

Policy Bifurcation in Safe Reinforcement Learning

Wenjun Zou¹, Yao Lyu¹, Jie Li¹, Yujie Yang¹, Shengbo Eben Li^{1*}, Jingliang Duan², Xianyuan Zhan^{3,4},
Jingjing Liu³, Yaqin Zhang³, Keqiang Li¹

Abstract—Safe reinforcement learning (RL) offers advanced solutions to constrained optimal control problems. Existing studies in safe RL implicitly assume continuity in policy functions, where policies map states to actions in a smooth, uninterrupted manner. However, our research finds that in some scenarios, the feasible policy should be discontinuous or multi-valued, and a continuous policy can inevitably lead to constraint violations. We are the first to identify the generating mechanism of such a phenomenon, and employ topological analysis to rigorously prove the existence of policy bifurcation in safe RL, which corresponds to the contractibility of the reachable tuple. Our theorem reveals that in scenarios where the obstacle-free state space is non-simply connected, a feasible policy is required to be bifurcated, meaning its output action needs to change abruptly in response to the varying state. To train such a bifurcated policy, we propose a safe RL algorithm called multimodal policy optimization (MUPO), which utilizes a Gaussian mixture distribution as the policy output. The bifurcated behavior can be achieved by selecting the Gaussian component with the highest mixing coefficient. Experiments with vehicle control tasks show that our algorithm successfully learns the bifurcated policy and ensures satisfying safety, while a continuous policy suffers from inevitable constraint violations.

Reinforcement learning (RL) has made significant advancements in solving complex optimal control problems (OCPs), such as robotic manipulation [1], navigation [2], and autonomous driving [3]–[5]. The primary objective of RL is to learn an optimal policy by maximizing the expected total rewards. In many practical control tasks, the policy needs to not only aim for reward maximization but also ensure safety by avoiding undesirable states. In an OCP, safety guarantees are often modeled by a set of hard state constraints, which cannot be violated. The corresponding RL algorithms are called safe RL [6], [7], which has become an effective tool for solving constrained OCPs.

Existing research in safe RL primarily focuses on integrating constrained optimization methods with RL algorithms. J. Achiam et al. [8] propose constrained policy optimization (CPO), employing trust region methods to achieve near-constraint satisfaction at each iteration theoretically. Alex Ray et al. [9] incorporate the Lagrangian multiplier method into proximal policy optimization (PPO) [10], resulting in PPO-Lagrangian (PPO-L), which enables constraint satisfaction and return maximization. Yu et al. [11] and Yang et al.

[12] tackle the problem by exploiting feasible sets in safe RL, defining them as state space regions where policies can consistently operate without violating safety constraints. Their work enhances the understanding and application of feasible sets in safe RL.

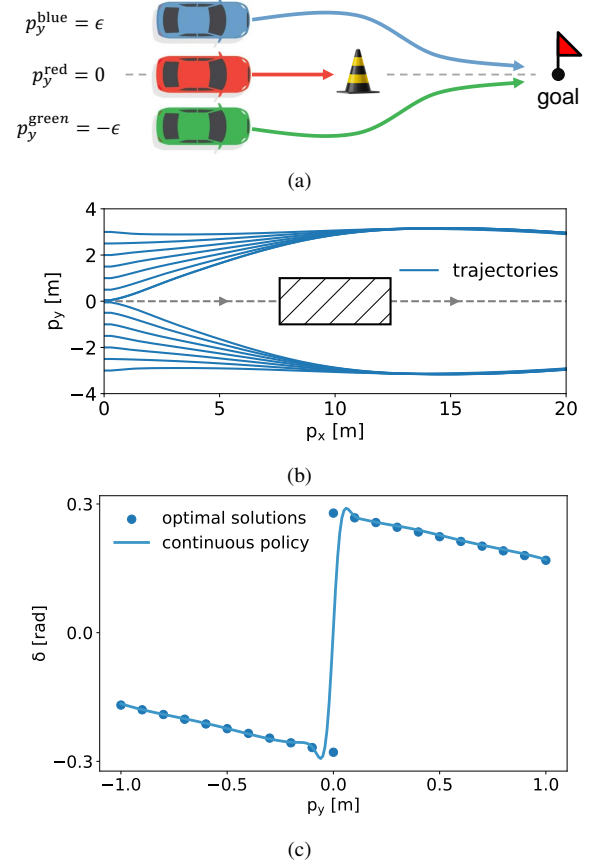


Fig. 1. Limitations of continuous policies in vehicle control problems. (a) Illustration of a control problem for autonomous vehicles that encounter an obstacle and must reach a goal region. (b) The trajectories of a vehicle that start from different lateral positions p_y (in meters) under the optimal policy. There is a bifurcation around $p_y = 0$ where the vehicle steers around the obstacle in different directions. (c) Correlation between the front wheel steering angle δ (in radians) and lateral position p_y for an autonomous vehicle starts at a fixed longitudinal position p_x and speed. The dots show optimal steering responses, with jumps for varying lateral positions at $p_y = 0$. The curve denotes a continuous policy that fails to initiate necessary avoidance maneuvers when starting from the road's center.

In terms of policy structure, existing methods apply deep neural networks to approximate policies [8], [9], [11]–[14]. For tasks involving continuous state and action spaces, the inherent continuity of neural networks results in policies that are

*All correspondence should be sent to Shengbo Eben Li. Email: lishbo@mail.tsinghua.edu.cn.

¹State Key Laboratory of Intelligent Green Vehicle and Mobility, School of Vehicle and Mobility, Tsinghua University, Beijing, China.

²School of Mechanical Engineering, University of Science and Technology Beijing, China.

³Institute for AI Industry Research (AIR), Tsinghua University, Beijing, China.

⁴Shanghai Artificial Intelligence Laboratory, Shanghai, China.

also continuous, mapping system states to actions smoothly. However, existing research overlooks a serious issue: in many cases, no feasible continuous policy solution may exist for constrained OCPs. In such situations, any continuous policy learned through any safe RL methods will inevitably violate the constraints for some initial states, which is extremely detrimental in practical applications.

Considering the scenario shown in Fig. 1a, a safe RL algorithm aims to learn a policy that ensures the ego vehicle reaches the goal without colliding with the obstacle, which is in the center of the road. The ego vehicle starts initially at different positions: blue on the left, red in the center, and green on the right side of the road. The blue vehicle should learn to detour to the left, while the green vehicle should move to the right. As observed, a challenge arises for the red vehicle in the center, as shown in Fig. 1c. If a continuous policy is adopted, the red vehicle may perform an interpolated action based on the actions taken by the blue and green vehicles, potentially leading to a direct collision with the obstacle.

This issue is closely related to multimodal state-action distributions, each of which corresponds to a local optimum of the non-convex OCP. The problem of action multimodality has been noted and studied in imitation learning [15] and offline RL [16]–[18], where data sets are derived from different experts. Applying a continuous policy in these contexts can result in interpolation between different modes, which is often low-value and risky. This issue has received little attention in online RL under the assumption that agents can explore and learn high-value modes while avoiding low-value interpolations. However, as we demonstrate in Fig. 1a, a continuous policy can result in undesirable intermediate behaviors and lead to constraint violation regardless of the algorithm used. In such scenarios, feasible policies should be bifurcated, where “bifurcated” signifies that the policy’s actions change abruptly in response to states. We term this phenomenon “policy bifurcation”. Revealing the underlying mechanisms of this phenomenon remains an unexplored area in current research.

To the best of our knowledge, this work is the first to identify the policy bifurcation mechanism. We established a set of sufficient conditions under which there is no feasible continuous policy for constrained OCPs. These conditions are commonly met in the problem formulations of many robotic and autonomous driving tasks, revealing an overlooked flaw of the policy structure in existing safe RL research.

In this work, we analyze constrained OCPs from a topological perspective. We introduce topological concepts such as paths, loops, and homotopy to define the critical property of contractibility, which plays a pivotal role in our analysis. Our findings reveal that for constrained OCPs with non-simply connected constraints, the set of reachable states and times—termed the reachable tuple—corresponding to a feasible continuous policy must exhibit this contractibility. This requirement often leads to suboptimal policies. Furthermore, when the initial state set is inherently noncontractible, a common occurrence, the reachable tuple cannot be contractible, indicating the impossibility of a continuous policy to be feasible. This highlights the necessity of adopting bifurcated policies in safe RL, especially for tasks with complex safety

constraints.

As a solution to the problem revealed by our theoretical findings, we propose the multimodal policy optimization (MUPO) algorithm. It utilizes a Gaussian mixture distribution to realize a bifurcated policy structure. Additionally, it incorporates spectral normalization and forward Kullback-Leibler (KL) divergence, enabling the policy to better capture multimodal action distributions, thereby facilitating the accurate learning of bifurcated policies. Our experimental results demonstrate the ability of MUPO to efficiently learn bifurcated policies that ensure safety in challenging vehicle control tasks.

