2020 American Control Conference
Denver, CO, USA, July 1-3, 2020

Density Functions for Guaranteed Safety on Robotic Systems

Yuxiao Chen, Andrew W. Singletary and Aaron D. Ames

Abstract— **The recent study on density functions as the dual of value functions for optimal control gives a new method for synthesizing safe controllers. A density function describes the state distribution in the state space, and its evolution follows the Liouville Partial Differential Equation (PDE). The duality between the density function and the value function in optimal control can be utilized to solve constrained optimal control problems with a primal-dual algorithm. This paper focuses on the application of the method on robotic systems and proposes an implementation of the primal-dual algorithm that is less computationally demanding than the method used in the literature. To be specific, we use kernel density estimation to estimate the density function, which scales better than the ODE approach in the literature and only requires a simulator instead of a dynamic model. The Hamilton Jacobi Bellman (HJB) PDE is solved with the finite element method in an implicit form, which accelerates the value iteration process. We show an application of the safe control synthesis with density functions on a segway control problem demonstrated experimentally.**

I. INTRODUCTION

Safety constraints such as collision avoidance are present in many robotic applications, the classic methods for dealing with safety constraints include motion planning [12], [15] and artificial potential field [14], both of which are effective methods yet not able to guarantee safety. Control Barrier Function (CBF) was first proposed in [1], and have seen success in obstacle avoidance for robotics [6], [10], [24]. With properly computed control invariant set, it is able to guarantee that the safety constraint is always satisfied. However, the computation of a control invariant set [4], [11] or a winning set in the differential game setting [19] is not trivial and the myopic nature of CBF makes it suboptimal.

For optimality, one needs to pose the problem as a constrained optimal control problem. There exist methods that solve the problem for a single initial condition; the standard solution would be the Pontryagin Maximum Principle (PMP) method with costate [13]. However, PMP is typically hard to solve except for some simple cases. Other methods such as model predictive control [18] and shooting method [5] use various types of approximation of the constrained optimal control problem to simplify the computation so that it can be handled by numerical optimization tools. For polynomial systems, occupation measure [16], [17], [26], [27] is capable of synthesizing safe controllers by relaxing the infinite dimensional linear programming to a semidefinite programming problem, but the computation of semidefinite programming do not scale well.

Yuxiao Chen, Andrew Singletary, and Aaron Ames are with the Department of Mechanical and Civil Engineering, California Institute of Technology, Pasadena, CA, 91106, USA. Emails: {chenyx, asinglet, ames}@caltech.edu

Density functions have been studied as the dual of the Lyapunov function in [21], [22]. We showed in [9] that when an optimal control problem has a solution, the density function under the optimal control strategy is the dual of the value function. Furthermore, the constrained optimal control problem can be cast as an optimization over the density function, which makes the representation of the safety constraints straightforward. As a result, the constrained optimal control problem can be solved with a primal-dual algorithm.

The contributions of this paper focus on the implementation of density functions on robotic systems: (i) We propose to apply numerical methods such as finite element method and kernel density estimation to approximate the solution of the HJB PDE and the Liouville PDE, which accelerates the computation. Moreover, comparing to the ODE approach in the literature, the kernel density estimation only requires a simulator instead of an explicit dynamic model. (ii) We apply the density approach on a practical robotic application, the segway, and presents experiment result, as shown in Fig. 1

Fig. 1: Segway experiment

Nomenclature For the remainder of the paper, $\mathbb{N}$ denotes the set of natural numbers, $\mathbb{N}_+$ denotes the positive natural number, $\mathbb{R}$ denotes the set of real numbers, $\mathbb{R}_{\geq 0}$ denotes the nonnegative real numbers. Given a dynamic equation $\dot{x} = f(x)$, $\Phi_f(x_0, T)$ denotes the flow map of the dynamics with initial state x_0 and horizon T. $\langle a, b\rangle_{\mathcal{X}} = \int_{\mathcal{X}} a(x) \cdot b(x) dx$ denotes the inner product of two functions a and b. $\mathbf{0}$ denotes a vector of all zeros or a function that is always zero, depending on the context. $\mathbb{1}_S$ denotes the indicator function of a set S.

II. REVIEW OF DENSITY FUNCTIONS

In this section, we review the concept of density function and value function and the duality between them.

Optimal control and value function Consider the following discounted optimal control problem with infinite horizon.

$$\min_{u} \int_{0}^{\infty} e^{-\kappa t} C\left(x\left(t\right), u\left(t\right)\right) dt \text{ s.t. } \dot{x} = F\left(x, u\right), \quad (1)$$

where $x \in \mathcal{X} \subseteq \mathbb{R}^n$ is the state, $u \in \mathcal{U} \subseteq \mathbb{R}^m$ is the control input, $F : \mathcal{X} \times \mathcal{U} \rightarrow \mathcal{X}$ is the dynamic equation described as an ODE, $C : \mathcal{X} \times \mathcal{U} \rightarrow \mathbb{R}$ is the running cost function, and κ is the discount factor.

