  $S$  is the (possibly infinite) set of states; (2)  $A$  is the (possibly infinite) set of actions; (3)  $P : S \times A \times S \mapsto [0, 1]$  is the transition probability with  $P(s, a, s')$  the probability of transitioning from state  $s$  to  $s'$  when action  $a \in A$  is taken; (4)  $R : S \times A \times S \mapsto \mathbb{R}$  is the reward associated with the transition  $P$  under the action  $a \in A$ ; (5)  $\gamma \in [0, 1]$  is a discount factor.

A policy  $\pi$  maps states to a probability distribution over actions where  $\pi(a|s)$  denotes the probability of choosing action  $a$  at state  $s$ . Let a function  $\phi : S \mapsto \mathbb{R}$

specify the initial state distribution. The objective of an MDP optimization is to find the optimal policy that maximizes the overall discounted reward  $J_p = \int_S \phi(s)V^\pi(s)ds$ , where  $V^\pi(s)$  is called the value function and satisfies  $V^\pi(s) = r^\pi(s) + \gamma \int_A \pi(a|s) \int_S P(s, a, s')V^\pi(s')ds'da$ , and  $r^\pi(s) = \int_A \pi(a|s) \int_S P(s, a, s')R(s, a, s')ds'da$  is the one-step reward from state  $s$  following policy  $\pi$ . For every state  $s$  with occurring as an initial state with probability  $\phi(s)$ , it incurs a expected cumulative discounted reward of  $V^\pi(s)$ . Therefore the overall reward is  $\int_S \phi(s)V^\pi(s)ds$ . The formulation of value functions in MDP typically cannot handle constraints on state distribution, which motivates the density functions.

**Density Functions.** State density functions  $\rho : S \mapsto \mathbb{R}_{\geq 0}$  measure the state concentration in the state space (Chen & Ames, 2019; Rantzer, 2001).<sup>1</sup> We will show later that generic duality relationship exists between density functions and Q functions, which allows us to directly impose density constraints in RL problems. For infinite horizon MDPs, given a policy  $\pi$  and an initial state distribution  $\phi$ , the stationary density of state  $s$  is expressed as:

$$\rho^\pi(s) = \sum_{t=0}^{\infty} \gamma^t P(s^t = s | \pi, s^0 \sim \phi),$$

which is the discounted sum of the probability of visiting  $s$ . We prove in the supplementary material that the density has an equivalent expression:

$$\rho^\pi(s) = \phi(s) + \gamma \int_S \int_A \pi(a|s')P(s', a, s)\rho^\pi(s')dad s',$$

## 4. Density Constrained Reinforcement Learning

In this section, we first formally define the novel DCRL problem and clarify its benefits compared to value function-based constraints. To solve DCRL, we will establish the duality between Q functions and density functions, then propose a general algorithm for DCRL with convergence guarantees.

### 4.1. Problem Statement

Given an MDP  $\mathcal{M} = \langle S, A, P, R, \gamma \rangle$  and an initial state distribution  $\phi$ , DCRL finds the optimal policy  $\pi^*$  to the following optimization problem:

$$\begin{aligned} \max \int_S \phi(s)V^\pi(s)ds \\ \text{s.t. } \rho_{min}(s) \leq \rho^\pi(s) \leq \rho_{max}(s), \forall s \in S, \end{aligned} \quad (1)$$

<sup>1</sup> $\rho$  is not necessarily a probability density function, which means  $\int_S \rho(s) = 1$  is not enforced.

where  $\int_S \phi(s)V^\pi(s)ds$  is the expected cumulative reward. The density constraints  $\rho_{min}(s)$  and  $\rho_{max}(s)$  are functions of states, and different states can have different lower and upper bound constraints on their densities.

The formulation in (1) is different from most of the previous work (Achiam et al., 2017; Tessler et al., 2019; Yang et al., 2020), which did not consider the density constraint but instead used the cost (or reward) value as constraint:

$$\begin{aligned} \max \int_S \phi(s)V^\pi(s)ds \\ \text{s.t. } \int_S \phi(s)V_C^\pi(s)ds \leq \eta, \end{aligned} \quad (2)$$

where  $V_C^\pi$  is the cost value function and  $\eta$  is the threshold for the expected cumulative cost.

**Benefits of the Density Constraint.** The density constraint in (1) can bring at least two benefits. First, density has a clear physical and mathematical interpretation as a measurement of state concentration in the state space. A wide range of real-world constraints can be conveniently expressed as density constraints (e.g., the vehicle density in certain areas, and the frequency of agents entering undesirable states). Second, the value function constraint in (2) requires the time-consuming process of designing and tuning cost functions, which are completely avoided by the density constraint since no cost function tuning is needed.

### 4.2. Duality of Density Functions and Q Functions

To solve the DCRL problem in (1), we cannot directly apply existing RL algorithms (Lillicrap et al., 2016; Schulman et al., 2015; 2017), which are designed to optimize the expected cumulative reward or cost, but not the state density function. However, we will show that density functions are actually dual to Q functions, and the density constraints can be realized by modifying the formulation of Q functions. Then the off-the-shelf RL algorithms can be used to optimize the modified Q functions to enforce the density constraints.

