

Risk-Averse Control via CVaR Barrier Functions: Application to Bipedal Robot Locomotion

Mohamadreza Ahmadi, Xiaobin Xiong, and Aaron D. Ames

Abstract—Enforcing safety in the presence of stochastic uncertainty is a challenging problem. Traditionally, researchers have proposed safety in the statistical mean as a safety measure for systems subject to stochastic uncertainty. However, ensuring safety in the statistical mean is only reasonable if system’s safe behavior in the large number of runs is of interest, which precludes the use of mean safety in practical scenarios. In this paper, we propose a risk sensitive notion of safety called conditional-value-at-risk (CVaR) safety. We introduce CVaR barrier functions as a tool to enforce CVaR-safety and propose conditions for their Boolean compositions. Given a legacy controller, we show that we can design a *minimally interfering CVaR-safe controller* via solving difference convex programs (DCPs). We elucidate the proposed method by applying it to a bipedal robot locomotion case study.

Index Terms—Autonomous systems, Robotics, Stochastic systems, Uncertain systems.

I. INTRODUCTION

With the rise of autonomous systems being deployed in real-world settings, the associated risk that stems from uncertain and unforeseen circumstances is correspondingly on the rise. For instance, there are several inherent sources of uncertainty in robotics systems, such as modeling uncertainty, sensor range and resolution limitations, dynamic and uncertain environments, noise and wear-and-tear in robot actuation [1], that lead to higher risk during deployment.

Mathematically speaking, risk can be quantified in numerous ways, such as chance constraints [2], [3]. However, applications in autonomy and robotics require more “nuanced assessments of risk” [4]. Artzner *et. al.* [5] characterized a set of natural properties that are desirable for a risk measure, called a coherent risk measure, and have obtained widespread acceptance in finance and operations research, among other fields. An important example of a coherent risk measure is the conditional value-at-risk (CVaR) that has received significant attention in decision making problems, such as Markov decision processes (MDPs) [6], [7], [8]. For stochastic discrete-time dynamical systems, a model predictive control technique with coherent risk objectives was proposed in [9], wherein the authors also proposed Lyapunov conditions for risk-sensitive exponential stability. Moreover, a method based on stochastic reachability analysis was proposed in [10] to estimate a CVaR-safe set of initial conditions via the solution to an MDP.

The authors are with Control & Dynamical Systems (CDS), California Institute of Technology, Pasadena, CA 91125. email: ({mrahmadi, xxiong, ames}@caltech.edu).




Fig. 1. Risk-averse obstacle avoidance using CVaR barrier functions (robot behavior and barrier function evolution). The shaded yellow area denotes safe regions. (a) safety violation with no barrier function; (b) safety violation with risk-neutral barrier function; (c) safe behavior with CVaR barrier function. Plots on the right side show the values of the barrier functions.

In this work, we use a special class of barrier functions as a tool for enforcing risk-sensitive safety. Control barrier functions have been proposed in [11] and have been used for designing safe controllers (in the absence of a legacy controller, *i.e.*, a desired controller that may be unsafe) and safety filters (in the presence of a legacy controller) for continuous-time dynamical systems, such as bipedal robots [12], with guaranteed robustness [13], [14]. For discrete-time systems, discrete-time barrier functions have been formulated in [15], [16] and applied to the multi-robot coordination problem [17]. Recently, for a class of stochastic (Ito) differential equations, safety in probability and statistical mean were also studied in [18], [19], [20].

In this paper, we go beyond the conventional notions of safety in probability and statistical mean for discrete-time systems subject to stochastic uncertainty. To this end, we define safety in the risk-sensitive CVaR sense, which is concerned with safety in the worst possible scenarios. We then propose CVaR barrier functions as a tool to enforce CVaR-safety and formulate conditions for their Boolean compositions. We propose a computational method based on difference convex programs (DCPs) to synthesize CVaR-safe controllers for stochastic linear discrete-time systems. These CVaR-safe controllers are designed such that they minimally interfere with a given legacy controller. We show the efficacy of our proposed method on collision avoidance scenarios involving a bipedal robot subject to modeling uncertainty (see Figure 1).




Fig. 2. Comparison of the mean, VaR, and CVaR for a given confidence $\beta \in (0, 1)$. The axes denote the values of the stochastic variable h and its pdf $p(h)$. The shaded area denotes the $\% \beta$ of the area under $p(h)$. If $h \geq 0$ represents a safe behavior, using $\mathbb{E}(h)$ as a performance measure is misleading (note that $\mathbb{E}(h)$ is positive). VaR gives the value of h at the β -tail of the distribution. But, it ignores the values of h with probability below β . CVaR is the average of the values of VaR with probability less than β (average of the worst-case values of h in the β tail of the distribution). Note that CVaR_β is negative indicating unsafe behavior.

The rest of the paper is organized as follows. In the next section, we introduce CVaR-safety and formulate CVaR barrier functions as a tool to synthesize risk-averse safe controllers. In Section III, we discuss our robot bipedal locomotion case study and present the obtained results. Finally, in Section V, we conclude the paper and give directions for future research.