The Pontryagin Maximum Principle [20] gives necessary conditions for the optimality of the solution, and the optimal control problem can be solved with Hamilton-Jacobi-Bellman PDE, which utilize the principle of optimality [3]. In the infinite horizon case, the HJB equation becomes

$$\min_{u \in \mathcal{U}} \left\{\nabla V\left(x\right) \cdot F\left(x, u\right) + C\left(x, u\right)\right\} - \kappa V = 0, \quad (2)$$

where

$$\begin{aligned} V\left(x_0\right) = \min_{u(t) \in \mathcal{U}} \int_{0}^{\infty} e^{-\kappa t} C\left(x\left(t\right), u\left(t\right)\right) dt, \\ \text{s.t. } x(0) = x_0, \dot{x} = F(x, u) \end{aligned} \quad (3)$$

is the value function of the optimal control problem, which represents the optimal cost-to-go for an initial condition x_0. Once the value function is known, the optimal policy is then

$$\mathbf{u}^{\star}\left(x\right) = \underset{u \in \mathcal{U}}{\arg\min} \left\{\nabla V\left(x\right) \cdot F\left(x, u\right) + C\left(x, u\right)\right\}. \quad (4)$$

Density function for dynamic systems Taking a different perspective, a density function $\rho : [0, \infty] \times \mathcal{X} \rightarrow \mathbb{R}$ can be understood as the measure of state concentration (or presence) in the state space. Given the dynamics $\dot{x} = f(x)$, the evolution of density function follows the Liouville PDE:

$$\begin{aligned} \frac{\partial \rho}{\partial t} + \nabla \cdot \left(\rho \cdot f\right) = \phi(t, x, \rho) \\ \rho\left(0, x\right) = \rho_0\left(x\right), \end{aligned} \quad (5)$$

where $\phi : [0, \infty) \times \mathcal{X} \times \mathbb{R} \rightarrow \mathbb{R}$ is the supply function, indicating the intensity that new states enter or leave the state space. $\phi(t, x_0, \rho(x_0, t)) > 0$ denotes a source, and $\phi(t, x_0, \rho(x_0, t)) < 0$ denotes a sink. We allow ϕ to depend on ρ to allow more flexible characterization of the supply.

The Liouville PDE can be solved as an ODE since

$$\frac{\partial \rho}{\partial t} + \nabla \cdot \left(\rho \cdot f\right) = \left.\frac{d\rho}{dt}\right|_{\dot{x} = f(x)} + (\nabla \cdot f)\rho = \phi. \quad (6)$$

This implies that we can integrate the following ODE to get the density function alone the trajectory of the dynamic system $\dot{x} = f(x)$ as

$$\left[\begin{array}{c} \dot{x} \\ \dot{\rho} \end{array}\right] = \left[\begin{array}{c} f\left(x\right) \\ \phi\left(t, x, \rho\right) - \nabla \cdot f \cdot \rho \end{array}\right]. \quad (7)$$

With this, we can evaluate the density function at any state x_t, any time $t > 0$ with the following two step procedure:

- First, solve the reverse ODE of $\dot{x} = -f(x)$ with initial condition x_t to get $\Phi_{-f}(x_t, t) = \Phi_f(x_t, -t)$.
- Then, solve the extended ODE in (7) with initial condition $[\Phi_f(x_t, -t)^\intercal, \rho_0(\Phi_f(x_t, -t))]^\intercal$ to time T to obtain $[x_t^\intercal, \rho(t, x_t)]^\intercal$.

Assumption 1. *$\mathcal{X}$ is forward invariant under all possible dynamics considered in the optimal control problem.*

Remark 1. *Assumption 1 is clearly true when $\mathcal{X} = \mathbb{R}^n$. For a compact $\mathcal{X}$, invariance of $\mathcal{X}$ can be achieved if some barrier function intervention is implemented on $\partial\mathcal{X}$, see [2], [11] for example.*

For a stationary supply function, i.e. a ϕ only depending on x and ρ, one would hope that there exists a stationary density function that any initial condition ρ_0 converges to, i.e.,

$$\phi\left(x, \rho_s\right) - \nabla \cdot \left(\rho_s \cdot f\right) = 0. \quad (8)$$

This is not always the case, but we provide sufficient condition for the convergence in Lemma 1.

Define the extended dynamics in (7) as $\overline{f}$. Given a stationary supply function ϕ and an initial density function ρ_0, from the two step procedure shown above, we have

$$\rho\left(x, t\right) = \Phi_{\overline{f}}([\Phi_f\left(x, -t\right), \rho_0\left(\Phi_f\left(x, -t\right)\right)]^\intercal, t)_{\downarrow\rho}, \quad (9)$$

where $\downarrow \rho$ means the projection of $[x^\intercal, \rho]^\intercal$ to ρ.