We extend the stationary density  $\rho^\pi$  to consider the action taken at each state. Let  $\bar{\rho} : S \times A \rightarrow \mathbb{R}_{\geq 0}$  be a stationary state-action density function, which represents the amount of state instances taking action  $a$  at state  $s$ .  $\bar{\rho}$  is related to  $\rho$  via marginalization:  $\rho(s) = \int_A \bar{\rho}(s, a)da$ . Under a policy  $\pi$ , we also have  $\bar{\rho}^\pi(s, a) = \rho^\pi(s)\pi(a|s)$ . Let  $r(s, a) = \int_S P(s, a, s')R(s, a, s')ds'$ . Consider the density function optimization:

$$\begin{aligned} J_d = \max_{\bar{\rho}, \pi} \int_S \int_A \bar{\rho}^\pi(s, a)r(s, a)dads \\ \text{s.t. } \bar{\rho}^\pi(s, a) = \pi(a|s) \left( \phi(s) + \right. \\ \left. \gamma \int_S \int_A P(s', a', s)\bar{\rho}^\pi(s', a')da'ds' \right) \\ \rho_{min}(s) \leq \rho^\pi(s) \leq \rho_{max}(s) \end{aligned} \quad (3)$$

**Lemma 1.** *The optimization problems in (1) and (3) share the same optimal policy  $\pi^*$ .*

The proof of Lemma 1 is provided in the supplementary material. It shows that solving (1) is equivalent to solving (3). Now we are ready to present the duality between density functions and Q functions, which reveals that (3) can be solved via standard RL algorithms by modifying the Q function. We use the Lagrangian method for (3) and denote the Lagrange multipliers for  $\rho^\pi \geq \rho_{min}$  and  $\rho^\pi \leq \rho_{max}$  as  $\sigma_- : S \mapsto \mathbb{R}_{\geq 0}$  and  $\sigma_+ : S \mapsto \mathbb{R}_{\geq 0}$ . Then we consider the following optimization problem:

$$\begin{aligned} J_p &= \max_{Q, \pi} \int_S \phi(s) \int_A Q^\pi(s, a) \pi(a|s) da ds \\ \text{s.t. } Q^\pi(s, a) &= r(s, a) + \sigma_-(s) - \sigma_+(s) + \\ &\quad \gamma \int_S P(s, a, s') \int_A \pi(a'|s') Q^\pi(s', a') da' ds' \end{aligned} \quad (4)$$

The difference between the Q function in (4) and the standard Q function is that the reward  $r(s, a)$  is modified to  $r(s, a) + \sigma_-(s) - \sigma_+(s)$ .

**Theorem 1.** *The density constrained optimization objectives  $J_d$  in (3) and  $J_p$  in (4) are dual to each other. If both are feasible and the KKT conditions are satisfied, then they share the same optimal policy  $\pi^*$ .*

The proof of Theorem 1 is provided in the appendix. Theorem 1 reveals that when the KKT conditions are satisfied, the optimal solution to the adjusted unconstrained primal problem (4) is exactly the optimal solution to the dual problem (3) with density constraints. Such an optimal solution not only satisfies the state density constraints, but also maximizes the total reward  $J_d$ . Thus it is the optimal solution to the DCRL problem. Note that (4) is solvable using standard RL algorithms (Lillicrap et al., 2016; Schulman et al., 2015; 2017). However, the update of policy  $\pi$  will also result in an update of state density function  $\rho^\pi$ , which will change the Lagrangian multipliers  $\sigma_+$  and  $\sigma_-$  in order to satisfy the KKT conditions of (3). This motivates us to develop an algorithm that iteratively update  $\pi, \rho^\pi, \sigma_+$  and  $\sigma_-$ , which will be detailed in Section 4.3.

**Remark 1.** While the duality between density functions and value functions has already been studied (Chen & Ames, 2019), we take a step further to show the duality property between the density functions and Q functions. We prove and use the duality between density functions and Q functions over continuous state space to solve density constrained RL from a novel perspective, which has not been explored and utilized by published RL literature.

**Remark 2.** The dual variables in Linear Programming are different from ours and do not have a clear physical interpretation. Technically, any non-negative dual variable satisfying the conservation law (Liouville PDE in the continuous

case, see Chen et al. (2019)) is a valid dual. However, among all valid dual variables, the state density is associated with a clear physical interpretation as the concentration of states, and we are able to directly apply constraints on the density in RL.

**Remark 3.** The value function-based constraint in CMDPs can be viewed as a special case of density constraint.  $\int_S \phi(s) V_C^\pi(s) ds \leq \eta$  is equivalent to  $\int_S \rho^\pi(s) r_c(s) ds \leq \eta$ . Thus the value function-based CRL problems can also be solved by the DCRL framework.

**Remark 4.** The general DCRL problem cannot be solved by value function-based RL methods. One may argue that for a single state  $s_k$ , its visiting frequency can be constrained by defining a cost function  $c_k$  that returns 1 when  $s_k$  is visited and 0 otherwise. Then value function-based methods can be used to constrain the expected value  $V_{C_k}$  of the cost, as is in  $\int_S \phi(s) V_{C_k}(s) ds \leq \eta_k$  (\*). However, when every state  $s_j$  has unique density constraints  $\rho_{min}(s_j)$  and  $\rho_{max}(s_j)$ , every state  $s_j$  needs a value function  $V_{C_j}$  (for the cost of visiting  $s_j$ ) and an inequality (\*). In order to cover all states, it would require an extremely large and even infinite number (when the state space is continuous) of value functions and inequalities, which can make the problem intractable for value function-based methods. Furthermore, we have explicitly mentioned in Remark 3 that the value function-based constraint in CMDP is equivalent to a special case of density constraint, and our general formulation of state-wise density constraint is more expressive and cannot be trivially converted into value function-based constraints.