Notation: We denote by \mathbb{R}^n the n -dimensional Euclidean space and $\mathbb{N}_{\geq 0}$ the set of non-negative integers. For a finite set \mathcal{A} , we denote by $|\mathcal{A}|$ the number of elements of \mathcal{A} . For a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, $\mathcal{L}_1(\Omega, \mathcal{F}, \mathbb{P})$. The Boolean operators are denoted by \neg (negation), \vee (conjunction), and \wedge (disjunction).

II. CVAR BARRIER FUNCTIONS FOR RISK-AVERSE CONTROL

In this section, we formulate the risk-averse safety problem and propose a solution based on a special class of barrier functions. We begin by describing the class of systems under study.

A. System Description

We consider Markov control processes of the form

$$x^{t+1} = f(x^t, u^t, w^t), \quad x^0 = x_0, \quad (1)$$

where $t \in \mathbb{N}_{\geq 0}$ denotes the time index, $x^t \in \mathcal{X}$ is the state, $u^t \in \mathcal{U}$ is the control input, $w^t \in \mathcal{W}$ is the uncertainty/disturbance, and $f: \mathcal{X} \times \mathcal{U} \times \mathcal{W} \rightarrow \mathcal{X}$ is a (possibly nonlinear) Borel-measurable function. We assume that the sets \mathcal{X} and \mathcal{U} are compact Borel subsets of \mathbb{R}^n and \mathbb{R}^m , respectively. Moreover, the stochastic process $\{w^t\}$ is i.i.d. and is drawn from a finite sample space $\mathcal{W} = \{w_1, \dots, w_{|\mathcal{W}|}\}$ with the probability mass function $p(w_i) := \mathbb{P}(w^t = w_i)$, $i = 1, 2, \dots, |\mathcal{W}|$, and w^t is independent of the deterministic initial condition x_0 and the state-control pairs (x^τ, u^τ) , for all $\tau \leq t$.

We are interested in studying the properties of the solutions to (1) with respect to the compact set \mathcal{S} described as

$$\mathcal{S} := \{x \in \mathcal{X} \mid h(x) \geq 0\}, \quad (2)$$

where $h: \mathcal{X} \rightarrow \mathbb{R}$ is a continuous function. For instance, \mathcal{S} can represent robot constraints, e.g. joint limits, safe exploration regions, etc.

B. Conditional Value-at-Risk Safety

In the presence of the stochastic uncertainty w , assuring almost sure (with probability one) invariance or safety may not be feasible. Moreover, enforcing safety in expectation is only meaningful if the law of large numbers can be invoked and we are interested in the long term performance, independent of the realization fluctuations.

At any time step $t \neq 0$, from the compactness of \mathcal{X} and continuity of h , it follows that the random variable $h^t := h(x^t) \in \mathcal{H}^t \subset \mathcal{L}_1(\mathcal{W}^t, \mathcal{F}^t, \mathbb{P}^t)$ is bounded, where the sample space \mathcal{W}^t is the t times Cartesian product of \mathcal{W} with itself, and \mathcal{F}^t and \mathbb{P}^t are the corresponding filtration and probability measures. From the description of the safe set (2), we infer that h is a continuous random variable for which higher values are of interest (for example, greater values of h indicate safer performance).

For a given confidence level $\beta \in (0, 1)$, value-at-risk (VaR_β) denotes the β -quantile value of $h^t \in \mathcal{H}^t$ described as $\text{VaR}_\beta(h^t) = \sup_{\zeta \in \mathbb{R}} \{\zeta \mid \mathbb{P}^t(h^t \leq \zeta) \leq \beta\}$.

Unfortunately, working with VaR for non-normal stochastic variables is numerically unstable, optimizing models involving VaR are intractable in high dimensions, and VaR ignores the values of h^t with probability less than β [21].

In contrast, CVaR overcomes the shortcomings of VaR. CVaR with confidence level $\beta \in (0, 1)$ denoted CVaR_β measures the expected loss in the β -tail given that the particular threshold VaR_β has been crossed, i.e., $\text{CVaR}_\beta(h^t) = \mathbb{E}[h^t \mid h^t \leq \text{VaR}_\beta(h^t)]$. An optimization formulation for CVaR was proposed in [21] that we use in this paper. That is, CVaR_β is given by

$$\text{CVaR}_\beta(h^t) := \inf_{\zeta \in \mathbb{R}} \mathbb{E} \left[\zeta + \frac{(-h^t - \zeta)_+}{\beta} \right], \quad (3)$$

where $(\cdot)_+ = \max\{\cdot, 0\}$. A value of $\beta \rightarrow 1$ corresponds to a risk-neutral case, i.e., $\text{CVaR}_1(h^t) = \mathbb{E}(h^t)$; whereas, a value of $\beta \rightarrow 0$ is rather a risk-averse case, i.e., $\text{CVaR}_0(h^t) = \text{VaR}_0(h^t)$ [22, Proposition 6]. Figure 2 illustrates these notions for an example h variable with distribution $p(h)$.

Unlike VaR, CVaR is a coherent risk measure on \mathcal{H} [23], which satisfies the following properties [24].