Lemma 1. *Given a stationary supply function ϕ and an initial density function ρ_0, assume that there exists a $\rho_s : \mathcal{X} \rightarrow \mathbb{R}$ such that*

$$\forall x \in \mathcal{X}, \phi\left(x, \rho_s\right) - \nabla \cdot \left(\rho_s \cdot f\right) = 0.$$

For any $x \in \mathcal{X}$, if there exists $T \geq 0$ such that $\forall t \geq T, \rho_0(\Phi_f(x, -t)) = \rho_s(\Phi_f(x, -t))$, then $\forall t \geq T, \rho(x, t) = \rho_s(x)$.

This is Theorem 1 in [9], see the proof therein.

Corollary 1. *With Lemma 1 and Assumption 1, if there exists a $T > 0$ such that $\Phi_f(x, -T) \notin \mathcal{X}$, then clearly*

$$\rho_0(\Phi_f(x, -T)) = 0 = \rho_s(\Phi_f(x, -T)),$$

therefore the density at x converges to the stationary density ρ_s in finite time T.

Duality in optimal control In this section, we review the duality relationship between the value function and the density function in optimal control problems. We consider the infinite optimal control problem in (1).

Assumption 2. *We assume that given the dynamics $\dot{x} = F(x, u)$, state cost C, input range $\mathcal{U}$, and supply function ϕ, there exists a unique differentiable value function V that satisfies the HJB PDE in* (2)*; and there exists a differentiable stationary density function ρ_s that satisfies* (8)*.*

Remark 2. *The existence and uniqueness of the HJB PDE can be guaranteed under mild assumptions of the dynamics and cost function, such as Lipschitz continuity, see [7] for example.*

For the uniqueness of ρ_s, we give the following proposition.

Proposition 1. *Assume that the supply function is*

$$\phi(x,\rho)=\phi_+(x)-\kappa\rho, \tag{10}$$

and $\exists T>0, \varepsilon>0$ *such that* $\forall x\in\mathcal{X}, \forall t>T$, $-\kappa-\nabla\cdot f\mid_{\Phi_f(x,t)}<-\varepsilon$. *If there exists a stationary density* ρ_s *satisfying* (8)*, then* ρ_s *is unique and any differentiable density function* $\rho(t,x)$ *satisfying* (5) *will converge to* ρ_s *as* $t\to\infty$.

Proof. Given any two differentiable density functions ρ^1,ρ^2 satisfying (5), let $\bar\rho\doteq\rho^2-\rho^1$ and define the Lyapunov function $V(t)=\frac{1}{2}\int_{\mathcal{X}}\bar\rho(x)^2dx$. Then we have

$$\frac{\partial\bar\rho}{\partial t}=-\nabla\cdot(\bar\rho\cdot f)-\kappa\rho,$$
$$\dot V=\int_{\mathcal{X}}\bar\rho\cdot\frac{\partial\bar\rho}{\partial t}dx=\int_{\mathcal{X}}\bar\rho\cdot(-\nabla\cdot(\bar\rho\cdot f)-\kappa\bar\rho)dx$$

Integrating by parts and use the adjoint condition due to Assumption 1,

$$\dot V=\int_{\mathcal{X}}-\kappa\bar\rho^2-\bar\rho\nabla\bar\rho\cdot f dx$$

$\forall t>T$, at any state $x\in\mathcal{X}$, there are two possibilities. First, $\Phi_f(x,-t)\notin\mathcal{X}$. In this case, by Corollary 1, $\bar\rho(t,x)=0$. Second, $\Phi_f(x,-t)\in\mathcal{X}$, then by assumption, $-\kappa-\nabla\cdot f\mid_{\Phi_f(x,t)}<-\varepsilon$. Combining the two cases, we have

$$\dot V\le\int_{\mathcal{X}}-\varepsilon\bar\rho^2dx=-2\varepsilon V,$$

which implies that $\lim_{t\to\infty}V=0$. Furthermore, since ρ^1,ρ^2 are differentiable, $\bar\rho$ is differentiable, thus continuous, therefore $V(t)=0$ implies $\bar\rho(t,x)=0$ everywhere in $\mathcal{X}$ as t goes to infinity. Taking $\rho^1(t,x)=\rho_s(x),\forall t$ completes the proof. ■

Lemma 1 shows that when the rate of state aggregation is slower than the dissipation caused by the discount, the density functions converge to a unique stationary ρ_s.

The duality is between the value function V and the stationary density function ρ_s satisfying (8), both functions of only x. Consider the overall cost rate under the supply ϕ_+. By Bellman's principle of optimality, we know that there exists a pure state feedback law $\mathbf{u}^\star$ that minimizes the overall cost J, and is determined by the following equation:

$$\begin{aligned}J_p^\star=&\langle\phi_+,V\rangle_{\mathcal{X}}\\ \text{s.t. }&\mathbf{u}^\star(x)=\arg\min_{u\in\mathcal{U}}\ \nabla V\cdot F(x,u)+C(x,u)\\ &C_{\mathbf{u}^\star}+\nabla V\cdot F_{\mathbf{u}^\star}-\kappa V=0,\end{aligned}\tag{11}$$

where $C_{\mathbf{u}}(x)=C(x,\mathbf{u}(x))$, $F_{\mathbf{u}}(x)=F(x,\mathbf{u}(x))$. The overall cost $J^\star$ is simply the inner product of the optimal value function and the positive supply ϕ_+. The optimization sign is left out because $V^{\mathbf{u}^\star}$ is completely determined by the equality constraints. We denote (11) as the primal problem.