Problem setting

We focus on the constrained OCP within a continuous dynamical system, where the objective is to maximize the expected total reward while satisfying safety constraints and reaching the goal. The system is characterized by a continuous state space \mathcal{X} and an action space \mathcal{U} , with continuous dynamics $\dot{x}(t) = f(x(t), u(t))$, where $x(t) \in \mathcal{X}$ and $u(t) \in \mathcal{U}$. The control action $u(t)$ is determined by a policy π_θ parameterized by θ , such that $u(t) = \pi_\theta(x(t))$. Let X_{init} be the set of initial states from which the system can start, and let d_{init} be the probability distribution over X_{init} . The safety constraints are denoted by the sub-zero level set of a time-invariant function $h : \mathcal{X} \rightarrow \mathbb{R}$, defining the constrained set $X_{\text{cstr}} = \{x \in \mathcal{X} \mid h(x) \leq 0\}$. Conversely, the region where $h(x) > 0$ corresponds to the violation region X_{viol} , denoting the states that violate the safety constraints. The goal state set X_{goal} denotes the desired outcomes that the system aims to achieve within the time horizon T . The formal statement of the constrained OCP is:

$$\begin{aligned} & \max_{\theta} \mathbb{E}_{x(0) \sim d_{\text{init}}} \left\{ \int_0^T \gamma^t r(x(t), u(t)) dt \right\} \\ \text{s.t. } & \dot{x}(t) = f(x(t), u(t)), \\ & u(t) = \pi_\theta(x(t)), \\ & h(x(t)) \leq 0, \\ & \forall t \in [0, T], x(0) \in X_{\text{init}}. \end{aligned} \quad (1)$$

Safe RL offers a promising approach to address the constrained OCP by learning a policy that improves performance while ensuring safety. In standard safe RL, the policy π_θ is often parameterized by neural networks. Such a representation of policy usually leads to a continuous policy, which is popular in modern RL algorithms [4], [10], [19], [20]. However, the continuous policy is far from enough for real-world applications, as it will cause unavoidable violations of safety constraints. As shown in Fig. 1b and Fig. 1c, extensive scenarios require us to design a bifurcated policy to ensure safety. The mechanism of policy bifurcation for the constrained OCP remains an open problem, which is the main focus of this work.

RESULTS

Suboptimality and infeasibility of continuous policies

We introduce two key concepts in our analysis: the *reachable tuple* \mathcal{R} , and *contractibility*. The reachable tuple \mathcal{R}




Fig. 2. **Suboptimality of continuous policies.** (a) Illustration of a constrained OCP in a 2D state space, extended with a time dimension. Due to the fact that the safety constraints are active at all times, the violation region X_{viol} is depicted as a cylinder that extends along the time dimension. For the open-loop optimal solution, trajectories from initial states x_a and x_b to goal states x'_a and x'_b bypass the obstacle on different sides. The loop $\ell_{aa'b'ba}$ cannot be continuously contracted to a point within X_{cstr} without passing through X_{viol} , due to the non-simply connected property of X_{cstr} . According to our theoretical analysis, for any continuous policy acting in this manner, there must exist an initial state x_c , from which the trajectory must violate the constraints. (b) Illustration of a feasible continuous policy for the constrained OCP, showing the trajectory that avoids the obstacle from one side to ensure the feasibility of the continuous policy. Continuous policy is forced to take a significantly suboptimal trajectory compared to open-loop optimal solutions, leading to a substantial loss in optimality.

denotes the set of all state-time pairs (x, t) that the system can reach within a given time frame T , meaning the system can be in state x at time t . Contractibility refers to the property that a set can be continuously transformed into a single point without violating the constraints of the system. It is important to note that this definition differs from the traditional notion of contractibility in topology. When the constrained set X_{cstr} is non-simply connected, its subset could be noncontractible. This noncontractibility further leads to limitations of continuous policies in constrained OCPs, which are presented as follows:

Suboptimality of continuous policies: for a constrained OCP characterized by a Lipschitz continuous dynamic function f , if the optimal policy corresponds to a reachable tuple \mathcal{R} that is noncontractible, then the optimal policy cannot be achieved by continuous policy. Fig. 2 provides an illustrative example of this theorem.

Infeasibility of continuous policies: if the initial state set X_{init} is noncontractible, and the goal state set X_{goal} is contractible, then a constrained OCP has no feasible continuous policy. Fig. 3 provides an illustrative example of this theorem.

These theorems highlight the inherent limitations of continuous policies in constrained OCPs. Importantly, these limitations are intrinsic to the continuous policies and are not dependent on the specific learning algorithm. As long as the policy is continuous, these challenges in achieving optimality and feasibility under non-simply connected constraints will persist.

Experimental validation for learning bifurcated policy

To address the limitations of continuous policies, we construct a bifurcated policy that outputs the parameters of a Gaussian mixture distribution. By selecting the mean of the Gaussian component with the highest gate probability as the




Fig. 3. **Infeasibility of continuous policies.** This figure illustrates a constrained OCP within a 2D state space. The initial state set X_{init} is noncontractible; for instance, a closed curve ℓ_{init} within X_{init} encircles the violation region X_{viol} . Under a continuous policy, all points on ℓ_{init} after the same time period form a new closed curve, such as ℓ_{mid} in the figure. However, there exists no continuous deformation in topology that can transform ℓ_{init} into ℓ_{goal} without intersecting X_{viol} , thus demonstrating the infeasibility of the continuous policy.

action, this approach allows for discontinuous changes in the action with respect to changes in the state. We propose the MUPO algorithm, developed based on maximum entropy RL [19], to train this bifurcated policy. It utilizes the forward KL divergence to enhance its ability to learn multimodal distributions, avoiding convergence to suboptimal solutions. Our experimental validation demonstrates that the bifurcated policies learned by the MUPO algorithm can effectively ensure safety and optimality in vehicle control tasks. In contrast, continuous control policies are shown to lead to constraint violations in these scenarios.

1) *Simulation experiments:* We build two experimental tasks with a three-degree-of-freedom vehicle model [21]. The