Definition 1 (Coherent Risk Measure): We call the one-step conditional risk measures $\rho^t: \mathcal{H}^{t+1} \rightarrow \mathcal{H}^t$, $t = 1, \dots, N-1$ a *coherent risk measure*, if it satisfies the following conditions

- **Convexity:** $\rho^t(\lambda h + (1-\lambda)h') \leq \lambda \rho^t(h) + (1-\lambda)\rho^t(h')$, for all $\lambda \in (0, 1)$ and all $h, h' \in \mathcal{H}^{t+1}$;
- **Monotonicity:** If $h \leq h'$, then $\rho^t(h) \leq \rho^t(h')$ for all $h, h' \in \mathcal{H}^{t+1}$;
- **Translational Invariance:** $\rho^t(h' + h) = \rho^t(h') + h$ for all $h \in \mathcal{H}^t$ and $h' \in \mathcal{H}^{t+1}$;
- **Positive Homogeneity:** $\rho^t(\beta h) = \beta \rho^t(h)$ for all $h \in \mathcal{H}^{t+1}$ and $\beta \geq 0$.

In fact, we use the nice mathematical properties of CVaR given in Definition 1 in the proofs of our main results in Section II-C.

In this work, we propose safety in a dynamic coherent risk measure, namely, CVaR sense, with conditional expectation

(risk-neutral case) as a special case $\beta \rightarrow 1$. Let $\text{CVaR}_\beta^t : \mathcal{H}^{t+1} \rightarrow \mathcal{H}^t$ consistent with Definition 1. We then define the dynamic CVaR measure $\text{CVaR}_\beta^{0:t} := \text{CVaR}_\beta^0 \circ \dots \circ \text{CVaR}_\beta^t : \mathcal{H}^{t+1} \rightarrow \mathcal{H}^0$. Since x^0 is deterministic, we have $\mathcal{H}^0 = \mathbb{R}$.

Definition 2 (CVaR-safety): Given a safe set \mathcal{S} as given in (2) and a confidence level $\beta \in (0, 1)$, we call the solutions to (1) starting at $x_0 \in \mathcal{S}$ CVaR-safe if and only if

$$\text{CVaR}_\beta^{0:t}(h(x^t)) \geq 0, \quad \forall t \geq 0. \quad (4)$$

Note that $\text{CVaR}_\beta^{0:t}$ is also a time-consistent risk measure [24, Definition 3]. The time consistency property ensures that contradictory evaluations of safety risk at different points in time does not happen.

C. CVaR Barrier Functions

In order to check and enforce CVaR-safety, we define CVaR barrier functions.

Definition 3 (CVaR Barrier Function): For the discrete-time system (1) and a confidence level $\beta \in (0, 1)$, the continuous function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ is a CVaR barrier function for the set \mathcal{S} as defined in (2), if there exists a constant $\alpha \in (0, 1)$ such that

$$\text{CVaR}_\beta^t(h(x^{t+1})) \geq \alpha h(x^t), \quad \forall x^t \in \mathcal{X}. \quad (5)$$

In the next result, we demonstrate that the existence of a CVaR barrier function indeed implies CVaR-safety.

Theorem 1: Consider the discrete-time system (1) and the set \mathcal{S} as described in (2). Let $\beta \in (0, 1)$ be a given confidence level. Then, \mathcal{S} is CVaR-safe, if there exists an CVaR barrier function as defined in Definition 3.

Proof: The proof is carried out by induction and using the properties of CVaR as a coherent risk measure as outlined in Definition 1. If (5) holds, for $t = 0$, we have

$$\text{CVaR}_\beta^0(h(x^1)) \geq \alpha h(x_0). \quad (6)$$

Similarly, for $t = 1$, we have

$$\text{CVaR}_\beta^1(h(x^2)) \geq \alpha h(x_1). \quad (7)$$

Since CVaR_β^0 is monotone (because it is a coherent risk measure), composing both sides of (7) with CVaR_β^0 does not change the inequality and we obtain

$$\text{CVaR}_\beta^0 \circ \text{CVaR}_\beta^1(h(x^2)) \geq \text{CVaR}_\beta^0(\alpha h(x_1)). \quad (8)$$

Since $\alpha \in (0, 1)$, from positive homogeneity property of CVaR, we obtain $\text{CVaR}_\beta^0(\alpha h(x_1)) = \alpha \text{CVaR}_\beta^0(h(x_1))$. Thus, (8) simplifies to $\text{CVaR}_\beta^0 \circ \text{CVaR}_\beta^1(h(x^2)) \geq \alpha \text{CVaR}_\beta^0(h(x_1))$.

Then, using inequality (6), we have

$$\text{CVaR}_\beta^0 \circ \text{CVaR}_\beta^1(h(x^2)) \geq \alpha \text{CVaR}_\beta^0(h(x_1)) \geq \alpha^2 h(x_0).$$

Therefore, by induction, at time t , we can show that

$$\text{CVaR}_\beta^{0:t}(h(x^t)) \geq \alpha^t h(x_0). \quad (9)$$

If $x_0 \in \mathcal{S}$, from the definition of the set \mathcal{S} , we have $h(x_0) \geq 0$. Since $\alpha \in (0, 1)$, then we can infer that (4) holds. Thus, the system is CVaR $_\beta$ -safe. ■

In many practical path planning scenarios, we encounter multiple obstacles and safe sets composed of Boolean compositions of several barrier functions [25], [17], [26]. Next, we propose conditions for checking Boolean compositions of CVaR barrier functions.