Alternatively, when the density function reaches a stationary distribution ρ_s under control strategy $\mathbf{u}$, the overall cost rate can also be represented as

$$J_d=\langle C_{\mathbf{u}},\rho_s\rangle_{\mathcal{X}}, \tag{12}$$

which is interpreted as the inner product of the stationary state distribution and the state cost.

This means that instead of thinking of the value function for each x, we can think about the stationary density distribution ρ_s. The following optimization solves for the optimal overall cost:

$$\begin{aligned}J_d^\star=&\min_{\rho_s,\mathbf{u}}\langle\rho_s,C_{\mathbf{u}}\rangle_{\mathcal{X}}\\ &\text{s.t.}\nabla\cdot(\rho_s\cdot F_{\mathbf{u}}))=\phi_+-\kappa\rho_s,\\ &\forall x\in\mathcal{X},\mathbf{u}(x)\in\mathcal{U},\ \rho_s(x)\ge0,\end{aligned}\tag{13}$$

which we denote as the dual problem.

Lemma 2. *The optimization in* (13) *and* (11) *are dual to each other and if there exists optimal solutions to both problems, there is no duality gap, i.e.,* $J_p^\star=J_d^\star$.

Proof. This is Theorem 1 in [8], see the proof therein. ■

III. Constrained optimal control synthesis with density function

In this section, we review the density function approach for constrained optimal control, utilizing the duality relationship reviewed in Section II. The reason for using the density function is that most state constraints cannot be posed on the value function, while posing them on the density function is straightforward. The following constrained optimal control problem is considered:

$$\begin{aligned}&\min_{\mathbf{u}}\int_0^\infty e^{-\kappa t}C(x(t),\mathbf{u}(x))dt\\ &\text{s.t. }\dot x=F(x,u),\forall x\in\mathcal{X},\mathbf{u}(x)\in\mathcal{U}\\ &\quad\forall t, x(t)\notin\mathcal{X}_d,\end{aligned}\tag{14}$$

where $\mathcal{X}_d$ is the danger set. The safety constraint $x\notin\mathcal{X}_d$ is hard to pose in the primal optimization, but can be easily posed in the dual problem:

$$\begin{aligned}&\min_{\rho_s,\mathbf{u}}\langle\rho_s,C_{\mathbf{u}}\rangle_{\mathcal{X}}\\ &\text{s.t.}\nabla\cdot(\rho_s\cdot F_{\mathbf{u}}))=\phi_+-\kappa\rho_s,\\ &\quad\forall x\in\mathcal{X},\mathbf{u}(x)\in\mathcal{U},\ \rho_s(x)\ge0,\\ &\quad\forall x\in\mathcal{X}_d,\rho_s(x)\le\rho^{\max}.\end{aligned}\tag{15}$$

where $\rho^{\max}$ is the tolerance, and it takes the value 0 if the constraint is absolute.

We showed in [8] that the constrained optimal control problem can be solved with a primal-dual algorithm. The primal problem is a optimal control problem (posed as HJB) with an additional term corresponding to the safety constraint:

$$\begin{aligned}J_p^\star=&\left\langle\phi_+,V^{\mathbf{u}^\star}\right\rangle_{\mathcal{X}}\\ &\text{s.t.}\mathbf{u}^\star(x)=\arg\min_{u\in\mathcal{U}}\ \nabla V\cdot F+C\\ &C_{\mathbf{u}^\star}+\sigma\mathbb{1}_{\mathcal{X}_d}+\nabla V\cdot F_{\mathbf{u}^\star}-\kappa V=0.\end{aligned}\tag{16}$$

where $\sigma:\mathcal{X}\to\mathbb{R}_{\ge0}$ is the Lagrange multiplier associated with the safety constraint and $\mathbb{1}_{\mathcal{X}_d}$ is the indicator function of the danger set. The dual problem evaluates the density function under the current controller and updates the Lagrange multiplier σ. The primal-dual algorithm iterates between the primal and dual problem until it finds a safe controller. The procedure is shown in Algorithm 1, where $\alpha > 0$ is the step size and $\varepsilon > 0$ is the tolerance on the complementary slackness condition.

Algorithm 1 Primal-dual algorithm for safe optimal control

```
σ[0] ← 0, k = 0
do
    Solve (16) with σ[k], get u*.
    Compute the stationary density ρ_s under u*.
    σ[k+1] ← max{0, σ[k] + α((ρ_s − ρ^max) 1_{X_d})}.
    k ← k + 1
while ‖max(0, ρ_s − ρ^max) 1_{X_d}‖_∞ > ε
return u*, ρ_s, V
```

IV. Safe control synthesis for a robotic system with density

In this section, we present the application of the density approach on a segway control problem.