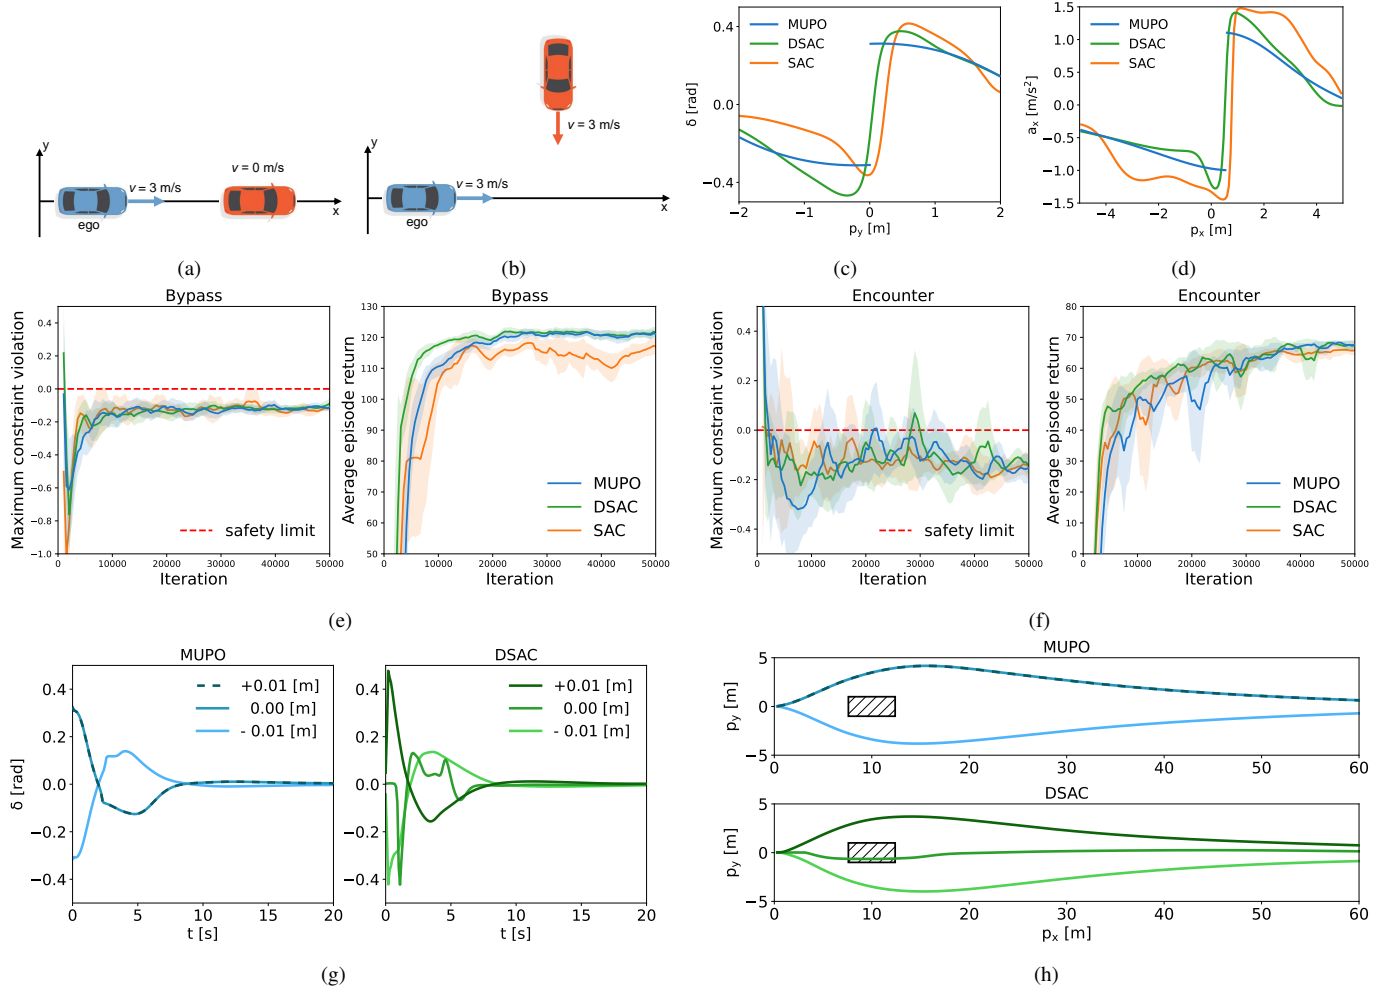


Fig. 4. **Experimental results and visualization.** (a) Bypass task illustration. (b) Encounter task illustration. (c) Open-loop control for the Bypass task. (d) Open-loop control for the Encounter task. (e) Training curves for the Bypass task. (f) Training curves for the Encounter task. (g) Front-wheel steering angle δ changes over time for the Bypass task where the vehicle starts from different lateral positions p_y . (h) Trajectories for the Bypass task where the vehicle starts from different lateral positions p_y .

primary objective of these tasks is to effectively maneuver the ego vehicle along a linear trajectory while adeptly circumventing any potential obstacles. The first task, referred to as “Bypass”, involves a stationary obstacle on the reference path. The objective of the ego vehicle is to bypass this obstacle and then return to the reference path. The second task, referred to as “Encounter”, involves dealing with a vehicle approaching at a constant speed from the left side. The goal for the ego vehicle in this scenario is to maintain on the reference path while ensuring there is no collision with the approaching vehicle.

The baselines for comparison are distributional soft actor-critic (DSAC) [22] and soft actor-critic (SAC) [19] with continuous policies. Detailed hyperparameters can be found in the supplementary materials. First, we evaluate the algorithms with two metrics: maximum constraint violation and average episode return. The training curves are shown in Fig. 4e and Fig. 4f. The results demonstrate that both MUPO and DSAC achieve comparable high average episode returns, surpassing the performance of SAC, while seemingly satisfying the safety constraints under the current evaluation setup. However, it is important to note that the evaluation based on a limited number

of random initializations cannot fully reflect the safety of the learned policies. Our next experiment will further show that only bifurcated policy can guarantee safety, while continuous policy cannot.

We plot the action variations against initial states for both the bifurcated policies of MUPO and the continuous policies of DSAC and SAC. In the Bypass task, Fig. 4c shows how the front-wheel steering angle δ changes with the initial lateral position p_y . All the policies tend to steer left (positive front-wheel angle δ) when initially positioned on the left side of the road, and steer right on the right side. However, near the road center, continuous policies interpolate between left and right steering—going straight, while bifurcated policies can exhibit discontinuous transitions. In the Encounter task, Fig. 4d shows how the longitudinal acceleration a_x changes with the initial longitudinal position p_x . All the policies accelerate when initially ahead, and decelerate when behind, to pass or yield, respectively. Again, only bifurcated policies exhibit discontinuous behavioral transitions. In both tasks, the MUPO learns bifurcated policies that exhibit discontinuous transitions in actions as the state changes, which helps avoid constraint

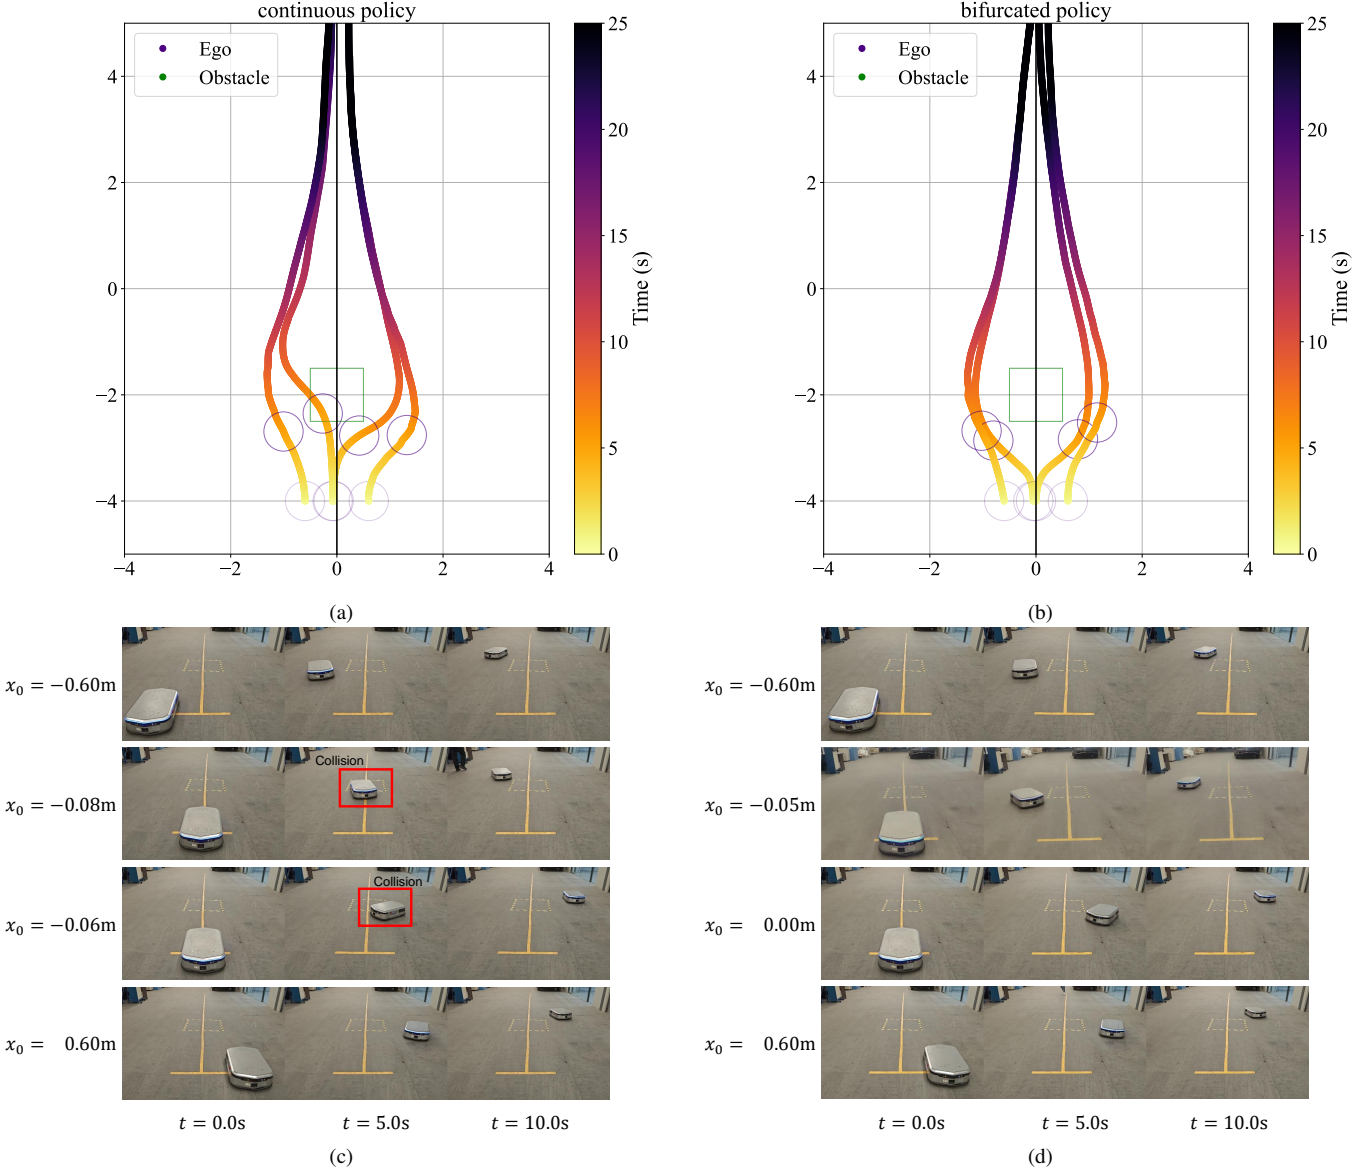


Fig. 5. Comparison of autonomous vehicle trajectories under continuous versus bifurcated policies in real-world execution. (a) Trajectory visualization of autonomous driving with continuous policy. (b) Trajectory visualization of autonomous driving with bifurcated policy. (c) Snapshot of autonomous vehicle executing continuous policy. (d) Snapshot of autonomous vehicle executing bifurcated policy.

violations.