Proposition 1: Let $\mathcal{S}_i = \{x \in \mathbb{R}^n \mid h_i(x) \geq 0\}$, $i = 1, \dots, k$ denote a family of safe sets with the boundaries and interior defined analogous to \mathcal{S} in (2). Consider the discrete-time system (1). If there exist a $\alpha \in (0, 1)$ such that

$$\text{CVaR}_\beta^t \left(\min_{i=1, \dots, k} h_i(x^{t+1}) \right) \geq \alpha \min_{i=1, \dots, k} h_i(x^t) \quad (10)$$

then the set $\{x \in \mathbb{R}^n \mid \bigwedge_{i=1, \dots, k} (h_i(x) \geq 0)\}$ is CVaR-safe. Similarly, if there exist a $\alpha \in (0, 1)$ such that

$$\text{CVaR}_\beta^t \left(\max_{i=1, \dots, k} h_i(x^{t+1}) \right) \geq \alpha \max_{i=1, \dots, k} h_i(x^t) \quad (11)$$

then the set $\{x \in \mathbb{R}^n \mid \bigvee_{i=1, \dots, k} (h_i(x) \geq 0)\}$ is CVaR-safe.

Proof: If (10) holds from the proof of Theorem 1, we can infer that

$$\text{CVaR}_\beta^{0:t} \left(\min_{i=1, \dots, k} h_i(x^t) \right) \geq \alpha^t \min_{i=1, \dots, k} h_i(x^0).$$

That is, if $x^0 \in \{x \in \mathbb{R}^n \mid \min_{i=1, \dots, k} h_i(x) \geq 0\}$, then $\text{CVaR}_\beta^{0:t}(\min_{i=1, \dots, k} h_i(x^t)) \geq 0$ for all $t \in \mathbb{N}_{\geq 0}$. Let $h_{i^*}(x^t)$ be the smallest among $h_i(x^t)$, $i = 1, 2, \dots, k$, i.e., it satisfies $h_j(x^t) \geq \dots \geq h_{i^*}(x^t)$, $\forall j \neq i^*$. Because CVaR is monotone (see Definition 1), the latter inequality implies $\text{CVaR}_\beta^{0:t}(h_j(x^t)) \geq \dots \geq \text{CVaR}_\beta^{0:t}(h_{i^*}(x^t))$, $\forall j \neq i^*$ (because the composition of monotone functions remains monotone). Since $\text{CVaR}_\beta^{0:t}(\min_{i=1, \dots, k} h_i(x^t)) = \text{CVaR}_\beta^{0:t}(h_{i^*}(x^t)) \geq 0$ for all $t \in \mathbb{N}_{\geq 0}$, we have

$$\text{CVaR}_\beta^{0:t}(h_j(x^t)) \geq \dots \geq \text{CVaR}_\beta^{0:t}(h_{i^*}(x^t)) \geq 0, \quad j \neq i^*.$$

Thus, $\text{CVaR}_\beta^{0:t}(h_i(x)) \geq 0$ for all $i \in \{1, \dots, k\}$.

Similarly, if (11) holds, we can infer that

$$\text{CVaR}_\beta^{0:t} \left(\max_{i=1, \dots, k} h_i(x^t) \right) \geq \alpha^t \max_{i=1, \dots, k} h_i(x^0).$$

Hence, using similar arguments as the proof of the conjunction case, $\text{CVaR}_\beta^{0:t}(\max_{i=1, \dots, k} h_i(x^t)) \geq 0$ for all $t \in \mathbb{N}_{\geq 0}$. That is, there exists at least an $i \in \{1, \dots, k\}$ for which $\text{CVaR}_\beta^{0:t}(\max_{i=1, \dots, k} h_i(x^t)) \geq 0$. ■

The negation operator is trivial and can be shown by checking if $-h$ satisfies inequality (5).

In the next section, we demonstrate how a sequence $\{u^t\}_{t>0}$ can be designed such that system (1) becomes CVaR-safe based on optimization techniques.

D. CVaR-Safe Controller Synthesis

Inspired by the quadratic programming formulations of conventional control barrier functions in the continuous-time case [11], we pose the controller synthesis problem as an optimization.

CVaR Control Barrier Function Optimization: At every time step t , given x^t , a set \mathcal{S} as described in (2), a confidence level $\beta \in (0, 1)$, a parameter $\alpha \in (0, 1)$, control upper bounds \bar{u} , lower bounds \underline{u} , and a legacy controller u_{legacy}^t , solve

$$u_*^t = \underset{u^t}{\operatorname{argmin}} (u^t - u_{\text{legacy}}^t)^T (u^t - u_{\text{legacy}}^t) \quad (12a)$$

subject to $\underline{u} \leq u^t \leq \bar{u}$,

$$\operatorname{CVaR}_\beta^t(h(f(x^t, u^t, w))) \geq \alpha h(x^t). \quad (12b)$$

Note that instantaneous controls u^t are the only variables in the optimization. The cost function $(u^t - u_{\text{legacy}}^t)^T (u^t - u_{\text{legacy}}^t) = \|u^t - u_{\text{legacy}}^t\|^2$ ensures that u^t remains as close as possible to the legacy controller u_{legacy}^t in the Euclidean 2-norm; hence, it guarantees the *minimal interference*.

