Risk-Aware Robotics

TAIL RISK MEASURES IN PLANNING, CONTROL, AND VERIFICATION

PRITHVI AKELLA, ANUSHRI DIXIT, MOHAMADREZA AHMADI, LARS LINDEMANN, MARGARET P. CHAPMAN, GEORGE J. PAPPAS, AARON D. AMES, and JOEL W. BURDICK



©SHUTTERSTOCK.COM/VS148

Why is the consideration of risk integral to robotics? Consider the scenario of a rescue robot designed to navigate through a constrained environment to locate and rescue victims. Despite limited visibility and terrain that is potentially difficult to traverse, the robot is required to assess whether to execute a planned route or seek an alternative route. Delays in finding an alternative route could negatively impact the victim's well-being, and opting for the planned route could lead to the robot getting stuck. How should the robot gauge these competing risks? The issue of risk assessment is not new to roboticists and control engineers, and in practical applications, it is usually tackled using heuristics. For example, the robot constructs a map of the environment using computer vision. As this map is uncertain, one may attempt to decrease the robot's risk by tightening the area that can safely be traversed. This tightening is often based on previous experimental data. However, these experimental data may not accurately represent real-world disaster scenarios. Additionally, an overly conservative tightening might lead the robot to waste time seeking alternate routes, while too little tightening could place the robot at risk of entrapment. Consequently, a systematic approach is needed to assess the risks and rewards associated with different actions in the face of uncertainty. (see "Summary").

The need for a systematic approach to risk assessment has only increased in recent years due to the ubiquity of autonomous systems that alter our day-to-day experiences and their need for safety, for example, for self-driving vehicles, mobile service robots, and bipedal robots. These systems are expected to function safely in unpredictable

Digital Object Identifier 10.1109/MCS.2025.3577050
Date of current version: 21 July 2025

1066-033X © 2025 IEEE. All rights reserved, including rights for text and data mining, and training of artificial intelligence and similar technologies.
Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

environments and interact seamlessly with humans whose behavior is notably challenging to forecast. To reason about risk in such settings, the fields of systems science and control engineering have a long history and a rich literature on analyzing and designing systems under uncertainty. Existing methodologies can be broadly classified into the three paradigms: worst-case, risk-neutral, and risk-aware approaches, as classified in [1]. In the *worst-case paradigm*, a system's ability to remain safe or perform satisfactorily is judged by examining its most severe safety violation or worst performance. This paradigm forms the basis for robust control [2], [3], [4] and robust safety analysis [5]. For instance, if a system is supposed to track a planned trajectory, the largest discrepancy between the system's realized trajectory and the planned trajectory can serve as a measure of the system's performance.

In contrast, in the *risk-neutral paradigm*, a system's capacity to stay safe or perform satisfactorily is evaluated on average or probabilistically.[1] This paradigm is often used in stochastic control and reinforcement learning (RL) [4], [6], [7] as well as verification [8]. In the case of trajectory tracking, the average discrepancy between the system's realized trajectory and the planned trajectory is used to gauge system performance when using a risk-neutral approach. If the system aims to reach a designated target while avoiding an unsafe region, the probability that the system reaches these goals can be used to assess the system's performance [9]. However, the worst-case paradigm may result in an overly conservative risk assessment and impractical solutions, while the risk-neutral paradigm is not designed to account for harmful but less likely events. We also note that the developments in the worst-case and risk-neutral paradigms were often mathematically motivated since one was able to derive closed-form analytical solutions, for example, as in linear robust control and linear-quadratic control [3], [10]. The *risk-aware paradigm*, on the other hand, lies conceptually between these two extremes and aims to evaluate a system's performance by giving attention to system outcomes that do not correspond to the worst case or the average. Historically, the risk-aware paradigm has focused on the application of mean-variance approximations, for example, for controller design [11], [12], [13], which is derived from the Markowitz model to evaluate the risk of financial portfolios [14].

A contemporary risk-aware approach adopted in this survey to mitigate rare and detrimental outcomes employs *tail risk measures*, a concept borrowed from the financial literature [15]. As such, this survey will introduce these measures and explain their relevance in the context of robotic systems. However, in preparation for this explanation, we will first provide a deeper overview of the aforementioned paradigms, including an emphasis on the limitations of the *worst-case* and *risk-neutral* paradigms, which led to the rise of the *risk-aware* study.

[1]Note that the probability of an undesirable event can be expressed as the expected value over an indicator function, making probabilistic reasoning conceptually similar to average reasoning.

Summary

Often, control theorists and roboticists expect systems to function as reliably and predictably as the equations we use to represent them. Sadly, reality is often more random than our equations. For example, take a robot navigating in two similar but unstructured environments. Random perturbations in terrain and scenery could cause the robot to take wildly different paths. In another example, take a perfectly orchestrated robotic swarm that finds itself in dissonance moments later due to network connectivity going down and package loss. Such randomness arises because our equations are imperfect models of reality. So, perhaps we should find a way to account for such randomness in our equations themselves. This article delves into how tail risk measures—formal mathematical concepts of risk traditionally used in the financial community—facilitate accounting for this randomness in planning, control, and verification. The exposition to follow both defines these measures and includes multiple examples of their use in prescribing risk-aware control across all levels of the modern control stack. Finally, we end with a brief survey of existing and open problems in the field.

Benefits of the Risk-Aware Paradigm. Many systems are affected by uncertainties that can be well modeled by random noise. For example, uneven terrain can alter the planned trajectory of a robot [16], [17], wind can destabilize an aerial vehicle [18], and smoke can interfere with perception during an autonomous rescue mission [19]. However, in practice, designing operating rules to optimize a system's performance on average does not always produce trustworthy or reliable performance [1], [20], [21], [22]. For example, consider the disaster scenario presented in [19]. The robot must be able to assess risks in debris-laden zones while taking into account uncertainties derived from sensor and state estimation inaccuracies. If the robot operates under a worst-case scenario strategy, it may not find a feasible path through the rubble because of its aversion to any risk. On the other hand, a risk-neutral path-planning approach might expose the robot to unsafe mission-jeopardizing situations. This example is among many that underline the limitations of the risk-neutral and worst-case paradigms in the context of robotic applications. Another example is how humans may be risk-averse when collaborating with robots, which requires the robot to interpret the risk-averse human model for successful collaborations [23]. Case studies from our previous work will be presented later to further illustrate these points, although we will start by introducing the tail risk measures that have helped foster the more recent risk-aware paradigm.

Tail Risk Measures. Risk measures are functions of scalar-valued random variables designed to identify characteristics of interest of the random variable in question [25]. By

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

tail risk measures, we imply that the risk measures of interest assess the upper right tail of the random variable's distribution. Typically, the random variable indicates a cost, so tail risk measures capture the risk of incurring a high cost. Figure 1 depicts a few examples of such tail risk measures, namely, value at risk (VaR), conditional VaR (CvaR), and entropic VaR (EVaR). The VaR at level $\beta \in (0,1)$ corresponds to the cutoff value for which a fraction β of the outcomes of the random variable lies to the right of this cutoff value. A more formal description of these risk measures will follow in the "Tail Risk Measures: Definitions and Notation" section. We focus on these measures, though, as they provide a systematic way of assessing rare and unsafe (costly) outcomes that must be limited during planning, control, and verification. For example, consider a robot traversing through an environment with uncertain information about the obstacles' positions due to sensor noise. In this case, we could define a cost random variable by negating the minimum distance to all obstacles. As such, more positive, cf., more costly outcomes, would correspond to more unsafe behavior as the system is not maintaining the required distance to the measured obstacle. Then, as we have uncertain knowledge of the obstacle's location, we aim to minimize the tail risk incurred by this random cost in an effort to realize safe risk-aware control actions. This intuitive approach to risk-aware decision selection can be applied to several facets of an autonomy stack (Figure 2), that is, planning, control, and verification, as has recently been done in both the controls and robotics communities.

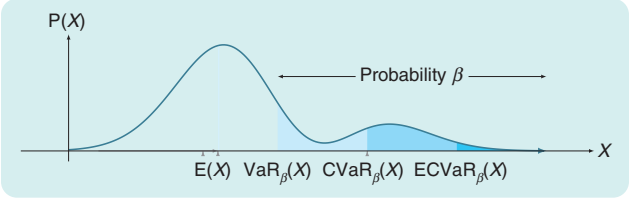


FIGURE 1 A visualization of common tail risk measures. The tail risk measures referenced in the article are shown in the figure, applied to the random variable X with probability density P. X could represent any cost random variable, for example—negative distance from an obstacle, distance from a goal, or energy used by the robot. (Source: Figure adapted from [24].)

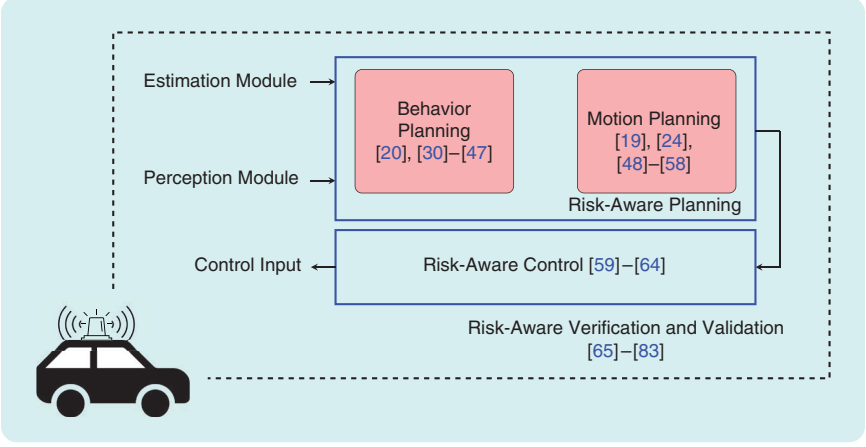


FIGURE 2 An overview of a typical (risk-aware) planning and verification pipeline in an autonomy stack. (See [19], [20], [24], [30], [31], [32], [33], [34], [35], [36], [37], [38], [39], [40], [41], [42], [43], [44], [45], [46], [47], [48], [49], [50], [51], [52], [53], [54], [55], [56], [57], [58], [59], [60], [61], [62], [63], [64], [65], [66], [67], [68], [69], [70], [71], [72], [73], [74], [75], [76], [77], [78], [79], [80], [81], [82], and [83] for a review.)

Organization. As such, this survey serves as an introduction to risk-aware planning, control, and verification in robotics, employing tail risk measures—an emerging field in the literature. Our focus areas include the following:

- A summary of risk-measure theory with an emphasis on tail risk measures is provided in the "Tail Risk Measures: Definitions and Notation" section. We highlight how these measures offer a consistent and intuitive means to adjust the system's risk aversion levels.
- A discussion of the key principles of risk-aware planning and control, an introduction to algorithms, and multiple presentations of real-world case studies, such as planning and control in subterranean environments, are covered in the "Risk-Aware Planning" and "Risk-Aware Control" sections.
- A review of temporal logic as a mathematical formalism for articulating complex robotic system specifications is provided in the "Verification" section. Additionally, we explore why it is important to consider the tail risk of system trajectories evaluated against these specifications.
- An introduction to a tail risk method for safety-critical controller verification is provided in the "Verification" section. We demonstrate this risk-aware verification approach's ability to identify potential mission issues while validating the probabilistic verification statements made within the described framework.
- We end the survey with open problems and future research directions.

Related Work. Despite the crucial need for systematic risk evaluation in robotics applications, recent surveys do not emphasize risk-aware planning, control, and verification [26], [27], [28]. For instance, Schwarting et al. [26] examine planning and decision-making methods for autonomous driving, Karpas and Magazzeni [27] discuss strategies that enable robots to automatically combine smaller tasks to achieve broader goals, and Brunke et al. [28] outline the interplay between control theory and RL for safety in robotic applications. However, these works do not focus on risk measures. Moreover, Hobbs et al. [29] provide a comprehensive introduction to nonstochastic run-time

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

assurance systems, such as a system that overrides an existing controller when a hazard is detected, without an emphasis on risk measures. Two closely related works to our survey are Majumdar and Pavone [21] and Wang and Chapman [1]. Majumdar and Pavone [21] propose axioms for risk measures suitable for robots and provide intuitive explanations for these axioms. However, their work does not present concrete algorithms for risk-aware planning, control, and verification as we do. On the other hand, Wang and Chapman [1] present historical and modern research on risk-aware autonomous systems, but they do not focus on robotics applications, temporal logic, or specific tail risk measures, such as EVaR and total variation distance-based risk measures. Our survey draws inspiration from Majumdar and Pavone [21] and Wang and Chapman [1] and aims to elucidate the concept of tail risk measures for the control systems community and to showcase their utility for planning, control, and verification of robotic systems. We provide the much-needed emphasis on these risk measures in robotics, helping to ensure that the development and application of autonomous systems remain safe, effective, and mindful of potential risks. As we move through this survey, we will provide additional links to the relevant literature.

TAIL RISK MEASURES: DEFINITIONS AND NOTATION

A Pedagogical Note on the Use of Measure Theory. A formal definition of tail risk measures, such as Definition 1, requires some concepts from measure theory, namely distributional equivalence, that is, $\overset{d}{=}$. As such, we begin this section in a more measure-theoretic context, first introducing the probability spaces and defining random variables over these spaces. For readers interested in a deeper study of probability and measure theory, please refer to [84]. For those unfamiliar, however, it is sufficient to consider the to-be-mentioned scalar random variables X as functions assigning real numbers to the random outcomes of interest, for example, the canonical example of assigning the numbers $1-6$ to the corresponding rolls of a six-sided die. Likewise, the distribution P of this random variable X denotes the probability of each sampleable event. Continuing with the six-sided die example and assuming a fair die, the probability of realizing any specific roll—the probability of sampling any number $1-6$ under the random variable X—would all be $1/6$. With this intuition, we will begin with a formal definition of tail risk measures.

Consider a probability space $(\Omega, \mathcal{F}, P)$, where Ω, $\mathcal{F}$, and P are the sample space, the σ-algebra over Ω, and the probability measure over $\mathcal{F}$, respectively. A random variable $X : \Omega \to \mathbb{R}$ denotes the cost of each outcome, and $\mathbb{X}$ is the set of all such random variables defined on Ω. For any random variable $X \in \mathbb{X}$, $F_X(x)$ refers to the cumulative distribution function with inverse $F_X^{-1}(\beta) = \inf\{x \in \mathbb{R} \mid F_X(x) \geq 1-\beta\}$. For any two random variables $X, X' \in \mathbb{X}$, the expression $X \overset{d}{=} X'$ indicates that the random variables X, X' have the same distribution under P. Similarly, we use $X \leq X'$ as a shorthand notation to indicate that $X(\omega) \leq X'(\omega)$ for almost all $\omega \in \Omega$. Finally, U denotes the uniform random variable between [0, 1].

A risk measure ρ is a function that maps a cost random variable to a real number, that is, $\rho : \mathbb{X} \to \mathbb{R}$. Informally, tail risk measures refer to the behavior of the cost X in the tail of its distribution. Mathematically, consider the random variable $X \in \mathbb{X}$ and the risk-level $\beta \in (0,1)$. As seen in [85], we define X_β to be the tail risk of X such that

$$X_\beta = F_X^{-1}(1-\beta+\beta \mathrm{U}). \tag{1}$$

In other words, the distribution of the random variable X_β is the distribution of X in its β-quantile (or tail) normalized to sum to one. It is also clear that as $\beta \to 0$, $X_\beta \to \operatorname{esssup}(X)$. We are now ready to define tail risk measures formally.

Definition 1 (Tail Risk Measures)

For $\beta \in (0,1)$, a risk measure ρ is a β-tail risk measure (adapted from [85]) (or simply tail risk measure) if $\rho(X) = \rho(X')$ for all $X, X' \in \mathbb{X}$ satisfying $X_\beta \overset{d}{=} X'_\beta$.

We defined tail risk measures with respect to the right tail of the distribution of the corresponding random variable; see also Figure 1 for an illustration. One could also consider the left tail of the distribution, which would correspond to the same quantity when considering the negated random variable. In this article, we sometimes consider the left tail instead of the right tail for simplicity. The aforementioned definition is a formal description of tail risk measures, three of which have featured more prominently in recent literature—VaR, CvaR, and EVaR. Their definitions will follow.

VaR

Chance constraints can be reformulated using a commonly used risk measure called the VaR. For a given confidence level $\beta \in (0,1)$, VaR_β denotes the β-quantile value of the random variable X and is defined as

$$\mathrm{VaR}_\beta(X) := \inf\{z \in \mathbb{R} \mid F_X(z) \geq 1-\beta\}.$$

Therefore, $\mathrm{VaR}_\beta(X) = F_X^{-1}(\beta)$. It follows that $\mathrm{VaR}_\beta(X) \leq 0 \Rightarrow P(X \leq 0) \geq 1-\beta$.

CvaR

The CVaR_β measures the expected loss in the β-tail of the random variable X. Formally, for some $\beta \in (0,1]$, CVaR_β is defined as follows per [15]:

$$\mathrm{CVaR}_\beta(X) := \inf_{z \in \mathbb{R}} \mathbb{E}\left[z + \frac{(X-z)^+}{\beta}\right] \tag{2}$$

where we use the notation $(X-z)^+ = \max(0, X-z)$. A value of $\beta \simeq 1$ corresponds to a risk-neutral case. A value

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

of $\beta \to 0$ is rather a risk-averse case.[2] Importantly, CVaR is a coherent risk measure [86] (see the "Coherent Risk Measures" section), prompting its widespread use in the recent literature. Furthermore, CVaR is a loose upper bound on VaR, that is

$$\mathrm{VaR}_\beta(X) \le \mathrm{CVaR}_\beta(X) \le 0 \Rightarrow P(X \le 0) \ge 1 - \beta. \tag{3}$$

EVaR

The EVaR_β, derived using the Chernoff inequality for the random variable in question, is the tightest upper bound for VaR and CVaR. It was shown in [87] that EVaR_β and CVaR_β are equal when $\mathrm{VaR}_t(X) = -\infty$ for all $\beta < t \le 1$. For some $\beta \in (0,1]$

$$\mathrm{EVaR}_\beta(X) := \inf_{z>0}\left[z^{-1}\ln\frac{\mathbb{E}[e^{Xz}]}{\beta}\right]. \tag{4}$$

Similar to CVaR_β, for EVaR_β, the limit $\beta \to 1$ corresponds to a risk-neutral case, whereas $\beta \to 0$ corresponds to a risk-averse case. It was shown in [88, Proposition 3.2] that $\lim_{\beta\to 0}\mathrm{EVaR}_\beta(X) = \mathrm{esssup}(X)$, where $\mathrm{esssup}(X)$ is the essential supremum of the random variable X. Finally, EVaR is also a coherent risk measure and is an upper bound for both VaR and CVaR

$$\begin{aligned}\mathrm{VaR}_\beta(X) \le \mathrm{CVaR}_\beta(X) \le \mathrm{EVaR}_\beta(X) \le 0,\\ \Rightarrow P(X \le 0) \ge 1 - \beta.\end{aligned}$$

EVaR gets its name from its dual representation as a risk measure that provides distributional robustness using relative entropy.

Coherent Risk Measures

Coherent risk measures are a prominent class of risk measures that are well regarded for their robust mathematical properties, their inherent intuitiveness for risk analysis, and their use in robotics [21]. Introduced by Artzner et al. [86] in the context of financial risk management, coherent risk measures satisfy four axiomatic properties: monotonicity, subadditivity, positive homogeneity, and translational invariance. Monotonicity indicates that adding a less risky outcome to a portfolio should not increase its risk. Subadditivity implies that diversifying a risk portfolio should not increase its overall risk. Positive homogeneity signifies that scaling all outcomes in a portfolio should proportionally scale its risk. Translational invariance denotes that adding a risk-free asset to a portfolio should decrease its risk correspondingly. Coherent risk measures offer a rich theoretical foundation to quantify and manage risk systematically. They enable us to capture extreme—but rare—high-consequence events and provide a mechanism to evaluate and compare different risk scenarios, making them particularly valuable for risk-aware planning, control, and verification in robotics. We are now ready to describe coherent risk measures.

Definition 2 (Coherent Risk Measure)

We call the risk measure $\rho : \mathbb{X} \to \mathbb{R}$, a coherent risk measure if it satisfies the following conditions:

- *Subadditivity:* $\rho(X + X') \le \rho(X) + \rho(X')$, *for all* $X, X' \in \mathbb{X}$.
- *Monotonicity: If* $X \le X'$ *then* $\rho(X) \le \rho(X')$ *for all* $X, X' \in \mathbb{X}$.
- *Translational Invariance:* $\rho(X + c) = \rho(X) + c$ *for all* $X \in \mathbb{X}$ *and* $c \in \mathbb{R}$.
- *Positive Homogeneity:* $\rho(\beta X) = \beta\rho(X)$ *for all* $X \in \mathbb{X}$ *and* $\beta \ge 0$.

Note that the subadditivity and positive homogeneity properties imply that coherent risk measures are also convex. CVaR and EVaR have been recognized as coherent risk measures. Conversely, VaR does not qualify as a coherent risk measure as it does not satisfy the subadditivity property. This fact limits its utility in situations where a joint assessment of risks is required, which is often the case in risk-aware robotic planning, control, and verification. While VaR is not a coherent risk measure, each of the measures defined previously has its place in risk evaluation, depending on the specific requirements and constraints of the task.

Coherent risk measures can also be written as the worst-case expectation over a convex and closed set of probability mass (or density) functions. This distributionally robust representation of coherent risk measures is defined next.

Definition 3

Every coherent risk measure can be represented in its dual form as the following (adapted from [89]):

$$\rho(X) := \sup_{Q \in \mathcal{Q}} E_Q(X),$$

where the ambiguity set (or risk envelope), $\mathcal{Q}$*, is convex and closed, and the probability density function* $Q(X)$ *is absolutely continuous with respect to the probability density function* $P(X)$; that is, $P(X) = 0 \Rightarrow Q(X) = 0$.

Remark 1

In this review, we limit our focus to tail risk measures for robotics, but we note that distributionally robust optimization is another widely used measure of risk [90], [91], [92], [93], [94]. Some tail risk measures, like CVaR and EVaR, also have a distributionally robust interpretation. Some methods that consider distributionally robust formulations of tail risk measures are also discussed in this review.

Coherent risk measures provide a snapshot of potential risks based on current conditions. These measures fall short when applied to dynamic systems, such as those common in robotics, where risks and conditions vary with time. *Dynamic coherent risk measures* extend beyond the static approach, taking into account the time-varying nature of

[2] Another, more intuitive, way to think about the widely used CVaR metric is that it is the expectation of the random variable X conditioned on $X \ge \mathrm{VaR}_\beta(X)$, that is, $\mathrm{CVaR}_\beta(X) = \mathbb{E}[X \mid X \ge \mathrm{VaR}_\beta(X)]$. For example, the 5% CVaR risk of a portfolio is equivalent to the expected (mean) return on a portfolio in the worst 5% of scenarios over a specified time horizon (in this definition, we assume that X is larger for worse returns).

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

risk. They are particularly adept at characterizing the evolving risk landscape in dynamic environments. Dynamic risk measures continuously monitor and update risk assessments in response to changes in the system and its environment. This capability to adapt and provide a comprehensive understanding of risk in a fluctuating context aligns well with the realities of robotic operations in unstructured environments.

To define these measures, we will first index a sequence of random variables X_t, $t = 0,\dots,N$, where $N \in \mathbb{N}_{\geq 0} \cup \{\infty\}$. Note that prior to that, each random variable X was defined over a corresponding probability space, and $\mathbb{X}$ denoted the set of all random variables defined in that space $(\Omega, \mathcal{F}, P)$. However, in the dynamic context, each sequenced random variable X_t belongs to (perhaps) different spaces of random variables $\mathbb{X}_t$. As such, denote $\mathbb{X}_{t:N} = \mathbb{X}_t \times \cdots \times \mathbb{X}_N$ as the space in which the corresponding sequence $X_{t:N}$ lives, and $\mathbb{X} = \mathbb{X}_0 \times \mathbb{X}_1 \times \cdots$. Finally, we assume that the sequence $\boldsymbol{X} \in \mathbb{X}$ is almost surely bounded (with exceptions having probability zero), that is, $\max_t \operatorname{esssup} \mid X_t(\omega) \mid < \infty$. To describe how one can evaluate the risk of subsequence $X_t,\dots,X_N$ from the perspective of stage t, we require the following definitions.

Definition 4 (Conditional Risk Measure)

A mapping $\rho_{t:N} : \mathbb{X}_{t:N} \to \mathbb{X}_t$, where $0 \leq t \leq N$, is called a conditional risk measure if it has the following monotonicity property:

$$\rho_{t:N}(\boldsymbol{X}) \leq \rho_{t:N}(\boldsymbol{X}'), \quad \forall \boldsymbol{X}, \forall \boldsymbol{X}' \in \mathbb{X}_{t:N} \text{ such that } \boldsymbol{X} \preceq \boldsymbol{X}'.$$

Definition 5 (Dynamic Risk Measure)

A dynamic risk measure is a sequence of conditional risk measures $\rho_{t:N} : \mathbb{X}_{t:N} \to \mathbb{X}_t$, $t = 0,\dots,N$.

A key attribute of dynamic risk measures is their temporal consistency [95, Definition 3]; if two scenarios X and X' are identical for a time interval $[\tau, \theta]$, and if X is evaluated to be as favorable as X' at some future time point θ, then it stands to reason that X should not be viewed as more risky than X' at the earlier time point τ. This principle ensures that the risk assessment remains coherent and consistent across different points in time. The definition of a dynamic coherent risk measure [96, p. 298] then follows from its static counterpart.