### 4.3. The DCRL Algorithm

The density function optimization in continuous state space is an infinite dimensional linear program and there is no convenient approximation method. We thus turn to the equivalent Q function optimization where standard RL tools are applicable. By utilizing the duality between density function and Q function (Theorem 1) in the density constrained optimization, the DCRL problem can be solved by alternating between the primal problem (4) and dual problem (3), as is illustrated in Figure 2. In the primal domain, we solve the adjusted primal problem (reward adjusted by Lagrange multipliers) in (4) using off-the-shelf unconstrained RL methods such as TRPO (Schulman et al., 2015) and DDPG (Lillicrap et al., 2016). Note that the density constraints are enforced in dual domain and the primal domain is still an unconstrained problem, which means we can make use of existing RL methods to solve the primal problem. In the dual domain, the policy is used to evaluate the state density function, which is described in details in Section 4.5. If the KKT conditions  $\sigma_+ \cdot (\rho^\pi - \rho_{max}) = 0$ ,  $\sigma_- \cdot (\rho_{min} - \rho^\pi) = 0$  and  $\rho_{min} \leq \rho^\pi \leq \rho_{max}$  are not satisfied, the Lagrange multipliers are updated and enter the next loop. The key insight






Figure 2. Illustration of the iterative optimization in DCRL. In the primal domain, we solve the adjusted primal problem in (4) to obtain the policy  $\pi$ . Then in the dual domain, the  $\pi$  is used to evaluate the state density. The Lagrange multipliers  $\sigma_+$  and  $\sigma_-$  are updated as  $\sigma_+ \leftarrow \max(\mathbf{0}, \sigma_+ + \alpha(\rho^\pi - \rho_{max}))$  and  $\sigma_- \leftarrow \max(\mathbf{0}, \sigma_- + \alpha(\rho_{min} - \rho^\pi))$ . In the next loop, since the reward  $r(s, a) + \sigma_-(s) - \sigma_+(s)$  is updated, the primal optimization solves for the new  $\pi$  under the updated reward. The loop stops when the KKT conditions are satisfied.

is that the density constraints can be enforced in the dual problem, and we can solve the dual problem by solving the equivalent primal problem using existing algorithms. Alternating between the primal and dual optimization can gradually adjust the Lagrange multipliers until the KKT conditions are satisfied.

#### Algorithm 1 Template of the DCRL algorithm

- 1: **Input** MDP  $\mathcal{M}$ , initial condition distribution  $\phi$ , constraints on the state density  $\rho_{max}$  and  $\rho_{min}$
- 2: Initialize  $\pi$  randomly,  $\sigma_+ \leftarrow \mathbf{0}$ ,  $\sigma_- \leftarrow \mathbf{0}$
- 3: Generate experience  $D_\pi \subset \{(s, a, r, s') \mid s^0 \sim \phi, a \sim \pi(s) \text{ then } r \text{ and } s' \text{ are observed}\}$
- 4: **repeat**
- 5:  $(s, a, r, s') \leftarrow (s, a, r + \sigma_-(s) - \sigma_+(s), s')$  for each  $(s, a, r, s')$  in  $D_\pi$
- 6: Solve for  $\pi$  of (4) using  $D_\pi$  via standard RL
- 7: Generate experience  $D_\pi \subset \{(s, a, r, s') \mid s^0 \sim \phi, a \sim \pi(s) \text{ then } r \text{ and } s' \text{ are observed}\}$
- 8: Compute stationary density  $\rho^\pi$  using  $D_\pi$
- 9:  $\sigma_+ \leftarrow \max(\mathbf{0}, \sigma_+ + \alpha(\rho^\pi - \rho_{max}))$
- 10:  $\sigma_- \leftarrow \max(\mathbf{0}, \sigma_- + \alpha(\rho_{min} - \rho^\pi))$
- 11: **until**  $\sigma_+ \cdot (\rho^\pi - \rho_{max}) = 0$ ,  $\sigma_- \cdot (\rho_{min} - \rho^\pi) = 0$  **and**  $\rho_{min} \leq \rho^\pi \leq \rho_{max}$
- 12: **Return**  $\pi$ ,  $\rho^\pi$

A general template of the density constraint policy optimization is provided in Algorithm 1. In Algorithm 1, the Lagrange multipliers  $\sigma_+$  and  $\sigma_-$  are used to adjust rewards, which lead to an update of the policy  $\pi$ . Then the policy is used to evaluate the stationary density, and the Lagrange multipliers are updated following dual ascent. The iteration stops when all the KKT conditions are satisfied.

#### 4.4. Convergence Guarantee under Imperfect Updates

If every loop of Algorithm 1 gives the perfect (optimal) policy  $\pi$  of (4) under the current  $\sigma_-$  and  $\sigma_+$ , then the convergence can be derived from the subgradient methods (Boyd et al., 2003) and will not be detailed here. But in reality, the policy update in each loop can be imperfect, which is unavoidable for standard RL algorithms.

In light of this, we are interested in the case where the policy solved at each iteration is imperfect, and will provide conver-

gence guarantees of Algorithm 1. We will also identify how the error in policy updates propagates to the solution found by Algorithm 1 and derive an upper bound of its distance to the optimal solution. Denote the negative objective function of (3) as  $f(\rho^\pi) = -\int_S \int_A \rho^\pi(s) \pi(a|s) r(s, a) da ds$ . We assume  $f$  is strongly convex modulus  $\mu$ . If this is not the case, one can add a strongly convex regularization term to the (already convex) objective function. Then (3) minimizes  $f(\rho^\pi)$  subject to the density constraints. For brevity, we only consider the density upper bound in this convergence analysis, and the same result exists for the density lower bound. The Lagrangian dual function is formulated as

$$d(\sigma_+) = \min_{\rho^\pi} f(\rho^\pi) + \int_S \sigma_+(s) (\rho^\pi(s) - \rho_{max}(s)) ds \quad (5)$$