For general nonlinear h , optimization problem (12) is a nonlinear program in the decision variable u^t (note that CVaR is a convex function in h since it is a coherent risk measure). Indeed, this was the case for optimization problems designed for synthesizing discrete control barrier functions for discrete-time systems even without stochastic uncertainty [16], as well. MATLAB functions such as `fmincon` can be used to solve the nonlinear program.

Next, we show that under some assumptions the search over CVaR-safe controls u^t can be carried out by solving DCPs. For the remainder of this section, we restrict our attention to the case when system (1) is a linear system. That is,

$$f(x^t, u^t, w^t) = A(w^t)x^t + B(w^t)u^t + G(w^t),$$

where $A : \mathcal{W} \rightarrow \mathbb{R}^{n \times n}$, $B : \mathcal{W} \rightarrow \mathbb{R}^{n \times m}$ and $G : \mathcal{W} \rightarrow \mathbb{R}^n$.

For such systems, we assume the CVaR barrier function takes the form of a linear function

$$h(x^t) = Hx^t + l, \quad (13)$$

where $H \in \mathbb{R}^{1 \times n}$ and $l \in \mathbb{R}$. Then, the term $\operatorname{CVaR}_\beta^t(h(f(x^t, u^t, w)))$ in constraints (12b) changes to

$$\operatorname{CVaR}_\beta^t(HA(w)x^t + HB(w)u^t + HG(w) + l). \quad (14)$$

Since $\operatorname{CVaR}_\beta$ is a convex function, the above term is a convex function in u^t , *i.e.*, the control variable.

Re-writing optimization problem (12) for linear discrete time systems with stochastic uncertainty and CVaR barrier function (13) gives the following optimization problem

$$u_*^t = \underset{u^t}{\operatorname{argmin}} (u^t - u_{\text{legacy}}^t)^T (u^t - u_{\text{legacy}}^t) \quad (15a)$$

subject to $\underline{u} \leq u^t \leq \bar{u}$,

$$(14) \geq \alpha h(x^t), \quad (15b)$$

Substituting the expression for CVaR (3) in (14) for uncertainty w with finite $|\mathcal{W}|$ yields

$$\inf_{\zeta \in \mathbb{R}} \left\{ \zeta + \frac{1}{\beta} \sum_{i=1}^{|\mathcal{W}|} (-HA(w_i)x^t - HB(w_i)u^t - HG(w_i) - l - \zeta)_+ p(w_i) \right\}, \quad (16)$$

which introduces the extra decision variable $\zeta \in \mathbb{R}$.

Hence, (15) can be rewritten in the standard DCP form

$$u_*^t = \underset{u^t, \zeta}{\operatorname{argmin}} q_0(u^t) \quad (17a)$$

subject to $q_1(u^t) \leq 0$ and $q_2(u^t) \leq 0$,

$$q_3 - \hat{q}_4(\zeta, u^t) \leq 0, \quad (17b)$$

where $q_0(u^t) = (u^t - u_{\text{legacy}}^t)^T (u^t - u_{\text{legacy}}^t)$ is a convex (quadratic) function, $q_1(u^t) = \underline{u} - u^t$ is a convex (linear) function, $q_2(u^t) = u^t - \bar{u}$ is a convex (linear) function, and $q_3 = \alpha Hx^t + \alpha l$ is a convex (constant) function. The expression for $\hat{q}_4(\zeta, u^t)$ is given inside the braces in (16) which is a convex function in u^t and ζ since $\hat{q}_4(\zeta, u^t)$ is convex in ζ [21, Theorem 1] because the function $(\cdot)_+$ is increasing and convex [27, Lemma A.1.].

DCPs like (17) arise in many applications, such as risk-averse MDPs [28] and inverse covariance estimation in statistics [29]. In order to solve DCPs, we use a variant of the convex-concave procedure [30], wherein the concave terms are replaced by a convex upper bound and solved. In fact, the disciplined convex-concave programming (DCCP) [30] technique linearizes DCP problems into a (disciplined) convex program (carried out automatically via the DCCP package [30]). Then, the cone program can be solved readily by available solvers, such as YALMIP [31].

We remark that solving (17) via the DCCP method, finds the (local) saddle points to optimization problem (17). Nonetheless, every such local u^t guarantees CVaR-safety.

III. CASE STUDY: BIPEDAL ROBOT LOCOMOTION

Planning and controlling of bipedal walking are challenging problems, and there has been various related approaches [32] in the literature. In this paper, we apply the approach in [33] to approximate the bipedal walking dynamics via a discrete linear system subject to stochastic uncertainty.