Definition 6 (Dynamic Coherent Risk Measure)

We call the one-step conditional risk measures $\rho_t : \mathbb{X}_{t+1} \to \mathbb{X}_t$, $t = 1,\dots,N-1$ a coherent risk measure if it satisfies the following conditions:

- *Convexity: $\rho_t(X + (1-\lambda)X') \leq \lambda\rho_t(X) + (1-\lambda)\rho_t(X')$, for all $\lambda \in (0,1)$ and all $X, X' \in \mathbb{X}_{t+1}$.*
- *Monotonicity: If $X \leq X'$ then $\rho_t(X) \leq \rho_t(X')$ for all $X, X' \in \mathbb{X}_{t+1}$.*
- *Translational Invariance: $\rho_t(X + X') = X + \rho_t(X')$ for all $X \in \mathbb{X}_t$ and $X' \in \mathbb{X}_{t+1}$.*
- *Positive Homogeneity: $\rho_t(\beta X) = \beta\rho_t(X)$ for all $X \in \mathbb{X}_{t+1}$ and $\beta \geq 0$.*

AN OVERVIEW OF TAIL RISK MEASURES IN ROBOTICS

The definitions in the previous section now allow us to provide a general overview of how tail risk measures are used to generate risk-aware behavior or validate systems in a risk-aware setting. For such an overview, let us define ρ to be a risk measure that maps a random variable X to a real number that indicates the risk associated with X. For applications in robotics, this random variable X denotes any random cost associated with (perhaps) stochastically evolving system trajectories. To be more specific, consider a (perhaps) nonlinear discrete-time control system at time t with state $\boldsymbol{x}(t) \in \mathcal{X} \subseteq \mathbb{R}^n$, input $\boldsymbol{u}(t) \in \mathcal{U}$, and system disturbances $\boldsymbol{d}(t) \sim \xi(\boldsymbol{x}(t), \boldsymbol{u}(t), t)$ where $\xi(\boldsymbol{x}(t), \boldsymbol{u}(t), t)$ is a (perhaps) state-, input-, and time-dependent probability distribution over $\mathbb{R}^n$

$$\boldsymbol{x}(t+1) = f(\boldsymbol{x}(t), \boldsymbol{u}(t), \boldsymbol{d}(t)). \tag{5}$$

Given a feedback controller $U : \mathcal{X} \to \mathcal{U}$, which prescribes $\boldsymbol{u}(t) = U(\boldsymbol{x}(t))$, we can construct a closed-loop stochastically evolving dynamical system. To be clear, we first define $\mathcal{S}^{\mathcal{X}}$ as the set of all signals mapping time to the state space $\mathcal{X}$, for example, $\mathcal{S}^{\mathcal{X}} = \{s : \mathbb{Z} \to \mathcal{X}\}$. As such, provided an initial condition $\boldsymbol{x}_0 \in \mathcal{X}$, this feedback controller U, and a randomly sampled disturbance sequence $\{\boldsymbol{d}(0), \boldsymbol{d}(1), \dots\}$, the corresponding state trajectory $\boldsymbol{x} = \{\boldsymbol{x}(0) = \boldsymbol{x}_0, \boldsymbol{x}(1), \dots\}$ following the update in (5) is a signal that lives in $\mathcal{S}^{\mathcal{X}}$, that is, $\boldsymbol{x} \in \mathcal{S}^{\mathcal{X}}$. Notably, this trajectory $\boldsymbol{x}$ is also a randomly sampled trajectory conditioned on the initial state x_0 due to the randomly sampled disturbance sequence. To be more specific in a tail risk context, we define $\Sigma(\boldsymbol{x}_0)$ as the random variable conditioned on the initial condition $\boldsymbol{x}_0$, whose samples are these random trajectories $\boldsymbol{x}$ living in $\mathcal{S}^{\mathcal{X}}$, that is

$$\text{a sample of } \Sigma(\boldsymbol{x}_0) \text{ is } \boldsymbol{x} \triangleq \{\boldsymbol{x}(0) = \boldsymbol{x}_0, \boldsymbol{x}(1), \dots\}. \tag{6}$$

To analyze the tail risk of this trajectory random variable $\Sigma(x_0)$, we require a mapping from trajectory samples $\boldsymbol{x}$ to the real line

$$C : \mathcal{S}^{\mathcal{X}} \to \mathbb{R}.$$

The function C may denote the inverse distance of a robot's trajectory to an obstacle, thereby encoding the system's robustness to a collision. These cost functions can be constructed from principled approaches when the system's desired behavior is expressed as a temporal logic specification. Applying this cost C to the random variable corresponding to the system trajectories $\Sigma(x_0)$ defines the scalar random variable X over which we can apply our tail risk measure ρ, that is, $X = C(\Sigma(x_0))$. For robotic planning, control, and verification, the risk of this cost can then be minimized $\rho(X)$, or it can be used as a risk constraint–$\rho(X) \leq r$ for some risk threshold r.

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

In what follows, we present the two main sections, "Risk-Aware Planning and Control" and "Risk-Aware Verification and Validation," in this order. We chose this specific organization as planning and control synthesis are usually carried out by making simplifying assumptions on the system and by considering low-fidelity models. The design using low-fidelity models is motivated by the complexity of the problem and computational considerations. The designed planners and controllers are, hence, not guaranteed to be correct for high-fidelity models or real robotic systems, emphasizing the need for system verification and validation (V&V) techniques that apply in realistic settings.

RISK-AWARE PLANNING AND CONTROL

The robotic behavior, motion planning, and control problems focus on designing algorithms that allow a robot to interact intelligently and safely with its surroundings. This complex process involves deciding on possible actions, constructing trajectories that a robot can take to achieve specific high-level goals, for example, safely navigating through unstructured environments, and controlling the instantaneous robot motions to track the trajectory. Behavior planning concerns the higher-level decision making needed to achieve higher-level robot objectives, while motion planning plans the details of robot movement, determining how a robot should move from one location to another while avoiding obstacles and factoring in the robot's kinematics and dynamics. The control level manages the details of motion execution by continually computing the system control inputs that minimize tracking error while accounting for uncertainty and unexpected events.

The process of designing robot plans and controls should critically consider potential risks and their overall system implications. These risks include physical risks to humans close to a robot and risks to the robot itself due to its unpredictable environment and imperfect robot sensing and perception systems. Notions of risk are fundamental in high-stakes robot tasks, such as search and rescue or exploration of hazardous sites. To properly manage these risks, this article advocates the use of tail risk measures. Importantly, tail risk measures focus on more extreme but rarer events, providing a systematic and principled approach to quantifying and managing risk in robotics. Integrating tail risk measures into robotic planning and control modules can lead to more robust and safer robot operation, effectively balancing performance and safety under uncertainty.

Historical Remark on Exponential Utility. (This remark is adapted from [97].) Risk-aware control has been in development for at least 50 years. The earliest contributions concern the exponential utility measure (that is, the entropic risk measure)

$$\rho_{\mathrm{EU},\theta}(X) := \frac{-2}{\theta}\log\mathbb{E}\left[\exp\left(\frac{-\theta}{2}X\right)\right] \tag{7}$$

where $\theta \neq 0$ is a risk parameter, and X is a nonnegative random variable, which can be interpreted as a mean-variance approximation [13]. When $\theta < 0$, the robotic system is risk averse, while the robot will exhibit more risk-tolerant behavior when $\theta > 0$. When $\theta = 0$, the system is risk neutral. Notably, the exponential utility is *not* a tail risk measure. To our knowledge, the first article in the area of risk-aware control was a 1972 study about finite-state Markov decision processes (MDPs) by Howard and Matheson [11], in which performance was assessed by using an exponential utility criterion. The authors took inspiration from game theory [99]. One year later, Jacobson [12] investigated the exponential utility criterion in the classical linear-quadratic setting (linear dynamics, quadratic costs, Euclidean spaces, and additive Gaussian noise). Whittle [13], [100] developed key contributions regarding risk-aware control in the linear-quadratic setting using the exponential utility measure, including showing close analogies to optimal control and state estimation results in the risk-neutral case. While exponential utility continues to be investigated (for example, see [97], [101], and [102]), recent attention focuses on different types of risk-aware performance and safety criteria, motivating this survey on tail risk.

Risk-Aware Behavior Planning

Behavior planning for robotics has been studied in many mathematical frameworks, though the most commonly utilized framework has been to reframe planning as a (partially observed) MDP (POMDP), which we will discuss in this section. Specifically, an MDP is a triple[3]

$$\mathcal{M} = (\mathcal{X}, \mathcal{U}, T) \tag{8}$$

where $\mathcal{X} = \{x_1, \ldots, x_{|\mathcal{X}|}\}$ represents the states of the autonomous agent(s) and the world model, $\mathcal{U} = \{u_1, \ldots, u_{|\mathcal{U}|}\}$ represents the actions available to the agent, and $T(x_j|x_i, u)$ defines the likelihood of transitioning to state x_j from state x_i when taking action $u \in \mathcal{U}$. Furthermore, the transition function must be such that

$$\sum_{x \in \mathcal{X}} T(x|x_i, u) = 1, \forall\, x_i \in \mathcal{X}, \forall\, u \in \mathcal{U} \tag{9}$$

and for a *finite* MDP, the state and action spaces must also be finite (see Figure 3 for an example of MDP). Finally, a policy π maps states to actions, that is, $\pi : \mathcal{X} \to \mathcal{U}$.

As such, risk-aware behavior planning aims to construct a policy π that minimizes a risk-measure evaluation over either each state or over some time horizon of states. Specifically, consider the cost c associated with transitioning from an initial state x_0 by taking an action u to a final state x_f, that is, $c : \mathcal{X} \times \mathcal{U} \times \mathcal{X} \to \mathbb{R}$ where $c(x_0, u, x_f)$ denotes the cost of the action-specific transition. As state evolution for an MDP is stochastic, the exact cost of an action is unknown a

[3]Here, we note that MDPs are typically defined with an associated objective function, and we will define one for risk-aware behavior planning in a section to follow.

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

priori as the final state x_f will be a sample from a random variable $X_f(x_0,u)$ conditioned on the initial state and action choice with distribution provided by the transition function T. As a result, the corresponding cost of an initial state and action pair is a scalar random variable corresponding to the costs of the random final state samples, that is, $C(x_0,u) = c(x_0,u,X_f(x_0,u))$. As such, the one-step optimal policy is to choose actions per state that minimize the tail risk ρ associated with this cost random variable $C(x_0,u)$, that is

$$\pi^*(x) = \operatorname{argmin}_{u\in\mathcal{U}} \rho(C(x,u)). \tag{10}$$

We could similarly define policies by considering not just the one-step risk but the risk over a horizon. To do so, we note that for any horizon length $T \in \{1,2,\ldots\} \triangleq= \mathbb{Z}_+$, an initial state x_0, and a policy π, the state evolution $S = \{x_0,x_1,\ldots,x_T\}$ is random due to probabilistic state transitions. Then, provided a similar sequence of risk measures $\hat{\rho}_T = \{\rho_i\}_{i=0}^{T-1}$, we can define a cost function J corresponding to sequential risk evaluations of each cost random variable over the state horizon, with perhaps some discount factor $\gamma \in (0,1]$

$$\begin{aligned} J(x_0,\pi,\gamma) = \; & \rho_0(C(x_0,\pi(x_0)) + \gamma\rho_1(C(x_1,\pi(x_1)) \\ & + \gamma^2\rho_2(C(x_2,\pi(x_2)) + \ldots). \end{aligned} \tag{11}$$

Here, we note that, while we have defined the interior states $x_1,x_2,\ldots,x_{T-1}$, they are random based on the probabilistic transitions from the previous state and the provided action. As such, this randomness must be taken into account during risk-measure evaluation, resulting in the nested risk measures in the cost mentioned previously. Similarly, the optimal policy minimizes this cost

$$\pi^* = \operatorname{argmin}_{\pi} J(x_0,\pi,\gamma). \tag{12}$$

Finally, for $\gamma < 1$, this is called the *discounted MDP problem*, and for $\gamma = 1$, this is called the *undiscounted MDP problem*, and for either case where $T \to \infty$, this is called the *infinite-horizon MDP problem*, all of which have been applied to risk-aware behavior planning, as will be discussed in the sections to follow.

The last extension to the MDP framework concerns the case where the state is not exactly known, that is, the system is partially observed at every time step, and as such, we do not know the exact state that the system occupies. Rather, we have a belief, a probability distribution, over the states that encodes our understanding of where the system is at any given time. We'll refrain from rigorously defining this POMDP as, at a high level, the policy selection pursuit is similar in spirit to what was shown in the prior two optimization problems (for more information on POMDPs, though, please reference [103] and [104]). The primary change arises in the consideration of the distribution over states when accounting for risk, whereas the state was known exactly in the prior MDP case.

Examples

Perhaps the most prevalent use of POMDPs in planning arises in navigation problems wherein the finite POMDP states correspond to a discretization of the robot's state space $\mathcal{X}$ into a finite set, for example, discrete cells on a 2D grid for planar navigation combined with the enumeration of a finite list of environmental phenomena, for example, the occupancy values of those same grid cells [105]. Actions then correspond to a minor navigation task, for example, moving from one discretized grid cell in the state space to another discretized grid cell. Finally, stochasticity between these transitions arises as either the environmental phenomena are partially observed and might frustrate the immediate application of minor navigation actions, faulty sensors might cause some drift in the robot's ability to achieve the navigation task, etc. Then, the goal is to identify an optimal policy against a risk measure, typically VaR or CVaR. Indeed, this paradigm underlies many works in this vein, a sampling of which is provided here and in the citations within [42], [45], [105], [106], [107], [108], and [109]. These works are by no means exhaustive; however, due to the prevalence of MDPs for behavior planning, both control theorists and roboticists alike have also progressed fundamental research in POMDPs in pursuit of advancements in robotic behavioral path planning. A discussion of those advances will follow.

Prior Work on Discounted MDPs: In an early work in this vein, the authors of [95] presented techniques for incorporating this measure into dynamic programming. This work resulted in a wave of new work evaluating risk measures in dynamic programming problems [21], [30], [31], [32], [110]. For example, in [33] and [34], the authors identified locally optimal solutions via gradient descent to MDP problems with CVaR constraints and total expected costs. In particular, [33] provides a convergence guarantee, whereas [34] does not. The authors of [35] and [36] extended these prior notions by developing sample-based saddle point algorithms to identify policies for MDPs whose cost is a coherent risk measure, though not specifically CVaR. Other relevant works include [37], [38], [39], and [40].

One question posed by the authors in [41] has caused a renewal of work in this vein. Specifically, the authors show

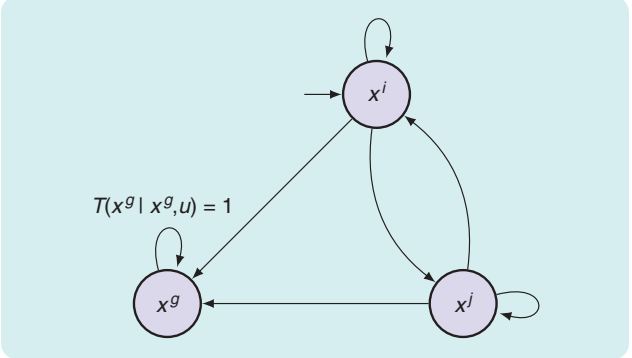


FIGURE 3 The transition graph of a transient MDP. The goal state $x = x^g$ is cost-free and absorbing.

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

that most risk-level dynamic programs cannot guarantee the recovery of a globally optimal value function despite a discretized state space. In [42] and [47], the authors generate optimal risk-aware policies for MDPs with dynamic coherent risk objectives and constraints to partially address this concern. Solutions can also be rapidly identified by phrasing policy generation as a difference convex program. Despite these advances, the field of risk-aware discounted MDPs still holds numerous avenues for future exploration. New algorithms and techniques that can handle an expansive range of coherent risk measures and can effectively manage constraints in MDPs are needed.

Prior Work on POMDPs: While POMDPs can be difficult to design and solve, significant strides have been made in addressing coherent risk-measure objectives. For example, [46] explored POMDPs with coherent risk-measure objectives. However, their noteworthy theoretical contributions fell short of providing a computational method for designing policies applicable to general coherent risk measures. Ahmadi et al. [20] aimed to address this gap by proposing a method for finding finite-state controllers for POMDPs with objectives defined in terms of coherent risk measures. Their novel approach took advantage of convex optimization techniques, showcasing the potential of mathematical optimization in policy design. Nevertheless, their method has its limitations; it can be applied only when the risk transition mapping is affine in the policy.

Recognizing this limitation, Ahmadi et al. [47] extended their previous work [20] to incorporate a broader set of coherent risk measures. They proposed an innovative approach, a bounded policy iteration method that identifies finite-state risk-averse policies. This methodology breaks the problem down into manageable pieces, tackling convex optimization problems at each policy iteration step. This approach substantially ameliorates the computational tractability of synthesizing risk-averse policies for POMDPs. By iteratively solving these convex optimization problems, the policy synthesis process becomes markedly more feasible.

However, the methodology outlined in [47] has its limitations. One notable constraint is that the technique can currently be applied only to problems involving hundreds of states due to the computational limitations inherent to convex optimization. Despite existing limitations, exploring POMDPs in the context of coherent risk measures presents a promising field of study. As our understanding deepens and computational methods evolve, we can anticipate the development of more pragmatic solutions for planning under uncertainty under a broader range of coherent risk measures.

Risk-Aware Motion Planning and Control

The behavior planning layer, as described in the previous section, is only the top layer of most controllers. Below this layer, controllers typically comprise a motion planning layer that translates the behavior into a lower-level sequence of commands, which are then tracked by the lower-level controller. This section will detail risk-aware work on the remaining two layers—motion planning and control. Optimization-based motion planning methods are increasingly popular because they provide optimized system behaviors that respect the system's dynamics while readily incorporating state and control constraints. Specifically, sudden system changes or disturbances must be taken into account to ensure system safety.

Our risk-aware motion planning review considers discrete-time controlled stochastic systems, with the form

$$\boldsymbol{x}(t+1) = f(\boldsymbol{x}(t), \boldsymbol{u}(t), \boldsymbol{d}(t)). \tag{13}$$

Here, $\boldsymbol{x}(t) \in \mathbb{R}^{n_x}$ and $\boldsymbol{u}(t) \in \mathbb{R}^{n_u}$ are the system state and controls at time t, respectively. The system is affected by a stochastic process noise $\boldsymbol{d}(t) \in \mathbb{R}^{n_d}$ and $f : \mathbb{R}^{n_x} \times \mathbb{R}^{n_u} \times \mathbb{R}^{n_d} \to \mathbb{R}^{n_x}$. An optimization-based planner seeks to minimize a system cost $J(\boldsymbol{x}, \boldsymbol{u}) \in \mathbb{R}$ for initial condition $\boldsymbol{x}(0) = \boldsymbol{x}_0$ at time $t = 0$. The optimal controller $U = [\boldsymbol{u}(0), \ldots, \boldsymbol{u}(N-1)]$ is the solution to the following optimization problem:

$$J^*(\boldsymbol{x}(0)) = \min_{U} \ \rho\left(\sum_{t=0}^{N-1} J(\boldsymbol{x}(t), \boldsymbol{u}(t))\right) \tag{14a}$$

$$\text{s.t. } \boldsymbol{x}(t+1) = f(\boldsymbol{x}(t), \boldsymbol{u}(t), \boldsymbol{d}(t)) \tag{14b}$$

$$\boldsymbol{x}(0) = \boldsymbol{x}_0, \quad \forall t \in \{0, \ldots N-1\}. \tag{14c}$$

Risk-Aware Model Predictive Control

Model predictive control (MPC) applies the finite-horizon controller (14) in a receding-horizon fashion. Uncertainty can arise for many reasons. Uncertainty in the robot's dynamics model causes the true system motion to differ from the predicted one. Such effects are typically accounted for via process noise, $\boldsymbol{d}(t)$. Sensor noise and imprecise robot localization or estimation of environment states are other common sources of uncertainty. There are many ways to account for these uncertainties in an MPC framework. For example, Robust MPC accounts for worst-case disturbances in a set of bounded uncertainties [111]. Robust approaches are often too conservative because they focus on worst-case events. Conversely, stochastic MPC [112] accounts only for the average realization of the cost while respecting a bound on the probability of violating the state and control constraints, see "MPC With Uncertainty."

Risk-aware MPC optimizes risk-averse behavioral policies; they are not as conservative as in the robust case. But, since they account for "risky" outcomes in the tail of the uncertainty distribution, they perform better in practice. In [48], the authors provide an MPC scheme for a discrete-time dynamical system with process noise whose objective was a CVaR measure. They further provided new Lyapunov conditions for risk-sensitive exponential stability. In [49], the authors devised an MPC scheme that expressed a distributionally robust chance constraint along with a risk-aware cost in terms of a CVaR reformulation. Optimal control using distributionally robust CVaR constraints with

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

MPC With Uncertainty

Consider a linear discrete-time system given by

$$\boldsymbol{x}(t+1) = A\boldsymbol{x}(t) + B\boldsymbol{u}(t) + D\boldsymbol{d}(t) \tag{S1}$$

where $\boldsymbol{x}(t) \in \mathbb{R}^{n_x}$ and $\boldsymbol{u}(t) \in \mathbb{R}^{n_u}$ are the system state and controls at time t, respectively. The system is affected by a stochastic, additive, process noise $\boldsymbol{d}(t) \in \mathbb{R}^{n_d}$.

Consider r_x state constraints of the form

$$\mathcal{X}_S := \{\boldsymbol{x} \in \mathbb{R}^{n_x} \mid F_x x \le g_x\}, F_x \in \mathbb{R}^{r_x \times n_x}, g_x \in \mathbb{R}^{r_x}.$$

We also assume r_u control constraints having the form

$$\mathcal{U} := \{\boldsymbol{u} \in \mathbb{R}^{n_u} \mid F_u \boldsymbol{u} \le g_u\}, F_u \in \mathbb{R}^{r_u \times n_x}, g_u \in \mathbb{R}^{r_u}.$$

The goal is to steer the system to a set

$$\mathcal{X}_F := \{\boldsymbol{x} \in \mathbb{R}^{n_x} \mid F_f \boldsymbol{x} \le g_f\}, F_f \in \mathbb{R}^{n_f \times n_x}, g_x \in \mathbb{R}^{n_f}$$

while minimizing the control effort and deviation from the desired trajectory, that is, we want to minimize the following cost:

$$J(\boldsymbol{x}, \boldsymbol{u}) := \boldsymbol{x}^T Q \boldsymbol{x} + \boldsymbol{u}^T R \boldsymbol{u}$$

where $Q \in \mathbb{R}^{n_x \times n_x}$ and $R \in \mathbb{R}^{n_u \times n_u}$ are weights on the state and control costs. MPC provides an optimization-based framework to compute the best N-step control input while satisfying the state and control constraints. The MPC optimization is given by

$$J_t^*(\boldsymbol{x}(t)) = \min_{U_t} \quad \mathbb{E}[x_{t+N|t}^T P \boldsymbol{x}_{t+N|t} + \tag{S2a}$$

$$\sum_{k=t}^{t+N-1} (\boldsymbol{x}_{k|t}^T Q \boldsymbol{x}_{k|t} + \boldsymbol{u}_{k|t}^T R \boldsymbol{u}_{k|t})] \tag{S2b}$$

$$\text{subject to } \boldsymbol{x}_{k+1|t} = A\boldsymbol{x}_{k|t} + B\boldsymbol{u}_{k|t} + D\boldsymbol{d}_{k|t} \tag{S2c}$$

$$\text{Prob}(\boldsymbol{x}_{k|t} \notin \mathcal{X}_S) \le \beta, \boldsymbol{u}_{k|t} \in \mathcal{U} \tag{S2d}$$

$$\text{Prob}(\boldsymbol{x}_{t+N|t} \notin \mathcal{X}_F) \le \beta \tag{S2e}$$

$$\boldsymbol{x}_{t|t} = \boldsymbol{x}(t) \quad \forall k \in \{t, \ldots t+N-1\} \tag{S2f}$$

where $\boldsymbol{x}_{k|t}$ is the state at time k as predicted at the time t while starting from the current state $\boldsymbol{x}_{t|t} = \boldsymbol{x}(t)$, and β is the user chosen risk level. Uncertainty is propagated through the states as

$$\boldsymbol{x}_{k+1|t} = A^k \boldsymbol{x}_{t|t} + \sum_{i=t}^{k} (A^{(k-i)} B \boldsymbol{u}_{i|t} + A^{(k-i)} D \boldsymbol{d}_{i|t}).$$

If the uncertainty is independent identically distributed Gaussian with $\boldsymbol{d}(t) \sim \mathcal{N}(0, \Sigma)$, the states $\boldsymbol{x}(t)$ are also Gaussian $\boldsymbol{x}_{k+1|t} \sim \mathcal{N}(\hat{x}_{k|t}, \Sigma_{k|t})$ where, $\hat{x}_{k|t} = A^k \boldsymbol{x}_{t|t} + \Sigma_{i=t}^{k} A^{(k-i)} B \boldsymbol{u}_{i|t}$ and $\Sigma_{k|t} = \Sigma_{i=t}^{k} D^T A^{(k-i)T} \Sigma A^{(k-i)} D$ (the family of normal distributions is closed under linear transformations). Hence, we can rewrite the aforementioned uncertain MPC optimization as the following deterministic quadratic program:

$$\begin{aligned} J_t^*(\boldsymbol{x}(t)) = \min_{U_t} \quad & \mathbb{E}[\boldsymbol{x}_{t+N|t}^T P \boldsymbol{x}_{t+N|t} + \sum_{k=t}^{t+N-1} (\boldsymbol{x}_{k|t}^T Q \boldsymbol{x}_{k|t} + \boldsymbol{u}_{k|t}^T R \boldsymbol{u}_{k|t})] \\ \text{subject to } \quad & \boldsymbol{x}_{k+1|t} = A\boldsymbol{x}_{k|t} + B\boldsymbol{u}_{k|t} + D\boldsymbol{d}_{k|t} \\ & F_x \hat{\boldsymbol{x}}_{k|t} + F_x \Phi^{-1}(1-\beta)\Sigma_{k|t} \le g_x \\ & F_u \boldsymbol{u}_{k|t} \le g_u, \\ & F_f \hat{\boldsymbol{x}}_{t+N|t} + F_f \Phi^{-1}(1-\beta)\Sigma_{t+N|t} \le g_f \\ & \boldsymbol{x}_{t|t} = \boldsymbol{x}(t) \, \forall k \in \{t, \ldots t+N-1\}. \end{aligned}$$

However, if the uncertainty distribution is non-Gaussian, the uncertain MPC (S2) is not easily reformulated into a convex optimization program. In this case, one possible solution is to sample the uncertainty distribution and reformulate the MPC optimization as a much more computationally expensive mixed-integer program [S1], [S2]. Many tail risk measures, such as CVaR and EVaR, provide intuitive convex inner approximations of chance constraints regardless of the uncertainty distribution. Hence, we advocate that risk-aware MPC formulations not only better account for uncertainty but also provide an efficient convex reformulation without making assumptions about the nature of the uncertainty. The resulting deterministic risk-aware MPC formulation is given by

$$\begin{aligned} J_t^*(\boldsymbol{x}(t)) = \min_{U_t} \quad & \rho[\boldsymbol{x}_{t+N|t}^T P \boldsymbol{x}_{t+N|t} + \sum_{k=t}^{t+N-1} (x_{k|t}^T Q x_{k|t} + u_{k|t}^T R u_{k|t})] \\ \text{subject to } \quad & \boldsymbol{x}_{k+1|t} = A\boldsymbol{x}_{k|t} + B\boldsymbol{u}_{k|t} + D\boldsymbol{d}_{k|t} \\ & \rho_\beta(F_x \boldsymbol{x}_{k|t} - g_x) \le 0, \boldsymbol{u}_{k|t} \in \mathcal{U} \\ & \rho_\beta(F_f \boldsymbol{x}_{t+N|t} - g_f) \le 0 \\ & \boldsymbol{x}_{t|t} = \boldsymbol{x}(t) \quad \forall k \in \{t, \ldots t+N-1\}. \end{aligned}$$

REFERENCES

[S1] L. Blackmore, "A probabilistic particle control approach to optimal, robust predictive control," in *Proc. AIAA Guid., Navigation, Control Conf. Exhib.*, 2006, p. 6240, doi: 10.2514/6.2006-6240.