Denote the feasible set of Lagrangian multipliers as  $\mathcal{P}$ . The dual ascent is to find  $\max_{\sigma_+ \in \mathcal{P}} d(\sigma_+)$ . We assume the policy update is imperfect and the suboptimality of the solution  $\hat{\rho}^\pi$  to (3) is upper bounded as:

$$f(\hat{\rho}^\pi) + \int_S \sigma_+(s) (\hat{\rho}^\pi(s) - \rho_{max}(s)) ds - d(\sigma_+) \leq \epsilon$$

And the imperfect dual ascent takes the form  $\sigma_+ \leftarrow \max(\mathbf{0}, \sigma_+ + \alpha \nabla \hat{d}(\sigma_+))$ , where  $\nabla \hat{d}(\sigma_+) = \hat{\rho}^\pi - \rho_{max}$ . Let  $\hat{g}(\sigma_+, \alpha) = \frac{1}{\alpha} (\max(\mathbf{0}, \sigma_+ + \alpha \nabla \hat{d}(\sigma_+)) - \sigma_+)$ , which is the residual of  $\sigma_+$  before and after the imperfect update in a single loop of Algorithm 1.

**Lemma 2.** *Following the imperfect dual ascent with step size  $\alpha \leq \mu$ , we have*

$$d(\sigma_+^{k+1}) \geq d(\sigma_+^k) + \frac{\alpha}{2} \|\hat{g}(\sigma_+^k, \alpha)\|^2 - \sqrt{\frac{2\epsilon}{\mu}} \|\hat{g}(\sigma_+^k, \alpha)\|$$

The superscript  $k$  denotes the  $k^{\text{th}}$  loop of Algorithm 1. The proof of Lemma 2 is provided in the supplementary material. Let  $\mathcal{P}^* = \{\sigma_+ \mid d(\sigma_+) = \max_{\sigma_+ \in \mathcal{P}} d(\sigma_+)\}$  be the set of optimal solutions to (5).

**Theorem 2.** *There exists constants  $\lambda > 0$  and  $\xi > 0$  such that Algorithm 1 with imperfect policy updates converges to a dual solution  $\hat{\sigma}_+$  that satisfies*

$$\min_{\sigma_+ \in \mathcal{P}^*} \|\hat{\sigma}_+(s) - \sigma_+' \| \leq \lambda \sqrt{\frac{\epsilon}{\mu}} \quad (6)$$

The dual function also converges to a bounded neighbourhood of its optimal value:

$$\min_{\sigma'_+ \in \mathcal{P}^*} \|d(\hat{\sigma}_+) - d(\sigma'_+)\| \leq \xi \frac{\epsilon}{\mu^2} \quad (7)$$

The proof of Theorem 2 is provided in the supplementary material. Theorem 2 shows that Algorithm 1 is guaranteed to converge to a near-optimal solution even under imperfect policy updates.

#### 4.5. Computational Approaches

Algorithm 1 requires computing the policy  $\pi$ , density  $\rho^\pi$ , Lagrange multipliers  $\sigma_+$  and  $\sigma_-$ . For  $\pi$ , there are well-developed representations such as neural networks and tabular methods. Solving  $\pi$  from experience  $D_\pi$  is also straightforward via standard model-free RL. By contrast, the computation of  $\rho^\pi$ ,  $\sigma_+$  and  $\sigma_-$  need to be addressed.

**Density Functions.** In the discrete state case,  $\rho^\pi$  is represented by a vector where each element corresponds to a state. To compute  $\rho^\pi$  from experience  $D_\pi$  (line 1 of Algorithm 1), let  $D_\pi$  contain  $N$  episodes where episode  $i$  ends at time  $T_i$ . Let  $s^{ij}$  represent the state reached at time  $j$  in the  $i^{th}$  episode. Initialize  $\rho^\pi \leftarrow \mathbf{0}$ . For all  $i \in \{1, \dots, N\}$  and  $j \in \{0, 1, \dots, T_i\}$ , do the update  $\rho^\pi(s^{ij}) \leftarrow \rho^\pi(s^{ij}) + \frac{1}{N} \gamma^j$ . The resulting vector  $\rho^\pi$  approximates the stationary state density. In the continuous state space,  $\rho^\pi$  cannot be represented as a vector since there are infinitely many states. We utilize the kernel density estimation method (Chen, 2017; Chen & Ames, 2019) that computes  $\rho^\pi(s)$  at state  $s$  using the samples in  $D_\pi$  with  $\rho^\pi(s) = \frac{1}{N} \sum_{i=1}^N \sum_{j=0}^{T_i} \gamma^j K_h(s - s^{ij})$  where  $K_h$  is the kernel function satisfying  $\forall s \in \mathcal{S}, K_h(s) \geq 0$  and  $\int_{\mathcal{S}} K_h(s) ds = 1$ . There are multiple choices of the kernel  $K_h$ , e.g. Gaussian, Spheric, and Epanechnikov kernels (Chen, 2017), and probabilistic guarantee of accuracy can be derived (Wasserman, 2019).