Let $\mathbf{x}_h = [c, p, v]^T$ denote the horizontal state, where c is the horizontal position of the center of mass (COM) of the robot relative to the inertia frame, p is the horizontal position of the COM relative to its stance foot, and v is the horizontal velocities of the COM. Then, the horizontal step-to-step (S2S) dynamics [33] of the horizontal COM state is represented as

$$\mathbf{x}_h^{t+1} = \mathcal{P}^h(x^t, \tau(t)), \quad (18)$$

where x is the full robot state and τ is the input torque. However, the S2S dynamics \mathcal{P}^h (18) cannot be obtained in analytical form due to the nonlinear and hybrid dynamics of the robot walking.

Our previous work [33] suggests that the S2S dynamics of the walking of the Hybrid-Linear Inverted Pendulum (H-LIP)

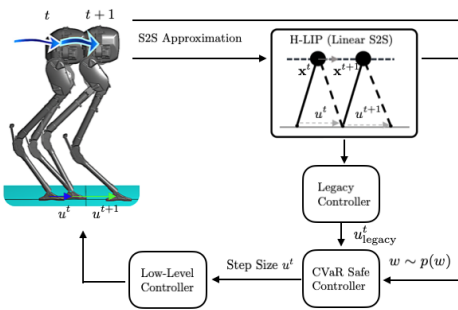


Fig. 3. A schematic diagram of our proposed risk-averse control method based on CVaR barrier functions for bipedal robots.

can be used to approximate the actual horizontal S2S dynamics of the walking of the robot in Eq. (18). The S2S dynamics of the H-LIP [33] is:

$$\mathbf{x}_{\text{H-LIP}}^{t+1} = A\mathbf{x}_{\text{H-LIP}}^t + B u_{\text{H-LIP}}^t \quad (19)$$

where $\mathbf{x}_{\text{H-LIP}}^{t+1} = [c_{\text{H-LIP}}, p_{\text{H-LIP}}, v_{\text{H-LIP}}]^T$ is the discrete pre-impact state of the H-LIP, and $u_{\text{H-LIP}}^t$ is the step size. The expressions of A, B can be found in [33]. By approximation, (18) can be rewritten as

$$\mathbf{x}_h^{t+1} = A\mathbf{x}_h^t + B u^t + w^t, \quad (20)$$

where $w^t := \mathcal{P}^h(x^t, \tau(t)) - A\mathbf{x}_h^t - B u^t \in \mathcal{W}$ can be treated as a stochastic disturbance to the linear system in (19).

For application of 3D bipedal walking, the H-LIP model is applied in each plane of walking: the sagittal and the lateral planes. The H-LIP based planning provides the desired step sizes for the robot, which become the desired outputs for the low-level controller to track [33].

A. Risk-Averse Bipedal Robot Control

We apply the CVaR barrier function based risk-averse control presented in Section II to the 3D bipedal walking as described in Fig. 3. Model discrepancy w is treated as stochastic uncertainty and as a risk factor that can lead to undesired behavior on the generated walking.

To circumvent this issue, we synthesize CVaR barrier functions based controllers to filter the H-LIP based stepping controller on the robot. The barrier functions are designed to represent the safe regions, where there are no obstacles. We are then interested in keeping the robot inside the safe (obstacle-free) regions during walking.

The uncertainty w is numerically calculated by several hours of simulations of different walking behaviors, which provided a polytopic set that bounds w . We took $|\mathcal{W}|$ random samples from the latter polytopic set. Since w is sparse in nature, we assumed a uniform distribution of w inside \mathcal{W} , i.e., $p(w) = 1/|\mathcal{W}|$. To design the risk-averse safe controllers, we then solve DCP (17), where A, B , and $G(w) = w$ are given by the approximated S2S dynamics (20).

B. Simulation Results

We apply the proposed approach in high-fidelity simulation on the under-actuated bipedal robot Cassie. DCP (17) is solved in YALMIP using MOSEK solver at each step. The

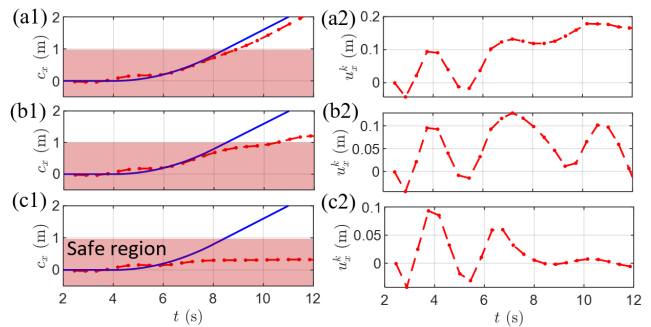


Fig. 4. Case 1: Trajectories of the positions (blue is the desired trajectory) and step length of the robot in the sagittal plane for (a) walking without CVaR barrier function, where (a2) shows the legacy controller input, (b1) walking with risk neutral CVaR barrier function with $\beta = 0.999$, where (b2) shows the output of the CVaR-safe controller with $\beta = 0.999$, (c) walking with the risk-averse CVaR barrier function with $\beta = 0.1$, where (c2) shows the output of the CVaR-safe controller with $\beta = 0.1$. The red shaded area indicates the safe region for the robot.

optimization typically takes 100 ~ 700 steps under 10 seconds to solve on a laptop with the processor intel(R) Core(TM) i7-7700HQ@2.8GHz. The low-level controller on the robot is solved at 1kHz. The legacy controller used in our experiments is a model predictive controller. The simulation video of all the experiments can be found at <https://youtu.be/QNMW1zey3cI>.