[S2] L. Blackmore, M. Ono, A. Bektassov, and B. C. Williams, "A probabilistic particle-control approximation of chance-constrained stochastic predictive control," *IEEE Trans. Robot.*, vol. 26, no. 3, pp. 502–517, Jun. 2010, doi: 10.1109/TRO.2010.2044948.

second-order moment ambiguity sets is posed as a semidefinite program in [50]. A tree-based approach for MPC that enumerates all possible extreme disturbance signals and searches for feedback policies that account for a tradeoff between robustness and performance through CVaR metrics was proposed in [51]. In [52], the authors considered multistage risk-averse and risk-constrained optimal control for general coherent risk measures with conic representations. A sampling-based approach, using a model predictive path integral to solve nonlinear optimal control problems using CVaR risk constraints, was considered in [113]. Other sampling-based approaches to solving the optimal control problem with state estimation uncertainty involve using a robust version of the RRT* algorithm [114].

Data-driven MPC that uses samples from the uncertainty distribution is becoming increasingly popular. Risk-aware MPC approaches provide the required robustness to account for the gap between enforcing the sample-based chance

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

constraints for the empirical distribution and the true chance constraint for the actual uncertainty distribution. In [53] and [54], the authors propose a distributionally robust data-enabled predictive control (DeePC) algorithm that uses finite samples of an unknown system to make trajectory predictions. Instead of learning the system dynamics model, the authors propose an *equivalent* formulation using these data-driven trajectory predictions that enjoys strong out-of-sample guarantees using Wasserstein distributionally robust CVaR constraints. Reference [55] considers a learning MPC framework whose infinite horizon CVaR-constrained optimal control solution is approximated iteratively given a finite number of safe states and uncertainty samples. Through this iterative method, the authors construct a data-driven terminal set for distributionally robust CVaR-constrained iterative MPC with safety and feasibility guarantees.

MPC is also useful for obstacle avoidance in motion planning tasks. Risk-aware MPC accounts for varied obstacle behaviors and sensor and process noise, not limited to Gaussian distributions. The MPC scheme in [56] avoids moving obstacles using a CVaR risk metric. Similar results were obtained in [57] on the EVaR metric for obstacle avoidance with additional guarantees of recursive feasibility and finite-time task completion while following a set of waypoints. In [24], these results were extended to general coherent risk measures for systems with process noise to obtain a disturbance feedback policy. The authors of [24] propose various constraint-tightening techniques to make the risk-aware obstacle avoidance MPC computationally tractable for motion planning while providing probabilistic guarantees on recursive feasibility and task completion in finite time. However, recursive feasibility for general nonlinear systems with guarantees that hold *at all times* remains an open problem.[4] A risk-constrained MPC formulation was also studied in [19] and [58], in which the authors calculated a risk map to traverse rough terrain using CVaR. They incorporated this CVaR terrain map into MPC constraints to account for obstacles and terrain hazards. All of the aforementioned methods, including the formulation presented in "MPC With Uncertainty," consider pointwise-in-time constraints for safety. Instead, the authors of [117] consider a CVaR-based optimal control formulation for obstacle avoidance such that the risk constraints hold for the entire trajectory by accounting for the timewise supremum safety cost inside the CVaR risk.

Risk-Aware Safety-Critical Control

In the previous sections, we considered behavior and motion planning algorithms through the lens of a risk-aware approach. Risk-aware feedback control methodologies, like the MPC techniques discussed in the previous section, can strike a balance between worst-case and nominal operating conditions and thus account for rare events by enhancing robustness while achieving high performance under nominal operation. Additionally, we must also consider risk with a view to safety-critical autonomous systems, such as those found in aerospace and human–robot applications. The specific risks here are often associated with the uncertainty of modeling intricate nonlinear dynamics, for example, bipedal robots [120], and/or sensing extreme unstructured environments, for example, subterranean or extraterrestrial exploration [121]. Safety in the feedback control layer is often formulated in terms of set-theoretic properties of dynamical systems [122], for example, reachability and invariance. Safety verification then involves ensuring that the system solutions remain within a predefined safe set or steer clear of a predetermined unsafe set. A common technique for this is to calculate the reachable set of a system under disturbances and controls [8], [123], [124]. However, for intricate high-dimensional systems, these methods may be impractical or excessively conservative.

Historically, alternative methods for assessing reachability date back to Nagumo's seminal research [125] on the set invariance of ordinary differential equations (ODEs). This work was later expanded to include ODEs with inputs by Aubin et al. [126] within the framework of viability theory. The rise in interest in hybrid systems in the 2000s led to the development of barrier certificates for safety verification [127]. However, the creation of these certificates involves solving complex polynomial optimization problems that are challenging for high-dimensional systems despite some progress made in the last decade [128]. The newer concept of barrier functions [129] offers a solution to the computational difficulties faced by barrier certificates. These functions can be formulated directly from the safe set's definition, simplifying the process. Using this attribute, barrier functions have been effectively applied to design safe controllers (without an existing controller) and safety filters (with an existing controller) for dynamic systems, such as biped robots [130] and trucks [131]. These applications have demonstrated assured performance and robustness [132].

CVaR is a useful measure for assessing how far a realized trajectory may deviate from a safe region of operation [50], [59], [60], [61], [62], [63]. Defining safety in terms of CVaR is well motivated when constraint violations may be unavoidable; the magnitude of the risk should be minimized during the undesired excursion [50], [62]. The sets of initial conditions whose safety is characterized by notions of CVaR can be estimated using dynamic programming [59], [60], [61], [62], [63]. Pointwise CVaR constraints have also been studied in [59]. The problem of optimizing the CVaR of a maximum random cost has been studied in different settings, such as deriving an upper-bound approximation [61], a finite-horizon solution [62], and an infinite-horizon solution [63].

[4]Please see [115] and [116] for existing approaches to ensure recursive feasibility for stochastic linear MPC.

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

RCBFs

Consider a discrete-time stochastic system given by

$$\boldsymbol{x}(t+1) = f(\boldsymbol{x}(t), \boldsymbol{u}(t), \boldsymbol{d}(t)), \quad x(0) = x_0 \tag{S3}$$

where at time $t \in \mathbb{N}_{\geq 0}$, $\boldsymbol{x}(t) \in \mathcal{X} \subset \mathbb{R}^n$ is the state, $\boldsymbol{u}(t) \in \mathcal{U} \subset \mathbb{R}^m$ is the control input, $\boldsymbol{d}(t) \in \mathcal{D}$ is the stochastic uncertainty/disturbance, and $f : \mathbb{R}^n \times \mathcal{U} \times \mathcal{D} \to \mathbb{R}^n$. We assume that the initial condition $\boldsymbol{x}_0$ is deterministic and that $|\mathcal{D}|$ is finite, that is, $\mathcal{D} = \{v_1, \ldots, v_{|\mathcal{D}|}\}$. At every time step t, for a state-control pair $(\boldsymbol{x}(t), \boldsymbol{u}(t))$, the process disturbance $\boldsymbol{d}(t)$ is drawn from set $\mathcal{D}$ according to the probabilities $p = [p_1, \ldots, p_{|\mathcal{D}|}]^T$, where $p_i := \mathbb{P}(\boldsymbol{d}(t) = v_i)$, $i = 1, 2, \ldots, |\mathcal{D}|$. Note that the probability mass function for the process disturbance is time invariant and that the process disturbance is independent of the process history and of the state-control pair $(\boldsymbol{x}(t), \boldsymbol{u}(t))$. Note that, in particular, the system (S3) can capture stochastic hybrid systems, such as Markovian jump systems.

In risk-aware safety analysis, we are interested in studying the properties of the solutions to (S3) with respect to the compact set $\mathcal{S}$ described by

$$\begin{aligned} \mathcal{S} &:= \{\boldsymbol{x} \in \mathcal{X} \mid h(\boldsymbol{x}) \geq 0\} \\ \text{Int}(\mathcal{S}) &:= \{\boldsymbol{x} \in \mathcal{X} \mid h(\boldsymbol{x}) > 0\} \\ \partial\mathcal{S} &:= \{\boldsymbol{x} \in \mathcal{X} \mid h(\boldsymbol{x}) = 0\} \end{aligned} \tag{S4}$$

where $h : \mathcal{X} \to \mathbb{R}$ is a continuous function.

In the presence of stochastic uncertainties $\boldsymbol{d}$, assuring almost sure (with probability one) invariance or safety may not be feasible. Moreover, enforcing safety in expectation is meaningful only if the law of large numbers can be invoked, and we are interested in the long-term performance, independent of the actual fluctuations. RCBFs focus on safety in the dynamic coherent risk-measure sense with conditional expectation as a special case, allowing for more robust measures of safety.

DEFINITION S1 (ρ-SAFETY)

Given a "safe set" $\mathcal{S}$ in (S4) and a time-consistent dynamic coherent risk measure $\rho_{0:t}$, the solutions to (S3), starting at $\boldsymbol{x}_0 \in \mathcal{S}$, are ρ-safe if and only if

$$\rho_{0,t}(0, 0, \ldots, h(\boldsymbol{x}(t))) \geq 0, \quad \forall t \in \mathbb{N}_{\geq 0}. \tag{S5}$$

When $\boldsymbol{x}_0 \in \mathcal{X} \backslash \mathcal{S}$, we often want to know if $\mathcal{S}$ can be reached in finite time.

DEFINITION S2 (ρ-REACHABILITY)

Consider system (S3) with initial condition $\boldsymbol{x}_0 \in \mathcal{X} \backslash \mathcal{S}$. Given a set $\mathcal{S}$ as in (S4) and a time-consistent dynamic coherent risk measure $\rho_{0:t}$, $\mathcal{S}$ is ρ-reachable if and only if there exists a constant t^* such that

$$\rho_{0,t^*}(0, 0, \ldots, h(\boldsymbol{x}(t^*))) \geq 0. \tag{S6}$$

Safety requirements can also be encoded and enforced via control barrier functions (CBFs), which were proposed in [129]. CBFs have been used to design safe controllers for continuous-time dynamical systems, such as bipedal robots [130] and trucks [131], with guaranteed robustness [132], [133] (see survey [134] and references therein). For discrete-time systems, discrete-time barrier functions were formulated in [135] and [136] and applied to multirobot coordination [137]. For a class of stochastic (Ito) differential equations, safety in probability and statistical mean was studied in [138], [139], [140], [141], and [142] via stochastic barrier functions.

The first attempt to formulate risk-aware CBFs (RCBFs) was carried out in [143], where the authors proposed CVaR CBFs as a composition of a dynamic CVaR metric with a CBF to study safety, in the CVaR sense, for a discrete-time dynamical system subject to stochastic uncertainty. (See "RCBFs" for more details.) A computational method based on difference convex programs (DCPs) was also proposed to synthesize CVaR-safe controllers. The method was applied to collision avoidance scenarios involving a bipedal robot subject to modeling uncertainty. The CVaR CBFs were generalized to RCBFs with general coherent risk measures in [44] and [64], where it was shown that the existence of such barrier functions implies invariance in a coherent risk sense. Furthermore, conditions based on finite-time RCBFs were proposed to guarantee finite-time reachability to a desired set. In recent work [145], sampling-based underapproximations of the CVaR for belief states were used to define risk CBFs.

Definition 7 (RCBF)

For the discrete-time system (S3) and a dynamic coherent risk measure ρ, the continuous function $h : \mathbb{R}^n \to \mathbb{R}$ is an RCBF for the set $\mathcal{S}$ as defined in (S4), if there exists a convex class-$\mathcal{K}$ function α[5] satisfying $\alpha(r) < r$ for all $r > 0$ such that

$$\rho(h(\boldsymbol{x}(t+1))) \geq \alpha(h(\boldsymbol{x}(t))), \quad \forall \boldsymbol{x}(t) \in \mathcal{X}. \tag{15}$$

In [64], the authors demonstrated that the existence of an RCBF implies invariance/safety in the coherent risk measure.

Theorem 1

For the discrete-time system (S3) and the set $\mathcal{S}$ as described in (S4), let ρ be a coherent risk measure. Then, $\mathcal{S}$ is ρ-safe if there exists an RCBF as defined in Definition 7.

[5]A class-$\mathcal{K}$ function is a continuous scalar function $\alpha(r)$ defined for $r \in [0, a)$ that is strictly increasing and satisfies $\alpha(0) = 0$.

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

Note that the most common choice for function α is a constant $\alpha = \alpha_0$, where $\alpha_0 \in (0,1)$ as $\alpha_0 r < r$, $\forall r > 0$. To study risk-aware reachability, we require the following.

Definition 8 (Finite-Time RCBF)

For discrete-time system (S3) and dynamic coherent risk measure ρ, the continuous function $h : \mathcal{X} \to \mathbb{R}$ is a finite-time RCBF for set $\mathcal{S}$, as defined in (S4), if there exist constants $0 < \gamma < 1$ and $\varepsilon > 0$ such that

$$\rho(h(\boldsymbol{x}(t+1))) - \gamma h(\boldsymbol{x}(t)) \geq \varepsilon(1-\gamma), \quad \forall \boldsymbol{x}(t) \in \mathcal{X}. \tag{16}$$

It was also shown in [64] that the existence of a finite-time RCBF implies ρ-reachability.

Theorem 2

Consider the discrete-time system (S3) and a dynamic coherent risk measure ρ. Let $\mathcal{S} \subset \mathcal{X}$ be as described in (S4). If there exists a finite-time RCBF $h : \mathcal{X} \to \mathbb{R}$ as in Definition 8, then for all $\boldsymbol{x}(0) \in \mathcal{X} \setminus \mathcal{S}$, there exists a $t^* \in \mathbb{N}_{\geq 0}$ such that $\mathcal{S}$ is ρ-reachable, that is, inequality (S6) holds. Furthermore

$$t^* \leq \log\left(\frac{\varepsilon - h(\boldsymbol{x}(0))}{\varepsilon}\right) / \log\left(\frac{1}{\gamma}\right) \tag{17}$$

where the constants γ and ε are as defined in Definition 8.

Case Study: Risk-Aware Robotic Motion Planning in Subterranean Environments

This case study examines a hierarchical risk-aware *traversability* and planning methodology that can be used for autonomous robot (legged or wheeled robot) traversal over extreme terrain [19], [58], as motivated by the DARPA Subterranean challenge. This is the first framework to integrate a CVaR-based risk-aware planning and control pipeline into a fully autonomous robotic system. We briefly describe the risk-aware traversability and planning pipeline.

We first need to interpret which parts of the environment the robot can traverse. The evaluation of a natural terrain's *traversability* is difficult due to uncertainties arising from sensor noise and robot localization errors. Furthermore, there are multiple sources of terrain hazards, such as steep slopes, loose surface material, sudden elevation drops, and physical obstacles. To account for these different sources of uncertainty systematically, we evaluate the CVaR of the terrain hazards to obtain a *risk map* that can be used in the planning and control pipeline. The traversability estimate is given by the random variable R that is constructed jointly from the grid map of the terrain, the robot state, and the applied control; see [19] and [58] for details on how to construct the map and associated traversability estimate. R provides a cost of traversing over each grid point on the terrain that we use to assess the CVaR value. This CVaR risk evaluation, $\text{CVaR}_\beta(R)$, enables a robot engineer to define the allowable traversability risk level based on the mission criteria. Furthermore, one can dynamically adjust the risk level, β, online based on 1) the mission-level states, that is, based on the robot's capabilities and the environment, and 2) whether the robot is stuck in a situation wherein there is no feasible path and decreasing the risk level (and consequently being less conservative) might allow the robot to find a feasible but possibly riskier path. The geometric planner and the kinodynamic MPC controller then utilize the risk evaluation, $\text{CVaR}_\beta(R)$, to obtain a risk-aware control policy.

After computing the risk-aware traversability cost, the authors utilize this cost for high-level geometric path planning using an A* algorithm to obtain waypoints for navigation. This risk-aware geometric planner uses the dynamic risk measure in the cost of the geometric planner optimization. The waypoints are then passed into an MPC optimization with CVaR constraints for safety for generating obstacle-free trajectories. The statistical performance of the aforementioned risk-aware controllers is evaluated in simulation with randomly generated environment maps and goals. This study illustrates the tradeoff between the risk taken by the robot to reach the goal versus the total distance traversed by the robot for different allowable risk levels, β (see Figure 4). When $\beta \to 1$ (or $\alpha \to 0$), the framework is risk neutral and mimics the deterministic setting, where we use only the mean value of the traversability estimate. We can see from Figure 4 that the robot takes paths with moderate to high risk more often (that is, paths with maximum risk $>$ ~0.3). Ultimately, the risk level desired is left up to the user, but the main takeaway is that the user can change the risk level to mimic deterministic or robust baselines as desired. The robot uses longer low-risk paths when the robot is risk averse (low

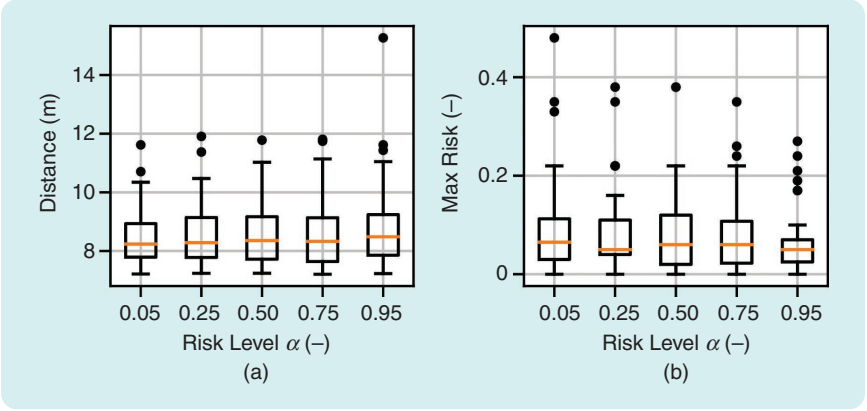


FIGURE 4 (a) The tradeoff between the distance traversed by the robot and different risk levels. (b) The tradeoff between maximum risk taken along the traversed path and different risk levels. Note that $\beta = 1 - \alpha$. (Source: Figure taken from [19].)

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

β) and shorter higher-risk paths when the robot is risk neutral (high β), see Figure 5.

This risk-aware traversability evaluation and planning framework was experimentally tested during the DARPA Subterranean Challenge and in other real-world subterranean environments. The final competition course of the DARPA Subterranean Challenge was composed of tunnel, urban, and cave environments, for which the traversability evaluation and navigation results are provided in Figure 6. The following list describes the difficult terrain hazards found within that environment:

- *Region A:* This is an office-like area with narrow corridors and small rooms where it is tough to find a feasible path if the maps are overinflated to avoid obstacles.
- *Region B*: This is a mock post-earthquake warehouse whose shelving and clutter are difficult to navigate around.
- *Region C*: This consists of a door connecting the urban and tunnel part of the course via stairs. The stairs act as a potential sudden drop-off (that is, a negative obstacle) for wheeled robots. The drop is hard to detect because of the narrow doorway.
- *Region D:* This is a narrow passage littered with debris, vertical pipes along the walls, and ceiling obstacles. The robot must correctly identify the pipes as obstacles.
- *Region E*: This consists of a small cave opening that mimics real caves wherein humans must crawl through the small openings to reach another cave chamber. The upward-sloping cave floor and downward-sloping ceiling make it difficult to differentiate between the ceiling and the ground. The ceiling height at the opening is very close to the ground height at the end of the opening.
- *Region F*: This is a small limestone cave with rubble and loose rock piles. The robot must distinguish between traversable and non-traversable rubble.

A statistical analysis of the simulations and the qualitative experimental results from the field show that a risk-aware traversability and planning pipeline provides a framework where the risk of the *entire* system can be adjusted by changing the risk-level β despite there being multiple risk sources, such as slopes, obstacles, low-ceilings, and mud. This framework is agnostic to the kind of ground robot utilized; it has been tested on wheeled robots (Clearpath Husky) and legged robots (Boston Dynamics Spot quadruped).

Case Study: Risk-Aware Bipedal Walking

Control of bipedal walking presents significant challenges, as evidenced by the variety of approaches taken in the literature to handle the nonlinearity and complexity of bipedal robot dynamics [146]. In practice, bipedal walking dynamics are often simplified by approximate models subject to stochastic uncertainty [147]. The horizontal robot state at time t is denoted by $\boldsymbol{x}_h(t) = [c, p, v]^T$, where c represents the horizontal position of the robot's center of mass (COM) relative to an inertial frame, p denotes the horizontal COM position with respect to the stance foot, and v denotes the horizontal COM velocity. The step-to-step (S2S) dynamics of the horizontal COM state are expressed as $\boldsymbol{x}_h(t+1) = \mathcal{P}^h(\boldsymbol{x}(t), \tau(t))$, where τ represents the joints' input torques. In practice, deriving the S2S dynamics analytically is challenging due to the robot's nonlinear and hybrid dynamics.

The authors in [147] and [148] suggested that a Hybrid-Linear Inverted Pendulum (H-LIP) walking model [148] provides an apt approximation for the actual horizontal S2S dynamics of robot walking. The H-LIP dynamics are represented as

$$\boldsymbol{x}_{\text{H-LIP}}(t+1) = A\boldsymbol{x}_{\text{H-LIP}}(t) + B\boldsymbol{u}_{\text{H-LIP}}(t) \tag{18}$$

where $\boldsymbol{x}_{\text{H-LIP}}(t+1)$ is the H-LIP's $[c_{\text{H-LIP}}, p_{\text{H-LIP}}, v_{\text{H-LIP}}]^T$ discrete preimpact state, and $\boldsymbol{u}_{\text{H-LIP}}(t)$ denotes the step size. The specific expressions for A and B can be found in [147]. With this approximation, the S2S dynamics can be rewritten as

$$\boldsymbol{x}_h(t+1) = A\boldsymbol{x}_h(t) + B\boldsymbol{u}(t) + \boldsymbol{d}(t) \tag{19}$$

where $\boldsymbol{d}(t) := \mathcal{P}^h(x(t), \tau(t)) - A\boldsymbol{x}_h(t) - \ -B\boldsymbol{u}(t) \in \mathcal{D}$ can be viewed as a stochastic disturbance to the linear system.

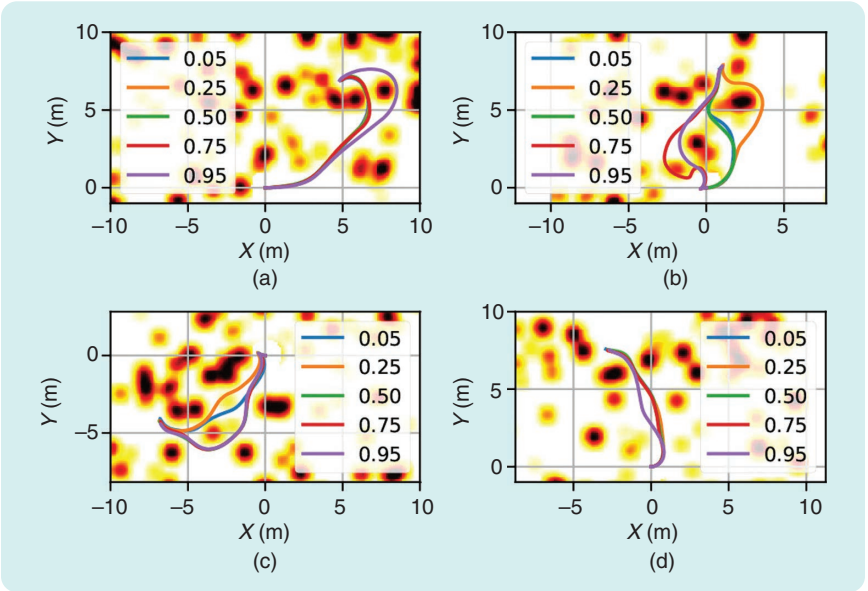


FIGURE 5 (a)–(d) Four instances from a Monte-Carlo simulation illustrate how different choices of risk levels, $1-\beta$, affect the paths taken by the robot. (Source: Figure taken from [19]).