Solving Hamilton Jacobi Bellman PDE Algorithm 1 requires solving the HJB PDE in (16) multiple times. We use a finite difference method (FDM) to turn the HJB into a difference equation by dividing the state space into a uniform grid. Let x_g denote the array of grid points and V_g denote the array of V values on x_g. The partial differentials become differences between neighbors in V_g. Upwind scheme is used to improve the stability of the FDM [23]. To be specific, let $I = (i_1, ..., i_n)$ be the index of a grid point, and $I_{-,k} = (i_1, ..., i_k - 1, ..., i_n)$, $I_{+,k} = (i_1, ..., i_k + 1, ..., i_n)$ are the indices of the two neighbors on the k-th dimension. Then define the backward and forward differences as

$$\nabla^+ V(x_g^I) = \begin{bmatrix} \frac{V_g^I - V_g^{I_{-,1}}}{\Delta x_1} \\ \cdots \\ \frac{V_g^I - V_g^{I_{-,n}}}{\Delta x_n} \end{bmatrix}, \nabla^- V(x_g^I) = \begin{bmatrix} \frac{V_g^{I_{+,1}} - V_g^I}{\Delta x_1} \\ \cdots \\ \frac{V_g^{I_{+,n}} - V_g^I}{\Delta x_n} \end{bmatrix} \tag{17}$$

When $I_{+,k}$ or $I_{-,k}$ are out of bounds of the grid on the boundary, the corresponding entries of the backward and forward differences are simply zero. Then under the upwind scheme, $\nabla V \cdot f(x_g^I)$ is computed as

$$\nabla V \cdot f(x_g^I) \approx \max\{\mathbf{0}, f(x_g^I)\} \cdot \nabla^+ V(x_g^I) + \min\{\mathbf{0}, f(x_g^I)\} \cdot \nabla^- V(x_g^I), \tag{18}$$

where the $\max$ and $\min$ are taken entry-wise.

To accelerate the computation, the HJB PDE is written in an implicit form and becomes a linear equation of V_g. When fixing the control strategy $\mathbf{u}$, the value function satisfies

$$C_{\mathbf{u}} + \sigma \mathbb{1}_{\mathcal{X}_d} + \nabla V \cdot F_{\mathbf{u}} = 0,$$

which is turned into a difference equation of V_g:

$$C_{\mathbf{u}}(x_g) + \sigma \mathbb{1}_{\mathcal{X}_d}(x_g) + \nabla V_g F_{\mathbf{u}}(x_g) = 0. \tag{19}$$

This is merely a linear algebraic equation of V_g, and can be solved with linear solvers. Although direct methods such as LU decomposition are more accurate, they are too slow and requires too much memory. Therefore, we use indirect methods such as gradient descent instead. As a comparison, to solve the HJB equation with 148176 grid points, the direct method takes 218s and the indirect method takes 12.6s with 10^{-5} accuracy. Once V_g is solved, the control input at each grid point x_g^I can be evaluated as

$$\mathbf{u}(x_g^I) = \underset{u \in \mathcal{U}}{\arg\min}\ \nabla V(x_g) \cdot F(x_g^I, u) + C(x_g^I, u) \tag{20}$$

The algorithm alternates between solving (19) for V_g and updating the controller $\mathbf{u}$ based on (20) until convergence.

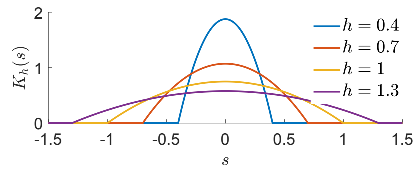

Fig. 2: Epanechnikov kernel

Kernel density estimation Since we are using a grid to compute HJB, σ at each grid point is required for the primal-dual algorithm, which in turn depends on ρ_s at each grid point. The ODE approach for computing ρ_s introduced in Section II is a convenient way to obtain $\rho_s(x)$, but it suffers from two shortcomings: (i) the computation load is heavy with a big grid (number of ODEs to solve grows exponentially with the grid dimension) (ii) it requires that an accurate model is known. We propose to use kernel density estimation, a method widely used in probability estimation, to approximate the density function.