**Lagrange Multipliers.** If the state space is discrete, both  $\sigma_+$  and  $\sigma_-$  are vectors whose length equals to the number of states. In each loop of Algorithm 1, after the stationary density is computed,  $\sigma_+$  and  $\sigma_-$  are updated following Line 1 and Line 1 respectively in Algorithm 1. If the state space is continuous, we construct Lagrange multiplier functions  $\sigma_+$  and  $\sigma_-$  from samples in the state space leveraging linear interpolation. Let  $\bar{s} = [s_1, s_2, \dots]$  represent the samples in the state space. In every loop of Algorithm 1, denote the Lagrange function computed by the previous loop as  $\sigma_+^o$  and  $\sigma_-^o$ . We compute the updated Lagrange multipliers at states  $\bar{s}$  as  $\bar{\sigma}_+ = [\max(0, \sigma_+^o(s_1) + \alpha(\rho^\pi(s_1) - \rho_{max}(s_1))), \max(0, \sigma_+^o(s_2) + \alpha(\rho^\pi(s_2) - \rho_{max}(s_2))), \dots]$  and  $\bar{\sigma}_- = [\max(0, \sigma_-^o(s_1) + \alpha(\rho_{min}(s_1) - \rho^\pi(s_1))), \max(0, \sigma_-^o(s_2) + \alpha(\rho_{min}(s_2) - \rho^\pi(s_2))), \dots]$ . Then the new  $\sigma_+$  and  $\sigma_-$  are obtained by

linearly interpolating  $\bar{\sigma}_+$  and  $\bar{\sigma}_-$  respectively.

## 5. Experiment

We consider a wide variety of CRL benchmarks and demonstrate how they can be effectively solve by the proposed DCRL approach. The density constrained benchmarks include autonomous electrical vehicle routing (Blahoudek et al., 2020), safe motor control (Traue et al., 2019) and agricultural drone control. The standard CRL benchmarks are from MuJoCo and Safety-Gym (Ray et al., 2019). The definition of reward and constraint vary from task to task and will be explained when each task is introduced.

**Baseline Approaches.** Three CRL baselines are compared. PCPO (Yang et al., 2020) first performs an unconstrained policy update then project the action to the constrained set. CPO (Achiam et al., 2017) maximizes the reward in a small neighbourhood that enforces the constraints. RCPO (Tessler et al., 2019) incorporates the cost terms and Lagrange multipliers with the reward function to encourage the satisfaction of the constraints. We used the original implementation of CPO and PCPO with KL-projection that leads to the best performance. For RCPO, since the official implementation is not available, we re-implemented RCPO and made sure it matches the original performance. All the three baseline approaches and our DCRL have the same number of neural network parameters. Note that the baseline approaches enforce the constraints by restricting *values functions* while our DCRL restricts the *state density functions*. When we train the value function-based CRL methods on density constrained tasks, we convert the density threshold to the value function threshold by duality between density functions and value functions (Chen & Ames, 2019).

### 5.1. Autonomous Electrical Vehicle Routing

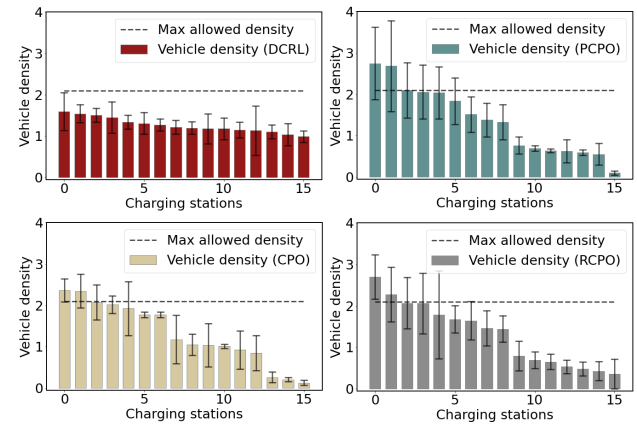


Figure 3. Vehicle densities at the charging stations. All results are averaged over 10 independent trials. 16 charging stations with the highest vehicle densities are kept for visualization.




Figure 4. Average reward and remaining energy. All results are averaged over 10 independent trials.

Our first case study is about controlling autonomous electric vehicles (EV) in the middle of Manhattan, New York. It is adopted from Blahoudek et al. (2020) and is shown in Figure 1. While EVs drive to their destinations, they can avoid running out of power by recharging at the fast charging stations along the roads. At the same time, the vehicles should not stay at the charging stations for too long in order to save resources and avoid congestion. An road intersection is called a node. In each episode, an autonomous EV starts from a random node and drives to the goals. At each node, the EV chooses a direction and reaches the next node along that direction at the next step. The consumed electric energy is assumed to be proportional to traveling distance. There are 1024 nodes and 137 charging stations in total. Denote the full energy capacity of the vehicle as  $c_f$  and the remaining energy as  $c_r$ . When arriving at a charging station, the EV chooses a charging time  $\tau \in [0, 1]$ , then its energy increases to  $\min(c_f, c_r + \tau c_f)$ . The action space includes the EV direction and the charging time  $\tau$ . The state space  $S \subset \mathbb{R}^3$  is consisted of the current 2D location and the remaining energy  $c_r$ . The agent receives a negative reward proportional to its traveling distance to restrict energy consumption, and a +10 reward when it reaches the goal. In one episode, a single agent starts from a random initial location and reaches the goal. The density is accumulated for 2000 episodes, which is equivalent to having 2000 agents.