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

As seen in Figure 7, a CVaR RCBF-based controller can ensure safe and risk-aware 3D bipedal walking. The model discrepancy w is treated as a stochastic uncertainty and risk factor that could elicit undesirable walking behavior. To mitigate this risk, CVaR RCBF-based controllers are synthesized and act as safety filters for the H-LIP-based stepping controller. We will omit the details of how these barrier functions were generated for this specific problem to keep the discussion at a high level. For those interested in the details, please reference [143]. Abstractly, however, these barrier functions delineate safe regions within which the bipedal robot can navigate by constraining the horizontal

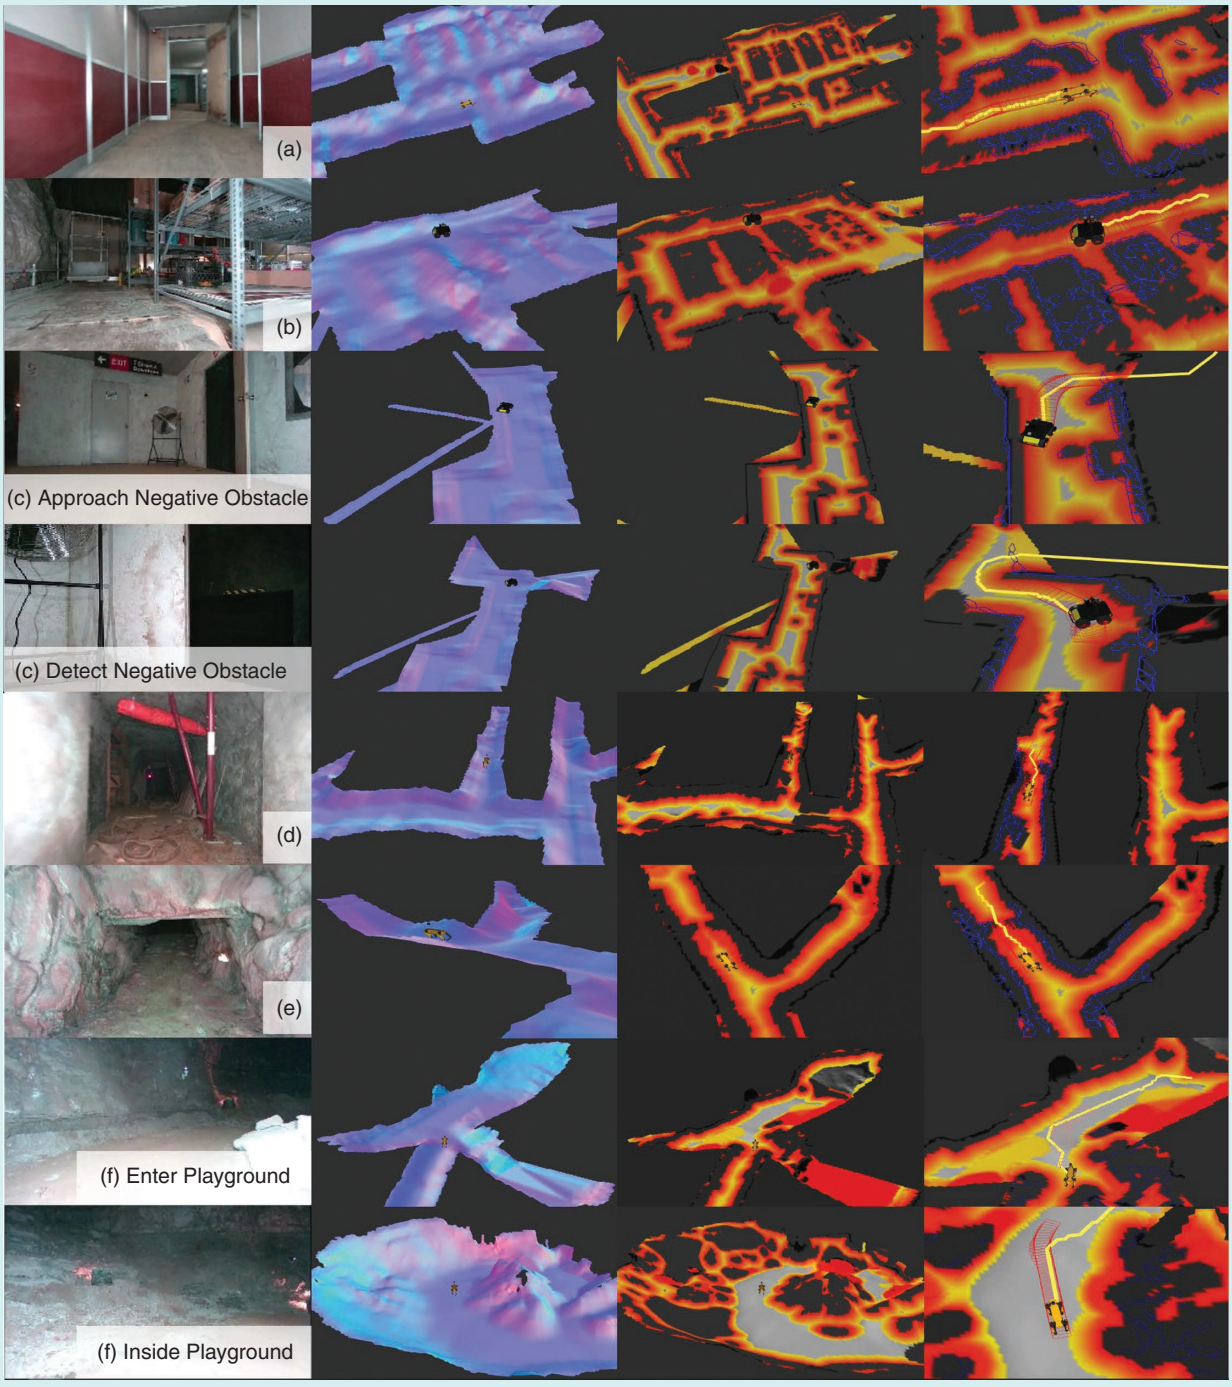


FIGURE 6 A risk-aware traversability analysis of the DARPA Subterranean Challenge final course. Columns, from left to right: robot front camera view, elevation map (colored by normal direction), risk map (colored by risk level—white: safe, yellow to red: moderate, black: risky), and planned geometric/kinodynamic paths (yellow lines/red boxes). (Source: Figure taken from [19].)

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

position of the robot to lie between the obstacles it is navigating between. To construct this CVaR RCBF, the authors carried out extensive simulations over a wide range of walking behaviors to numerically define an uncertainty set. Then, to calculate the CVaR of the RCBF evaluation at the subsequent time steps—as required in the CVaR RCBF condition expressed in [143]—the authors approximated the true distribution via a uniform sampling over the previously calculated uncertainty set and calculated the CVaR against this distribution. Even still, the corresponding walking experiments still represented safe behavior, further supporting the use of RCBFs for risk-aware safe controller synthesis.

The simulation results shown in Figure 8 showcase the effectiveness of the CVaR-based CBF, especially when taking into account the inherent uncertainties in 3D bipedal dynamics. Figure 9 presents snapshots from an experiment conducted at the Caltech AMBER Lab using the Agility Robotics' Cassie bipedal robot.

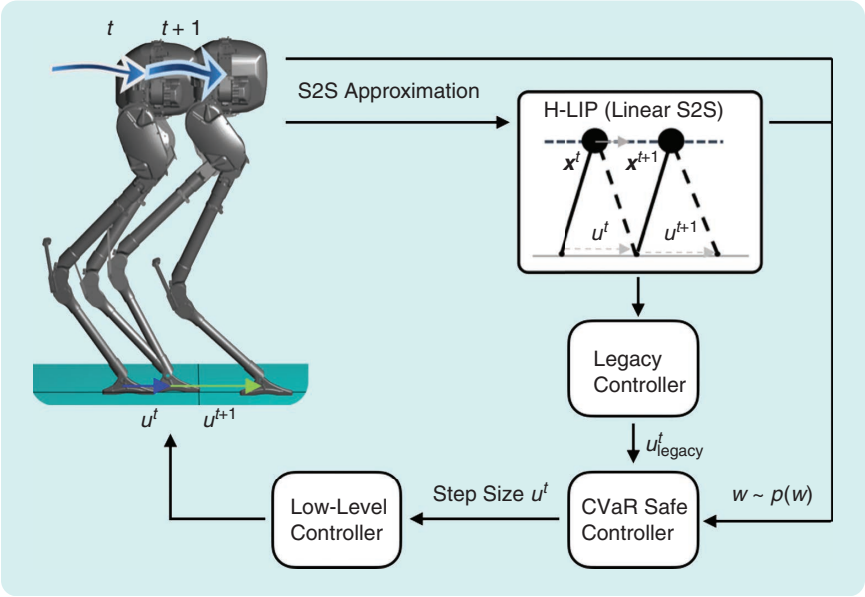


FIGURE 7 A schematic diagram of a risk-aware bipedal robot path planning method based on RCBFs with CVaR risk measure [143]. The diagram features the sequential flow of processes starting with the S2S approximation of robot dynamics, which leverages the dynamics of the walking robot, modeled after the H-LIP system, where $\boldsymbol{x}^t \equiv \boldsymbol{x}(t)$ denotes the horizontal position of the COM of the robot relative to the inertia frame. The H-LIP approximation is controlled via a legacy controller, such as a MPC, which outputs step size commands u_{legacy}. The difference in terms of the horizontal position of the COM between the H-LIP model and simulation/experimental is measured offline and used to construct the discrete distribution over the uncertainty $p(w)$. The distribution over uncertainty is used to tune a CVaR-safe controller based on RCBFs. The outputs of the legacy controller are then amended online using the CVaR-safe controller to ensure risk-aware safety in the presence of uncertainty $p(w)$, which adjusts the robot's locomotion parameters (in particular, step size u) in real time.

Open Questions and Future Directions

This section outlines recent advances in risk-aware planning and control. We applied these ideas to bipedal walking and terrain traversability analysis for wheeled and legged robots. Many future research directions are suggested by current work in risk-aware planning and control.

Computation. The best choice of a risk measure for a specific problem remains an open question. The popular CVaR risk measure is computationally attractive as it can be formed into linear programs when the original stochastic optimal control program is also linear, that is, when the cost, dynamics, and constraints are linear. Other risk measures, such as the Kullback–Leibler divergence-based EVaR metric or the Wasserstein metric, provide rich expressions of the uncertainty but are computationally expensive. Sampling-based methods for risk computation can also require many samples to guarantee the correctness of the resulting risk-aware plan, which remains challenging with computationally limited hardware, that is, without a GPU.

Risk-Aware Planning With Nonlinear Dynamics. Keeping with the computation theme, forward computations of nonlinear dynamics with complex policies are already difficult. Adding risk-aware planning further complicates the issue, making such a procedure currently intractable for nonsample-based approaches. Even still, real-time applications of these methods further limit the number of samples that can be taken to inform planning or control input selection, making efficient estimation of risk measures in these scenarios an open problem. Recent work indicates that parallelizing computation over a sample-based scenario approach might improve efficiency, but such an approach is still computationally intensive [149].

Multiagent Interactions. Our discussion of risk-aware planning and control only considered a single agent. However, real-world dynamic agents may react to the motion of the controlled agent, and these effects could potentially cause unmodeled uncertainty distribution shifts. An important open problem is how to account for the interactions between dynamic agents in a risk-aware manner.

Approximations of Risk. Many risk-aware planning and control techniques either assume that the uncertainty is discrete or use approximation techniques, like sample average approximation (SAA), for continuous distributions. How can we guarantee the correctness of risk evaluation and control design when using continuous distribution approximations? The next section introduces methods to verify the risk-aware behavior of autonomous systems.

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

Constructing CBFs. In the previous "Risk-Aware Safety-Critical Control" section, we introduced CBFs as a potential method by which a controller can ensure safety even in a risk-aware context. However, it is worth noting that constructing CBFs—let alone the risk-aware formulations—remains an open problem for complex systems. There has been recent work aimed at learning CBFs [150], [151], [152], [153] and constructing them numerically via methods such as sum of squares [154], [155], [156] (see [157] for a tutorial on the sum of squares method). As such, doing so even in a risk-aware context remains an open problem.

RISK-AWARE V&V

The previous section provided a high-level summary of risk-aware planning and control, with a more in-depth review of the existing literature. We highlighted the importance of analyzing the inherent uncertainties and risks in robotic operations, particularly when navigating through unstructured environments. This section briefly summarizes the important companion problem of "risk-aware V&V" in robotic systems. The verification process determines whether a given system exhibits its desired behavior in the environments in which it is required to operate. This crucial process ensures that the integrated system performs safely, reliably, and as intended under a broad range of operating conditions.

The integration of risk awareness into the V&V process allows for a more comprehensive evaluation of a robotic system and its potential to properly respond to risky situations. Specifically, V&V aims to quantify robotic reliability and safety, explicitly considering interactions with uncertain and unstructured environments [158], [159], [160], [161]. As such, risk-aware V&V requires probabilistic risk assessments, stochastic models, and rigorous testing methods that cover a wide range of

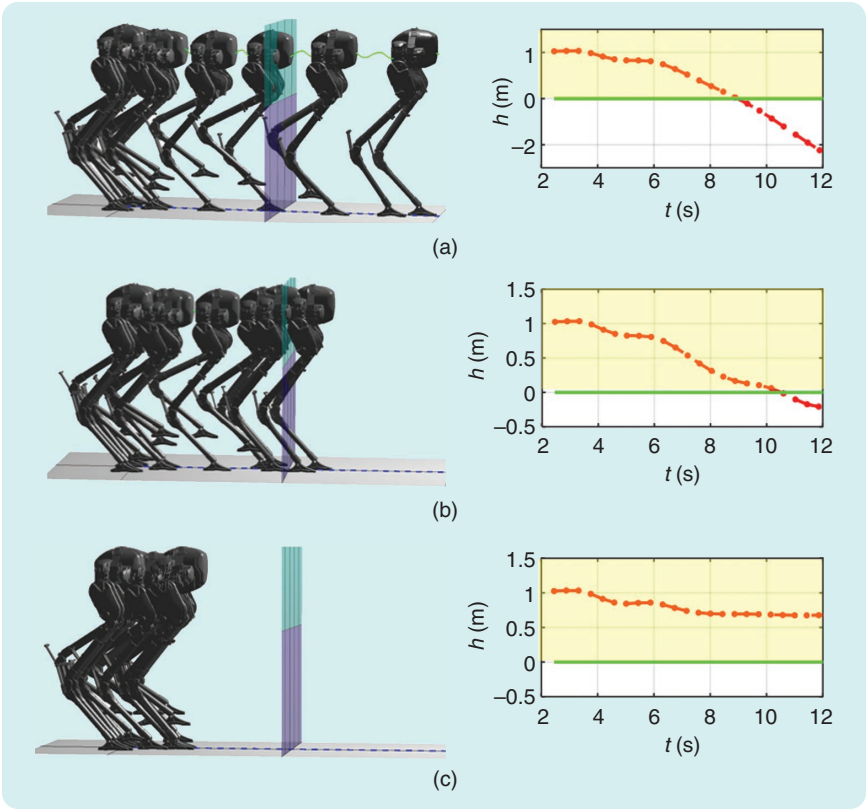


FIGURE 8 Risk-averse obstacle avoidance using CVaR barrier functions (robot behavior and barrier function evolution). The shaded yellow area denotes safe regions. (a) Safety violation with no barrier function. (b) Safety violation with a risk-neutral barrier function. (c) Safe behavior with CVaR barrier function. Plots on the right side show the values of the barrier functions [143].

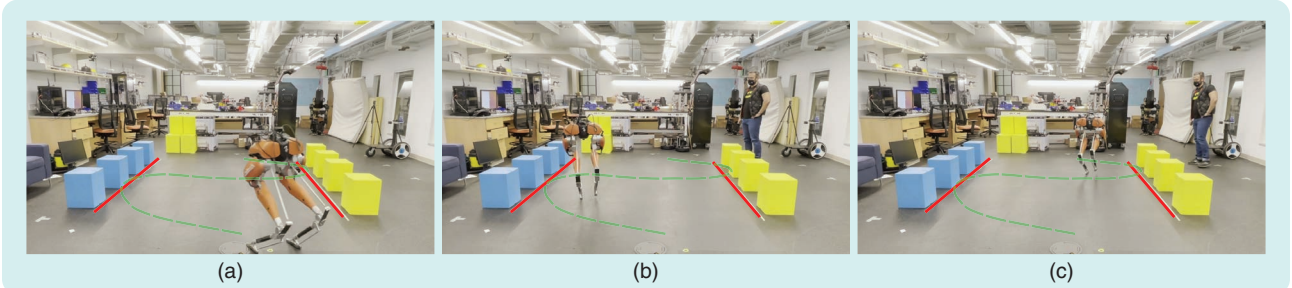


FIGURE 9 Snapshots from an experiment conducted at Caltech's AMBER Lab, featuring Agility Robotics' Cassie Robot. The OptiTrack system that was used for localization tracks reflective markers on the robot to determine its position and orientation with high precision and provides these data to the robot's onboard computer for dynamic movement updates. (a) The initial setup, where the robot's trajectory is aligned with a sinusoidal path, represented by a dashed green line. (b) Midcourse navigation highlighting the effectiveness of the risk CBF-based safety filter. This filter is designed to allow the robot to dynamically avoid obstacles and unsafe regions, which are indicated by the red solid lines, representing the boundaries of areas the robot should not enter. (c) Successfully following a trajectory that has been adjusted by the risk-aware safety filter. This demonstrates the practical application of the risk-aware control method outlined in Figure 7, where a robot not only plans its path in consideration of potential risks but also dynamically adjusts its course in real time to maintain safe navigation.

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

potential scenarios. For instance, recent work in human–robot interaction has developed procedures to quantify the riskiness of actions taken by systems in a collaborative context [162], [163], [164], [165]. These methods typically formalize the risk assessment against existing international standards, for example, ISO 14121 [166] and ISO 12100 [167]. Risk assessment of autonomous systems in other fields has also emerged [168], [169], [170], [171].

As mentioned, previous notions of risk have typically been defined against a corresponding standard and are developed on a case-by-case basis. This observation has prompted recent work in risk-aware verification to adopt the same formal treatment of risk—tail risk measures—as used by the synthesis community [21], [172]. This section delves into the key methodologies and approaches employed for risk-aware V&V in robotics, starting with a brief overview of its theoretical foundations and moving toward practical applications and case studies. As works in this area typically exploit the quantifiable semantics of temporal logics to quantify (un)safe system behavior to make risk-aware verification statements, we begin with a brief overview of temporal logics [173], [174], [175].

Temporal Logic

Temporal logics can be used to express complex system specifications and were originally developed for the analysis and design of reactive systems, that is, systems with external inputs, such as control systems [176], [177], [178]. Temporal logics are extensions of Boolean logic (for example, propositions, negations, conjunctions, and disjunctions) by adding temporal operators (for example, until, eventually, and always) to reason about the temporal properties of a system. One can make a distinction between temporal logics that reason over qualitative and quantitative temporal properties in their temporal operators. Linear temporal logic, arguably the most common temporal logic, reasons only over qualitative temporal properties, while real-time temporal logics, such as metric temporal logic [179], can reason over quantitative temporal properties. This distinction is best exemplified where a qualitative temporal property is "eventually reach the goal region" while its quantitative counterpart could be "eventually within the next 5 min reach the goal region." We focus on temporal logic specifications with quantitative temporal reasoning that can encode combinations of timed reachability ("reach region A within 30 s"), timed surveillance ("visit regions B, C, and D every 10–60 s while agents form a triangular formation"), timed safety ("always between 5 and 25 s stay at least 1 m away from region E"), and many others.

Temporal logics are formally defined by their syntax and semantics, where the syntax defines the rules to construct a temporal logic specification ϕ while the semantics define when a temporal logic specification ϕ is satisfied (or violated). Spatiotemporal logics, as opposed to temporal logics, also permit reasoning about spatial properties, for example, allowing a system designer to quantify to what extent (with what safety margin) an obstacle is avoided by a robot. It is this property that enables us to quantify how well a specification is satisfied, cf., how severely a specification is violated, which in turn helps define risk for system verification. Signal temporal logic (STL) is a commonly used spatiotemporal logic introduced in [180], and we provide a brief introduction to its syntax and semantics in "Signal Temporal Logic: Syntax and Semantics" and "Signal Temporal Logic: Robust Semantics." For the remainder of the exposition in this section, though, assume that ϕ denotes a system specification.

Concerning the use of these logics in V&V, over the past decades, the formal methods community proposed and studied a broad range of system verification techniques. Existing techniques focus on the verification of 1) deterministic systems or 2) uncertain systems, with the two previously discussed viewpoints of the risk-neutral and the worst-case paradigms. In fact, automated verification tools have been developed for deterministic systems, for example, model checking [177], [178] and theorem proving [181], [182]. Verification of uncertain systems was particularly studied in the risk-neutral paradigm using probabilistic model checking [183], [184], [185], an extension of deterministic model checking, and statistical model checking (SMC) [186], [187], [188], [189], which are sampling-based techniques for probabilistic system verification. System verification techniques, however, should not only be able to reason about the probability of violating a specification but also be able to reason about the severity of a violation in terms of rare, harmful outcomes. As briefly discussed previously, spatiotemporal logics enable us to quantify how well (severely) a specification is satisfied (violated), and in turn, allow us to define risk in terms of this quantitative measure for stochastic systems.

Motivations for Tail Risk Measures in Verification

As mentioned previously, the existence of these spatiotemporal logics makes tail risk measures uniquely suited to serve as the risk measure of choice for V&V, and this section will provide an example supporting that claim using the notation offered in "Signal Temporal Logic: Robust Semantics." Consider, for the sake of argument, that we have two controlled systems Σ_1, Σ_2, a signal temporal logic specification ϕ denoting the desired behavior required of both systems, and a robustness measure ρ^ϕ for the same specification ϕ. For context, every signal temporal logic specification ϕ comes equipped with a quantitative measure. Let's further assume that every time we query either system Σ_1 or Σ_2, we receive a random trajectory $\boldsymbol{x}_1$ or $\boldsymbol{x}_2$, respectively. Let R_i, $i = 1,2$, be a random variable whose

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

samples r_i are the robustness of the trajectories sampled from system Σ_i, that is, $r_i = \rho^{\phi}(\boldsymbol{x}_i)$. Now, let's assume that an oracle tells us that for both systems, with probability $1-\beta$ for some $\beta \in (0,1)$, the random robustness exceeds a cutoff value $\epsilon > 0$. Furthermore, with probability β, Σ_1 exhibits robustness values between ϵ and 0, whereas Σ_2 exhibits robustness values strictly less than zero. In other words, Σ_1 will always realize the desired behavior, while Σ_2 will, in some rare cases, be unable to realize the desired behavior.

Put into the context of tail risk measures, both systems exhibit a robustness VaR at risk level β that is positive. This fact arises as the oracle mentioned that with minimum probability $1-\beta$, both systems exhibit robustnesses exceeding $\epsilon > 0$, that is, $\mathrm{VaR}_{\beta}(R_i) = \epsilon > 0$. Ending the analysis here results in the correct conclusion that both systems exhibit some minimum probability of realizing satisfactory behavior, and this is a traditional probabilistic V&V statement. However, by considering CVaR, we can further discriminate between the two systems as system 1 is expected to exhibit satisfactory behavior even in the worst $100\beta\%$ of cases, whereas system 2 is expected to produce unsatisfactory behavior in the same cases. This conclusion arises as even in the worst $100\beta\%$ of cases, the oracle mentioned that system 1 exhibits robustness values $r_1 \in [0,\epsilon]$, whereas system 2 exhibits robustness values $r_2 < 0$. Taking into account the expected value over those cases—the definition of CvaR—we conclude that $\mathrm{CVaR}_{\beta}(R_1) \geq 0$, whereas $\mathrm{CVaR}_{\beta}(R_2) < 0$. Therefore, even if both systems exhibit similar minimum probabilities of specification satisfaction, system 1 is "better" than system 2 as it is still expected to exhibit satisfactory behavior in the worst $100\beta\%$ of cases.

The example described previously highlights the utility of tail risk measures in risk-aware V&V. By considering the robustness VaR, we can make statements on the minimum probability with which a system exhibits a desired behavior in its operating environment(s)—this is the traditional notion of probabilistic V&V. Additionally, we can utilize tail risk measures also to make statements on expected worst-case robustness using the CVaR, lower-bound such expected worst-case robustness using EVaR, and calculate these values without exact distribution knowledge, as will be described in sections to follow.

A General Overview of Risk-Aware V&V

For most relevant problems, the system to be verified can be recast as a discrete-time system with known state and input spaces and (perhaps) known dynamics and disturbance spaces. Formally, at some time $t \in \mathbb{Z}_+ = \{0,1,2,\ldots\}$, let $\boldsymbol{x}(t) \in \mathcal{X}$ be the system state, $\boldsymbol{u}(t) \in \mathcal{U}$ be the system control input, and $\boldsymbol{d}(t) \in \mathcal{D}$ be a

Signal Temporal Logic: Syntax and Semantics

STL specifications are interpreted over continuous-time signals $\boldsymbol{x}: \mathbb{R}_{\geq 0} \to \mathbb{R}^n$. An STL specification ϕ is recursively constructed from atomic predicates by using Boolean operators and temporal operators. These atomic predicates are Boolean-valued functions $\mu: \mathbb{R}^n \to \{1,0\}$ whose truth value is obtained after the evaluation of a real-valued function $b: \mathbb{R}^n \to \mathbb{R}$. At time t, we obtain the truth value of μ as

$$\mu(x(t)) := \begin{cases} 1 & \text{if } b(x(t)) \geq 0 \\ 0 & \text{otherwise.} \end{cases} \tag{S7}$$

Predicate functions h can encode relationships between state variables, such as relative and absolute distances. The syntax of STL defines a set of rules according to which well-defined STL specifications can be constructed and is given as

$$\phi ::= 1 \mid \mu \mid \neg\phi' \mid \phi' \wedge \phi'' \mid \phi' U_I \phi'' \tag{S8}$$

where the operators $\neg$, $\wedge$, and U_I encode negations, conjunctions, and the until over the time interval $I \subseteq \mathbb{R}_{\geq 0}$, respectively. The syntax in (S8) can be understood as follows; the symbol $::=$ assigns one of the expressions from the right-hand side, which are separated by vertical bars, to the free variable ϕ on the left-hand side. The variables ϕ' and ϕ'' on the right-hand side are already well-defined STL specifications. While the meaning of negations and conjunctions is clear, the until operator $\phi' U_I \phi''$ encodes that ϕ' has to hold until ϕ'' holds, which has to happen within the time interval I. We can now use logical equivalences to derive the Boolean disjunction, implication, and equivalence operators and the temporal eventually and always operators. In what follows, $\top$ denotes *true* in the corresponding logical specification:

$$\begin{aligned}
\phi' \vee \phi'' &:= \neg(\neg\phi' \wedge \neg\phi'') && \text{(disjunction)} \\
\phi' \Rightarrow \phi'' &:= \neg\phi' \vee \phi'' && \text{(implication)} \\
\phi' \Leftrightarrow \phi'' &:= (\phi' \Rightarrow \phi'') \wedge (\phi'' \Rightarrow \phi') && \text{(equivalence)} \\
F_I \phi' &:= \top U_I \phi' && \text{(eventually)} \\
G_I \phi &:= \neg F_I \neg\phi' && \text{(always).}
\end{aligned}$$

The semantics of STL define when a signal x satisfies an STL specification ϕ. These semantics are formally defined as a relation $\vDash$ between x and ϕ, and $(x,t) \vDash \phi$ means that the signal x satisfies the specification ϕ at time t. We recursively define the semantics as

$$\begin{aligned}
(x,t) &\vDash 1 && \text{iff} && \text{holds by definition} \\
(x,t) &\vDash \mu && \text{iff} && h(x(t)) \geq 0 \\
(x,t) &\vDash \neg\phi' && \text{iff} && (x,t) \nvDash \phi' \\
(x,t) &\vDash \phi' \wedge \phi'' && \text{iff} && (x,t) \vDash \phi' \text{ and } (x,t) \vDash \phi'' \\
(x,t) &\vDash \phi' U_I \phi'' && \text{iff} && \exists t'' \in t \oplus I \text{ s.t. } (x,t'') \vDash \phi'' \text{ and} \\
& && && \forall t' \in (t,t''),\ (x,t') \vDash \phi'.
\end{aligned}$$

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

randomly sampled disturbance. Then, for some distribution function $\xi : \mathcal{X} \times \mathcal{U} \times \mathbb{Z}_+ \times \mathcal{D} \to [0,1]$ over $\mathcal{D}$

$$x(t+1) = f(x(t), u(t), d(t)) \tag{20a}$$
$$\text{s.t. } d(t) \sim \pi(x(t), u(t), t). \tag{20b}$$

Finally, let $\mathbb{U} : \mathcal{X} \times \Theta \times \mathbb{Z}_+ \to \mathcal{U}$ be the controller closing the loop for the system to be verified. Note that the controller $\mathbb{U}$ is parameterized with a parameter $\theta \in \Theta$ to account for exogenous user-specific inputs that may influence controller behavior, for example, parameterized three-space locations for packages in a warehouse that a warehouse robot receives on the fly from a central command station when a package is required to be collected.