More specifically, given ϕ_+, we sample N initial conditions x_0 according to ϕ_+ and for each x_0, simulate the system under $\mathbf{u}$, resulting in state sequence $\{x(t)\}, t = 0, t_s, 2t_s, ..., Mt_s$. M is picked to be large enough so that $e^{-\kappa M t_s} \ll 1$. Each sample then bears mass

$$m(x(t)) = e^{-\kappa t} \frac{t_s \Phi_+}{N}, \tag{21}$$

where $\Phi_+ = \int_{\mathcal{X}} \phi_+(x) dx$ denotes the total supply rate. To get the density at x, one simply compute

$$\hat{\rho}_s(x) = \sum_{i=1}^{N} \sum_{j=1}^{M} K_h(x - x^i(jt_s)) m(x^i(jt_s)), \tag{22}$$

where K_h is the kernel with bandwidth h and $x^i(jt_s)$ denotes the j-th state sample in the i-th trial. Here we choose the Epanechnikov kernel, which is defined as

$$K_h(s) = \begin{cases} \frac{3}{4h}\left(1 - \frac{s^2}{h^2}\right), & |s| \le h \\ 0 & |s| > h \end{cases}. \tag{23}$$

It satisfies the requirement for a kernel, i.e., $\int_{-\infty}^{\infty} K_h(s) ds = 1$ for all $h > 0$, as shown in Fig. 2. For $s \in \mathbb{R}^n$, one can simply take the product of n Epanechnikov kernels to get an n-dim kernel:

$$K_{\mathbf{h}}(s) = \prod_{i=1}^{n} K_{h_i}(s_i),$$

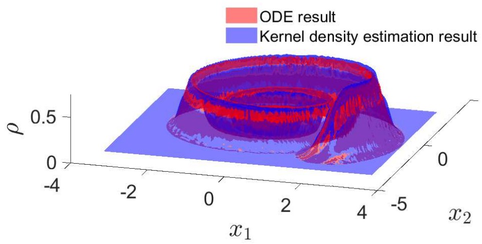

Fig. 3: Kernel density estimation

where $\mathbf{h} \in \mathbb{R}^n_{>0}$ is the bandwidth vector. Moreover, the Epanechnikov kernel is a kernel with finite support, which implies

$$\forall s \notin \{s \mid |s_i| < h_i\}, K_{\mathbf{h}}(s) = 0.$$

This is particularly useful for density function estimation for safe control synthesis, since it gives a compact neighborhood of the query point x such that any sample outside the neighborhood will not affect the density estimation of x. Suppose we use a kernel with infinite support, such as the Gaussian kernel, the density inside the danger set would never be zero and the primal-dual algorithm would not be implementable.

There exist many results about the accuracy of kernel density estimation, we quote the following lemma from [25]:

Lemma 3. *For simplicity, let $\mathbf{h} = [h, ..., h]$. Let $\hat{\rho}_s$ be the kernel density estimation of ρ_s, then for a $\rho_s \in \Sigma(\beta, L)$, for any fixed $\delta > 0$,*

$$\mathbb{P}\left\{\sup_{x\in\mathbb{R}^n} |\hat{\rho}_s(x) - \rho_s(x)| > \Phi_+\sqrt{\frac{C\log n}{nh^d}} + \Phi_+ch^\beta\right\} \leq \delta$$

for some constants c and C where C depends on δ.

$$\Sigma(\beta, L) = \left\{ \begin{array}{c} g : |D^s g(x) - D^s g(y)| \leq L\,\|x-y\|, \\ \forall\,|s| = \beta - 1, \forall x, y \end{array} \right\},$$

where $D^s g(x) = \frac{\partial^{s_1+\ldots+s_d}}{\partial x_1^{s_1}\ldots\partial x_d^{s_d}}$ and $|s| = \sum_{i=1}^d s_i$.

This is adapted from Theorem 9 in [25], see more detail and other bias analysis results therein. Lemma 3 essentially says that if ρ_s is smooth enough and one takes enough samples, the bias of the kernel density estimator can be bounded uniformly in the state space.

To demonstrate the performance of kernel density estimation, the following toy example is considered.

$$x = \begin{bmatrix} -0.05 & 0.6 \\ -0.6 & -0.05 \end{bmatrix} x, \phi_+ \sim \mathcal{N}(\mu, \Sigma), \kappa = 0.15,$$

with $\mu = [-2, 2]^\intercal$, $\Sigma = 0.1I_2$. The discount factor κ is picked such that the condition for Lemma 1 is satisfied since $-\kappa - \nabla \cdot f(x) = -0.05$ everywhere. Fig. 3 shows the comparison between the ODE computation and the result of kernel density estimation. As the theorem predicts, the bias is almost negligible.

V. Application on a Segway

To demonstrate the proposed method, we apply the proposed method on a segway control example. The control objective of the segway is to move from the initial position to the destination while avoiding an obstacle, which has a round shape. The state of the segway is $x = [s, v, \theta, \dot{\theta}]^\intercal$, where s and v are the longitudinal position and velocity, θ and $\dot{\theta}$ are the pitch angle and angular velocity. The dynamics of the segway can be obtained via the Lagrangian method:

$$\begin{bmatrix} \dot{v} \\ \ddot{\theta} \end{bmatrix} = \begin{bmatrix} M+m & ml\cos(\theta) \\ ml\cos(\theta) & ml^2 + J_e \end{bmatrix}^{-1} \begin{bmatrix} u/r \\ mlg\theta - u \end{bmatrix}, \quad (24)$$

where m is the cart mass, l is the CG height of the cart, r is the wheel radius, g is the gravitational acceleration, M is the equivalent mass of the wheels (accounting for the moment of inertia) and J_e is the moment of inertia of the cart.