Case 1: We consider a scenario where the robot is following a straight path and an obstacle is placed in this path. The results are shown in Fig. 1 (a) and Fig. 4. The legacy controller is not aware of this wall, which results in collision that in practice would cause hardware failure. Then, we apply a CVaR barrier function to filter the output of the legacy controller. The safe set is defined as

$$h(c_x) = p_x - c_x \geq 0,$$

where $p_x = 1$ is the position of the obstacle, c_x denotes the position of the robot in the forward direction. We first apply the CVaR barrier function with $\beta = 0.999$ (risk neutral) risk-averse controller. The result is shown in Fig. 1 (b): the robot walks and stop at the location of the obstacle. However, due the stochastic uncertainty w , the risk-neutral controller violates the safety requirement. Lastly, we apply the CVaR barrier function with $\beta = 0.1$ (risk-averse case), which generates the walking in Fig. 1 (c). The legacy controller directs the robot forward, but the CVaR-safe controller keeps the robot away from the obstacle.

Case 2: In this scenario, we consider the robot following a forward reference path. However, there is a wall at an angle, which does not completely prevent the robot from walking forward. The safe set is defined as

$$h(c_x, c_y) = c_y + k(c_x - p) \geq 0,$$

where k indicates the angle of the wall, p indicates the location of the wall in forward direction, and c_y is the position of the robot in the lateral plane. Here $k = -0.5$ and $p = 2$. Fig. 5 (a) shows the generated walking behavior. With the CVaR barrier function with $\beta = 0.5$, the robot keeps a distance from the wall and maintains its original forward walking behavior in its sagittal plane, which is similar to the walking in Fig.




Fig. 5. The generated walking for Case 2 (a1, a2) and Case 3 (b).

4 (a1). As a result, the robot also walks laterally as well to assure CVaR-safety.

Case 3: We consider a scenario with multiple barrier functions. The robot is supposed to follow a sinusoidal path. We add two walls on its way. The safe set is then defined as $\min(h_1, h_2) \geq 0$, where

$$h_1(c_y) = c_y + p_1 \geq 0, \quad h_2(c_y) = -c_y + p_2 \geq 0,$$

with $p_1 = 2$ and $p_2 = 0$. Fig. 5 (b) illustrates the walking with the CVaR barrier function $\beta = 0.5$, where the robot successfully avoided the collision with the walls.

IV. CONCLUSION

We proposed a method based on CVaR barrier functions to verify and enforce risk-averse safety for discrete-time stochastic systems. We proposed a computational method for synthesizing CVaR-safe controllers in the case of linear dynamics. The method was applied to enforce risk-averse safety of a bipedal robot. Future work will extend the CVaR barrier functions to other coherent risk measures, continuous-time systems, and applications involving cooperative human-robot teams and imperfect sensor measurements [34] and convex/polytopic approximations of barrier functions leading to convex programs.