In our closed-loop simulation, we analyze the behavior of the ego vehicle initialized at three specific lateral positions: $+0.01m$, $0.00m$, and $-0.01m$. The control curves of the front-wheel steering angle are displayed in Fig. 4g, and the trajectories of the ego vehicle are shown in Fig. 4h, with the obstacle represented by a black rectangle. The policy of MUPO demonstrates a bifurcation characteristic, where infinitesimal variations in the initial state induce significant shifts in behavioral modes. On the other hand, the DSAC adopts a continuous policy, adjusting behavior smoothly in response to initial state changes. However, between left and right bypass maneuvers around an obstacle, there exists no continuous transition that satisfies constraints. Consequently, employing the continuous policy inevitably leads to constraint violations. It is worth noting that the initial states leading to

constraint violations constitute a minute fraction of the overall state space, as evidenced by position deviation presented at a sub-centimeter scale. This explains why previous studies have not identified this issue: the probability of constraint violations is extremely low when initial states are randomly selected. However, the non-zero probability of such violations occurring cannot be overlooked. In scenarios where safety is critical, the adoption of bifurcated policies is essential.

2) *Real-world experiments*: We explore the potential for constraint violations under continuous policies and evaluate MUPO's capacity to ensure safety in the real world. We utilize the Geekplus M200 mobile robot in our experiments, which is programmed to follow a linear trajectory at $0.3m/s$ while avoiding a centrally placed square obstacle, as illustrated in Fig. 5. The robot is initialized randomly from various positions, as depicted in Fig. 5c and Fig. 5d. Our observations,

as depicted in Fig. 5a, confirm that the robot could avoid the obstacle when initiated from the sides under continuous policy control. However, initiating from a central position led to a collision. In contrast, as shown in Fig. 5b, the bifurcated policy effectively facilitates obstacle avoidance for the robot, regardless of its initial position. This empirical observation aligns with our theoretical predictions, showcasing that the bifurcation property of the policy is crucial in meeting constraints and approaching the global optimum.

DISCUSSION

Our study has revealed the inherent limitations of continuous policies in constrained OCPs. Specifically, our main theoretical results establish that continuous policies fail to achieve optimality and feasibility when the system’s constrained set is non-simply connected. If the initial state set is noncontractible while the goal state set is contractible, no feasible continuous policy exists for the constrained OCP. Although our analysis concentrates on these topological conditions, they are not rare exceptions but rather common occurrences in a wide range of real-world control tasks. This underscores the need for alternative approaches in the design of control policies for constrained OCPs.

The implications of this finding are profound for the field of safe RL, challenging the prevailing use of continuous policy functions in mainstream RL approaches. Our study introduces a novel perspective by highlighting the necessity of adopting bifurcated policies to address the topological challenges in constrained OCPs, a consideration that has been largely overlooked in existing literature. The proposed MUPO algorithm represents a significant advancement in this direction, learning a bifurcated policy that ensures safety and optimality in vehicle control tasks. Experimental validation confirms that MUPO’s bifurcated policies can ensure safety by exhibiting behavior patterns that abruptly change with the system’s state, whereas continuous control policies often result in constraint violations. However, it is important to note that MUPO applies a mixture of Gaussian distributions, which can not approximate arbitrary distributions, leading to potential limitations in performance. Furthermore, while existing methods like diffusion models [23]–[25] can fit arbitrary continuous distributions, they often involve high computational complexity to obtain the most probable actions. Future work will aim to address these challenges to enhance the applicability and effectiveness of the MUPO algorithm.

In conclusion, our study identifies a key limitation in existing RL methods for handling safety control tasks with continuous state and action spaces. This suggests a need to move away from traditional continuous policy functions and towards policy network structures with bifurcation properties. The MUPO algorithm represents a significant advancement in formulating safe and effective control policies for such challenges, with potential applications extending to autonomous driving and other safety-critical domains.

METHODS

Continuous state transition in control system

In this work, we focus on constrained OCP with continuous state space, action space, and dynamic functions, as introduced in the Problem setting section. When dealing with such continuous control problems, existing RL methods primarily adopt continuous policy functions, based on the assumption that continuous functions can satisfy both feasibility and optimality. However, this work challenges that assumption. To begin with, such continuous control systems result in continuous state transitions.

Lemma 1. *Given the control system defined by the dynamic function f and the control policy π , we define the continuous-time state transition function as $F_\pi : \mathcal{X} \times \mathbb{R}_{\geq 0} \rightarrow \mathcal{X}$. $F_\pi(x_{\text{init}}, t)$, which maps an initial state x_{init} to the subsequent state at a later time t under the policy π . It is determined by solving the initial value problem:*

$$\frac{dx}{dt} = f(x, \pi(x)), \quad x(0) = x_{\text{init}}.$$

Here, the relationship between F_π and the solution of the differential equation is given by:

$$F_\pi(x_{\text{init}}, t) = x(t)|_{x(0)=x_{\text{init}}}.$$

Under the assumption of Lipschitz continuity for f and π , the solution $F_\pi(x, t)$ is guaranteed to exist uniquely and be continuous with respect to both x and t .

Proof. Given that f and π are Lipschitz continuous functions, let us define a composite function $g(x) = f(x, \pi(x))$ that denotes the right-hand side of the state’s differential equation. By the properties of Lipschitz functions, the function g is also Lipschitz continuous since for any two points $x_1, x_2 \in \mathcal{X}$, the Lipschitz condition gives:

$$\begin{aligned} \|g(x_1) - g(x_2)\| &= \|f(x_1, \pi(x_1)) - f(x_2, \pi(x_2))\| \\ &\leq L_f(\|x_1 - x_2\| + \|\pi(x_1) - \pi(x_2)\|) \\ &\leq L_f(1 + L_\pi)\|x_1 - x_2\|, \end{aligned}$$

where L_f and L_π are the Lipschitz constants of f and π , respectively. The Picard-Lindelöf theorem [26] asserts that for a Lipschitz continuous function g , there exists a unique solution to the initial value problem $\dot{x}(t) = g(x(t))$ for each initial condition $x(0)$ in \mathcal{X} . Therefore, the state transition function F_π , which describes the state of the system at any time t , exists uniquely for each initial state x and is continuous with respect to both x and t . \square

Time-augmented state dynamics

Definition 1 (Augmented state). To facilitate subsequent proofs in the context of control systems, we consider augmenting the system’s state by incorporating time. Specifically, time t is introduced as an additional dimension to the system state, leading to the augmented state denoted as $\tilde{x} = (x, t)$. The augmented state space is defined as

$$\tilde{\mathcal{X}} = \mathcal{X} \times \mathbb{R}_{\geq 0}.$$

The constraint set X_{cstr} is also augmented to align with $\tilde{\mathcal{X}}$, resulting in the augmented constraint set defined as

$$\tilde{X}_{\text{cstr}} = \{(x, t) \in \tilde{\mathcal{X}} \mid h(x) \leq 0\}.$$

The augmented initial set \tilde{X}_{init} is defined as

$$\tilde{X}_{\text{init}} = \{(x, 0) \mid x \in X_{\text{init}}\}.$$

To clearly delineate the components of the augmented state \tilde{x} , we introduce two projection functions denoted by ϕ . Specifically, $\phi_x(\tilde{x}) = x$ extracts the spatial component of the augmented state, whereas $\phi_t(\tilde{x}) = t$ extracts the temporal component. These functions allow us to refer to and manipulate the individual elements of \tilde{x} in subsequent discussions and proofs, thus simplifying the exposition and analysis of the system's dynamics.

Definition 2 (Truncated augmented state transition function). We define the truncated augmented state transition function, denoted as $\tilde{F}_\pi^T : \tilde{\mathcal{X}} \times \mathbb{R}_{\geq 0} \rightarrow \tilde{\mathcal{X}}$. This function includes an additional constraint that the initial time t does not exceed the upper time limit T . The function is specifically defined as

$$\tilde{F}_\pi^T(\tilde{x}, t) = \begin{cases} (F_\pi(\phi_x(\tilde{x}), t), \phi_t(\tilde{x}) + t) & \text{if } \phi_t(\tilde{x}) + t \leq T \\ (F_\pi(\phi_x(\tilde{x}), T - \phi_t(\tilde{x})), T) & \text{if } \phi_t(\tilde{x}) + t > T \end{cases}.$$

This ensures that the state transition respects the time horizon bound T , truncating the transition at T when $\phi_t(\tilde{x}) + t$ exceeds it. Given Lemma 1, which establishes the continuity of F_π , and considering that the time component is continuous, it follows that \tilde{F}_π^T is also continuous.