Remark. For systems that operate in adversarial environments or in the presence of obstacles, the disturbance distribution ξ can be defined as the singleton distribution over the adversarial choice at the given state $x(t)$, input $u(t)$, and time t. For more information, see Dirac distributions [190] and adversarial testing works such as [191], [192], [193], and [194]. We can consider both cases—the case with nonadversarial randomized disturbances and the case with adversarial or otherwise known disturbances—in the same stochastic setting. Furthermore, the state and input spaces have been left arbitrary to allow both the continuous space definitions in the controls community and the finite-state and input definitions in the POMDP computer science literature.

Closing the loop between (20) and the assumed controller $\mathbb{U}$ generates a system Σ that, when queried with a specific initial condition $x_0 \in \mathcal{X}_0 \subseteq \mathcal{X}$ and controller parameter $\theta \in \Theta$, produces a (perhaps) different state trajectory x, which is an element of the signal space $\mathcal{S} = \{s : \mathbb{Z}_+ \to \mathcal{X}\}$

$$\text{a sample of } \Sigma(x_0, \theta) \text{ is } x \triangleq \{x_0 \equiv x(0), x(1), \ldots\} \tag{21a}$$
$$\text{subject to } x(t+1) = f(x(t), \mathbb{U}(x(t), \theta), d(t)). \tag{21b}$$

Finally, let C be a classifier function mapping from the signal and parameter spaces to the real line, that is, $C : \mathcal{S} \times \Theta \to [-a, b]$ where parameters $a, b \in \mathbb{R}_{++}$ are finite. The classifier C delineates between satisfactory behavior—trajectory and parameter pairs (x, θ) that evaluate to a positive value, that is, $C(x, \theta) \geq 0$—and unsatisfactory behavior—pairs that evaluate to a negative value. Examples of such a classifier could be the robustness functions ρ from signal temporal logic, the minimum value of a barrier function h over time [129], etc. For a description of robustness measures in signal temporal logic, please see "Signal Temporal Logic: Robust Semantics." Generally speaking, we define the outcome of function C to be the *robustness* of the corresponding trajectory and parameter pair, that is, for $r = C(x, \theta)$, where r is the trajectory and parameter pair's robustness value. More positive values of $C(x, \theta)$ indicate a better and *more robust* realization of the desired behavior.

Remark. The rationale to analyze multiple nonunique trajectories x as defined in (21) arises from the fact that disturbances d in (20) are sampled randomly at each time t from the distribution function ξ. If the distribution function ξ were the singleton distribution corresponding to a specific disturbance $d \in \mathcal{D}$ for each state, input, and time $(x, u, k) \in \mathcal{X} \times \mathcal{U} \times \mathbb{Z}_+$, then $\Sigma(x_0, \theta)$ would always produce the same trajectory x upon successive queries at (x_0, θ), and this argument holds $\forall\ (x_0, \theta) \in \mathcal{X}_0 \times \Theta$.

The goal of risk-aware verification, then, is to bound the risk-measure evaluation of the classifications of these (perhaps) stochastically evolving trajectories. Stated formally, let χ be a tail risk measure, for example, VaR, CVaR, EVaR, etc.[6] Then, at some risk-level $\beta \in [0,1]$, determine an upper or lower bound to $\chi(C(\Sigma(x_0, \theta)))$ for some $(x_0, \theta) \in \mathcal{X}_0 \times \Theta$.

Signal Temporal Logic: Robust Semantics

While the STL semantics tell us if a signal $x : \mathbb{R}_{\geq 0} \to \mathbb{R}^n$ satisfies an STL specification ϕ, it does not give us any information about the quality of satisfaction. To obtain such information, one can define robust semantics that quantify how robustly the signal x satisfies the specification ϕ. If x satisfies ϕ, we would hence like to quantify how different a signal $x^* : \mathbb{R} \to \mathbb{R}^n$ can be from x while still satisfying ϕ. To quantify this, we first define the closeness of two signals $x, x^* : \mathbb{R} \to \mathbb{R}^n$ as

$$d(x, x^*) := \sup_{t \in \mathbb{R}_{\geq 0}} \| x(t) - x^*(t) \|.$$

We now want to compute a value ρ^ϕ such that all signals x^* that are such that $d(x, x^*) < \rho^\phi$ will also satisfy ϕ, that is, $(x^*, t) \vDash \phi$. To do so, we first define the signed distance of the signal value $x(t)$ to the set of states that satisfy a predicate μ, denoted by $O^\mu : \{x \in \mathbb{R}^n \mid b(x) \geq 0\}$, as

$$\text{Dist}(x(t), O^\mu) := \begin{cases} \inf_{x^* \in \text{cl}(\mathbb{R}^n \setminus O^\mu)} \| x^* - x(t) \| & \text{if } x \in O^\mu \\ -\inf_{x^* \in \text{cl}(O^\mu)} \| x^* - x(t) \| & \text{otherwise.} \end{cases}$$

Note that $\text{Dist}(x(t), O^\mu)$ quantifies the extent to which μ is satisfied (if $\text{Dist}(x(t), O^\mu) > 0$) or violated (if $\text{Dist}(x(t), O^\mu) < 0$). We can now recursively define the robust semantics of ϕ as a real-valued function $\rho^\phi(x, t)$ as follows:

$$\begin{aligned} \rho^\top(x, t) &:= \infty \\ \rho^\mu(x, t) &:= \text{Dist}(x(t), O^\mu) \\ \rho^{\neg\phi'}(x, t) &:= -\rho^{\phi'}(x, t) \\ \rho^{\phi' \wedge \phi''}(x, t) &:= \min(\rho^{\phi'}(x, t), \rho^{\phi}(x, t)) \\ \rho^{\phi' U_I \phi''}(x, t) &:= \sup_{t'' \in t \oplus I} \min(\rho^{\phi''}(x, t''), \inf_{t' \in (t, t'')} \rho^{\phi'}(x, t')). \end{aligned}$$

Finally, it holds that $(x, t) \vDash \phi$ if and only if $(x^*, t) \vDash \phi$ for all signals $x^* : \mathbb{R} \to \mathbb{R}^n$ that are such that $d(x, x^*) < |\rho^\phi(x, t)|$.

[6] We select χ here to denote risk measures to avoid confusion with the notation of the robust semantics ρ^ϕ of an STL formula ϕ.

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

Figure 10 depicts this generic risk-aware verification pipeline, and the following section will specify how this pipeline has been implemented in a variety of recent works. To facilitate that discussion, we will define the *robustness* $R(x_0,\theta)$ to be the random variable whose samples $r = \boldsymbol{C}(\boldsymbol{x},\theta)$, where $\boldsymbol{x}$ is a sample of the random variable $\Sigma(x_0,\theta)$. In other words, $R(x_0,\theta) = \boldsymbol{C}(\Sigma(x_0,\theta),\theta)$—this term was first defined in [66], [73], and [75]. We provide details on the sampling-based estimation of $\chi(\boldsymbol{C}(\Sigma(x_0,\theta)))$ for tail risk measures χ in "Concentration Inequalities: Risk-Measure Estimation."

Examples

Perhaps the most prevalent examples of risk-aware verification arise from a reframing of traditional work in the SMC community [189], [195], [196], [197]. With respect to the aforementioned pipeline, SMC assumes the ability to collect system traces—trajectories $\boldsymbol{x}$—and evaluate their satisfaction of a desired behavior. These behaviors are typically expressed as a specification ϕ in Probabilistic Computational Tree Logic [198], which is a form of temporal logic (see "Signal Temporal Logic: Robust Semantics"). As such, each of these behaviors has satisfiability metrics—classifier functions $\boldsymbol{C}$ in our overarching methodology—with which to determine the trace satisfaction of the desired behavior ϕ.

SMC consists of two different analyses. The first, hypothesis testing, asserts that the system Σ realizes the behavior ϕ with minimum probability p and determines the minimum number of system traces that have to be evaluated to accept or reject this hypothesis. The second, estimation, exploits either the Chernoff bound or Hoeffding's inequality to estimate the probability p with which Σ realizes ϕ within some tolerance bounds that are a function of the number of trajectories sampled and evaluated. In both cases, however, the probability of satisfaction p has a one-to-one correspondence with the VaR of the random variable $R(x_0)$ (we omit θ in the notation here as SMC typically does not consider parameterized trajectories). More specifically, p is such that $\mathrm{VaR}_{1-p}(R(x_0)) \geq 0$.

These are not the only works that take a VaR approach to system verification. In [65], the authors use scenario optimization to lower-bound the VaR of the robustness random variable $R(x_0,\theta)$ for a user-defined $\beta \in [0,1]$. Similarly, the authors of [66] use a sample-average-approximation procedure to estimate the VaR of the same robustness variable for any user-defined $\beta \in [0,1]$. In [67], the authors go one step further and express policy or controller synthesis as an optimization problem over a general class of risk measures for verification purposes. They show numerical examples of the success of a convex-concave procedure in identifying such policies for an MDP. In this case, the policies optimize for a certain risk sensitivity as expressed by VaR among other risk measures expressed in Cumulative Prospect Theory [199]. Finally, in [68], the authors modify a learned controller online whenever the learned controller outputs an infeasible trajectory. Via a gradient-descent method, they update controller parameters until the resulting trajectory passes an intermediary risk-aware verification step before the implementation of the modified controller. In general, however, any of the aforementioned risk-aware works could also be conceived of as VaR-based verification insofar as the classifier functions $\boldsymbol{C}$ were developed against specific standards for their respective applications [162], [163], [164], [165], [168], [169], [170], [171].

However, VaR verification represents a smaller fraction of risk-aware verification efforts as compared to works using coherent risk measures, such as CVaR. For example, in a similar paradigm as in [67], in [69], the author proves that there exist polynomial time algorithms to determine policies for an MDP that are verifiable by default. Verification arises as the policies are synthesized to achieve a minimum CvaR with respect to objective satisfaction. Similarly, in [70], the authors utilize a CVaR constraint for their optimal controller and verify that the system remains within a risk-sensitive safe set defined by the same CVaR

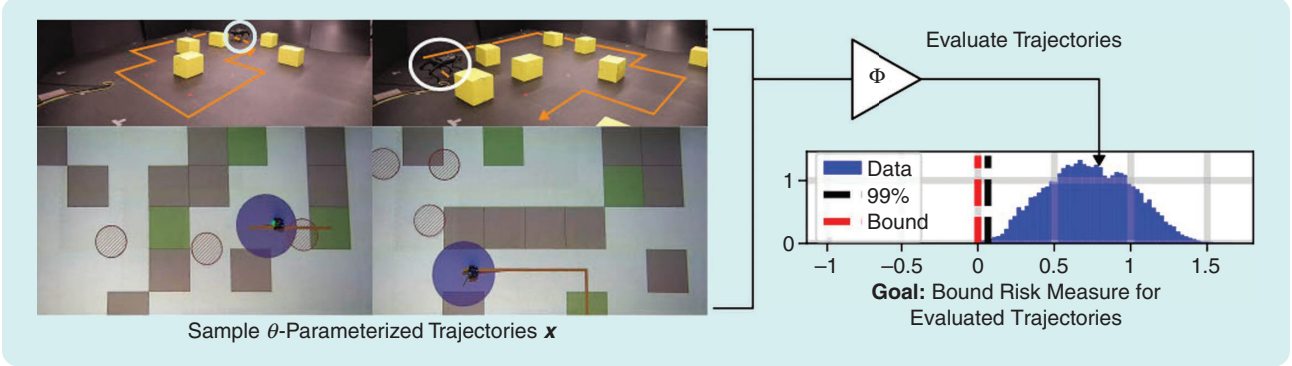


FIGURE 10 A flowchart for a general risk-aware verification pipeline. In the figure, the parameters θ correspond to obstacle locations and waypoints for the robots highlighted by white and blue circles [83]. Risk-aware verification bounds the risk-measure evaluation of evaluated trajectories—the blue data shown in the right figure.

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

constraint. In [71], the authors develop a procedure for learning a controller to tackle simultaneous performance and safety tradeoffs for nonlinear systems and verify the learned controller by estimating the CVaR of a corresponding robustness random variable. In [72], the authors constrain the optimal design of supersonic aircraft bodies against both VaR and CVaR requirements to ensure verifiable risk-aware performance despite uncertainties arising from the transition from laminar to turbulent flow, manufacturing uncertainties, etc. Examples from a larger body of work, including those by the authors, can be found in references [61], [62], [63], [66], [70], [73], [74], [75], [76], [77], [78], [79], and [80].

Case Study: Risk-Aware Verification of Lane-Keeping Controllers

In this case study, the goal is to find the least risky controller among two neural network lane-keeping controllers in the autonomous driving simulator CARLA during a left turn [81]; see Figure 11(a) and (b). Lane-keeping is realized by tracking a set of predefined waypoints. For verification, we consider the cross-track error c_e and the orientation error θ_e with respect to the current and next waypoint; see Figure 11(c). We consider an imitation learning (IL) controller [200] and a CBF controller [201]. The car model is stochastic as the control inputs are subject to normally distributed noise, and we uniformly sample the car's initial position from the set $(c_e, \theta_e) \in [-1,1] \times [-0.4, 0.4]$. We collected $N := 1{,}000$ trajectories $\boldsymbol{x}_i$ in a validation set D_{val} for each controller.

We are first concerned with cross-track error and consider the specifications $\phi_1 := G_{[0,\infty)}(|c_e| \leq 2.25)$ where 2.25 is an empirically obtained threshold indicating that the car stays within the lane. For the following analysis, recall that a negative value of $-\rho^{\phi_1}$ indicates satisfaction of ϕ_1, and positive values indicate a failure to lane-keep. Upper bounds on $\text{VaR}_{0.05}(-\rho^{\phi_1}(\boldsymbol{x}))$, $\text{CVaR}_{0.15}(-\rho^{\phi_1}(\boldsymbol{x}))$, and $E(-\rho^{\phi_1}(\boldsymbol{x}))$ are reported in Table 1, along with the empirical satisfaction rate $P_{\phi_1} := |\{\boldsymbol{x}_i \in D_{\text{val}} \mid \boldsymbol{x}_i \text{ satisfies } \phi_1\}|/N$. The computation of these bounds follows the SAA method discussed in "Concentration Inequalities: Risk-Measure Estimation." Clearly, the IL controller is the least risky one in terms of $\text{VaR}_{0.05}$ and $\text{CVaR}_{0.15}$, and it also has the highest empirical satisfaction rate. Interestingly, though, the CBF controller performs better on average. This result can also be seen in the empirical histograms of Figure 12(a). We hypothesize that this behavior arises from the long tail of risky behavior for the CBF controller, which corresponds to transient system behavior. We also analyzed the controllers' behavior more closely by looking at the cross-track error during the steady-state and transient phases for the specifications $\phi_2 := G_{[10,\infty)}(|c_e| \leq 2.25)$ and $\phi_3 := F_{[0,5]}G_{[0,5]}(|c_e| \leq 1.25)$, respectively. The upper bounds of the $\text{VaR}_{0.05}(-\rho^{\phi_i}(\boldsymbol{x}))$ and $E(-\rho^{\phi_i}(\boldsymbol{x}))$ for ϕ_2 and ϕ_2 are shown in Table 1 as well.

Interestingly, the IL controller is the least risky one only during the transient phase, while the CBF controller is the least risky one in the steady state. The corresponding empirical distributions are shown in Figure 12(b) and (c). Finally, let us verify the controller risk in terms of the orientation error θ_e. Consider $\phi_4 := G_{[0,\infty)}((c_e \geq 1.25) \Rightarrow F_{[0,2]}G_{[0,1]}(\theta_e \leq 0)) \wedge (c_e \leq -1.25) \Rightarrow F_{[0,2]}G_{[0,1]}(\theta_e \geq 0))$, which expresses the need for the controller to react to large cross-track errors c_e using the right orientation adjustment.

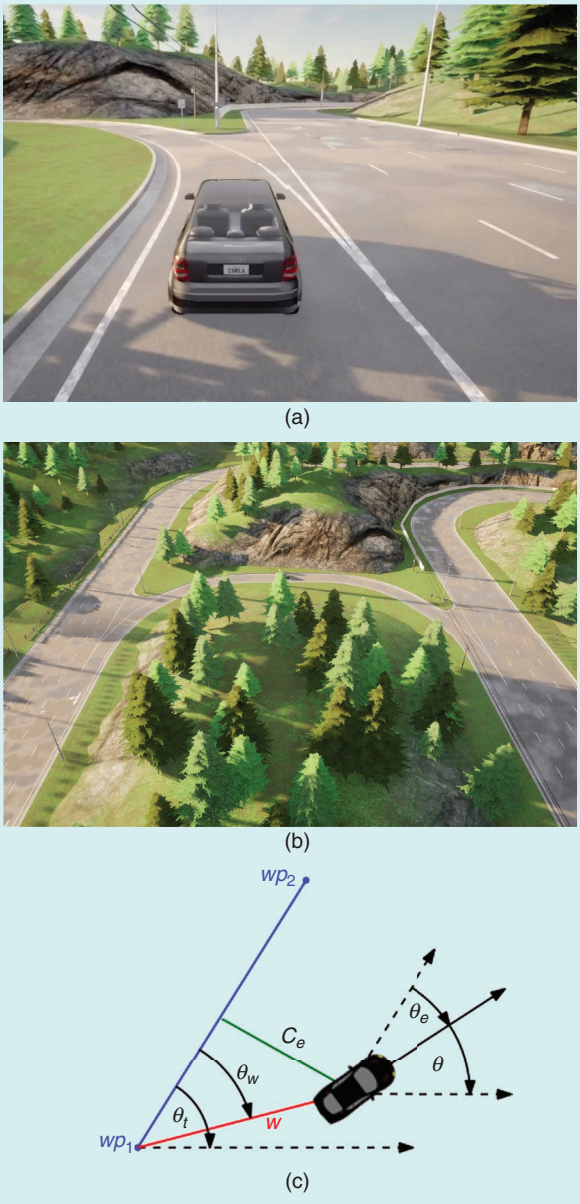


FIGURE 11 (a) Simulation environment in the CARLA autonomous driving simulator. (b) Left turn on which we evaluate two neural network lane-keeping controllers. (c) The car's cross-track error c_e and orientation error θ_e with respect to waypoints. (Source: Figures taken from [73].)

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

For this specification, the CBF controller is the least risky controller, which aligns with our observation that it is a better controller during the steady state. It can further be observed that both controllers have the same empirical satisfaction probability, while our risk analysis better quantifies a controller's quality. The empirical distribution of both controllers is shown in Figure 12(d).

TABLE 1 **Tabulated data from [73] for the case study on risk-aware lane-keeping of the IL and CBF controllers.**

R \ Controller	IL	CBF
$\overline{\text{VaR}}_{0.05}(-\rho^{\phi_1}(\boldsymbol{x}))$	0.462	1.125
$\overline{\text{CVaR}}_{0.15}(-\rho^{\phi_1}(\boldsymbol{x}))$	1.436	1.818
$\overline{E}(-\rho^{\phi_1}(\boldsymbol{x}))$	−0.248	−0.375
P_{ϕ_1}	0.975	0.863
$\overline{\text{VaR}}_{0.05}(-\rho^{\phi_2}(\boldsymbol{x}))$	0.462	−0.794
$\overline{E}(-\rho^{\phi_2}(\boldsymbol{x}))$	−0.254	−0.81
$\overline{\text{VaR}}_{0.05}(-\rho^{\phi_3}(\boldsymbol{x}))$	−0.324	0.063
$\overline{E}(-\rho^{\phi_3}(\boldsymbol{x}))$	−0.652	−0.297
$\overline{\text{VaR}}_{0.05}(-\rho^{\phi_4}(\boldsymbol{x}))$	−0.13	−0.32
$\overline{E}(-\rho^{\phi_4}(\boldsymbol{x}))$	−0.517	−0.533
P_{ϕ_5}	1	1

The green highlighted values indicate the risk evaluation for the controller that "scored" better for that corresponding risk metric.

Case Study: Risk-Aware Verification of Quadrupedal Locomotion

This case study aims to verify a quadruped's ability to render positive a collision-avoidance barrier function h for at least $T = 150$ time steps with a time step $\triangle t = 0.1$ [82]. Keeping consistent with our notation for the general overview for risk-aware V&V, our parameters θ include the locations of four randomly placed static obstacles in a 5 × 5-m grid and the center coordinates of a goal region g in the same grid. Hence, $\theta \in [-5,5]^{10} \triangleq \Theta$. To simplify our analysis, we represent the quadruped as a unicycle system, and as such, we assume that we can initialize the quadruped at a random planar position and angular orientation in the grid, that is, $x_0 \in [-5,5]^2 \times [0,2\pi] \triangleq \mathcal{X}_0$. Our classifier function $\boldsymbol{C}$ evaluates the discrete-time fractional difference of a candidate barrier function h that the quadruped is to keep positive. As such, the classifier outputs the minimum value over all time steps k of $h(\boldsymbol{x}(t+1)/h(\boldsymbol{x}(t))$, as realized by the quadruped over one trajectory $\boldsymbol{x} = \{\boldsymbol{x}(0), \boldsymbol{x}(1), \ldots, \boldsymbol{x}(150)\}$. Therefore, $\boldsymbol{C}(\boldsymbol{x}) < 0$ is equivalent to stating that there existed an interior time step $x_j \in \boldsymbol{x}$ such that $h(x_j) < 0$ and the quadruped failed to remain safe.

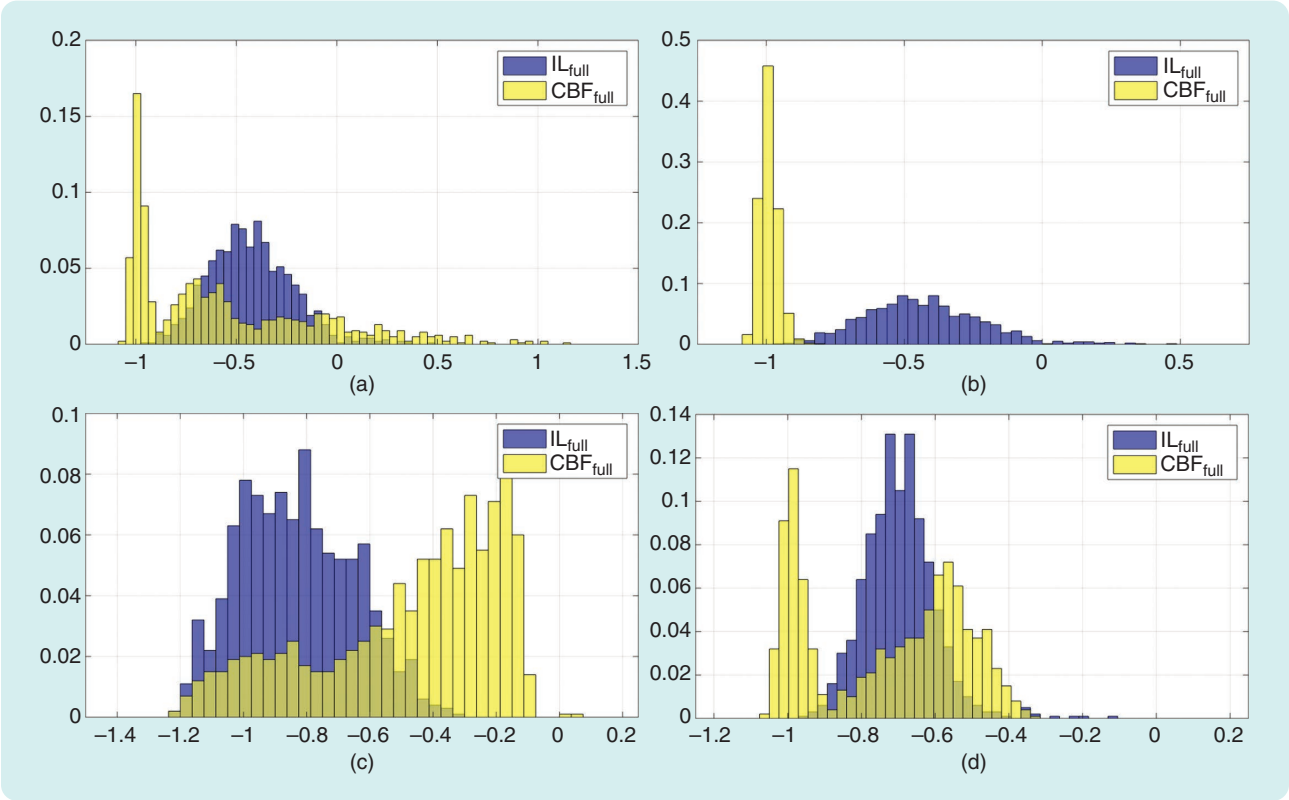


FIGURE 12 (a)–(d) Empirical distributions for the IL and CBF-based controllers and for the specifications $\phi_1 - \phi_4$ as described in the case study on risk-aware lane-keeping. (Source: Figures taken from [73].)

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

Slightly different from the general overview, however, instead of aiming to determine a lower bound on the VaR level $\beta = 0.9$ of the robustness random variable $R(x_0,\theta)$, we also randomize over initial conditions and parameters (x_0,θ) from their combined space. As such, the evaluation r of a sampled trajectory $\boldsymbol{x}$ generated by first sampling $(x_0,\theta) \sim \mathrm{U}[\mathcal{X}_0 \times \Theta]$ is a sample of the *holistic robustness random variable* R—this term was first defined in [75]. That being said, we still aim to lower bound $\mathrm{VaR}_{\beta=0.9}(R)$, which, according to the sample-based methods detailed in "Concentration Inequalities: Risk-Measure Estimation," has a known sample complexity (number N of trajectories to be evaluated) to determine such a lower bound. To be specific, risk-aware verification in this context amounts to identifying a high-probability lower bound on the VaR of the holistic robustness random variable R through a sample-based approach predicated on scenario optimization [202]. Therefore, after taking $N = 50$ trajectories, we can state with ≈99.5% confidence that the quadruped will realize positive value trajectories with 90% probability as the identified lower bound on $\mathrm{VaR}_{\beta=0.9}(R)$ was positive. The associated safety information is depicted in Figure 13.

We can also show that the reported probabilistic bounds are accurate. In [83], we employed the same risk-aware verification procedure as described previously to validate the controllers for a quadruped and a multiagent robotic system [203]. The expressions for the classifier function were the same in both cases, though we will refrain from reproducing them here for the sake of brevity. Suffice it to say that any trajectory that evaluates to a positive value under $\boldsymbol{C}$ would have made nontrivial progress toward a goal while avoiding static/moving obstacles within 10 s. To that end, we implemented controllers on hardware systems only once they exhibited a positive lower bound for $\mathrm{VaR}_{0.99}(R)$. To determine such a lower bound, we sampled 300 trajectories for both controllers and calculated their robustness under the aforementioned classifier $\boldsymbol{C}$. Doing so for our chosen controllers indicated positive lower bounds, and we can verify the accuracy of these lower bounds by taking 20,000 trajectories, evaluating them, numerically approximating the distribution of the holistic robustness random variable R, and reporting the numeric $\mathrm{VaR}_{0.99}(R)$ as our approximation. Figure 14 showcases the validity of the reported lower bounds, overlaid on the numeric approximation of the distribution for R, and Figure 15 shows the tiles of the controllers implemented on their respective hardware systems. Note that in all of the randomized cases, the controllers steered the systems successfully to their goals while avoiding all obstacles. This controller reliability is the primary reason we take a tail risk approach to verification, as the purpose of the procedure is to identify rare unsafe phenomena and ensure that even in those rare cases, the system still performs admirably.