REFERENCES

- [1] S. Thrun, W. Burgard, and D. Fox, *Probabilistic robotics*. Cambridge, Mass.: MIT Press, 2005.
- [2] M. Ono, M. Pavone, Y. Kuwata, and J. Balam, "Chance-constrained dynamic programming with application to risk-aware robotic space exploration," *Autonomous Robots*, vol. 39, no. 4, pp. 555–571, 2015.
- [3] A. Wang, A. M. Jasour, and B. Williams, "Non-Gaussian chance-constrained trajectory planning for autonomous vehicles under agent uncertainty," *IEEE Robotics and Automation Letters*, 2020.
- [4] A. Majumdar and M. Pavone, "How should a robot assess risk? towards an axiomatic theory of risk in robotics," in *Robotics Research*. Springer, 2020, pp. 75–84.
- [5] P. Artzner, F. Delbaen, J. Eber, and D. Heath, "Coherent measures of risk," *Mathematical finance*, vol. 9, no. 3, pp. 203–228, 1999.
- [6] Y. Chow, A. Tamar, S. Mannor, and M. Pavone, "Risk-sensitive and robust decision-making: a CVaR optimization approach," in *Advances in Neural Information Processing Systems*, 2015, pp. 1522–1530.
- [7] Y. Chow and M. Ghavamzadeh, "Algorithms for CVaR optimization in MDPs," in *Advances in neural information processing systems*, 2014, pp. 3509–3517.
- [8] N. Bäuerle and J. Ott, "Markov decision processes with average-value-at-risk criteria," *Mathematical Methods of Operations Research*, vol. 74, no. 3, pp. 361–379, 2011.
- [9] S. Singh, Y. Chow, A. Majumdar, and M. Pavone, "A framework for time-consistent, risk-sensitive model predictive control: Theory and algorithms," *IEEE Transactions on Automatic Control*, 2018.
- [10] M. P. Chapman, J. Lacotte, A. Tamar, D. Lee, K. M. Smith, V. Cheng, J. F. Fisac, S. Jha, M. Pavone, and C. J. Tomlin, "A risk-sensitive finite-time reachability approach for safety of stochastic dynamic systems," in *2019 American Control Conference (ACC)*. IEEE, 2019, pp. 2958–2963.
- [11] 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*, vol. 62, no. 8, pp. 3861–3876, 2016.
- [12] Q. Nguyen, A. Hereid, J. W. Grizzle, A. D. Ames, and K. Sreenath, "3d dynamic walking on stepping stones with control barrier functions," in *2016 IEEE 55th Conference on Decision and Control (CDC)*. IEEE, 2016, pp. 827–834.
- [13] X. Xu, P. Tabuada, J. W. Grizzle, and A. D. Ames, "Robustness of control barrier functions for safety critical control," *IFAC-PapersOnLine*, vol. 48, no. 27, pp. 54–61, 2015.
- [14] S. Kolathaya and A. D. Ames, "Input-to-state safety with control barrier functions," *IEEE control systems letters*, vol. 3, no. 1, pp. 108–113, 2018.
- [15] M. Ahmadi, A. Singletary, J. W. Burdick, and A. D. Ames, "Safe policy synthesis in multi-agent POMDPs via discrete-time barrier functions," in *2019 IEEE 58th Conference on Decision and Control (CDC)*. IEEE, 2019, pp. 4797–4803.
- [16] A. Agrawal and K. Sreenath, "Discrete Control Barrier Functions for Safety-Critical Control of Discrete Systems with Application to Bipedal Robot Navigation," in *Robotics: Science and Systems*, 2017.
- [17] M. Ahmadi, A. Singletary, J. W. Burdick, and A. D. Ames, "Barrier functions for multiagent-POMDPs with DTL specifications," in *The 59th IEEE Conference on Decision and Control*, 2020.
- [18] A. Clark, "Control barrier functions for complete and incomplete information stochastic systems," in *2019 American Control Conference (ACC)*. IEEE, 2019, pp. 2928–2935.
- [19] C. Santoyo, M. Dutreix, and S. Coogan, "A barrier function approach to finite-time stochastic system verification and control," *arXiv preprint arXiv:1909.05109*, 2019.
- [20] W. Luo, W. Sun, and A. Kapoor, "Multi-robot collision avoidance under uncertainty with probabilistic safety barrier certificates," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [21] R. T. Rockafellar, S. Uryasev *et al.*, "Optimization of conditional value-at-risk," *Journal of risk*, vol. 2, pp. 21–42, 2000.
- [22] R. T. Rockafellar and S. Uryasev, "Conditional value-at-risk for general loss distributions," *Journal of banking & finance*, vol. 26, no. 7, pp. 1443–1471, 2002.
- [23] A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on stochastic programming: modeling and theory*. SIAM, 2014.
- [24] A. Ruszczyński, "Risk-averse dynamic programming for markov decision processes," *Mathematical programming*, vol. 125, no. 2, pp. 235–261, 2010.
- [25] P. Glotfelter, J. Cortés, and M. Egerstedt, "Nonsmooth barrier functions with applications to multi-robot systems," *IEEE control systems letters*, vol. 1, no. 2, pp. 310–315, 2017.
- [26] M. Ahmadi, A. Israel, and U. Topcu, "Safe controller synthesis for data-driven differential inclusions," *IEEE Transactions on Automatic Control*, 2020.
- [27] J. T. Ott, *A Markov decision model for a surveillance application and risk-sensitive Markov decision processes*, 2010.
- [28] M. Ahmadi, U. Rosolia, M. D. Ingham, R. M. Murray, and A. D. Ames, "Constrained risk-averse Markov decision processes," in *The 35th AAAI Conference on Artificial Intelligence (AAAI-21)*, 2021.
- [29] J. Thai, T. Hunter, A. K. Akametalu, C. J. Tomlin, and A. M. Bayen, "Inverse covariance estimation from data with missing values using the concave-convex procedure," in *53rd IEEE Conference on Decision and Control*. IEEE, 2014, pp. 5736–5742.
- [30] X. Shen, S. Diamond, Y. Gu, and S. Boyd, "Disciplined convex-concave programming," in *2016 IEEE 55th Conference on Decision and Control (CDC)*. IEEE, 2016, pp. 1009–1014.
- [31] J. Lofberg, "Yalmip: A toolbox for modeling and optimization in matlab," in *Computer Aided Control Systems Design, 2004 IEEE International Symposium on*. IEEE, 2004, pp. 284–289.
- [32] J. W. Grizzle, C. Chevallereau, R. W. Sinnet, and A. D. Ames, "Models, feedback control, and open problems of 3D bipedal robotic walking," *Automatica*, vol. 50, no. 8, pp. 1955–1988, 2014.
- [33] X. Xiong and A. D. Ames, "Dynamic and versatile humanoid walking via embedding 3D actuated SLIP model with hybrid LIP based stepping," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 6286–6293, 2020.
- [34] M. Ahmadi, M. Ono, M. D. Ingham, R. M. Murray, and A. D. Ames, "Risk-averse planning under uncertainty," in *2020 American Control Conference (ACC)*. IEEE, 2020, pp. 3305–3312.