Two types of constraints are considered: (1) the minimum remaining energy should keep close to a required threshold and (2) the vehicle density at charging stations should be less than a given threshold. Apparently, if the EV chooses a larger  $\tau$ , then the constraint (1) is more likely to be satisfied, while (2) is more likely to be violated, since a larger  $\tau$  will increase the vehicle density at the charging station. These contradictory constraints pose a greater challenge to the RL algorithms. Both constraints can be naturally expressed as density constraints. For constraint (1), we can limit the density of low-energy states. For constraint (2), it is straightforward to limit the EV density (a function of  $\mathbb{E}[\tau]$ ) at charging stations. The threshold of density constraints are transformed to the threshold of value functions to be used by the baseline methods. The conversion is based on the duality of density functions and value functions (Chen & Ames, 2019). Figure 3 demonstrates that our DCRL can avoid

the vehicle densities from exceeding the thresholds, while the baseline methods suffer from constraint violation. This is because DCRL allows us to explicitly set state density thresholds. Figure 4 shows that all the compared methods can satisfy the constraint on remaining energy, and DCRL can reach the highest reward.

## 5.2. MuJoCo Benchmark and Safety-Gym

Experiments are conducted on three tasks adopted from CPO (Achiam et al., 2017) and two tasks from the Safety-Gym (Ray et al., 2019), built on top of the MuJoCo simulator. The tasks include Point-Gather, Point-Circle, Ant-Circle, Car-Goal and Car-Button. In the Point-Gather task, a point agent moves on a plane and tries to gather as many apples as possible while keeping the probability of gather bombs below the threshold. In the Point-Circle task, a point agent tries to follow a circular path that maximizes the reward, while constraining the probability of exiting an given area. The Ant-Circle task is similar to the Point-Circle task except that the agent is an ant instead of a point. Detailed configurations of the three benchmarks can be found in Achiam et al. (2017). In the Car-Goal task, a car agent tries to navigate to a goal while avoiding hazards. In the Car-Button task, a car agent tries to press the goal button while avoiding hazards, and not to press the wrong buttons. Detailed descriptions of the Safety-Gym benchmark can be found in Ray et al. (2019). The original value function constraints for these benchmarks are converted to state density constraints for DCRL by duality between density functions and value functions.

Figure 5 demonstrates the performance of the four methods. In general, DCRL is able to achieve higher reward than other methods while satisfying the constraint thresholds. In the Point-Gather and Point-Circle environments, all the four approaches exhibit stable performance with relatively small variances. In the Ant-Circle environment, the variances of reward and constraint values are significantly greater than that in Point environments, which is mainly due to the complexity of ant dynamics. In Ant-Circle, after 600 iterations of policy updates, the constraint values of the four approaches converge to the neighbourhood of the threshold. The reward of DCRL falls behind PCPO in the first 400 iterations of updates but outperforms PCPO thereafter.

## 5.3. Safe Motor Control

The safe motor control environment is adopted from Traue et al. (2019) and shown in Figure 6 (a). The objective is to control the direct current series motor and ensure its angular velocity follows a random reference trajectory and prevent the motor from overheating. The state space  $S \subset \mathbb{R}^6$  and consists of six variables: angular velocity, torque, current, voltage, temperature and the reference angular velocity. The

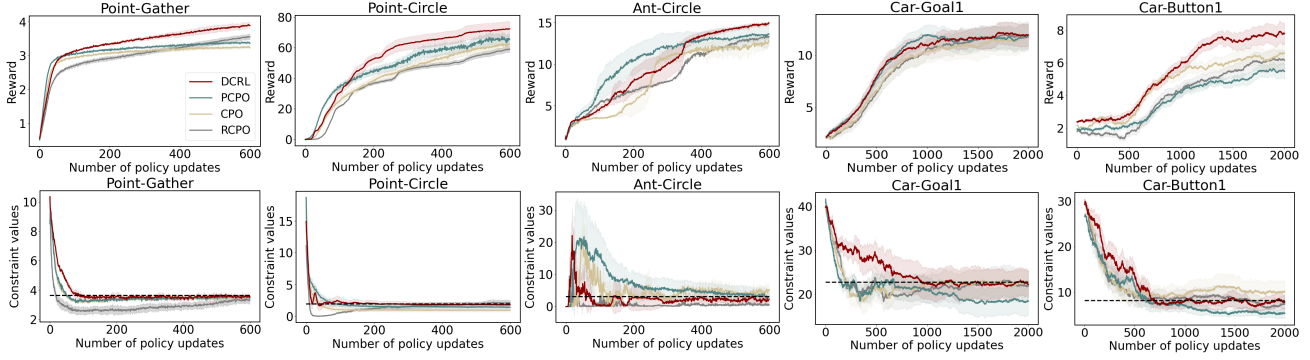


Figure 5. Performance on the constrained reinforcement learning tasks on the MuJoCo (Todorov et al., 2012) simulator. The first three tasks are from CPO (Achiam et al., 2017) and the last two tasks are from the Safety-Gym (Ray et al., 2019). All results are averaged over 10 independent trials. The constraint thresholds are all upper bounds.




Figure 6. Illustration of the safe motor control and drone application environments. (a) Control the rotor to follow a reference angular velocity trajectory while limiting the density of high-temperature states. (b) Control the drones to spray pesticide over a farmland which is divided into five parts and each requires different densities of pesticide.




Figure 7. Reward and state density constraint in the safe motor control task. Results are averaged over 10 independent trials.

action space  $A \subset \mathbb{R}$  is the electrical duty cycle that controls the motor power. The agent outputs a duty cycle at each time step to drive the angular velocity close to the reference. When the reference angular velocity increases, the required duty cycle will need to increase. As a result, the motor’s power and angular velocity will increase and cause the motor temperature to grow. The reward is defined as the negative distance between the measured angular velocity and the reference angular velocity. We consider the state density w.r.t. the motor temperature as constraint. High-temperature states should have a low density to protect the motor. When we train the baseline methods, the density con-




Figure 8. Result of the agricultural spraying problem. (a) Percentage of the entire area that satisfies the pesticide density requirement. (b) Time consumption in steps to reach the destination. Whiskers in the plots denote confidence intervals. Results are calculated with 10 independent trials after each method converges.

straints are converted to equivalent cost value constraints by duality between density functions and value functions. Figure 7 shows the reward and state density for each method. We find that DCRL is successful at reaching the highest reward among the compared methods while respecting the state density constraints. More experiments under different settings can be found in the supplementary material.