Open Questions and Future Directions in V&V

Sample Complexity

There exist several open questions in risk-aware verification, though the most notable one concerns the small tail probability requirements for industrial applications. More specifically, for most product-level robotic systems requiring a verification statement, that is, autonomous cars, factory robots, flight software, etc., current standards require these systems to be verified to arbitrarily high probabilities, that is, $1-10^{-4}, 1-10^{-6}, 1-10^{-9}$, or even higher. If we assume that the underlying distributions are known, that is, Gaussian as is typically done, then this analysis can be carried out in a tractable—even analytic—fashion. However, if we follow the philosophy underlying the sample-based works that have recently become more popular as they do not assume underlying distributional knowledge for verification, then verifying systems to these probabilities could require hundreds of thousands of samples or even more. If each of those samples constitutes even one experimental run of the system, then this makes the direct application of these theoretical concepts exceedingly costly or time intensive. As such, reducing this sample complexity, whether via intelligent test design or by leveraging partial system knowledge, would go a long way to facilitate the widespread industrial adoption of these currently theoretical pipelines.

Compositional Verification

In a similar vein to that of the previous question, this second open question is also primarily of industrial interest. Namely, typical large-scale systems are composed of a variety of moving parts, each of which has to satisfy its own component specifications such that the larger architectured system satisfies a grander objective. Using the prior pipeline, each component could be verified separately in a probabilistic fashion, but in systems with potentially hundreds of separately engineered parts, separate verification procedures could be potentially prohibitive. On the other hand, the system could be verified as a large-scale closed-box system, though this could similarly fall under the sample-complexity questions that arose in the previous section. As such, determining an optimal way to break down these larger-scale system-level specifications into easily verifiable subcomponents for their respective systems remains an open problem. Indeed, the satisfiability of a given signal temporal logic specification itself remains a challenging problem. Determining the minimum number of such subcomponents would also mitigate any further sample-complexity issues arising from separate verification procedures as well. On the other hand, perhaps via smart instrumentation, all verification procedures for all subcomponents could be performed simultaneously.

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

Concentration Inequalities: Risk-Measure Estimation

This sidebar will briefly describe two sampling-based methods to estimate the tail risk measures expounded upon in this article. We will describe these methods as they are applied to arbitrary scalar random variables X.

SAA. This first method estimates $\mathrm{VaR}_\beta(X)$, $\mathrm{CVaR}_\beta(X)$ for any $\beta \in (0,1)$ and any scalar random variable X. Let $\{x_i\}_{i=1}^N$ be a set of N independently drawn samples of X. The empirical distribution function $\hat{\mathrm{F}}_N(x)$ for X based on this set of samples is

$$\hat{\mathrm{F}}_N(x) = \frac{1}{N}\sum_{i=1}^{N} \mathbf{1}(x \le x_i), \ \forall\, x_i \in \{x_i\}_{i=1}^N. \tag{S9}$$

with $\mathbf{1}$ being the indicator function. The SAA exploits the Dvoretzky-Kiefer-Wolfowitz Inequality [S3] built upon by Paul Massart in [S4], which proves that the empirical distribution has bounded deviation with respect to the true cumulative distribution function F for X to within some probability $\delta \in (0,1)$

$$\hat{\mathrm{F}}_N(x) - \sqrt{\frac{1}{2N}\ln\left(\frac{2}{\delta}\right)} \le F(x) \tag{S10a}$$

$$\le \hat{\mathrm{F}}_N(x) + \sqrt{\frac{1}{2N}\ln\left(\frac{2}{\delta}\right)}, \text{ w.p. } \ge 1-\delta. \tag{S10b}$$

The tail risk $\mathrm{VaR}_\beta(X)$ can be lower- and upper-bounded for any $\beta \in (0,1)$ using (S10). Define upper bound $\overline{\mathrm{VaR}}_\beta(X)$ as

$$\overline{\mathrm{VaR}}_\beta(X,\delta) = \inf\left\{x \in \mathbb{R} \,\middle|\, \hat{\mathrm{F}}_N(x) - \sqrt{\frac{1}{2N}\ln\left(\frac{2}{\delta}\right)} \ge 1-\beta\right\}$$

and let the lower bound $\underline{\mathrm{VaR}}_\beta(X)$ be defined as

$$\underline{\mathrm{VaR}}_\beta(X,\delta) = \inf\left\{x \in \mathbb{R} \,\middle|\, \hat{\mathrm{F}}_N(x) + \sqrt{\frac{1}{2N}\ln\left(\frac{2}{\delta}\right)} \ge 1-\beta\right\}.$$

Then, using (S10), the following result holds $\forall\, \beta, \delta \in (0,1)$

$$\underline{\mathrm{VaR}}_\beta(X,\delta) \le \mathrm{VaR}_\beta(X) \le \overline{\mathrm{VaR}}_\beta(X,\delta) \text{ w.p. } \ge 1-\delta.$$

Note that as the number of samples, N, of the random variable X increases, the gap between the upper and lower bounds shrinks as the bounds converge to the true value $\mathrm{VaR}_\beta(X)$. Similar methods exist to estimate $\mathrm{CVaR}_\beta(X)$ as well [S5].

Scenario Bounds. The second method upper bounds $\mathrm{VaR}_\beta(X), \mathrm{CVaR}_\beta(X), \mathrm{EVaR}_\beta(X)$ for any $\beta \in (0,1)$. As before, let $\{x_i\}_{i=1}^N$ be a set of N independently drawn samples of the scalar random variable X. Consider the following optimization problem, termed a *scenario program* [202]:

$$\begin{aligned} \zeta_N^\star = \underset{\zeta \in \mathbb{R}}{\operatorname{argmin}} \ & \zeta, \\ \text{subject to} \ & \zeta \ge x_i, \ \forall\, x_i \in \{x_k\}_{k=1}^N. \end{aligned} \tag{S11}$$

The theory of scenario optimization states that the solution to this optimization problem is an upper bound on $\mathrm{VaR}_\beta(X)$ with minimum probability $1-(1-\beta)^N$, that is, if X has the probability density function π, then

$$\mathbb{P}_\pi^N[\zeta_N^\star \ge \mathrm{VaR}_\beta(X)] \ge 1-(1-\beta)^N. \tag{S12}$$

The aforementioned result was proven in [65]. Note that (S12) does not need the density function π for X to be known. It just requires an ability to take N independent samples of X. Therefore, if we have a constant $c \in \mathbb{R}$ such that $\mathbb{P}_\pi[x \le c] = 1$, then we can exploit this inequality (S12) to similarly upper bound $\mathrm{CVaR}_\beta(X)$ and $\mathrm{EVaR}_\beta(X)$. Details on this approach can be found in [75, Sec. 3]. Besides scenario bounds, we refer to [S6] for VaR and CVaR bounds via linear and second-order cone programming.

Utilization in Controller Synthesis and Verification. As will be mentioned in various sections in this article, utilizing these concentration inequalities to estimate risk measures assumes that any taken samples are identically distributed, and the risk-aware result then corresponds to this same distribution. This may pose a problem for controller synthesis as oftentimes the randomness prompting a risk-aware analysis arises from random disturbances whose distributions are conditioned on the control input to be chosen. As such, any successive sampled disturbances over multiple control inputs would likely be samples from different distributions. Similarly, for verification, one must take care to ensure that the method used to randomly sample trajectories to be verified consistently samples these random trajectories from the same distribution. This oftentimes requires "resetting" the system to a predetermined starting configuration to ensure repeatability. As ensuring that these samples are identically distributed oftentimes requires considerable effort or may be impossible, recent work aims to estimate these risk measures over nonstationary distributions, though this remains an open problem.

REFERENCES

[S3] A. Dvoretzky, J. Kiefer, and J. Wolfowitz, "Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator," *Ann. Math. Statist.*, vol. 27, no. 3, pp. 642–669, 1956, doi: 10.1214/aoms/1177728174.

[S4] P. Massart, "The tight constant in the Dvoretzky-Kiefer-Wolfowitz inequality," *Ann. Probability*, vol. 18, no. 3, pp. 1269–1283, 1990, doi: 10.1214/aop/1176990746.

[S5] P. Thomas and E. Learned-Miller, "Concentration inequalities for conditional value at risk," in *Proc. 36th Int. Conf. Mach. Learn.*, vol. 97, K. Chaudhuri and R. Salakhutdinov, Eds., PMLR, 2019, pp. 6225–6233.

[S6] J. Miller, M. Tacchi, M. Sznaier, and A. Jasour, "Peak value-at-risk estimation for stochastic differential equations using occupation measures," in *Proc. 62nd IEEE Conf. Decis. Control (CDC)*, Piscataway, NJ, USA: IEEE Press, 2023, pp. 4836–4842, doi: 10.1109/CDC49753.2023.10383958.

OPEN PROBLEMS AND FUTURE DIRECTIONS

Our discourse up to this point has predominantly centered on the utilization of tail risk measures in the domains of planning, control, and verification within robotic systems. However, it is crucial to emphasize the broader applicability and potential impact of these notions. In this section, we detail several emerging areas that have gained substantial attention.

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

Risk-Aware Learning

There exists a rich body of literature exploring the integration of tail risk measures within learning paradigms such as RL, supervised learning, and unsupervised learning. These studies delve into diverse topics ranging from risk-sensitive reward functions and policy optimization to risk-aware feature learning and model training. Such research underscores the versatile role of tail risk measures in not only guiding robotic behavior in uncertain environments but also enabling robots to learn and adapt in a risk-aware manner over time. Thus, to provide a comprehensive overview of the role of tail risk measures in robotics, it is crucial to shed light on their applications in learning-based contexts as well.

The recent exploration into risk-averse RL is well encapsulated by the work of Greenberg et al. [208]. They emphasized the challenges of optimizing risk measures as conventional methods often overlook high-return strategies. To address this, they proposed a soft risk mechanism coupled with a cross-entropy module for efficient risk sampling. This innovative method, while maintaining risk aversion, demonstrated improved risk aversion across diverse benchmarks, setting a precedent for future exploration in this realm. In another interesting direction, Lacotte et al. [209] explored a risk-sensitive Generative Adversarial Imitation Learning (GAIL) approach aimed at performing as well as or better than the expert with respect to a risk profile.

Focusing on risk-constrained RL, Chow et al. [30] developed algorithms for risk-constrained MDPs, using chance constraints or CVaR as the risk representation. Their work represents an important step toward understanding and implementing risk constraints in RL and how these can be used for practical applications. Finally, Kose and Ruszczynski [210] proposed a novel RL methodology employing a Markov coherent dynamic risk measure. This work provided new risk-averse counterparts for basic and multistep methods of temporal differences, paving the way for future exploration in risk-averse learning methodologies.

Despite the notable strides made in these studies, there are still many open problems in risk-aware RL. One that specifically ties into the next section includes risk-aware RL when the underlying distribution changes over time. Such nonstationarity would frustrate, for example, immediate applications of distributional RL for policy development as such approaches aim to learn the underlying distribution to inform future risk-aware action choices. However, we expect such nonstationarity in real-world contexts, for example, changing wind patterns over time for a drone. Keeping with the realistic theme, most risk-aware RL approaches are sample complex. While this is fine in a simulation or lab-based setting, how do we efficiently learn such policies in real-world contexts where experiments are not cheap? Alternatively, how might we transition the learned simulation or lab-based policies to real-world contexts where disturbance distributions are almost certainly different?

Risk Awareness With Nonstationary and Dependent Data

This second area has garnered substantial interest, as it breaks the assumptions in the works discussed in this survey. Namely, the vast majority of the work discussed has all centered around risk-aware methodologies where either 1) the underlying distribution was known either exactly or in a distributionally robust sense, or 2) the distribution is queryable in a sample-based fashion, where the algorithm receives independent samples. What happens when either of these assumptions fails to hold? This is the central question in a new area of work that is just beginning in the risk-aware space—and for important reasons as well. Consider a robot that ambulates on uneven terrain. As the robot

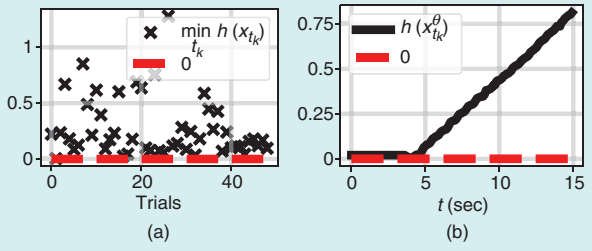


FIGURE 13 Worst-case safety information for the quadrupedal case study. Over the 50 sampled trials, the quadruped realizes a positive barrier value every time, which, according to the concentration inequality results in "Concentration Inequalities: Risk-Measure Estimation," implies that the system should always keep positive the barrier function h with minimum probability 90% and with confidence 99.5%. (a) Quadruped safety information. (b) Worst-case trial. (Source: Information for this figure comes from [82].)

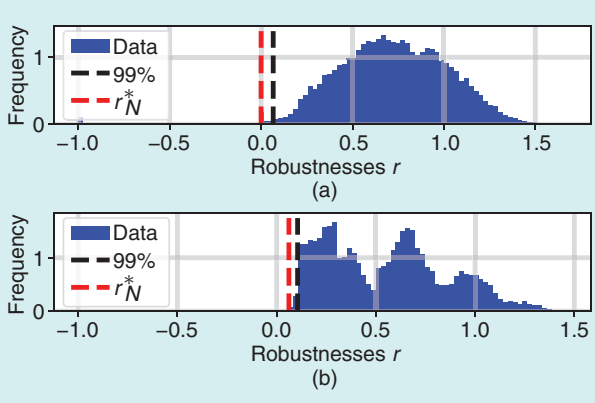


FIGURE 14 Validation data for probabilistic lower bounds reported on $\mathrm{VaR}_{0.99}(R)$ for controllers generated for (a) a quadruped and (b) a robotarium. As shown, the reported lower bounds (red) generated by the scenario approach mentioned in "Concentration Inequalities: Risk-Measure Estimation" are accurate as they lower bounds the true $\mathrm{VaR}_{0.99}(R)$ (black). (a) Quadruped controller verification data. (b) Robotarium controller verification data. (Source: Information for this figure comes from [82].)

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

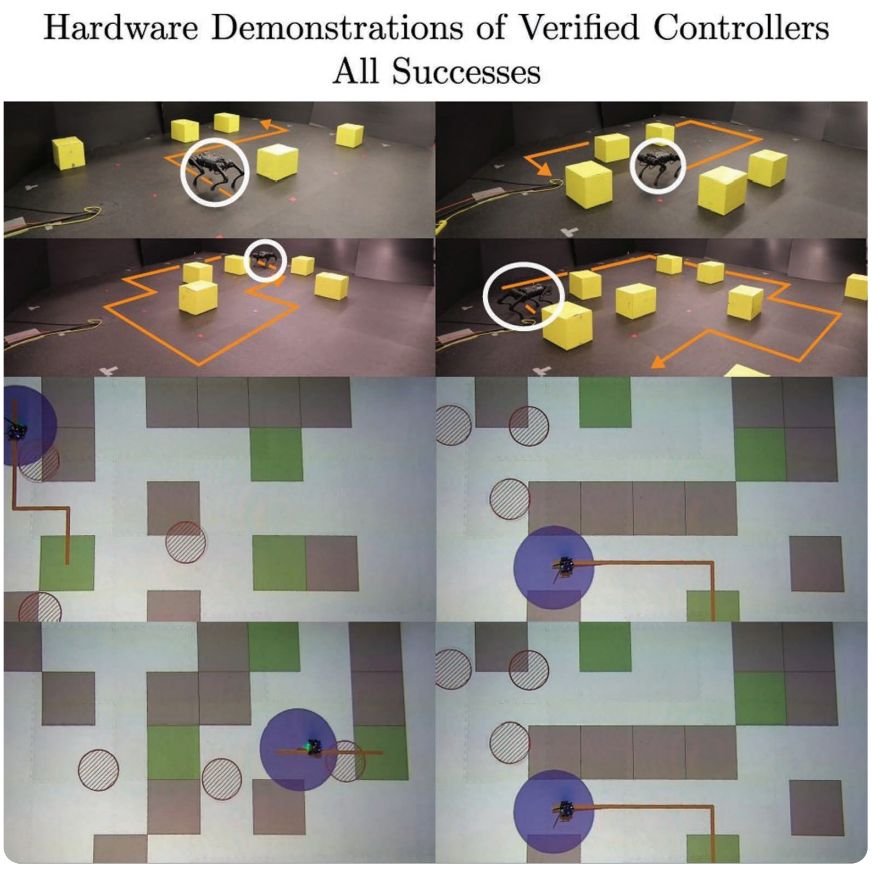


FIGURE 15 Hardware demonstration of controllers verified from a tail risk perspective. As the implemented controllers were "verified" since the reported lower bound on their robustness VaR was positive, we expect decent behavior in practice. Indeed, the verified controllers performed admirably despite randomized test cases generated for each system. (Source: Figure adapted from [83].)

traverses the space, any recording of the unevenness of the terrain would correspond to samples from a nonstationary distribution, and if the robot's controller builds a map of the terrain using this information and chooses actions predicated on this map, then successive data points are not necessarily independent.

When the underlying distribution of the uncertainty changes during a motion planning task, we would ideally like to understand the level of this shift so that we can account for it in our risk-aware planners. There has been a push toward identifying out-of-distribution (OOD) data for learning-based tasks [211], [212]. In [212], the authors study task-driven OOD detection using the Probably Approximately Correct (PAC)-Bayes theory to train the robot. The PAC-Bayes procedure provides a performance bound such that violating this bound signals that the robot is operating in an OOD environment.

In addition to the *detection* of OOD scenarios, we would also like to *respond* to such scenarios online. Here, distribution-free prediction schemes, like those offered by conformal prediction, are gaining more traction [213]. As these tools offer ways of generating probabilistically accurate predictors provided streams of potentially non-independent data, these predictors have been used for motion planning [118], [119], [204], [214], for confidence regions for learned classifiers [205], [206], and even for risk-aware decision making, though not in a tail risk sense [207]. If we can identify and adapt to distribution shifts in a risk-aware manner, we can enable robotic systems to react to data drift, unseen data, or spurious correlations [98]. By dynamically adjusting the risk level to adapt to the changing uncertainty distribution and guaranteeing the desired level of safety for the motion planning task, robots can operate in a wider array of unstructured environments while guaranteeing safety, task completion, and efficiency.

ACKNOWLEDGMENT

Prithvi Akella and Anushri Dixit contributed equally to this work.

AUTHOR INFORMATION

Prithvi Akella (prithvi.akella@gmail.com) received the B.S. degree in mechanical engineering from the University of California, Berkeley, in 2018, and the M.S. and Ph.D. degrees in mechanical engineering from the California Institute of Technology in 2020 and 2023, respectively. He is with Siemens Technology, Berkeley, CA 94704 USA. He is the recipient of the Bell Family Graduate Fellowship in Engineering and Applied Sciences. His research interests include the automated test and evaluation of cyberphysical systems.

Anushri Dixit received the Ph.D. degree in control and dynamical systems from the California Institute of Technology in 2023 and the B.S. degree in electrical engineering from the Georgia Institute of Technology in 2017. She is an assistant professor in the Department of Mechanical and Aerospace Engineering, University of California, Los Angeles, Los Angeles, CA 90024 USA. She was a postdoctoral researcher in the Department of Mechanical and Aerospace Engineering at Princeton University from 2023 to 2024. Her research interests include the motion planning and control of robots in unstructured environments while accounting for uncertainty in a principled manner. She has received the Outstanding Student Paper Award at the Conference on Decision and Control and the Best Student Paper Award at the Conference of Robot Learning (as a coauthor) and was selected as a Rising Star in Data Science by The University of Chicago.

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

Mohamadreza Ahmadi earned the DPhil (Ph.D.) degree in engineering science in November 2016 at the University of Oxford, UK, as a Clarendon Scholar. His Ph.D. was followed by research positions at the University of Texas, Austin, TX, the Center for Autonomous Systems and Technologies (CAST) at the California Institute of Technology, Pasadena, CA, NASA Jet Propulsion Laboratory, La Canada, CA, and TuSimple, La Jolla, CA. He is currently a technical lead in planning and prediction at Gatik AI, Mountain View, CA 94041 USA. He is the recipient of the Sloan-Robinson Engineering Fellowship, an Edgell-Sheppee Award, and an ICES Postdoctoral Fellowship. His research interests include planning and prediction under uncertainty with application to autonomous driving, in particular, short- to long-haul trucks.

Lars Lindemann received the Ph.D. degree in electrical engineering from the KTH Royal Institute of Technology in 2020. He is currently an assistant professor in the Thomas Lord Department of Computer Science, University of Southern California, Los Angeles, CA 90089 USA, where he is also a member of the Ming Hsieh Department of Electrical and Computer Engineering (courtesy), the Robotics and Autonomous Systems Center, and the Center for Autonomy and Artificial Intelligence. Between 2020 and 2022, he was a postdoctoral fellow in the Department of Electrical and Systems Engineering at the University of Pennsylvania. His research interests include systems and control theory, formal methods, and autonomous systems. He received the Outstanding Student Paper Award at the 58th IEEE Conference on Decision and Control and the Student Best Paper Award (as a coadvisor) at the 60th IEEE Conference on Decision and Control. He was a finalist for the Best Paper Award at the 2022 Conference on Hybrid Systems: Computation and Control and for the Best Student Paper Award at the 2018 American Control Conference. He is a Member of IEEE.

Margaret P. Chapman received the Ph.D. degree in electrical engineering and computer sciences (EECS) from the University of California, Berkeley (UC Berkeley) in May 2020. She is an assistant professor in the Edward S. Rogers Sr. Department of Electrical and Computer Engineering, University of Toronto, Toronto, ON M6G 2T6, Canada. Her research interests include risk-aware control theory, and she also investigates different control-theoretic applications, including those in leukemia treatment. In 2021, she received the Leon O. Chua Award for outstanding achievement in nonlinear science from EECS at UC Berkeley, and in 2023, she received the Connaught New Researcher Award from the Office of the Vice-President, Research and Innovation at the University of Toronto. She is a Member of IEEE.

George J. Pappas received the Ph.D. degree in electrical engineering and computer sciences from the University of California, Berkeley, Berkeley, CA, USA, in 1998. He is currently the Joseph Moore Professor in and the chair of the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104 USA. He also holds a secondary appointment with the Department of Computer and Information Sciences and the Department of Mechanical Engineering and Applied Mechanics. He is a member of the General Robotics, Automation, Sensing, and Perception Lab and the Penn Research in Embedded Computing and Integrated Systems Engineering Center. He was previously the deputy dean for research with the School of Engineering and Applied Science. His research interests include control theory and, in particular, hybrid systems, embedded systems, cyberphysical systems, and hierarchical and distributed control systems, with applications to autonomous aerial vehicles, distributed robotics, green buildings, and biomolecular networks. He was a recipient of various awards, such as the Antonio Ruberti Young Researcher Prize, the IEEE Control Systems Society George S. Axelby Award, the O. Hugo Schuck Best Paper Award, the George H. Heilmeier Award, the National Science Foundation Presidential Early Career Award for Scientists and Engineers, and numerous best student papers awards. He is a Fellow of IEEE.

Aaron D. Ames received the B.S. degree in mechanical engineering and the B.A. degree in mathematics from the University of St. Thomas, Saint Paul, MN, USA, in 2001 and the M.A. degree in mathematics and the Ph.D. degree in electrical engineering and computer sciences from the University of California, Berkeley, CA, USA, in 2006. From 2006 to 2008, he served as a postdoctoral scholar in control and dynamical systems at the California Institute of Technology (Caltech), Pasadena, CA, USA. In 2008, he began his faculty career at Texas A&M University, College Station, TX, USA. He was an associate professor in the Woodruff School of Mechanical Engineering and the School of Electrical and Computer Engineering at the Georgia Institute of Technology, Atlanta, GA, USA. Since 2017, he has been a Bren Professor of Mechanical and Civil Engineering and Control and Dynamical Systems at the Department of Mechanical and Civil Engineering, California Institute of Technology, Pasadena, CA 91125 USA. His research interests include the areas of robotics, nonlinear, safety-critical control, and hybrid systems, with a special focus on applications to dynamic robots, both formally and through experimental validation. He was a recipient of the 2005 Leon O. Chua Award for Achievement in Nonlinear Science and the 2006 Bernard Friedman Memorial Prize in Applied Mathematics from the University of California, Berkeley. He received the NSF CAREER Award in 2010, the 2015 Donald P. Eckman Award, and the 2019 IEEE CSS Antonio Ruberti Young Researcher Prize. He is a Fellow of IEEE.

Joel W. Burdick received undergraduate degrees in mechanical engineering and chemistry from Duke University and the M.S. and Ph.D. degrees in mechanical engineering from Stanford University. He is the Richard L. and Dorothy

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

M. Hayman Professor of Mechanical Engineering and Bioengineering at the Department of Mechanical and Civil Engineering, California Institute of Technology, Pasadena, CA 91125 USA. He has been with the Department of Mechanical Engineering at Caltech since May 1988, where he has been the recipient of the NSF Presidential Young Investigator award, the Office of Naval Research Young Investigator award, and the Feynman Fellowship. He has also received the ASCIT Award for Excellence in Undergraduate Teaching and the GSA Award for Excellence in Graduate Student Education, and he received the Popular Mechanics Breakthrough Award in 2011. In addition to mechanical engineering, he is a core faculty of the control and dynamical systems option as well as a faculty affiliate in the options of bioengineering) and computational and neural systems. His research interests mainly include the areas of robotics, kinematics, and mechanical systems. He is a Member of IEEE.

REFERENCES

[1] Y. Wang and M. P. Chapman, "Risk-averse autonomous systems: A brief history and recent developments from the perspective of optimal control," *Artif. Intell.*, vol. 311, Oct. 2022, Art. no. 103743, doi: 10.1016/j.artint.2022.103743.

[2] D. P. Bertsekas and I. B. Rhodes, "On the minimax reachability of target sets and target tubes," *Automatica*, vol. 7, no. 2, pp. 233–247, 1971, doi: 10.1016/0005-1098(71)90066-5.

[3] K. Zhou and J. C. Doyle, *Essentials of Robust Control*. Upper Saddle River, NJ, USA: Prentice Hall, 1998.

[4] J. Blake Rawlings, D. Q. Mayne, and M. M. Diehl, *Model Predictive Control: Theory, Computation, and Design*, 2nd ed. Madison, WI, USA: Nob Hill, 2017.

[5] M. Chen and C. J. Tomlin, "Hamilton–Jacobi reachability: Some recent theoretical advances and applications in unmanned airspace management," *Annu. Rev. Control, Robot., Auton. Syst.*, vol. 1, no. 1, pp. 333–358, 2018, doi: 10.1146/annurev-control-060117-104941.