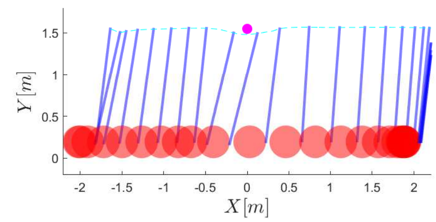

Fig. 4: Simulation of the safe controller

Initially, the segway is equipped with a legacy controller designed with LQR:

$$u_0 = K_1\mathbf{Sat}_\eta(s - s_{des}) + K_2 v + K_3\theta + K_4\dot{\theta}, \quad (25)$$

where K is obtained by solving the Riccati equation for the linearized dynamics, s_{des} is the position of the destination, set to be $s_{des} = 1.8m$, and $\mathbf{Sat}_\eta$ is the symmetric saturation function with range $[-\eta, \eta]$.

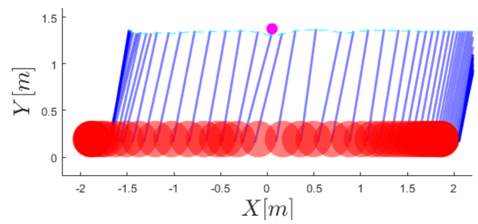

Fig. 5: Experiment run of the segway

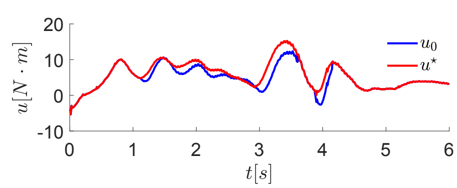

Fig. 6: Input signals during the experiment

The segway has a pole installed on one side, which would collide with the obstacle if following u_0. We follow the procedure presented in Section IV to synthesize a safe controller that avoids the collision with the obstacle, which uses linear equation solvers to solve the HJB PDE and kernel density estimation for the density computation. The cost function follows the same setting in [8]:

$$\mathcal{J} = \int_0^\infty e^{-\kappa t}\,\|u(t) - u_0(t)\|^2\,dt, \quad (26)$$

which penalizes the deviation from u_0.

Let $z(x) \doteq [s+L\sin(\theta+\theta_0)+r\cos(\theta), r+L\cos(\theta+\theta_0)-r\sin(\theta)]^\intercal$ be the position of the tip of the pole, where L is the length of the pole, θ_0 is the original pitch angle of the pole. The danger set $\mathcal{X}_d$ is defined as

$$\mathcal{X}_d = \{x \mid z(x) \in X_{obs} \vee |\theta| > \theta^{\max}\}, \qquad (27)$$

where $X_{obs} \subseteq \mathbb{R}^2$ is the circle area shown in Fig. 4, and $\theta^{\max}$ is the bound for the pitch angle. The set of initial condition is the following

$$\mathcal{X}_0 = [-1.9, 1.8]\times[-0.1, 0.1]\times[-0.01, 0.01]\times[-0.02, 0.02].$$

The initial condition appears with intensity proportional to a Gaussian distribution within $\mathcal{X}_0$

$$\phi_+(x) = \begin{cases} \nu \cdot p(x \mid \mu, \Sigma), & x \in \mathcal{X}_0 \\ 0, & x \notin \mathcal{X}_0 \end{cases},$$

where ν is the intensity, p is the probability density function of a multivariate Gaussian with mean and covariance $\mu = [-1.85, 0, 0, 0]^\intercal$, $\Sigma = diag([0.25, 0.01, 10^{-4}, 4\cdot10^{-4}])$. The primal-dual algorithm terminates after 223 iterations, outputting a safe controller. A simulation run of the obtained controller is shown in Fig. 4. The segway accelerates before reaching the obstacle so that the tip got tilted down to avoid the obstacle. The dotted line shows the trajectory of the pole tip, which indeed avoids the obstacle.

We conducted experiment on a customized segway platform in AMBER Lab at Caltech, as shown in Fig. 1. The Segway frame is a Ninebot Elite E+, with stock motors, and fully custom internals. The wheel encoders and a YOST LX Embedded IMU are connected to a Jetson TX2, which runs the controller on the ERIKA 3 RTOS at 500 hz. The torque commands are then sent through two Elmo Gold Solo Twitters to the motors.

Fig. 5 shows the result of the experiment, where the hardware experiment reproduced the behavior demonstrated in the simulation, avoiding the obstacle. Fig. 6 shows the input signals during the experiment. The red plot denotes the control input of the safe controller $\mathbf{u}^\star$, and the blue plot denotes the original LQR controller $\mathbf{u}_0$. The video of the experiment can be found at https://youtu.be/pdEtknFGu-A.