Definition 3 (Reachable tuple). Consider the set of initial states X_{init} in the context of a control system with a state transition function F_π . For a given finite time T , the *reachable tuple* \mathcal{R} is defined as the collection of all augmented states traversed within the time interval $[0, T]$ starting from any initial state in X_{init} , i.e.,

$$\mathcal{R} = \{(F_\pi(x_{\text{init}}, t), t) \mid 0 \leq t \leq T, x_{\text{init}} \in X_{\text{init}}\}.$$

Topological concepts in control system

In this work, we assume that the state space of the constrained OCP is a subset of Euclidean space. Since Euclidean spaces are topological spaces with the standard topology induced by the Euclidean metric, any subset of Euclidean space, including our state space, can be considered a topological space. Therefore, when we refer to a set of the state space, it is understood to be equipped with the topology inherited from the Euclidean space.

Definition 4 (Path and loop). In a set X , a path is defined as a continuous mapping $\ell : [0, 1] \rightarrow X$. A path is termed a loop if it satisfies $\ell(0) = \ell(1)$, indicating that the trajectory forms a closed curve in X . The reverse of a path ℓ is denoted as ℓ^{-1} and defined by $\ell^{-1}(s) = \ell(1 - s)$ for all $s \in [0, 1]$, denoting a path that traverses the same trajectory as ℓ but in the opposite direction.

The concept of path and loop is important for the concept of “non-simply connected” and “contractibility of a subset”.

“non-simply connected” is a fundamental concept in topology, referring to a connected space that contains loops that cannot be continuously contracted to a point, usually due to obstacles. The “contractibility of a subset” is further detailed in Definition 5. These concepts, illustrated in Fig. 6, are essential for analyzing the topological properties of the reachable tuple and assessing the feasibility of continuous policies.




Fig. 6. **Visualization of topological concepts.** This figure illustrates the concepts of contractibility as defined in Definition 5. The space X_{cstr} is non-simply connected due to the presence of X_{viol} . The initial state space X_{init} is noncontractible within X_{cstr} , as shown by the loop ℓ_{init} , which cannot be continuously contracted to a point within X_{cstr} . In contrast, the goal state space X_{goal} is an obstacle-free region, allowing every loop to be continuously contracted to a point, demonstrating its contractibility within X_{cstr} .

Definition 5 (Contractibility). Let X and Y be two sets. Y is said to be contractible within X if $Y \subseteq X$ and for any loop $\ell : [0, 1] \rightarrow Y$, there is a point $c \in X$ such that ℓ can be continuously deformed into c within X . Formally, there exists a continuous mapping $H : [0, 1] \times [0, T] \rightarrow X$ where $H(s, 0) = \ell(s)$, and $H(s, T) = c$ for all $s \in [0, 1]$. Otherwise, Y is said to be noncontractible within X .

For brevity, when we say a set of state or augmented state is contractible or noncontractible, we mean that it is within X_{cstr} or \tilde{X}_{cstr} . Based on the lemmas and definitions, the following theorem explores the contractibility of the reachable tuple \mathcal{R} .

Contractibility of the reachable tuple under constraints and continuous dynamics

Lemma 2 (Contractibility of the reachable tuple). *For a control system with dynamic function f and policy π , the reachable tuple \mathcal{R} is contractible when the following conditions are satisfied:*

- 1) **Policy continuity:** the policy $\pi : \mathcal{X} \rightarrow \mathcal{U}$ is Lipschitz continuous on \mathcal{X} .
- 2) **Dynamic function continuity:** the dynamic function $f : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}^n$ is Lipschitz continuous on $\mathcal{X} \times \mathcal{U}$, where \mathbb{R}^n denotes the space of state derivatives.
- 3) **Safety constraint:** $F_\pi(x(0), t) \in X_{\text{cstr}}$ for all $x(0) \in X_{\text{init}}$ and $t \in [0, T]$. This implies that $\mathcal{R} \subseteq \tilde{X}_{\text{cstr}}$.
- 4) **Finite time reachability:** $F_\pi(x(0), T) \in X_{\text{goal}}$ for all $x(0) \in X_{\text{init}}$.
- 5) **Goal set contractibility:** X_{goal} is contractible.

Proof. To begin with, from Lemma 1, we know that the system's state transition function $F_\pi(x, t)$ is continuous under

the conditions of *policy continuity* (condition 1) and *dynamic function continuity* (condition 2). Consequently, the truncated augmented state transition function \tilde{F}_π^T is continuous as well.

Let $\ell_{\mathcal{R}}$ denote any loop in \mathcal{R} , and consider a continuous mapping $H_{\mathcal{R}} : [0, 1] \times [0, T] \rightarrow \mathcal{R}$ where $H_{\mathcal{R}}(s, t) = \tilde{F}_\pi^T(\ell_{\mathcal{R}}(s), t)$. This mapping represents the continuous deformation of $\ell_{\mathcal{R}}$ under the state transition function until T .

Furthermore, since the *safety constraint* (condition 3) ensures $\mathcal{R} \subseteq \tilde{X}_{\text{cstr}}$, it follows that for every $s \in [0, 1]$ and $t \in [0, T]$, we have $H_{\mathcal{R}}(s, t) \in \tilde{X}_{\text{cstr}}$.

$H_{\mathcal{R}}$ leads to the formation of a loop $\ell_{\mathcal{R}, T}(s) = H_{\mathcal{R}}(s, T)$. Projecting $\ell_{\mathcal{R}, T}$ onto \mathcal{X} results in another loop, defined as $\ell_{\mathcal{R}, T}^x(s) = \phi_x(H_{\mathcal{R}}(s, T)) = \phi_x(\tilde{F}_\pi^T(\ell_{\mathcal{R}}(s), T)) = F_\pi(\phi_x(\ell_{\mathcal{R}}(s)), T - \phi_t(\ell_{\mathcal{R}}(s)))$.

Given the *finite time reachability* (condition 4) and the fact that $\ell_{\mathcal{R}}(s) \in \mathcal{R}$, it follows that $F_\pi(\phi_x(\ell_{\mathcal{R}}(s)), T - \phi_t(\ell_{\mathcal{R}}(s))) \in X_{\text{goal}}$. Consequently, $\ell_{\mathcal{R}, T}^x$ forms a loop in X_{goal} .

According to the *goal set contractibility* (condition 5), a continuous mapping $H_c : [0, 1] \times [0, T] \rightarrow X_{\text{cstr}}$ exists such that $H_c(s, 0) = \ell_{\mathcal{R}, T}^x(s)$, and there exists a constant $c \in X_{\text{goal}}$ such that $H_c(s, T) = c$ for all $s \in [0, 1]$.

Subsequently, we define $H_{\tilde{c}}(s, t) = (H_c(s, t), T)$, which contracts $\ell_{\mathcal{R}, T}(s)$ to a point $\tilde{c} = (c, T) \in \mathcal{R}$.

Therefore, a continuous mapping $H_{\mathcal{R}, \tilde{c}} : [0, 1] \times [0, T] \rightarrow \tilde{X}_{\text{cstr}}$ can be constructed, mapping $\ell_{\mathcal{R}}$ to a point in \mathcal{R} :

$$H_{\mathcal{R}, \tilde{c}}(s, t) = \begin{cases} H_{\mathcal{R}}(s, 2t) & \text{if } t \in [0, T/2] \\ H_{\tilde{c}}(s, 2t - T) & \text{if } t \in (T/2, T] \end{cases}.$$

In summary, with the fulfillment of conditions 1 to 5, any loop $\ell_{\mathcal{R}}$ in \mathcal{R} can be continuously transformed into a point \tilde{c} in \tilde{X}_{cstr} by $H_{\mathcal{R}, \tilde{c}}$, implying that \mathcal{R} is contractible. \square

Suboptimality and infeasibility of continuous policy in continuous control system

Lemma 2 demonstrates that for a continuous and feasible policy, the reachable tuple \mathcal{R} is contractible. This provides a key theoretical foundation for analyzing the performance of continuous policies in constrained OCPs. However, further examination of this theorem suggests that continuous policies may either significantly compromise optimality or lead to violations of constraints.