[6] D. P. Bertsekas and S. E. Shreve, *Stochastic Optimal Control: The Discrete-Time Case*. Belmont, MA, USA: Athena Scientific, 1996.

[7] Richard, S. Sutton, and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. Cambridge, MA, USA: MIT Press, 2014.

[8] A. Abate, M. Prandini, J. Lygeros, and S. Sastry, "Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems," *Automatica*, vol. 44, no. 11, pp. 2724–2734, 2008, doi: 10.1016/j.automatica.2008.03.027.

[9] S. Summers and J. Lygeros, "Verification of discrete time stochastic hybrid systems: A stochastic reach-avoid decision problem," *Automatica*, vol. 46, no. 12, pp. 1951–1961, 2010, doi: 10.1016/j.automatica.2010.08.006.

[10] S. Skogestad and I. Postlethwaite, *Multivariable Feedback Control: Analysis and Design*. Hoboken, NJ, USA: Wiley, 2005.

[11] R. A. Howard and J. E. Matheson, "Risk-sensitive Markov decision processes," *Manage. Sci.*, vol. 18, no. 7, pp. 356–369, 1972, doi: 10.1287/mnsc.18.7.356.

[12] D. Jacobson, "Optimal stochastic linear systems with exponential performance criteria and their relation to deterministic differential games," *IEEE Trans. Autom. Control*, vol. 18, no. 2, pp. 124–131, Apr, 1973, doi: 10.1109/TAC.1973.1100265.

[13] P. Whittle, "Risk-sensitive linear/quadratic/Gaussian control," *Adv. Appl. Probab.*, vol. 13, no. 4, pp. 764–777, 1981, doi: 10.2307/1426972.

[14] H. Markowitz, "Porfolio selection," *J. Finance*, vol. 7, no. 1, pp. 77–91, 1952.

[15] R. Tyrrell Rockafellar and S. Uryasev, "Optimization of conditional value-at-risk," *J. Risk*, vol. 2, pp. 21–42, 2000.

[16] J. Guzzi, R. O. Chavez-Garcia, L. M. Gambardella, and A. Giusti, "On the impact of uncertainty for path planning," in *Proc. Int. Conf. Robot. Automat. (ICRA)*, Piscataway, NJ, USA: IEEE Press, 2019, pp. 5929–5935, doi: 10.1109/ICRA.2019.8793782.

[17] S. Karlsson, "Risk aware path planning and dynamic obstacle avoidance towards enabling safe robotic missions," Ph.D. thesis, Luleå Univ. of Technology, Luleå, Sweden, 2023.

[18] P. Akella, S. X. Wei, J. W. Burdick, and A. Ames, "Learning disturbances online for risk-aware control: Risk-aware flight with less than one minute of data," in *Proc. Learn. Dyn. Control Conf.*, PMLR, 2023, pp. 665–678.

[19] A. Dixit, D. D. Fan, K. Otsu, S. Dey, A.-A. Agha-Mohammadi, and J. W. Burdick, "Step: Stochastic traversability evaluation and planning for risk-aware off-road navigation; results from the DARPA subterranean challenge," *IEEE Trans. Field Robot.*, vol. 4, pp. 182–210, 2024, doi: 10.55417/fr.2024006.

[20] M. Ahmadi, M. Ono, M. D. Ingham, R. M. Murray, and A. D. Ames, "Risk-averse planning under uncertainty," in *Proc. Amer. Control Conf. (ACC)*, Piscataway, NJ, USA: IEEE Press, 2020, pp. 3305–3312, doi: 10.23919/ACC45564.2020.9147792.

[21] A. Majumdar and M. Pavone, "How should a robot assess risk? Towards an axiomatic theory of risk in robotics," in *Robotics Research*, Cham, Switzerland: Springer-Verlag, 2020, pp. 75–84.

[22] K. M. Smith and M. P. Chapman, "On exponential utility and conditional value-at-risk as risk-averse performance criteria," *IEEE Trans. Control Syst. Technol.*, vol. 31, no. 6, pp. 2555–2570, Nov. 2023, doi: 10.1109/TCST.2023.3274843.

[23] M. Kwon, E. Biyik, A. Talati, K. Bhasin, D. P. Losey, and D. Sadigh, "When humans aren't optimal: Robots that collaborate with risk-aware humans," in *Proc. ACM/IEEE Int. Conf. Human-Robot Interact. (HRI)*, New York, NY, USA: ACM, 2020, pp. 43–52, doi: 10.1145/3319502.3374832.

[24] A. Dixit, M. Ahmadi, and J. W. Burdick, "Risk-averse receding horizon motion planning for obstacle avoidance using coherent risk measures," *Artif. Intell.*, vol. 325, Dec. 2023, Art. no. 104018, doi: 10.1016/j.artint.2023.104018.

[25] A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on Stochastic Programming: Modeling and Theory*. Philadelphia, PA, USA: MPS-SIAM, 2009.

[26] W. Schwarting, J. Alonso-Mora, and D. Rus, "Planning and decision-making for autonomous vehicles," *Annu. Rev. Control, Robot., Auton. Syst.*, vol. 1, no. 1, pp. 187–210, 2018, doi: 10.1146/annurev-control-060117-105157.

[27] E. Karpas and D. Magazzeni, "Automated planning for robotics," *Annu. Rev. Control, Robot., Auton. Syst.*, vol. 3, no. 1, pp. 417–439, 2020, doi: 10.1146/annurev-control-082619-100135.

[28] L. Brunke et al., "Safe learning in robotics: From learning-based control to safe reinforcement learning," *Annu. Rev. Control, Robot., Auton. Syst.*, vol. 5, no. 1, pp. 411–444, 2022, doi: 10.1146/annurev-control-042920-020211.

[29] K. L. Hobbs, M. L. Mote, M. C. L. Abate, S. D. Coogan, and E. M. Feron, "Run time assurance for safety-critical systems: An introduction to safety filtering approaches for complex control systems," *IEEE Control Syst. Mag.*, vol. 43, no. 2, pp. 28–65, Apr. 2023, doi: 10.1109/MCS.2023.3234380.

[30] Y. Chow, M. Ghavamzadeh, L. Janson, and M. Pavone, "Risk-constrained reinforcement learning with percentile risk criteria," *J. Mach. Learn. Res.*, vol. 18, no. 1, pp. 6070–6120, 2017.

[31] L. A. Prashanth and M. C. Fu, "Risk-sensitive reinforcement learning via policy gradient search," *Found. Trends® Mach. Learn.*, vol. 15, no. 5, pp. 537–693, 2022, doi: 10.1561/2200000091.

[32] N. Bäuerle and U. Rieder, "More risk-sensitive Markov decision processes," *Math. Oper. Res.*, vol. 39, no. 1, pp. 105–120, 2014, doi: 10.1287/moor.2013.0601.

[33] L. Prashanth, "Policy gradients for CVaR-constrained MDPs," in *Proc. Int. Conf. Algorithmic Learn. Theory*, Cham, Switzerland: Springer-Verlag, 2014, pp. 155–169.

[34] Y. Chow and M. Ghavamzadeh, "Algorithms for CVaR optimization in MDPs," in *Proc. 28th Int. Conf. Neural Inf. Process. Syst. (NIPS)*, 2014, vol. 2, pp. 3509–3517.

[35] A. Tamar, Y. Chow, M. Ghavamzadeh, and S. Mannor, "Sequential decision making with coherent risk," *IEEE Trans. Autom. Control*, vol. 62, no. 7, pp. 3323–3338, Jul. 2017, doi: 10.1109/TAC.2016.2644871.

[36] A. Tamar, Y. Chow, M. Ghavamzadeh, and S. Mannor, "Policy gradient for coherent risk measures," in *Proc. 29th Int. Conf. Neural Inf. Process. Syst. (NIPS)*, 2015, vol. 1, pp. 1468–1476.

[37] W. B. Haskell and R. Jain, "A convex analytic approach to risk-aware Markov decision processes," *SIAM J. Control Optim.*, vol. 53, no. 3, pp. 1569–1598, 2015, doi: 10.1137/140969221.

[38] S. Feyzabadi and S. Carpin, "Risk-aware path planning using hirerachical constrained markov decision processes," in *Proc. IEEE Int. Conf. Automat. Sci. Eng. (CASE)*, Piscataway, NJ, USA: IEEE Press, 2014, pp. 297–303, doi: 10.1109/CoASE.2014.6899341.

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

[39] A. A. Pereira, J. Binney, G. A. Hollinger, and G. S. Sukhatme, "Risk-aware path planning for autonomous underwater vehicles using predictive ocean models," *J. Field Rob.*, vol. 30, no. 5, pp. 741–762, 2013, doi: 10.1002/rob.21472.
[40] T. Lam, A. Verma, B. K. H. Low, and P. Jaillet, "Risk-aware reinforcement learning with coherent risk measures and non-linear function approximation," in *Proc. 11th Int. Conf. Learn. Representations*, 2023.
[41] J. L. Hau, E. Delage, M. Ghavamzadeh, and M. Petrik, "On dynamic program decompositions of static risk measures," 2023, *arXiv:2304.12477*.
[42] M. Ahmadi et al., "Constrained risk-averse Markov decision processes," in *Proc. AAAI Conf. Artif. Intell.*, 2021, vol. 35, no. 13, pp. 11,718–11,725, doi: 10.1609/aaai.v35i13.17393.
[43] S. Carpin, Y. Chow, and M. Pavone, "Risk aversion in finite Markov decision processes using total cost criteria and average value at risk," in *Proc. IEEE Int. Conf. Robot. Automat. (ICRA)*, Piscataway, NJ, USA: IEEE Press, 2016, pp. 335–342, doi: 10.1109/ICRA.2016.7487152.
[44] C. Gavriel, G. Hanasusanto, and D. Kuhn, "Risk-averse shortest path problems," in *Proc. IEEE 51st IEEE Conf. Decis. Control (CDC)*, Piscataway, NJ, USA: IEEE Press, 2012, pp. 2533–2538, doi: 10.1109/CDC.2012.6426188.
[45] M. Ahmadi, A. Dixit, J. W. Burdick, and A. D. Ames, "Risk-averse stochastic shortest path planning," in *Proc. 60th IEEE Conf. Decis. Control (CDC)*, Piscataway, NJ, USA: IEEE Press, 2021, pp. 5199–5204, doi: 10.1109/CDC45484.2021.9683527.
[46] J. Fan and A. Ruszczyński, "Risk measurement and risk-averse control of partially observable discrete-time Markov systems," *Math. Methods Operations Res.*, vol. 88, no. 2, pp. 161–184, 2018, doi: 10.1007/s00186-018-0633-5.
[47] M. Ahmadi, U. Rosolia, M. D. Ingham, R. M. Murray, and A. D. Ames, "Risk-averse decision making under uncertainty," *IEEE Trans. Autom. Control*, vol. 69, no. 1, pp. 55–68, Jan. 2023, doi: 10.1109/TAC.2023.3264178.
[48] S. Singh, Y. Chow, A. Majumdar, and M. Pavone, "A framework for time-consistent, risk-sensitive model predictive control: Theory and algorithms," *IEEE Trans. Autom. Control*, vol. 64, no. 7, pp. 2905–2912, Jul. 2019, doi: 10.1109/TAC.2018.2874704.
[49] A. Dixit, M. Ahmadi, and J. W. Burdick, "Distributionally robust model predictive control with total variation distance," *IEEE Contr. Syst. Lett.*, vol. 6, pp. 3325–3330, 2022, doi: 10.1109/LCSYS.2022.3184921.
[50] B. P. G. Van Parys, D. Kuhn, P. J. Goulart, and M. Morari, "Distributionally robust control of constrained stochastic systems," *IEEE Trans. Autom. Control*, vol. 61, no. 2, pp. 430–442, Feb. 2016, doi: 10.1109/TAC.2015.2444134.
[51] Y. Chen, U. Rosolia, W. Ubellacker, N. Csomay-Shanklin, and A. D. Ames, "Interactive multi-modal motion planning with branch model predictive control," *IEEE Robot. Autom. Lett.*, vol. 7, no. 2, pp. 5365–5372, Apr. 2022, doi: 10.1109/LRA.2022.3156648.
[52] P. Sopasakis, M. Schuurmans, and P. Patrinos, "Risk-averse risk-constrained optimal control," in *Proc. 18th Eur. Control Conf. (ECC)*, 2019, pp. 375–380, doi: 10.23919/ECC.2019.8796021.
[53] J. Coulson, J. Lygeros, and F. Dörfler, "Data-enabled predictive control: In the shallows of the DeePC," in *Proc. 18th Eur. Control Conf. (ECC)*, Piscataway, NJ, USA: IEEE Press, 2019, pp. 307–312, doi: 10.23919/ECC.2019.8795639.
[54] J. Coulson, J. Lygeros, and F. Dörfler, "Distributionally robust chance constrained data-enabled predictive control," *IEEE Trans. Autom. Control*, vol. 67, no. 7, pp. 3289–3304, Jul. 2022, doi: 10.1109/TAC.2021.3097706.
[55] A. Zolanvari and A. Cherukuri, "Data-driven distributionally robust iterative risk-constrained model predictive control," in *Proc. Eur. Control Conf. (ECC)*, 2022, pp. 1578–1583, doi: 10.23919/ECC55457.2022.9838319.
[56] A. Hakobyan, G. C. Kim, and I. Yang, "Risk-aware motion planning and control using CVaR-constrained optimization," *IEEE Robot. Autom. Lett.*, vol. 4, no. 4, pp. 3924–3931, Oct. 2019, doi: 10.1109/LRA.2019.2929980.
[57] A. Dixit, M. Ahmadi, and J. W. Burdick, "Risk-sensitive motion planning using entropic value-at-risk," in *Proc. Eur. Control Conf.*, 2021, pp. 1726–7132, doi: 10.23919/ECC54610.2021.9655104.
[58] D. D. Fan, K. Otsu, Y. Kubo, A. Dixit, J. Burdick, and A.-A. Agha-Mohammadi, "Step: Stochastic traversability evaluation and planning for risk-aware off-road navigation," in *Robotics: Science and Systems*, RSS Foundation, 2021, pp. 1–21.
[59] S. Samuelson and I. Yang, "Safety-aware optimal control of stochastic systems using conditional value-at-risk," in *Proc. Annu. Amer. Control Conf. (ACC)*, 2018, pp. 6285–6290, doi: 10.23919/ACC.2018.8430957.
[60] M. P. Chapman et al., "A risk-sensitive finite-time reachability approach for safety of stochastic dynamic systems," in *Proc. Amer. Control Conf. (ACC)*, Piscataway, NJ, USA: IEEE Press, 2019, pp. 2958–2963, doi: 10.23919/ACC.2019.8815169.
[61] M. P. Chapman, R. Bonalli, K. M. Smith, I. Yang, M. Pavone, and C. J. Tomlin, "Risk-sensitive safety analysis using conditional value-at-risk," *IEEE Trans. Autom. Control*, vol. 67, no. 12, pp. 6521–6536, Dec. 2022, doi: 10.1109/TAC.2021.3131149.
[62] M. P. Chapman, M. Fauß, and K. M. Smith, "On optimizing the conditional value-at-risk of a maximum cost for risk-averse safety analysis," *IEEE Trans. Autom. Control*, vol. 68, no. 6, pp. 3720–3727, Jun. 2023, doi: 10.1109/TAC.2022.3195381.
[63] C. Wei, M. Fauß, and M. P. Chapman, "CVaR-based safety analysis in the infinite time horizon setting," in *Proc. Amer. Control Conf. (ACC)*, 2022, pp. 2863–2870, doi: 10.23919/ACC53348.2022.9867285.
[64] A. W. Singletary, A. D. Ames, and M. Ahmadi, "Controlling a moveable device utilizing risk control barrier functions," U.S. Patent 18/324,718, Nov. 30, 2023.
[65] P. Akella, M. Ahmadi, and A. D. Ames, "A scenario approach to risk-aware safety-critical system verification," 2022, *arXiv:2203.02595*.
[66] L. Lindemann, N. Matni, and G. J. Pappas, "Stl robustness risk over discrete-time stochastic processes," in *Proc. 60th IEEE Conf. Decis. Control (CDC)*, Piscataway, NJ, USA: IEEE Press, 2021, pp. 1329–1335, doi: 10.1109/CDC45484.2021.9683305.
[67] M. Cubuktepe and U. Topcu, "Verification of Markov decision processes with risk-sensitive measures," in *Proc. Annu. Amer. Control Conf. (ACC)*, Piscataway, NJ, USA: IEEE Press, 2018, pp. 2371–2377, doi: 10.23919/ACC.2018.8430905.
[68] K. Leung and M. Pavone, "Semi-supervised trajectory-feedback controller synthesis for signal temporal logic specifications," in *Proc. Amer. Control Conf. (ACC)*, Atlanta, GA, USA, 2022, pp. 178–185, doi: 10.23919/ACC53348.2022.9867345.
[69] T. Meggendorfer, "Verification of discrete-time Markov decision processes," Ph.D. thesis, Technische Universität München, München, Germany, 2021.
[70] S. Samuelson and I. Yang, "Safety-aware optimal control of stochastic systems using conditional value-at-risk," in *Proc. Annu. Amer. Control Conf. (ACC)*, Piscataway, NJ, USA: IEEE Press, 2018, pp. 6285–6290, doi: 10.23919/ACC.2018.8430957.
[71] N. Hashemi et al., "Risk-awareness in learning neural controllers for temporal logic objectives," 2022, *arXiv:2210.07439*.
[72] D. Quagliarella and E. Iuliano, "Robust design of a supersonic natural laminar flow wing-body," *IEEE Comput. Intell. Mag.*, vol. 12, no. 4, pp. 14–27, Nov. 2017, doi: 10.1109/MCI.2017.2742718.
[73] L. Lindemann, L. Jiang, N. Matni, and G. J. Pappas, "Risk of stochastic systems for temporal logic specifications," *ACM Trans. Embedded Comput. Syst.*, vol. 22, no. 3, pp. 1–31, 2023, doi: 10.1145/3580490.
[74] L. Lindemann, A. Rodionova, and G. Pappas, "Temporal robustness of stochastic signals," in *Proc. 25th ACM Int. Conf. Hybrid Syst.: Comput. Control*, 2022, pp. 1–11.
[75] P. Akella, A. Dixit, M. Ahmadi, J. W. Burdick, and A. D. Ames, "Sample-based bounds for coherent risk measures: Applications to policy synthesis and verification," 2022, *arXiv:2204.09833*.
[76] M. Rigter, B. Lacerda, and N. Hawes, "Risk-averse bayes-adaptive reinforcement learning," in *Proc. 35th Int. Conf. Neural Inf. Process. Syst. (NIPS)*, 2021, pp. 1142–1154.
[77] M. Godbout, M. Heuillet, S. Chandra, R. Bhati, and A. D. Carl, "Conditional-value-at-risk adversarial reinforcement learning," 2021, *arXiv:2109.09470*.
[78] L. J. Hong and G. Liu, "Monte Carlo estimation of value-at-risk, conditional value-at-risk and their sensitivities," in *Proc. Winter Simul. Conf. (WSC)*, Piscataway, NJ, USA: IEEE Press, 2011, pp. 95–107, doi: 10.1109/WSC.2011.6147743.
[79] J.-H. Kim, J. Lee, and S.-K. Joo, "Conditional value-at-risk-based method for evaluating the economic risk of superconducting fault current limiter installation," *IEEE Trans. Appl. Supercond.*, vol. 25, no. 3, pp. 1–4, Jun. 2015, doi: 10.1109/TASC.2014.2386321.
[80] D. Kim and S. Oh, "Efficient off-policy safe reinforcement learning using trust region conditional value at risk," *IEEE Robot. Autom. Lett.*, vol. 7, no. 3, pp. 7644–7651, Jul. 2022, doi: 10.1109/LRA.2022.3184793.
[81] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "CARLA: An open urban driving simulator," in *Proc. Conf. Robot Learn.*, Mountain View, CA, USA, 2017, pp. 1–16.

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

[82] P. Akella and A. D. Ames, "A barrier-based scenario approach to verifying safety-critical systems," *IEEE Robot. Autom. Lett.*, vol. 7, no. 4, pp. 11,062–11,069, Oct. 2022, doi: 10.1109/LRA.2022.3192805.
[83] P. Akella, W. Ubellacker, and A. D. Ames, "Safety-critical controller verification via Sim2Real gap quantification," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2023, pp. 10,539–10,545, doi: 10.1109/ICRA48891.2023.10161126.
[84] R. B. Ash and C. A. Doléans-Dade, *Probability and Measure Theory*. San Diego, CA, USA: Academic Press, 2000.
[85] F. Liu and R. Wang, "A theory for measures of tail risk," *Math. Oper. Res.*, vol. 46, no. 3, pp. 1109–1128, Aug. 2021, doi: 10.1287/moor.2020.1072.
[86] P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath, "Coherent measures of risk," *Math. Finance*, vol. 9, no. 3, pp. 203–228, 1999, doi: 10.1111/1467-9965.00068.
[87] A. Ahmadi-Javid and A. Pichler, "An analytical study of norms and banach spaces induced by the entropic value-at-risk," *Math. Financial Econ.*, vol. 11, no. 4, pp. 527–550, 2017, doi: 10.1007/s11579-017-0197-9.
[88] A. Ahmadi-Javid, "Entropic value-at-risk: A new coherent risk measure," *J. Optim. Theory Appl.*, vol. 155, no. 3, pp. 1105–1123, 2012, doi: 10.1007/s10957-011-9968-2.
[89] F. Delbaen, "Coherent risk measures on general probability spaces," in *Proc. Adv. Finance Stochastics: Essays Honour Dieter Sondermann*, 2002, pp. 1–37.
[90] H. Rahimian and S. Mehrotra, "Distributionally robust optimization: A review," 2019, *arXiv:1908.05659*.
[91] A. Shapiro, "Tutorial on risk neutral, distributionally robust and risk averse multistage stochastic programming," *Eur. J. Oper. Res.*, vol. 288, no. 1, pp. 1–13, 2021, doi: 10.1016/j.ejor.2020.03.065.
[92] S. Safaoui, B. J. Gravell, V. Renganathan, and T. H. Summers, "Risk-averse RRT* planning with nonlinear steering and tracking controllers for nonlinear robotic systems under uncertainty," in *Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, 2021, pp. 3681–3688, doi: 10.1109/IROS51168.2021.9636834.
[93] V. Renganathan, S. Safaoui, A. Kothari, B. Gravell, I. Shames, and T. Summers, "Risk bounded nonlinear robot motion planning with integrated perception & control," *Artif. Intell.*, vol. 314, Jan. 2023, Art. no. 103812, doi: 10.1016/j.artint.2022.103812.
[94] V. Renganathan, I. Shames, and T. H. Summers, "Towards integrated perception and motion planning with distributionally robust risk constraints," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 15,530–15,536, 2020, doi: 10.1016/j.ifacol.2020.12.2396.
[95] A. Ruszczyński, "Risk-averse dynamic programming for Markov decision processes," *Math. Program.*, vol. 125, no. 2, pp. 235–261, 2010, doi: 10.1007/s10107-010-0393-3.
[96] A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on Stochastic Programming: Modeling and Theory*. SIAM, 2014.
[97] M. P. Chapman and K. M. Smith, "Classical risk-averse control for a finite-horizon Borel model," *IEEE Contr. Syst. Lett.*, vol. 6, pp. 1525–1530, 2021, doi: 10.1109/LCSYS.2021.3114126.
[98] O. Wiles et al., "A fine-grained analysis on distribution shift," 2021, *arXiv:2110.11328*.
[99] R. D. Luce and H. Raiffa, *Games and Decisions*. New York, NY, USA: Wiley, 1957.
[100] P. Whittle, *Risk-Sensitive Optimal Control*. New York, NY, USA: Wiley, 1990.
[101] N. Saldi, T. Başar, and M. Raginsky, "Approximate Markov-Nash equilibria for discrete-time risk-sensitive mean-field games," *Math. Oper. Res.*, vol. 45, no. 4, pp. 1596–1620, 2020, doi: 10.1287/moor.2019.1044.
[102] R. Blancas-Rivera, R. Cavazos-Cadena, and H. Cruz-Suárez, "Discounted approximations in risk-sensitive average markov cost chains with finite state space," *Math. Methods Operations Res.*, vol. 91, no. 2, pp. 241–268, 2020, doi: 10.1007/s00186-019-00689-3.
[103] V. Krishnamurthy, *Partially Observed Markov Decision Processes*. Cambridge, U.K.: Cambridge Univ. Press, 2016.
[104] M. Ahmadi, N. Jansen, B. Wu, and U. Topcu, "Control theory meets POMDPs: A hybrid systems approach," *IEEE Trans. Autom. Control*, vol. 66, no. 11, pp. 5191–5204, Nov. 2021, doi: 10.1109/TAC.2020.3035755.
[105] M. Ono, M. Pavone, Y. Kuwata, and J. Balaram, "Chance-constrained dynamic programming with application to risk-aware robotic space exploration," *Auton. Robots*, vol. 39, no. 4, pp. 555–571, 2015, doi: 10.1007/s10514-015-9467-7.
[106] L. Nardi and C. Stachniss, "Uncertainty-aware path planning for navigation on road networks using augmented MDPs," in *Proc. Int. Conf. Robot. Autom. (ICRA)*, Piscataway, NJ, USA: IEEE Press, 2019, pp. 5780–5786, doi: 10.1109/ICRA.2019.8794121.
[107] J. Wöhlke, F. Schmitt, and H. van Hoof, "Hierarchies of planning and reinforcement learning for robot navigation," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, Piscataway, NJ, USA: IEEE Press, 2021, pp. 10,682–10,688, doi: 10.1109/ICRA48506.2021.9561151.
[108] C. Wei, M. Fauß, and M. P. Chapman, "CVaR-based safety analysis in the infinite time horizon setting," in *Proc. Amer. Control Conf. (ACC)*, Piscataway, NJ, USA: IEEE Press, 2022, pp. 2863–2870, doi: 10.23919/ACC53348.2022.9867285.
[109] P. Stegmaier, K. Kurzer, and M. J. Zöllner, "Cooperative trajectory planning in uncertain environments with monte Carlo tree search and risk metrics," in *Proc. IEEE 25th Int. Conf. Intell. Transp. Syst. (ITSC)*, Piscataway, NJ, USA: IEEE Press, 2022, pp. 4109–4116, doi: 10.1109/ITSC55140.2022.9921861.
[110] A. Shapiro, D. Dentcheva, and A. Ruszczynski, *Lectures on Stochastic Programming: Modeling and Theory*. Philadelphia, PA, USA: SIAM, 2021.
[111] A. Bemporad and M. Morari, "Robust model predictive control: A survey," in *Robustness in Identification and Control*, London, U.K.: Springer-Verlag, 1999, pp. 207–226.
[112] A. Mesbah, "Stochastic model predictive control: An overview and perspectives for future research," *IEEE Control Syst. Mag.*, vol. 36, no. 6, pp. 30–44, Dec. 2016, doi: 10.1109/MCS.2016.2602087.
[113] J. Yin, Z. Zhang, and P. Tsiotras, "Risk-aware model predictive path integral control using conditional value-at-risk," in *Proc. IEEE Int. Conf. Robot. Automat. (ICRA)*, Piscataway, NJ, USA: IEEE Press, 2023, pp. 7937–7943, doi: 10.1109/ICRA48891.2023.10161100.
[114] L. Lindemann, M. Cleaveland, Y. Kantaros, and G. J. Pappas, "Robust motion planning in the presence of estimation uncertainty," in *Proc. 60th IEEE Conf. Decis. Control (CDC)*, Piscataway, NJ, USA: IEEE Press, 2021, pp. 5205–5212, doi: 10.1109/CDC45484.2021.9683651.
[115] J. Köhler, F. Geuss, and M. N. Zeilinger, "On stochastic MPC formulations with closed-loop guarantees: Analysis and a unifying framework," in *Proc. 62nd IEEE Conf. Decis. Control (CDC)*, 2023, pp. 6692–6699, doi: 10.1109/CDC49753.2023.10383313.
[116] L. Hewing, K. P. Wabersich, and M. N. Zeilinger, "Recursively feasible stochastic model predictive control using indirect feedback," *Automatica*, vol. 119, Sep. 2020, Art. no. 109095, doi: 10.1016/j.automatica.2020.109095. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0005109820302934
[117] T. Lew, R. Bonalli, and M. Pavone, "Risk-averse trajectory optimization via sample average approximation," *IEEE Robot. Autom. Lett.*, vol. 9, no. 2, pp. 1500–1507, Feb. 2024, doi: 10.1109/LRA.2023.3331889.
[118] L. Lindemann, M. Cleaveland, G. Shim, and G. J. Pappas, "Safe planning in dynamic environments using conformal prediction," *IEEE Trans. Robot. Autom.*, vol. 8, no. 8, pp. 5116–5123, Aug. 2023, doi: 10.1109/LRA.2023.3292071.
[119] R. Luo et al., "Sample-efficient safety assurances using conformal prediction," in *Proc. Int. Workshop Algorithmic Found. Robot.*, Cham, Switzerland: Springer-Verlag, 2022, pp. 149–169.
[120] J. Reher and A. D. Ames, "Dynamic walking: Toward agile and efficient bipedal robots," *Annu. Rev.*, vol. 4, pp. 535–572, Dec. 2020.
[121] T. Rouček et al., "DARPA subterranean challenge: Multi-robotic exploration of underground environments," in *Proc. Int. Conf. Model. Simul. Auton. Syst.*, Cham, Switzerland: Springer-Verlag, 2019, pp. 274–290.
[122] F. Blanchini and S. Miani, *Set-Theoretic Methods in Control*, vol. 78. Cham, Switzerland: Springer-Verlag, 2008.
[123] M. Althoff, O. Stursberg, and M. Buss, "Reachability analysis of nonlinear systems with uncertain parameters using conservative linearization," in *Proc. 47th IEEE Conf. Decis. Control*, Piscataway, NJ, USA: IEEE Press, 2008, pp. 4042–4048, doi: 10.1109/CDC.2008.4738704.
[124] I. M. Mitchell, A. M. Bayen, and C. J. Tomlin, "A time-dependent Hamilton-Jacobi formulation of reachable sets for continuous dynamic games," *IEEE Trans. Autom. Control*, vol. 50, no. 7, pp. 947–957, Jul. 2005, doi: 10.1109/TAC.2005.851439.
[125] M. Nagumo, "Über die lage der integralkurven gewöhnlicher differentialgleichungen," in *Proc. Physico-Math. Soc. Japan (3rd Ser.)*, 1942, vol. 24, pp. 551–559.
[126] J.-P. Aubin, A. M. Bayen, and P. Saint-Pierre, *Viability Theory: New Directions*. Berlin, Germany: Springer-Verlag, 2011.
[127] S. Prajna, A. Jadbabaie, and G. J. Pappas, "A framework for worst-case and stochastic safety verification using barrier certificates," *IEEE Trans. Autom. Control*, vol. 52, no. 8, pp. 1415–1428, Aug. 2007, doi: 10.1109/TAC.2007.902736.
[128] A. A. Ahmadi and A. Majumdar, "DSOS and SDSOS optimization: LP and SOCP-based alternatives to sum of squares optimization," in *Proc. 48th*