#### 5.4. Agricultural Spraying Drone

We consider the problem of using drones to spray pesticides over a farmland in simulation. The farmland is shown in Figure 6 (b), which is divided into 5 regions and different regions require different pesticide densities represented by a lower bound  $\rho_{min}$  and an upper bound  $\rho_{max}$ . The drone starts from the top-left corner, flies over the farmland spraying pesticides, and stops at the bottom-right corner. The state space constrains the position and velocity of the drone, as well as the required pesticide density of the current region. The action space contains the angular acceleration and the vertical thrust to control the drone. The drone sprays a constant volume of pesticide at each time step, and controls the pesticide density on the land by adjusting its position and velocity. A +1 reward will be given when the pesticide



density of the area below the drone enters the required range, and a -1 reward will be given when the density exists the range. A +10 reward will be given when the drone reaches the bottom-right corner that is the destination. Figure 8 shows the experimental results. With the DCRL method, we can obtain the highest percentage of area that satisfies the pesticide density constraint. Also, DCRL requires the least steps to reach the destination. Details of the experiment configurations and results under different settings can be found in the supplementary material.

## 6. Conclusion and future works

We study a novel class of constrained reinforcement learning problems where the constraints are imposed on state density functions, rather than value functions considered by previous literature. State densities have clear physical meanings and can express a variety of constraints of the environment and the system. We prove the duality between density functions and Q functions, then leverage the duality to develop a general algorithm to solve the DCRL problems. We also provide the convergence guarantee of our algorithm even under imperfect policy updates. Note that our algorithm does not guarantee the satisfaction of density constraints during training, and we do see occasional violations at the beginning iterations. In the future, we aim to improve the algorithm to enforce the density constraints throughout the entire training process.

## 7. Acknowledgement

The authors acknowledge support from the DARPA Assured Autonomy under contract FA8750-19-C-0089 and from the Defense Science and Technology Agency in Singapore. The views, opinions, and/or findings expressed are those of the authors and should not be interpreted as representing the official views or policies of the Department of Defense, the U.S. Government, DSTA Singapore, or the Singapore Government.

## References

- Abbeel, P., Coates, A., and Ng, A. Y. Autonomous helicopter aerobatics through apprenticeship learning. *The International Journal of Robotics Research*, 29(13):1608–1639, 2010.
- Achiam, J., Held, D., Tamar, A., and Abbeel, P. Constrained policy optimization. In *International Conference on Machine Learning (ICML)*, pp. 22–31, 2017.
- Altman, E. *Constrained Markov decision processes*, volume 7. CRC Press, 1999.
- Altman, E. and Shwartz, A. Constrained markov games: Nash equilibria. In *Advances in Dynamic Games and Applications*, pp. 213–221. Springer, 2000.
- Basu, A., Bhattacharyya, T., and Borkar, V. S. A learning algorithm for risk-sensitive cost. *Mathematics of operations research*, 33(4):880–898, 2008.
- Blahoudek, F., Brázdil, T., Novotný, P., Ornik, M., Thangeda, P., and Topcu, U. Qualitative controller synthesis for consumption markov decision processes. *International Conference on Computer-aided Verification*, 2020.
- Borkar, V. S. Q-learning for risk-sensitive control. *Mathematics of operations research*, 27(2):294–311, 2002.
- Bovopoulos, A. D. and Lazar, A. A. The effect of delayed feedback information on network performance. *Annals of Operations Research*, 36(1):101–124, 1992.
- Boyd, S., Xiao, L., and Mutapcic, A. Subgradient methods. *lecture notes of EE392o, Stanford University, Autumn Quarter*, 2004:2004–2005, 2003.
- Chen, Y. and Ames, A. D. Duality between density function and value function with applications in constrained optimal control and markov decision process. *arXiv preprint arXiv:1902.09583*, 2019.
- Chen, Y., Singletary, A., and Ames, A. Density functions for guaranteed safety on robotic systems. *American Control Conference*, 10 2019.
- Chen, Y.-C. A tutorial on kernel density estimation and recent advances. *Biostatistics & Epidemiology*, 1(1):161–187, 2017.
- Chow, Y., Nachum, O., Duenez-Guzman, E., and Ghavamzadeh, M. A lyapunov-based approach to safe reinforcement learning. In *Advances in neural information processing systems*, pp. 8092–8101, 2018.
- Clouse, J. A. and Utgoff, P. E. A teaching method for reinforcement learning. In *Machine Learning Proceedings 1992*, pp. 92–101. Elsevier, 1992.
- Dai, B., Shaw, A., He, N., Li, L., and Song, L. Boosting the actor with dual critic. *arXiv preprint arXiv:1712.10282*, 2017.
- Dalal, G., Dvijotham, K., Vecerik, M., Hester, T., Paduraru, C., and Tassa, Y. Safe exploration in continuous action spaces. *arXiv preprint arXiv:1801.08757*, 2018.
- Ding, D., Zhang, K., Basar, T., and Jovanovic, M. Natural policy gradient primal-dual method for constrained markov decision processes. *Advances in Neural Information Processing Systems*, 33, 2020.

