2021 60th IEEE Conference on Decision and Control (CDC)
December 13-15, 2021. Austin, Texas

STL Robustness Risk over Discrete-Time Stochastic Processes

Lars Lindemann, Nikolai Matni, and George J. Pappas

2021 60th IEEE Conference on Decision and Control (CDC) | 978-1-6654-3659-5/21/$31.00 ©2021 IEEE | DOI: 10.1109/CDC45484.2021.9683305

***Abstract*— We present a framework to interpret signal temporal logic (STL) formulas over discrete-time stochastic processes in terms of the induced risk. Each realization of a stochastic process either satisfies or violates an STL formula. In fact, we can assign a robustness value to each realization that indicates how robustly this realization satisfies an STL formula. We then define the risk of a stochastic process not satisfying an STL formula robustly, referred to as the *STL robustness risk*. In our definition, we permit general classes of risk measures such as, but not limited to, the conditional value-at-risk. While in general hard to compute, we propose an approximation of the STL robustness risk. This approximation has the desirable property of being an upper bound of the STL robustness risk when the chosen risk measure is monotone, a property satisfied by most risk measures. Motivated by the interest in data-driven approaches, we present a sampling-based method for estimating the approximate STL robustness risk from data for the value-at-risk. While we consider the value-at-risk, we highlight that such sampling-based methods are viable for other risk measures.**

I. Introduction

Consider the scenario in which an autonomous car equipped with noisy sensors navigates through urban traffic. As a consequence of imperfect sensing, the environment is not perfectly known. Instead, we can describe the scenario as a stochastic process that models each possible outcome along with the probability of an outcome. In this paper, we are interested in quantifying the associated risk in such safety-critical systems. In particular, we consider system specifications that are formulated in signal temporal logic (STL) [1] and, for the first time, propose a systematic way to assess the risk associated with such system specifications when evaluated over discrete-time stochastic processes.

Signal temporal logic has been introduced as a formalism to express a large variety of complex system specifications. STL particularly allows to express temporal and spatial system properties, e.g., surveillance ("visit regions A, B, and C every $10-60$ sec"), safety ("always between $5-25$ sec stay at least 1 m away from region D"), and many others. STL specifications are evaluated over deterministic signals and a given signal, for instance the trajectory of a robot, either satisfies or violates the STL specification at hand. Towards quantifying the robustness by which a signal satisfies an STL specification, the authors in [2] proposed the robustness degree as a tube around a nominal signal so that all signals in this tube satisfy (violate) the specification if the nominal signal satisfies (violates) the specification. In this way, the size of the tube indicates the robustness of the nominal signal with respect to the specification. As the robustness degree is in general hard to calculate, the authors in [2] proposed approximate yet easier to calculate robust semantics. Several other approximations have appeared such as the space and time robustness [3], the arithmetic-geometric mean robustness [4], the smooth cumulative robustness [5], averaged STL [6], or robustness metrics tailored for guiding reinforcement learning [7]. Also related is the work [8] where metrics for STL formulas are presented and [9] where a connection with linear time-invariant filtering is made.

For stochastic signals, the authors in [10]–[14] propose notions of probabilistic signal temporal logic in which chance constraints are defined over the atomic elements (called predicates) of STL, while the Boolean and temporal operators of STL are not altered. Similarly, notions of risk signal temporal logic have recently appeared in [15] and [16] by defining risk constraints over the atomic elements only. The work in [17] considers the probability of an STL specification being satisfied instead of applying chance or risk constraints on the atomic level. For the less expressive linear temporal logic, the authors in [18]–[20] consider control over belief spaces, while the authors in [21] consider probabilistic satisfaction over Markov decision processes. In contrast, in this work we quantify the risk of not satisfying an STL specification robustly. Probably closest to our paper are [22] and [23] in which the authors present a framework for the robustness of STL under stochastic models. Our work differs from these in several directions. Most importantly, we do not limit our attention to average satisfaction, termed average robustness degree and defined via the distribution of the approximate robustness degree. We instead allow for general risk measures towards an axiomatic risk theory for temporal logics. We also argue that the STL robustness risk should conceptually be defined differently than the average robustness degree in [22] and [23]. We further present an efficient way to reliably estimate the STL robustness risk for the value-at-risk.