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

Annu. Conf. Inf. Sci. Syst. (CISS), Piscataway, NJ, USA: IEEE Press, 2014, pp. 1–5, doi: 10.1109/CISS.2014.6814141.
[129] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Trans. Autom. Control*, vol. 62, no. 8, pp. 3861–3876, Aug. 2017, doi: 10.1109/TAC.2016.2638961.
[130] Q. Nguyen, A. Hereid, J. W. Grizzle, A. D. Ames, and K. Sreenath, "3D dynamic walking on stepping stones with control barrier functions," in *Proc. IEEE 55th Conf. Decis. Control (CDC)*, Piscataway, NJ, USA: IEEE Press, 2016, pp. 827–834, doi: 10.1109/CDC.2016.7798370.
[131] Y. Chen, A. Hereid, H. Peng, and J. Grizzle, "Enhancing the performance of a safe controller via supervised learning for truck lateral control," *J. Dyn. Syst. Meas. Control*, vol. 141, no. 10, 2019, Art. no. 101005, doi: 10.1115/1.4043487.
[132] 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.
[133] S. Kolathaya and A. D. Ames, "Input-to-state safety with control barrier functions," *IEEE Contr. Syst. Lett.*, vol. 3, no. 1, pp. 108–113, Jan. 2019, doi: 10.1109/LCSYS.2018.2853698.
[134] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *Proc. 18th Eur. Control Conf. (ECC)*, Piscataway, NJ, USA: IEEE Press, 2019, pp. 3420–3431, doi: 10.23919/ECC.2019.8796030.
[135] M. Ahmadi, A. Singletary, J. W. Burdick, and A. D. Ames, "Safe policy synthesis in multi-agent POMDPs via discrete-time barrier functions," in *Proc. IEEE 58th Conf. Decis. Control (CDC)*, Piscataway, NJ, USA: IEEE Press, 2019, pp. 4797–4803, doi: 10.1109/CDC40024.2019.9030241.
[136] A. Agrawal and K. Sreenath, "Discrete control barrier functions for safety-critical control of discrete systems with application to bipedal robot navigation," in *Proc. Robot.: Sci. Syst.*, 2017.
[137] M. Ahmadi, A. Singletary, J. W. Burdick, and A. D. Ames, "Barrier functions for multiagent-POMDPs with DTL specifications," in *Proc. 59th IEEE Conf. Decis. Control*, 2020, pp. 1380–1385, doi: 10.1109/CDC42340.2020.9304266.
[138] O. So, A. Clark, and C. Fan, "Almost-sure safety guarantees of stochastic zero-control barrier functions do not hold," 2023, *arXiv:2312.02430*.
[139] C. Santoyo, M. Dutreix, and S. Coogan, "A barrier function approach to finite-time stochastic system verification and control," *Automatica*, vol. 125, Mar. 2021, Art. no. 109439.
[140] S. Yaghoubi, K. Majd, G. Fainekos, T. Yamaguchi, D. Prokhorov, and B. Hoxha, "Risk-bounded control using stochastic barrier functions," *IEEE Contr. Syst. Lett.*, vol. 5, no. 5, pp. 1831–1836, Nov. 2021, doi: 10.1109/LCSYS.2020.3043287.
[141] S. Yaghoubi, G. Fainekos, T. Yamaguchi, D. Prokhorov, and B. Hoxha, "Risk-bounded control with Kalman filtering and stochastic barrier functions," in *Proc. 60th IEEE Conf. Decis. Control (CDC)*, Piscataway, NJ, USA: IEEE Press, 2021, pp. 5213–5219, doi: 10.1109/CDC45484.2021.9683756.
[142] M. Vahs, C. Pek, and J. Tumova, "Belief control barrier functions for risk-aware control," *IEEE Robot. Autom. Lett.*, vol. 8, no. 12, pp. 8565–8572, Dec. 2023, doi: 10.1109/LRA.2023.3330662.
[143] M. Ahmadi, X. Xiong, and A. D. Ames, "Risk-averse control via CVaR barrier functions: Application to bipedal robot locomotion," *IEEE Contr. Syst. Lett.*, vol. 6, pp. 878–883, 2022, doi: 10.1109/LCSYS.2021.3086854.
[144] A. Singletary, M. Ahmadi, and A. D. Ames, "Safe control for nonlinear systems with stochastic uncertainty via risk control barrier functions," *IEEE Contr. Syst. Lett.*, vol. 7, pp. 349–354, 2023, doi: 10.1109/LCSYS.2022.3187458.
[145] M. Vahs and J. Tumova, "Risk-aware control for robots with non-Gaussian belief spaces," 2023, *arXiv:2309.12857*.
[146] 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, doi: 10.1016/j.automatica.2014.04.021.
[147] X. Xiong and A. D. Ames, "Dynamic and versatile humanoid walking via embedding 3D actuated SLIP model with hybrid LIP based stepping," *IEEE Robot. Autom. Lett.*, vol. 5, no. 4, pp. 6286–6293, Oct. 2020, doi: 10.1109/LRA.2020.3013924.
[148] X. Xiong, J. Reher, and A. Ames, "Global position control on underactuated bipedal robots: Step-to-step dynamics approximation *for step planning*," in *Proc. IEEE Int. Conf. Robot. Automat. (ICRA)*, 2021, pp. 2825–2831, doi: 10.1109/ICRA48506.2021.9561961.
[149] P. Akella, W. Ubellacker, and A. D. Ames, "Probabilistic guarantees for nonlinear safety-critical optimal control," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Piscataway, NJ, USA: IEEE Press, 2023, pp. 8120–8126, doi: 10.1109/IROS55552.2023.10342474.
[150] A. Robey et al., "Learning control barrier functions from expert demonstrations," in *Proc. 59th IEEE Conf. Decis. Control (CDC)*, Piscataway, NJ, USA: IEEE Press, 2020, pp. 3717–3724, doi: 10.1109/CDC42340.2020.9303785.
[151] A. Taylor, A. Singletary, Y. Yue, and A. Ames, "Learning for safety-critical control with control barrier functions," in *Proc. Learn. Dyn. Control*, PMLR, 2020, pp. 708–717.
[152] C. Wang, Y. Meng, Y. Li, S. L. Smith, and J. Liu, "Learning control barrier functions with high relative degree for safety-critical control," in *Proc. Eur. Control Conf. (ECC)*, Piscataway, NJ, USA: IEEE Press, 2021, pp. 1459–1464, doi: 10.23919/ECC54610.2021.9655206.
[153] B. Dai, P. Krishnamurthy, and F. Khorrami, "Learning a better control barrier function," in *Proc. IEEE 61st Conf. Decis. Control (CDC)*, Piscataway, NJ, USA: IEEE Press, 2022, pp. 945–950, doi: 10.1109/CDC51059.2022.9993334.
[154] S. Prajna and A. Rantzer, "Convex programs for temporal verification of nonlinear dynamical systems," *SIAM J. Control Optim.*, vol. 46, no. 3, pp. 999–1021, 2007, doi: 10.1137/050645178.
[155] A. Clark, "Verification and synthesis of control barrier functions," in *Proc. 60th IEEE Conf. Decis. Control (CDC)*, Piscataway, NJ, USA: IEEE Press, 2021, pp. 6105–6112, doi: 10.1109/CDC45484.2021.9683520.
[156] H. Zhang, Z. Li, H. Dai, and A. Clark, "Efficient sum of squares-based verification and construction of control barrier functions by sampling on algebraic varieties," in *Proc. 62nd IEEE Conf. Decis. Control (CDC)*, Piscataway, NJ, USA: IEEE Press, 2023, pp. 5384–5391, doi: 10.1109/CDC49753.2023.10384199.
[157] A. Papachristodoulou and S. Prajna, "A tutorial on sum of squares techniques for systems analysis," in *Proc. Amer. Control Conf.*, vol. 4, Piscataway, NJ, USA: IEEE Press, 2005, pp. 2686–2700, doi: 10.1109/ACC.2005.1470374.
[158] J. Guiochet, M. Machin, and H. Waeselynck, "Safety-critical advanced robots: A survey," *Rob. Auton. Syst.*, vol. 94, pp. 43–52, Aug. 2017, doi: 10.1016/j.robot.2017.04.004.
[159] J. Yoo, E. Jee, and S. Cha, "Formal modeling and verification of safety-critical software," *IEEE Softw.*, vol. 26, no. 3, pp. 42–49, May/Jun. 2009, doi: 10.1109/MS.2009.67.
[160] A. Aniculaesei, D. Arnsberger, F. Howar, and A. Rausch, "Towards the verification of safety-critical autonomous systems in dynamic environments," 2016, *arXiv:1612.04977*.
[161] T. Grimm, D. Lettnin, and M. Hübner, "A survey on formal verification techniques for safety-critical systems-on-chip," *Electronics*, vol. 7, no. 6, 2018, Art. no. 81, doi: 10.3390/electronics7060081.
[162] F. Vicentini, M. Askarpour, M. G. Rossi, and D. Mandrioli, "Safety assessment of collaborative robotics through automated formal verification," *IEEE Trans. Robot.*, vol. 36, no. 1, pp. 42–61, Feb. 2020, doi: 10.1109/TRO.2019.2937471.
[163] X. Cheng, L. Zhou, W. Liu, Y. Li, M. Peng, and Y. Wang, "Construction and verification of risk predicting models to evaluate the possibility of venous thromboembolism after robot-assisted radical prostatectomy," *Ann. Surgical Oncol.*, vol. 29, no. 8, pp. 1–10, 2022, doi: 10.1245/s10434-022-11574-5.
[164] M. Askarpour, D. Mandrioli, M. Rossi, and F. Vicentini, "SAFER-HRC: Safety analysis through formal verification in human-robot collaboration," in *Proc. Int. Conf. Comput. Saf., Rel., Secur.*, Cham, Switzerland: Springer-Verlag, 2016, pp. 283–295.
[165] T. P. Huck, N. Münch, L. Hornung, C. Ledermann, and C. Wurll, "Risk assessment tools for industrial human-robot collaboration: Novel approaches and practical needs," *Saf. Sci.*, vol. 141, Sep. 2021, Art. no. 105288, doi: 10.1016/j.ssci.2021.105288.
[166] "Safety of machinery - Risk assessment - Part 2: Practical guidance and examples of methods," ISO, Geneva, Switzerland, ISO/TR 14121-2:2012, 2012. [Online]. Available: https://www.iso.org/obp/ui/#iso:std:iso:tr:14121:-2:ed-2:v1:en
[167] "Safety of machinery - General principles for design - Risk assessment and risk reduction," ISO, Geneva, Switzerland, ISO 12100:2010, 2010. [Online]. Available: https://www.iso.org/obp/ui/#iso:std:iso:12100:ed-1:v1:en
[168] M.-Y. Yu, R. Vasudevan, and M. Johnson-Roberson, "Occlusion-aware risk assessment for autonomous driving in urban environments," *IEEE Robot. Autom. Lett.*, vol. 4, no. 2, pp. 2235–2241, Apr. 2019, doi: 10.1109/LRA.2019.2900453.

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.

[169] M. Strickland, G. Fainekos, and H. Ben Amor, "Deep predictive models for collision risk assessment in autonomous driving," in *Proc. IEEE Int. Conf. Robot. Automat. (ICRA)*, Piscataway, NJ, USA: IEEE Press, 2018, pp. 4685–4692, doi: 10.1109/ICRA.2018.8461160.
[170] C. Ryan, F. Murphy, and M. Mullins, "End-to-end autonomous driving risk analysis: A behavioural anomaly detection approach," *IEEE Trans. Intell. Transp. Syst.*, vol. 22, no. 3, pp. 1650–1662, Mar. 2021, doi: 10.1109/TITS.2020.2975043.
[171] S. Lefèvre, D. Vasquez, and C. Laugier, "A survey on motion prediction and risk assessment for intelligent vehicles," *ROBOMECH J.*, vol. 1, no. 1, pp. 1–14, 2014, doi: 10.1186/s40648-014-0001-z.
[172] "Fundamental review of the trading book: A revised market risk framework," BCBS, 2013.
[173] E. M. Clarke and J. M. Wing, "Formal methods: State of the art and future directions," *ACM Comput. Surv.*, vol. 28, no. 4, pp. 626–643, 1996, doi: 10.1145/242223.242257.
[174] J. Woodcock, P. G. Larsen, J. Bicarregui, and J. Fitzgerald, "Formal methods: Practice and experience," *ACM Comput. Surv.*, vol. 41, no. 4, pp. 1–36, 2009, doi: 10.1145/1592434.1592436.
[175] "What is formal methods." NASA Langley Formal Methods (.gov). [Online]. Available: https://shemesh.larc.nasa.gov/fm/fm-what.html
[176] A. Pnueli, "The temporal logic of programs," in *Proc. 18th Annu. Symp. Found. Comput. Sci.*, Washington, DC, USA, 1977, pp. 46–57.
[177] C. Baier, B. Haverkort, H. Hermanns, and J.-P. Katoen, "Model checking continuous-time Markov chains by transient analysis," in *Proc. Int. Conf. Comput. Aided Verification*, Berlin, Germany: Springer-Verlag, 2000, pp. 358–372.
[178] E. M. Clarke, "Model checking," in *Proc. Int. Conf. Found. Softw. Technol. Theor. Comput. Sci.*, Berlin, Germany: Springer-Verlag, 1997, pp. 54–56.
[179] R. Koymans, "Specifying real-time properties with metric temporal logic," *Real-Time Syst.*, vol. 2, no. 4, pp. 255–299, 1990, doi: 10.1007/BF01995674.
[180] O. Maler and D. Nickovic, "Monitoring temporal properties of continuous signals," in *Formal Techniques, Modelling and Analysis of Timed and Fault-Tolerant Systems*, Berlin, Germany: Springer-Verlag, 2004, pp. 152–166.
[181] Y. Shoukry, P. Nuzzo, A. L. Sangiovanni-Vincentelli, S. A. Seshia, G. J. Pappas, and P. Tabuada, "SMC: Satisfiability modulo convex optimization," in *Proc. 20th Int. Conf. Hybrid Syst.: Comput. Control*, 2017, pp. 19–28.
[182] M. Sheeran, S. Singh, and G. Stålmarck, "Checking safety properties using induction and a SAT-solver," in *Proc. Int. Conf. Formal Methods Comput.-Aided Des.*, Berlin, Germany: Springer-Verlag, 2000, pp. 127–144.
[183] A. Bianco and L. d. Alfaro, "Model checking of probabilistic and nondeterministic systems," in *Proc. Int. Conf. Found. Softw. Technol. Theor. Comput. Sci.*, Berlin, Germany: Springer-Verlag, 1995, pp. 499–513.
[184] H. Hansson and B. Jonsson, "A logic for reasoning about time and reliability," *Formal Aspects Comput.*, vol. 6, no. 5, pp. 512–535, 1994, doi: 10.1007/BF01211866.
[185] M. Kwiatkowska, G. Norman, and D. Parker, "PRISM 4.0: Verification of probabilistic real-time systems," in *Proc. Int. Conf. Comput. Aided Verification*, Berlin, Germany: Springer-Verlag, 2011, pp. 585–591.
[186] H. L. Younes and R. G. Simmons, "Probabilistic verification of discrete event systems using acceptance sampling," in *Proc. Int. Conf. Comput. Aided Verification*, Berlin, Germany: Springer-Verlag, 2002, pp. 223–235.
[187] H. L. Younes and R. G. Simmons, "Statistical probabilistic model checking with a focus on time-bounded properties," *Inf. Comput.*, vol. 204, no. 9, pp. 1368–1409, 2006, doi: 10.1016/j.ic.2006.05.002.
[188] A. Legay, A. Lukina, L. M. Traonouez, J. Yang, S. A. Smolka, and R. Grosu, "Statistical model checking," in B. Steffen and G. Woeginger, *Computing and Software Science*, Cham, Switzerland: Springer-Verlag, 2019, pp. 478–504.
[189] A. Legay, B. Delahaye, and S. Bensalem, "Statistical model checking: An overview," in *Proc. Int. Conf. Runtime Verification*, Berlin, Germany: Springer-Verlag, 2010, pp. 122–135.
[190] R. P. Kanwal, *Generalized Functions Theory and Technique: Theory and Technique*. Boston, MA, USA: Birkhäuser, 1998.
[191] P. Akella, M. Ahmadi, R. M. Murray, and A. D. Ames, "Formal test synthesis for safety-critical autonomous systems based on control barrier functions," in *Proc. 59th IEEE Conf. Decis. Control (CDC)*, Piscataway, NJ, USA: IEEE Press, 2020, pp. 790–795, doi: 10.1109/CDC42340.2020.9303776.
[192] P. Akella, M. Ahmadi, R. M. Murray, and A. D. Ames, "Barrier-based test synthesis for safety-critical systems subject to timed reach-avoid specifications," 2023, *arXiv:2301.09622*.
[193] J. Zhang and J. Li, "Testing and verification of neural-network-based safety-critical control software: A systematic literature review," *Inf. Softw. Technol.*, vol. 123, Jul. 2020, Art. no. 106296, doi: 10.1016/j.infsof.2020.106296.
[194] S. Ghosh, F. Berkenkamp, G. Ranade, S. Qadeer, and A. Kapoor, "Verifying controllers against adversarial examples with Bayesian optimization," in *Proc. IEEE Int. Conf. Robot. Automat. (ICRA)*, Piscataway, NJ, USA: IEEE Press, 2018, pp. 7306–7313, doi: 10.1109/ICRA.2018.8460635.
[195] M. Kwiatkowska, G. Norman, and D. Parker, "Stochastic model checking," in *Proc. Int. School Formal Methods Des. Comput., Commun. Softw. Syst.*, Berlin, Germany: Springer-Verlag, 2007, pp. 220–270.
[196] K. Sen, M. Viswanathan, and G. Agha, "On statistical model checking of stochastic systems," in *Proc. Int. Conf. Comput. Aided Verification*, Berlin, Germany: Springer-Verlag, 2005, pp. 266–280.
[197] G. Agha and K. Palmskog, "A survey of statistical model checking," *ACM Trans. Model. Comput. Simul.*, vol. 28, no. 1, pp. 1–39, 2018, doi: 10.1145/3158668.
[198] F. Ciesinski and M. Größer, "On probabilistic computation tree logic," in *Validation of Stochastic Systems*, C. Baier, B. R. Haverkort, H. Hermanns, J. P. Katoen, and M. Siegle, Eds., Berlin, Germany: Springer-Verlag, 2004, pp. 147–188.
[199] P. Wakker and A. Tversky, "An axiomatization of cumulative prospect theory," *J. Risk Uncertainty*, vol. 7, no. 2, pp. 147–175, 1993, doi: 10.1007/BF01065812.
[200] S. Ross and D. Bagnell, "Efficient reductions for imitation learning," in *Proc. Int. Conf. Artif. Intell. Statist.*, Sardinia, Italy, 2010, pp. 661–668.
[201] L. Lindemann, A. Robey, L. Jiang, S. Tu, and N. Matni, "Learning robust output control barrier functions from safe expert demonstrations," *IEEE Open J. Control Syst.*, vol. 3, pp. 158–172, 2024, doi: 10.1109/OJCSYS.2024.3385348.
[202] M. C. Campi and S. Garatti, "The exact feasibility of randomized solutions of uncertain convex programs," *SIAM J. Optim.*, vol. 19, no. 3, pp. 1211–1230, 2008, doi: 10.1137/07069821X.
[203] S. Wilson et al., "The robotarium: Globally impactful opportunities, challenges, and lessons learned in remote-access, distributed control of multirobot systems," *IEEE Control Syst. Mag.*, vol. 40, no. 1, pp. 26–44, Feb. 2020, doi: 10.1109/MCS.2019.2949973.
[204] A. Z. Ren et al., "Robots that ask for help: Uncertainty alignment for large language model planners," 2023, *arXiv:2307.01928*.
[205] J. Vazquez and J. C. Facelli, "Conformal prediction in clinical medical sciences," *J. Healthcare Inform. Res.*, vol. 6, no. 3, pp. 241–252, 2022, doi: 10.1007/s41666-021-00113-8.
[206] H. Papadopoulos, A. Gammerman, and V. Vovk, "Reliable diagnosis of acute abdominal pain with conformal prediction," *Int. J. Eng. Intell. Syst. Elect. Eng. Commun.*, vol. 17, nos. 2–3, pp. 127–137, 2009.
[207] S. Ramakrishna, B. Luo, Y. Barve, G. Karsai, and A. Dubey, "Risk-aware scene sampling for dynamic assurance of autonomous systems," in *Proc. IEEE Int. Conf. Assured Autonomy (ICAA)*, Piscataway, NJ, USA: IEEE Press, 2022, pp. 107–116, doi: 10.1109/ICAA52185.2022.00022.
[208] I. Greenberg, Y. Chow, M. Ghavamzadeh, and S. Mannor, "Efficient risk-averse reinforcement learning," in *Proc. 36th Int. Conf. Neural Inf. Process. Syst. (NIPS)*, 2022, pp. 32,639–32,652.
[209] J. Lacotte, M. Ghavamzadeh, Y. Chow, and M. Pavone, "Risk-sensitive generative adversarial imitation learning," in *Proc. 22nd Int. Conf. Artif. Intell. Statist.*, PMLR, 2019, pp. 2154–2163.
[210] Ü. Köse and A. Ruszczyński, "Risk-averse learning by temporal difference methods with Markov risk measures," *J. Mach. Learn. Res.*, vol. 22, no. 38, pp. 1–34, 2021.
[211] R. Sinha et al., "A system-level view on out-of-distribution data in robotics," 2022, *arXiv:2212.14020*.
[212] A. Farid, S. Veer, and A. Majumdar, "Task-driven out-of-distribution detection with statistical guarantees for robot learning," in *Proc. 5th Conf. Robot Learn.*, vol. 164, A. Faust, D. Hsu, and G. Neumann, Eds., PMLR, 2022, pp. 970–980.
[213] G. Shafer and V. Vovk, "A tutorial on conformal prediction," *J. Mach. Learn. Res.*, vol. 9, no. 3, pp. 371–421, 2008.
[214] A. Dixit, L. Lindemann, S. X. Wei, M. Cleaveland, G. J. Pappas, and J. W. Burdick, "Adaptive conformal prediction for motion planning among dynamic agents," in *Proc. Learn. Dyn. Control Conf.*, PMLR, 2023, pp. 300–314.

Authorized licensed use limited to: University of Pennsylvania. Downloaded on October 29,2026 at 13:04:34 UTC from IEEE Xplore. Restrictions apply.