VI. Conclusion

This paper presents the implementation of the density function approach for safe control synthesis on robotic problems. The primal HJB PDE was solved with a finite element method in the implicit form, which is faster than solving it in the explicit form. Kernel density estimation is used to approximate the density function under a fixed controller, and we showed guarantee on the precision of the approximation. The density function safe synthesis algorithm is applied on a segway control example, and the simulation and experiment result shows the efficacy of the method.

References

[1] A. D. Ames, J. W. Grizzle, and P. Tabuada. Control barrier function based quadratic programs with application to adaptive cruise control. In *53rd IEEE Conference on Decision and Control*, pages 6271–6278. IEEE, 2014.

[2] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada. Control barrier function based quadratic programs for safety critical systems. *IEEE Transactions on Automatic Control*, 62(8):3861–3876, 2017.

[3] R. Bellman. *Dynamic programming*. Courier Corporation, 2013.

[4] F. Blanchini. Set invariance in control. *Automatica*, 35(11):1747–1767, 1999.

[5] H. G. Bock and K.-J. Plitt. A multiple shooting algorithm for direct solution of optimal control problems. *IFAC Proceedings Volumes*, 17(2):1603–1608, 1984.

[6] U. Borrmann, L. Wang, A. D. Ames, and M. Egerstedt. Control barrier certificates for safe swarm behavior. *IFAC-PapersOnLine*, 48(27):68–73, 2015.

[7] A. Bressan. Viscosity solutions of hamilton-jacobi equations and optimal control problems. *Lecture notes*, 2011.

[8] Y. Chen, M. Ahmadi, and A. D. Ames. Optimal safe controller synthesis: A density function approach. *arXiv preprint arXiv:1909.11798*, 2019.

[9] Y. Chen and A. D. Ames. Duality between density function and value function with applications in constrained optimal control and markov decision process. *arXiv preprint arXiv:1902.09583*, 2019.

[10] Y. Chen, H. Peng, and J. Grizzle. Obstacle avoidance for low-speed autonomous vehicles with barrier function. *IEEE Transactions on Control Systems Technology*, 26(1):194–206, 2017.

[11] Y. Chen, H. Peng, J. Grizzle, and N. Ozay. Data-driven computation of minimal robust control invariant set. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 4052–4058. IEEE, 2018.

[12] Y. Chen, H. Peng, and J. W. Grizzle. Fast trajectory planning and robust trajectory tracking for pedestrian avoidance. *Ieee Access*, 5:9304–9317, 2017.

[13] R. F. Hartl, S. P. Sethi, and R. G. Vickson. A survey of the maximum principles for optimal control problems with state constraints. *SIAM review*, 37(2):181–218, 1995.

[14] O. Khatib. Real-time obstacle avoidance for manipulators and mobile robots. In *Autonomous robot vehicles*, pages 396–404. Springer, 1986.

[15] J. J. Kuffner Jr and S. M. LaValle. Rrt-connect: An efficient approach to single-query path planning. In *ICRA*, volume 2, 2000.

[16] J. B. Lasserre, D. Henrion, C. Prieur, and E. Trélat. Nonlinear optimal control via occupation measures and lmi-relaxations. *SIAM journal on control and optimization*, 47(4):1643–1666, 2008.

[17] A. Majumdar, R. Vasudevan, M. M. Tobenkin, and R. Tedrake. Convex optimization of nonlinear feedback controllers via occupation measures. *The International Journal of Robotics Research*, 33(9):1209–1230, 2014.

[18] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, 2000.

[19] I. M. Mitchell, A. M. Bayen, and C. J. Tomlin. A time-dependent hamilton-jacobi formulation of reachable sets for continuous dynamic games. *IEEE Transactions on automatic control*, 50(7):947–957, 2005.

[20] L. S. Pontryagin. *Mathematical theory of optimal processes*. Routledge, 2018.

[21] S. Prajna, P. A. Parrilo, and A. Rantzer. Nonlinear control synthesis by convex optimization. *IEEE Transactions on Automatic Control*, 49(2):310–314, 2004.

[22] A. Rantzer. A dual to lyapunov's stability theorem. *Systems & Control Letters*, 42(3):161–168, 2001.

[23] B. Sun and B.-Z. Guo. Convergence of an upwind finite-difference scheme for hamilton–jacobi–bellman equation in optimal control. *IEEE Transactions on Automatic Control*, 60(11):3012–3017, 2015.

[24] L. Wang, A. D. Ames, and M. Egerstedt. Safety barrier certificates for collisions-free multirobot systems. *IEEE Transactions on Robotics*, 33(3):661–674, 2017.

[25] L. Wasserman. Lecture notes in statistical methods for machine learning. Online material, last visited on 2019/08/25.

[26] P. Zhao, S. Mohan, and R. Vasudevan. Control synthesis for nonlinear optimal control via convex relaxations. In *2017 American Control Conference (ACC)*, pages 2654–2661. IEEE, 2017.

[27] P. Zhao, S. Mohan, and R. Vasudevan. Optimal control for nonlinear hybrid systems via convex relaxations. *arXiv preprint arXiv:1702.04310*, 2017.