Theorem 1 (Suboptimality of continuous policies). *If the optimal solution of an open-loop constrained OCP corresponds to a reachable tuple \mathcal{R} that cannot be contracted within \tilde{X}_{cstr} , then this optimal solution will be discontinuous. Consequently, continuous policies fail to represent this optimal solution of the open-loop constrained OCP.*

Proof. We prove this theorem by contradiction. According to Lemma 2, if a continuous policy π_{conti} is feasible, then its corresponding reachable tuple \mathcal{R} is contractible. This implies that if the reachable tuple \mathcal{R} of π_{conti} is noncontractible, then π_{conti} is not feasible. Therefore, for an open-loop OCP whose optimal solution leads to a reachable tuple \mathcal{R} that is noncontractible, the optimal solution cannot be represented in the form of a continuous policy function. Consequently,

continuous policies fail to represent this optimal solution of the open-loop constrained OCP. \square

To illustrate the limitations of continuous policies in achieving optimal solutions in certain scenarios, consider the constrained OCP with a 2D state space depicted in Fig. 2a. This figure shows a scenario where the reachable tuple \mathcal{R} cannot be contracted within \tilde{X}_{cstr} . We observe two augmented initial states, $\tilde{x}_a = (x_a, 0)$ and $\tilde{x}_b = (x_b, 0)$, and a path ℓ_{ab} that connects these states and remains within \tilde{X}_{init} for all $s \in [0, 1]$.

The desired continuous policy, denoted as π_{des} , mandates the vehicle to (1) execute a right detour from \tilde{x}_a to reach $\tilde{x}_{a'}$ and a left detour from \tilde{x}_b to reach $\tilde{x}_{b'}$; (2) ensure collision avoidance with the obstacle and achieve goal set within finite time irrespective of initial state $x \in X_{\text{init}}$. Using Lemma 2, we will demonstrate that these two conditions present a contradiction for a continuous policy.

To formalize the trajectories under the desired policy π_{des} , we define paths $\ell_{aa'}(s)$ and $\ell_{bb'}(s)$. These paths represent the trajectories of the vehicle under the policy π_{des} from the initial states \tilde{x}_a and \tilde{x}_b to their respective goal states $\tilde{x}_{a'}$ and $\tilde{x}_{b'}$ over a normalized time scale $s \in [0, 1]$. Specifically, the paths are defined as

$$\begin{aligned} \ell_{aa'}(s) &= \tilde{F}_{\pi_{\text{des}}}^T(\tilde{x}_a, sT), \\ \ell_{bb'}(s) &= \tilde{F}_{\pi_{\text{des}}}^T(\tilde{x}_b, sT). \end{aligned}$$

ℓ_{ab} is a path in X_{init} , it starts from \tilde{x}_a , ends at \tilde{x}_b . $\ell_{a'b'}$ is a path in X_{goal} , it starts from $\tilde{x}_{a'}$, ends at $\tilde{x}_{b'}$. Consider the loop $\ell_{aa'b'ba}$ within \mathcal{R} , which transitions through the sequence $a \rightarrow a' \rightarrow b' \rightarrow b \rightarrow a$:

$$\ell_{aa'b'ba}(s) = \begin{cases} \ell_{aa'}(4s) & \text{for } 0 \leq s \leq \frac{1}{4}, \\ \ell_{a'b'}(4s - 1) & \text{for } \frac{1}{4} < s \leq \frac{1}{2}, \\ \ell_{bb'}^{-1}(4s - 2) & \text{for } \frac{1}{2} < s \leq \frac{3}{4}, \\ \ell_{ab}^{-1}(4s - 3) & \text{for } \frac{3}{4} < s \leq 1. \end{cases}$$

As depicted in Fig. 2a, the loop $\ell_{aa'b'ba}$ cannot be continuously contracted to a point within \tilde{X}_{cstr} due to the obstacle, suggesting that the reachable tuple \mathcal{R} is noncontractible. This noncontractibility contradicts Lemma 2. Consequently, as shown in Fig. 2b, a feasible continuous policy is limited to only circumventing the obstacle from one side, which allows its reachable tuple to be contractible but results in a lower total reward.

Moreover, when the initial state space X_{init} is noncontractible, the constrained OCP can become unsolvable, indicating the absence of any feasible continuous policy under such conditions. This concept is further explained in the following theorem, which specifically addresses the conditions leading to the non-existence of solutions in constrained OCPs.

Theorem 2 (Infeasibility of continuous policies). *Consider a constrained OCP characterized by a Lipschitz continuous dynamic function f and policy π . The problem has no feasible solution if the initial state set X_{init} is noncontractible, and the goal state set X_{goal} is contractible. Specifically, this constrained OCP is unsolvable when the following six conditions are simultaneously met:*

- 1) **Policy continuity:** the policy $\pi : \mathcal{X} \rightarrow \mathcal{U}$ is Lipschitz continuous on \mathcal{X} .

- 2) **Dynamic function continuity:** the dynamic function $f : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}^n$ is Lipschitz continuous on $\mathcal{X} \times \mathcal{U}$, where \mathbb{R}^n denotes the space of state derivatives.
- 3) **Safety constraint:** $F_\pi(x(0), t) \in X_{\text{cstr}}$ for all $x(0) \in X_{\text{init}}$ and $t \in [0, T]$. This implies that $\mathcal{R} \subseteq \tilde{X}_{\text{cstr}}$.
- 4) **Finite time reachability:** $F_\pi(x(0), T) \in X_{\text{goal}}$ for all $x(0) \in X_{\text{init}}$.
- 5) **Goal set contractibility:** X_{goal} is contractible.
- 6) **Initial state set noncontractibility:** X_{init} is noncontractible.

Proof. As established in Lemma 2, under conditions 1 to 5, the reachable tuple \mathcal{R} is contractible. However, by incorporating the *initial state set noncontractibility* (condition 6), we demonstrate that this contractibility does not hold for \mathcal{R} within \tilde{X}_{cstr} .

According to the definition of contractibility (see Definition 5), condition 6 implies the existence of a loop $\ell_{\text{init}} : [0, 1] \rightarrow X_{\text{init}}$ such that there is no such continuous mapping $H : [0, 1] \times [0, T] \rightarrow X_{\text{cstr}}$, where $H(s, 0) = \ell_{\text{init}}(s)$ and $H(s, T) = c$ for some $c \in X_{\text{cstr}}$ and all $s \in [0, 1]$. This indicates that ℓ_{init} cannot be continuously contracted to a point within X_{cstr} .

The augmented loop $\ell_{\text{init, aug}}(s) = (\ell_{\text{init}}(s), 0)$ forms a loop in \mathcal{R} . To demonstrate the non-existence of a continuous mapping capable of contracting $\ell_{\text{init, aug}}$ to a point within \tilde{X}_{cstr} , we initially assume that a continuous mapping $H_{\text{init, aug}}$ exists. Then, define $H_{\text{init}} : [0, 1] \times [0, T] \rightarrow X_{\text{cstr}}$ as $H_{\text{init}}(s, t) = \phi_x(H_{\text{init, aug}}(s, t))$. H_{init} contracts the loop ℓ_{init} to a point in X_{cstr} . However, this directly conflicts with the noncontractibility of X_{init} as stated in condition 6. Hence, this contradiction demonstrates that \mathcal{R} is noncontractible under condition 6, since $\ell_{\text{init, aug}}$ cannot be contracted to a point within \tilde{X}_{cstr} .

In summary, conditions 1 to 5 imply that the reachable tuple \mathcal{R} is contractible. However, condition 6 leads to the noncontractibility of \mathcal{R} . This contradiction underlines the unsolvability of the constrained OCP under these specific conditions. \square

Consequently, continuous policies fail to achieve optimality and feasibility in scenarios with complex constraints, indicating the need for an improved policy structure. The following sections explore Safe RL with bifurcated policies to address the challenges inherent in continuous policies.

Bifurcated policy construction with Gaussian mixture distribution

In this section, we introduce a construction for bifurcated policy using Gaussian mixture distribution. By selecting the action with the highest probability, the policy's output can abruptly change in response to minor variations in inputs, enabling it to meet state constraints.

Definition 6 (Bifurcated policy). A ‘‘bifurcated policy’’ is a control policy that exhibits distinct behavioral modes in response to continuous changes in state. This bifurcation is

marked by discontinuities at certain critical points x_c in the state space, formalized as

$$\lim_{x \rightarrow x_c, x \in \Omega^-} \pi(x) \neq \lim_{x \rightarrow x_c, x \in \Omega^+} \pi(x), \quad (2)$$

where Ω^- and Ω^+ are two disjoint regions in the state space that converge to x_c from different directions.

An example of such a critical point is the center line of a road, as illustrated in Fig. 1a. At these points, the policy behavior exhibits an abrupt change. This characteristic is particularly beneficial in environments with complex safety constraints, where different modes of action are required under varying conditions.