The theory of risk has a long history in finance [24], [25]. More recently, there has been an interest to also apply such risk measures in robotics and control applications [26]. Risk-aware control and estimation frameworks have recently appeared in [27]–[34] using various forms of risk. We remark that these frameworks are orthogonal to our work as they present design tools while we provide a generic framework for quantifying the risk of complex system specifications expressed in STL. We hope that such quantification will be useful to guide the design and analysis process in the future.

The authors are with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104, USA. Email: larsl@seas.upenn.edu (L. Lindemann), nmatni@seas.upenn.edu (N. Matni), and pappasg@seas.upenn.edu (G.J. Pappas)

This research was generously supported by NSF award CPS-2038873, NSF CAREER award ECCS-2045834, and AFOSR grant FA9550-19-1-0265 (Assured Autonomy in Contested Environments).

In this paper, we consider signal temporal logic specifications interpreted over discrete-time stochastic processes. Our contributions can be summarized as follows:

1) We show that the semantics, the robust semantics, and the robustness degree of STL are measurable functions so that these functions are well-defined and have a probability distribution.
2) We define the risk of a discrete-time stochastic process not satisfying an STL specification robustly and refer to this definition as the "STL robustness risk".
3) We argue that the robustness risk is in general hard to calculate and propose an approximation of the robustness risk that has the desireable property of being an upper bound of the STL robustness risk, i.e., more risk averse, if the risk measure is monotone.
4) We present a sampling-based estimate of the approximate robustness risk for the value-at-risk. We show that this estimate is an upper bound of the approximate robustness risk with high probability. We thereby establish an interesting connection between data-driven design approaches and the risk of an STL specification.

Sec. II presents background on STL, stochastic processes, and risk measures. In Sec. III, we define the STL robustness risk, while we show in Sec. IV how the approximate robustness risk can be estimated from data for the case of the value-at-risk. A case study is presented in Sec. V followed by conclusions in Sec. VI. An extended version of this paper containing all technical proofs can be found in [35].

II. Background

True and false are encoded as $\top := 1$ and $\bot := -1$, respectively, with the set $\mathbb{B} := \{\top, \bot\}$. Let $\mathbb{R}$ and $\mathbb{N}$ be the set of real and natural numbers. Let $\overline{\mathbb{R}} := \mathbb{R} \cup \{\infty, -\infty\}$ be the set of extended real numbers. Also let $\mathbb{R}_{\geq 0}$ be the set of non-negative real numbers and $\mathbb{R}^n$ be the real n-dimensional vector space. For a metric space (S, d), a point $s \in S$, and a nonempty set $S' \subseteq S$, let $\bar{d}(s, S') := \inf_{s' \in S'} d(s, s')$ be the distance of s to S'. We use the extended definition of the supremum and infimum operators, i.e., the supremum of the empty set is the smallest element of the corresponding domain and the infimum of the empty set is the largest element of the domain. For $t \in \mathbb{R}$ and $I \subseteq \mathbb{R}$, let $t \oplus I$ and $t \ominus I$ denote the Minkowski sum and the Minkowski difference of t and I, respectively. For $a, b \in \mathbb{R}$, let

$$\mathbb{I}(a \leq b) := \begin{cases} 1 & \text{if } a \leq b \\ 0 & \text{otherwise} \end{cases}$$

be the indicator function. Let $\mathfrak{F}(T, S)$ denote the set of all measurable functions mapping from the domain T into the domain S, i.e., $f \in \mathfrak{F}(T, S)$ is a function $f : T \to S$.

A. Signal Temporal Logic

Signal temporal logic [1] is based on deterministic signals $x : T \to \mathbb{R}^n$ where $T := \mathbb{N}$ is assumed throughout the paper. The atomic elements of STL are predicates that are functions $\mu : \mathbb{R}^n \to \mathbb{B}$. Let now M be a set of such predicates μ and let us associate an observation map $O^\mu \subseteq \mathbb{R}^n$ with μ. The observation map O^μ indicates regions where μ is true, i.e.,

$$O^\mu := \mu^{-1}(\top)$$

where $\mu^{-1}(\top)$ denotes the inverse image of $\top$ under μ. We assume throughout the paper that the sets O^μ and $O^{\neg\mu}$ are non-empty and measurable for any $\mu \in M$, i.e., O^μ and $O^{\neg\mu}$ are elements of the Borel σ-algebra $\mathcal{B}^n$ of $\mathbb{R}^n$.