- Dotan Di Castro, A. T. and Mannor, S. Policy gradients with variance related risk criteria. In *Proceedings of the 29th International Conference on Machine Learning, Edinburgh, Scotland, UK*, 2012.
- Garcia, J. and Fernández, F. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research*, 16(1):1437–1480, 2015.
- Geibel, P. and Wyzotzki, F. Risk-sensitive reinforcement learning applied to control under constraints. *Journal of Artificial Intelligence Research*, 24:81–108, 2005.
- Geramifard, A., Redding, J., and How, J. P. Intelligent cooperative control architecture: a framework for performance improvement using safe learning. *Journal of Intelligent & Robotic Systems*, 72(1):83–103, 2013.
- Gu, S., Holly, E., Lillicrap, T., and Levine, S. Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates. In *2017 IEEE international conference on robotics and automation (ICRA)*, pp. 3389–3396. IEEE, 2017.
- Heger, M. Consideration of risk in reinforcement learning. In *Machine Learning Proceedings 1994*, pp. 105–111. Elsevier, 1994.
- Hou, C. and Zhao, Q. Optimization of web service-based control system for balance between network traffic and delay. *IEEE Transactions on Automation Science and Engineering*, 15(3):1152–1162, 2017.
- Howard, R. A. and Matheson, J. E. Risk-sensitive markov decision processes. *Management science*, 18(7):356–369, 1972.
- Kadota, Y., Kurano, M., and Yasuda, M. Discounted markov decision processes with utility constraints. *Computers & Mathematics with Applications*, 51(2):279–284, 2006.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. Continuous control with deep reinforcement learning, 2016.
- Lötjens, B., Everett, M., and How, J. P. Safe reinforcement learning with model uncertainty estimates. In *2019 International Conference on Robotics and Automation (ICRA)*, pp. 8662–8668. IEEE, 2019.
- Luo, Z.-Q. and Tseng, P. On the convergence rate of dual ascent methods for linearly constrained convex minimization. *Mathematics of Operations Research*, 18(4):846–867, 1993.
- Moldovan, T. M. and Abbeel, P. Safe exploration in markov decision processes. *arXiv preprint arXiv:1205.4810*, 2012.
- Nachum, O. and Dai, B. Reinforcement learning via fenchel-rockafellar duality. *arXiv preprint arXiv:2001.01866*, 2020.
- Nachum, O., Chow, Y., Dai, B., and Li, L. Dualdice: Behavior-agnostic estimation of discounted stationary distribution corrections. In *Advances in Neural Information Processing Systems*, pp. 2318–2328, 2019.
- Nilim, A. and El Ghaoui, L. Robust control of markov decision processes with uncertain transition matrices. *Operations Research*, 53(5):780–798, 2005.
- Paternain, S., Chamon, L., Calvo-Fullana, M., and Ribeiro, A. Constrained reinforcement learning has zero duality gap. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 7553–7563, 2019.
- Pham, T.-H., De Magistris, G., and Tachibana, R. Optlayer-practical constrained optimization for deep reinforcement learning in the real world. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6236–6243. IEEE, 2018.
- Rantzer, A. A dual to lyapunov’s stability theorem. *Systems and Control Letters*, 42(3):161–168, 2001.
- Ray, A., Achiam, J., and Amodei, D. Benchmarking Safe Exploration in Deep Reinforcement Learning. 2019.
- Rockafellar, R. T. *Convex analysis*. Number 28. Princeton university press, 1970.
- Sato, M., Kimura, H., and Kobayashi, S. Td algorithm for the variance of return and mean-variance reinforcement learning. *Transactions of the Japanese Society for Artificial Intelligence*, 16(3):353–362, 2001.
- Schulman, J., Levine, S., Abbeel, P., Jordan, M., and Moritz, P. Trust region policy optimization. In *International conference on machine learning*, pp. 1889–1897, 2015.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Stooke, A., Achiam, J., and Abbeel, P. Responsive safety in reinforcement learning by pid lagrangian methods. In *International Conference on Machine Learning*, pp. 9133–9143. PMLR, 2020.
- Tang, J., Singh, A., Goehausen, N., and Abbeel, P. Parameterized maneuver learning for autonomous helicopter flight. In *2010 IEEE International Conference on Robotics and Automation*, pp. 1142–1148. IEEE, 2010.
- Tang, Z., Feng, Y., Li, L., Zhou, D., and Liu, Q. Doubly robust bias reduction in infinite horizon off-policy estimation. *arXiv preprint arXiv:1910.07186*, 2019.

- Tessler, C., Mankowitz, D. J., and Mannor, S. Reward constrained policy optimization. In *International Conference on Learning Representations (ICLR)*, 2019.
- Thomaz, A. L., Breazeal, C., et al. Reinforcement learning with human teachers: Evidence of feedback and guidance with implications for learning performance. In *Aaai*, volume 6, pp. 1000–1005. Boston, MA, 2006.
- Todorov, E., Erez, T., and Tassa, Y. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5026–5033. IEEE, 2012.
- Traue, A., Book, G., Kirchgässner, W., and Wallscheid, O. Towards a reinforcement learning environment toolbox for intelligent electric motor control, 2019.
- Wasserman, L. Lecture notes in statistical methods for machine learning, 2019. URL <http://www.stat.cmu.edu/~larry/=sml/densityestimation.pdf>. Online material, last visited on 2019/08/25.
- Yang, T.-Y., Rosca, J., Narasimhan, K., and Ramadge, P. J. Projection-based constrained policy optimization. In *International Conference on Learning Representations*, 2020.
- Zhang, Y., Vuong, Q., and Ross, K. First order constrained optimization in policy space. *Advances in Neural Information Processing Systems*, 33, 2020.