Our approach to constructing a bifurcated policy involves creating a stochastic policy that outputs a Gaussian mixture distribution:

$$\pi(\cdot|x) = \sum_{i=1}^k \pi_{\text{gate}}^i(x; \theta) \mathcal{N}(\mu_i(x; \theta), \sigma_i^2(x; \theta)), \quad (3)$$

where $\pi_{\text{gate}}^i(x; \theta)$ denotes the probability of selecting the i -th Gaussian component in the mixture distribution, i.e., gate probability. The policy $\pi(\cdot|x)$ is a mixture of k Gaussian components, with each component $\mathcal{N}(\mu_i(x; \theta), \sigma_i^2(x; \theta))$ characterized by its mean $\mu_i(x; \theta)$ and variance $\sigma_i^2(x; \theta)$. The parameters θ are used to parameterize the mixture distribution.

To obtain a deterministic action from this mixture distribution, we select the mean of the Gaussian component with the highest gate probability:

$$u = \mu_{\text{max}}(x; \theta), \quad \text{where } \text{max} = \arg \max_i \pi_{\text{gate}}^i(x; \theta). \quad (4)$$

Through this approach, the stochastic policy with Gaussian mixture distribution is transformed into a bifurcated policy, characterized by discontinuous action outputs that change abruptly depending on the state x . The training process, discussed in subsequent sections, involves learning the parameters θ to implement this policy in varying environments effectively.

It is worth noting that the discontinuity in our policy is derived from selecting the maximum gate probability, π_{gate} , instead of learning a very steep continuous function for the mean $\mu_i(x; \theta)$ of each component. Therefore, we propose the incorporation of spectral normalization [27] to prevent each component from overfitting on discontinuous segments, thereby promoting the learning of bifurcated policy. This process constrains the Lipschitz constant of the policy network to a predefined limit, denoted by L_π .

Bifurcated policy learning

In this section, we propose MUPO, a safe RL algorithm specifically designed for the efficient learning of bifurcated policy. Built upon the actor-critic architecture, MUPO comprises an actor learning a bifurcated policy π_θ and a critic evaluating the policy. The critic utilizes two distinct action-value distribution functions denoted as \mathcal{Z}_{ω_1} and \mathcal{Z}_{ω_2} , where ω_1 and ω_2 represent their parameters respectively.

We utilize the straightforward and effective reward penalty method [28] to handle state constraints. The modified reward function with the constraint $h(x) \leq 0$ is defined as follows:

$$\tilde{r}(x, u) = \begin{cases} r(x, u) & \text{if } h(x) \leq 0, \\ r(x, u) - C & \text{otherwise.} \end{cases} \quad (5)$$

As long as the penalty C is sufficiently large, it ensures that the RL algorithm converges to the policy that satisfies the constraints.

1) *Policy evaluation*: Our algorithm incorporates action-value distribution from the DSAC algorithm [22] within the critic component, enabling a comprehensive capture of potential future rewards. This integration is particularly crucial given the multimodal nature of our policy function, which often results in a significant variance in state-action values (i.e., Q-values). Compared to methods like SAC [19] and PPO [10], which utilize deterministic value functions, the action-value distribution in the DSAC algorithm effectively mitigates the overestimation issue in value functions, thereby enhancing policy performance. The loss function for the action-value distribution \mathcal{Z}_ω is defined as the negative expected log-likelihood of the target value given the current value distribution:

$$J_{\mathcal{Z}}(\omega) = - \mathbb{E}_{\substack{(x, u, r, x') \sim \mathcal{B}, u' \sim \pi_\theta \\ Z(x', u') \sim \mathcal{Z}_\omega(\cdot | x', u')}} [\log P(y_z | \mathcal{Z}_\omega(\cdot | x, u))], \quad (6)$$

where \mathcal{B} is the replay buffer containing historical samples. The target value y_z is computed by:

$$y_z = \tilde{r} + \gamma (Z(x', u') - \alpha \log \pi_\theta(u' | x')), \quad (7)$$

where \tilde{r} is the modified immediate reward, γ is the discount factor, and α is the temperature parameter that balances the exploration-exploitation trade-off.

Given that \mathcal{Z}_ω is assumed to be Gaussian, it can be expressed as $\mathcal{Z}_\omega(\cdot | x, u) = \mathcal{N}(Q_\omega(x, u), \sigma_\omega(x, u)^2)$, where $Q_\omega(x, u)$ and $\sigma_\omega(x, u)$ are the mean and standard deviation of the state-action value distribution, respectively.

2) *Policy improvement*: We adopt an energy-based policy approach in line with algorithms like DSAC and SAC. This approach is particularly effective in balancing exploration and exploitation, especially in high-dimensional action spaces. The energy-based policy $\pi_\omega(u|x)$ is defined as

$$\pi_\omega(u|x) = \frac{\exp(\alpha^{-1} Q_\omega(x, u))}{N_\omega(x)}, \quad (8)$$

where the partition function $N_\omega(x)$ normalizes the distribution.

Policy optimization aims at aligning the current policy π_θ as closely as possible with the energy-based policy π_ω . While various projection methods are feasible, utilizing the KL divergence proves to be particularly effective.

The KL divergence is a measure of the difference between two probability distributions, it comes in two forms: forward and reverse. The reverse KL divergence, defined as $\mathbb{E}_{x \sim \mathcal{B}} [D_{\text{KL}}(\pi_\theta(\cdot | x) \| \pi_\omega(\cdot | x))]$, is highly efficient for unimodal policies such as Gaussian distributions. However, in multi-modal distribution scenarios, reverse KL divergence tends to learn only one mode while neglecting others [29], limiting its effectiveness in capturing the full complexity of such distributions.

On the other hand, the forward KL divergence, expressed as $\mathbb{E}_{x \sim \mathcal{B}} [D_{\text{KL}}(\pi_\omega(\cdot | x) \| \pi_\theta(\cdot | x))]$, plays an important role in

learning multi-modal distributions. It fosters a wider coverage by the policy, which is essential for exploring diverse behavioral modes. However, the accurate computation of forward KL divergence relies on precise sampling from the energy-based policy π_ω , which is computationally challenging. Consequently, policies learned using forward KL divergence may not achieve high performance.

Our method combines these two forms of KL divergence, harnessing the advantages of both forward and reverse divergences. This combination enables the learned multimodal policies to achieve high performance.

Our policy loss function combines the reverse and forward KL divergences, aiming to balance reward exploitation with modal exploration. The loss function is denoted as

$$\begin{aligned} J_\pi(\theta) &= J_{\text{rev KL}} + \lambda J_{\text{fwd KL}}, \\ J_{\text{rev KL}} &= \mathbb{E}_{x \sim \mathcal{B}} [D_{\text{KL}}(\pi_\theta(\cdot | x) \| \pi_\omega(\cdot | x))], \\ J_{\text{fwd KL}} &= \mathbb{E}_{x \sim \mathcal{B}} [D_{\text{KL}}(\pi_\omega(\cdot | x) \| \pi_\theta(\cdot | x))], \end{aligned} \quad (9)$$

where λ modulates the contributions of forward KL divergences. This formulation efficiently exploits high-reward regions while enabling the effective exploration of new potential solutions.

Specifically, the reverse KL divergence is expressed as follows, where terms with no gradient to θ are omitted:

$$J_{\text{rev KL}} = \mathbb{E}_{x \sim \mathcal{B}, u \sim \pi_\theta} [\log \pi_\theta(u|x) - Q_\omega(x, u)], \quad (10)$$

and the forward KL divergence is given by:

$$J_{\text{fwd KL}} = \mathbb{E}_{x \sim \mathcal{B}, u \sim \pi_\omega} [-\log \pi_\theta(u|x)], \quad (11)$$

where we can use gradient-based Langevin sampling [30] to sample actions from π_ω .

The temperature parameter α is pivotal in moderating the balance between exploration and exploitation behaviors. Drawing on the approach delineated in [19], α is updated by:

$$\alpha \leftarrow \alpha - \beta_\alpha \mathbb{E}_{x \sim \mathcal{B}, u \sim \pi_\theta} [-\log \pi_\theta(u|x) - \mathcal{H}], \quad (12)$$

where \mathcal{H} is a given constant target entropy. The overall algorithm is detailed in the supplementary materials.

CODE AVAILABILITY

The custom codes developed for our training environment and the MUPO algorithm are available at the following GitHub repository: <https://github.com/THUzouwenjun/MUPO>.

REFERENCES

- [1] E. Kaufmann, L. Bauersfeld, A. Loquercio, M. Müller, V. Koltun, and D. Scaramuzza, "Champion-level drone racing using deep reinforcement learning," *Nature*, vol. 620, no. 7976, pp. 982–987, 2023.
- [2] M. G. Bellemare, S. Candido, P. S. Castro, J. Gong, M. C. Machado, S. Moitra, S. S. Ponda, and Z. Wang, "Autonomous navigation of stratospheric balloons using reinforcement learning," *Nature*, vol. 588, no. 7836, pp. 77–82, 2020.
- [3] Y. Guan, Y. Ren, Q. Sun, S. E. Li, H. Ma, J. Duan, Y. Dai, and B. Cheng, "Integrated decision and control: Toward interpretable and computationally efficient driving intelligence," *IEEE transactions on cybernetics*, vol. 53, no. 2, pp. 859–873, 2022.
- [4] P. R. Wurman, S. Barrett, K. Kawamoto, J. MacGlashan, K. Subramanian, T. J. Walsh, R. Capobianco, A. Devlic, F. Eckert, F. Fuchs, *et al.*, "Outracing champion gran turismo drivers with deep reinforcement learning," *Nature*, vol. 602, no. 7896, pp. 223–228, 2022.

- [5] S. Feng, H. Sun, X. Yan, H. Zhu, Z. Zou, S. Shen, and H. X. Liu, “Dense reinforcement learning for safety validation of autonomous vehicles,” *Nature*, vol. 615, no. 7953, pp. 620–627, 2023.
- [6] E. Altman, *Constrained Markov decision processes*. Routledge, 2021.
- [7] S. E. Li, *Reinforcement learning for sequential decision and optimal control*. Springer, 2023.
- [8] J. Achiam, D. Held, A. Tamar, and P. Abbeel, “Constrained policy optimization,” in *International conference on machine learning*, pp. 22–31, PMLR, 2017.
- [9] A. Ray, J. Achiam, and D. Amodei, “Benchmarking safe exploration in deep reinforcement learning,” *arXiv preprint arXiv:1910.01708*, vol. 7, no. 1, p. 2, 2019.
- [10] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” *arXiv preprint arXiv:1707.06347*, 2017.
- [11] D. Yu, W. Zou, Y. Yang, H. Ma, S. E. Li, Y. Yin, J. Chen, and J. Duan, “Safe model-based reinforcement learning with an uncertainty-aware reachability certificate,” *IEEE Transactions on Automation Science and Engineering*, 2023.
- [12] Y. Yang, Y. Zhang, W. Zou, J. Chen, Y. Yin, and S. E. Li, “Synthesizing control barrier functions with feasible region iteration for safe reinforcement learning,” *IEEE Transactions on Automatic Control*, 2023.
- [13] H. Ma, C. Liu, S. E. Li, S. Zheng, and J. Chen, “Joint synthesis of safety certificate and safe control policy using constrained reinforcement learning,” in *Learning for Dynamics and Control Conference*, pp. 97–109, PMLR, 2022.
- [14] D. Yu, H. Ma, S. Li, and J. Chen, “Reachability constrained reinforcement learning,” in *International Conference on Machine Learning*, pp. 25636–25655, PMLR, 2022.
- [15] P. Florence, C. Lynch, A. Zeng, O. A. Ramirez, A. Wahid, L. Downs, A. Wong, J. Lee, I. Mordatch, and J. Tompson, “Implicit behavioral cloning,” in *Conference on Robot Learning*, pp. 158–168, PMLR, 2022.
- [16] T. Osa, A. Hayashi, P. Deo, N. Morihira, and T. Yoshiike, “Offline reinforcement learning with mixture of deterministic policies,” *Transactions on Machine Learning Research*, 2023.
- [17] A. Ajay, Y. Du, A. Gupta, J. Tenenbaum, T. Jaakkola, and P. Agrawal, “Is conditional generative modeling all you need for decision-making?,” *arXiv preprint arXiv:2211.15657*, 2022.
- [18] C. Lu, H. Chen, J. Chen, H. Su, C. Li, and J. Zhu, “Contrastive energy prediction for exact energy-guided diffusion sampling in offline reinforcement learning,” *arXiv preprint arXiv:2304.12824*, 2023.
- [19] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, “Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor,” in *International conference on machine learning*, pp. 1861–1870, PMLR, 2018.
- [20] J. Hwangbo, J. Lee, A. Dosovitskiy, D. Bellicoso, V. Tsounis, V. Koltun, and M. Hutter, “Learning agile and dynamic motor skills for legged robots,” *Science Robotics*, vol. 4, no. 26, p. eaau5872, 2019.
- [21] Q. Ge, Q. Sun, S. E. Li, S. Zheng, W. Wu, and X. Chen, “Numerically stable dynamic bicycle model for discrete-time control,” in *2021 IEEE Intelligent Vehicles Symposium Workshops (IV Workshops)*, pp. 128–134, IEEE, 2021.
- [22] J. Duan, W. Wang, L. Xiao, J. Gao, and S. E. Li, “Dsac-t: Distributional soft actor-critic with three refinements,” *arXiv preprint arXiv:2310.05858*, 2023.
- [23] L. Yang, Z. Huang, F. Lei, Y. Zhong, Y. Yang, C. Fang, S. Wen, B. Zhou, and Z. Lin, “Policy representation via diffusion probability model for reinforcement learning,” *arXiv preprint arXiv:2305.13122*, 2023.
- [24] Y. Zheng, J. Li, D. Yu, Y. Yang, S. E. Li, X. Zhan, and J. Liu, “Safe offline reinforcement learning with feasibility-guided diffusion model,” *arXiv preprint arXiv:2401.10700*, 2024.
- [25] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song, “Diffusion policy: Visuomotor policy learning via action diffusion,” *arXiv preprint arXiv:2303.04137*, 2023.
- [26] P. Hartman, *Ordinary differential equations*. SIAM, 2002.
- [27] T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida, “Spectral normalization for generative adversarial networks,” *arXiv preprint arXiv:1802.05957*, 2018.
- [28] G. Thomas, Y. Luo, and T. Ma, “Safe reinforcement learning by imagining the near future,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 13859–13869, 2021.
- [29] A. Chan, H. Silva, S. Lim, T. Kozuno, A. R. Mahmood, and M. White, “Greedification operators for policy optimization: Investigating forward and reverse kl divergences,” *The Journal of Machine Learning Research*, vol. 23, no. 1, pp. 11474–11552, 2022.
- [30] M. Welling and Y. W. Teh, “Bayesian learning via stochastic gradient langevin dynamics,” in *Proceedings of the 28th international conference on machine learning (ICML-11)*, pp. 681–688, Citeseer, 2011.

ACKNOWLEDGEMENTS

This study is supported by National Key R&D Program of China with 2022YFB2502901, and NSF China with U20A20334 and 52072213. It is also partially supported by Tsinghua University-Didi Joint Research Center for Future Mobility.

AUTHOR CONTRIBUTIONS

W.Z. conceived the idea, conducted the theoretical analysis and experiments, and drafted the paper. Y.L. assisted with real-world experiments. J.D. contributed to theoretical proofs. All authors discussed and interpreted the results, and edited the paper.

COMPETING INTERESTS

The authors declare no competing interests.

ALGORITHMIC PSEUDO-CODE

Algorithm 1 Multimodal Policy Optimization (MUPO)

Input: $\theta, \omega_1, \omega_2, \alpha, \mathcal{B}, \beta_{\mathcal{Z}}, \beta_{\pi}, \beta_{\alpha}, t$
Initialize: target networks $\bar{\omega}_1 \leftarrow \omega_1, \bar{\omega}_2 \leftarrow \omega_2, \bar{\theta} \leftarrow \theta$
for each iteration **do**
 for each sampling step **do**
 Calculate action $u \sim \pi_{\theta}(u|x)$
 Get reward r and new state x'
 Store transition (x, u, r, x') in buffer \mathcal{B}
 end for
 for each update step **do**
 Sample a batch of data from \mathcal{B}
 Update critic using: $\omega_i \leftarrow \omega_i - \beta_{\mathcal{Z}} \nabla_{\omega_i} J_{\mathcal{Z}}(\omega_i)$,
 Update actor using: $\theta \leftarrow \theta - \beta_{\pi} \nabla_{\theta} J_{\pi}(\theta)$
 Update temperature α
 Update the target networks using:
 $\bar{\omega}_i \leftarrow (1 - t) \cdot \bar{\omega}_i + t \cdot \omega_i$
 $\bar{\theta} \leftarrow (1 - t) \cdot \bar{\theta} + t \cdot \theta$
 end for
end for

HYPERPARAMETERS

TABLE I

Hyperparameter	Value
<i>Shared</i>	
Discount factor (γ)	0.99
Number of hidden layers	3
Number of hidden neurons	256
Optimizer	Adam($\beta_1 = 0.99, \beta_2 = 0.999$)
Learning rate ($\beta_{\mathcal{Z}}$)	1e-3 \rightarrow 5e-5
Learning rate (β_V)	1e-3 \rightarrow 5e-5
Learning rate (β_{π})	1e-3 \rightarrow 5e-5
Learning rate (β_{α})	1e-3 \rightarrow 5e-5
Activation function	GeLU
Batch size	256
Target entropy	$-\dim(\mathcal{U})$
Initial temperature	1.0
<i>MUPO</i>	
Number of components (k)	2
forward KL weight (λ)	1.0 \rightarrow 0
Lipschitz constant (L_{π})	1.0