Remark 1: For convenience, the predicate μ is often defined via a predicate function $h : \mathbb{R}^n \to \mathbb{R}$ so that

$$\mu(\zeta) := \begin{cases} \top & \text{if } h(\zeta) \geq 0 \\ \bot & \text{otherwise} \end{cases}$$

for $\zeta \in \mathbb{R}^n$. In this case, we have $O^\mu = \{\zeta \in \mathbb{R}^n | h(\zeta) \geq 0\}$.

For $\mu \in M$, the syntax of STL is defined as

$$\phi \ ::= \ \top \mid \mu \mid \neg\phi \mid \phi' \wedge \phi'' \mid \phi' U_I \phi'' \mid \phi' \underline{U}_I \phi'' \qquad (1)$$

where ϕ' and ϕ'' are STL formulas and where U_I is the future until operator with $I \subseteq \mathbb{R}_{\geq 0}$, while $\underline{U}_I$ is the past until-operator. The operators $\neg$ and $\wedge$ encode negations and conjunctions. Also define the set of operators

$$\begin{aligned}
\phi' \vee \phi'' &:= \neg(\neg\phi' \wedge \neg\phi'') && \text{(disjunction operator)},\\
F_I \phi &:= \top U_I \phi && \text{(future eventually operator)},\\
\underline{F}_I \phi &:= \top \underline{U}_I \phi && \text{(past eventually operator)},\\
G_I \phi &:= \neg F_I \neg\phi && \text{(future always operator)},\\
\underline{G}_I \phi &:= \neg \underline{F}_I \neg\phi && \text{(past always operator)}.
\end{aligned}$$

1) Semantics: We give an STL formula ϕ as in (1) a meaning by the satisfaction function $\beta^\phi : \mathfrak{F}(T, \mathbb{R}^n) \times T \to \mathbb{B}$. In particular, $\beta^\phi(x, t) = \top$ indicates that the signal x satisfies ϕ at time t, while $\beta^\phi(x, t) = \bot$ indicates that x does not satisfy ϕ at time t. For a formal definition of $\beta^\phi(x, t)$, we refer to [35, Appendix A]. An STL formula ϕ is said to be satisfiable if $\exists x \in \mathfrak{F}(T, \mathbb{R}^n)$ such that $\beta^\phi(x, 0) = \top$.

Example 1: Consider a scenario in which a robot operates in a hospital environment. The robot needs to perform two time-critical sequential delivery tasks in regions A and B while avoiding areas C and D. In particular, consider

$$\phi := G_{[0,3]}(\neg\mu_C \wedge \neg\mu_D) \wedge F_{[1,2]}(\mu_A \wedge F_{[0,1]}\mu_B). \qquad (2)$$

To define μ_A, μ_B, μ_C, and μ_D, let a, b, c, and d denote the midpoints of the regions A, B, C, and D as

$$a := \begin{bmatrix} 4 & 5 \end{bmatrix}^T \ b := \begin{bmatrix} 7 & 2 \end{bmatrix}^T \ c \ := \begin{bmatrix} 2 & 3 \end{bmatrix}^T \ d := \begin{bmatrix} 6 & 4 \end{bmatrix}^T.$$

Define the state at time t as $x(t) := \begin{bmatrix} r(t) & a & b & c & d \end{bmatrix}^T \in \mathbb{R}^{10}$ where $r(t)$ is the robot position. The predicates μ_A, μ_B, μ_C, and μ_D are now described by the observation maps

$$\begin{aligned}
O^{\mu_A} &:= \{x \in \mathbb{R}^{10} | \|r - a\|_\infty < 0.5\},\\
O^{\mu_B} &:= \{x \in \mathbb{R}^{10} | \|r - b\|_2 < 0.7\},\\
O^{\mu_C} &:= \{x \in \mathbb{R}^{10} | \|r - c\|_\infty < 0.5\},\\
O^{\mu_D} &:= \{x \in \mathbb{R}^{10} | \|r - d\|_2 < 0.7\}.
\end{aligned}$$

where $\|\cdot\|_2$ is the Euclidean and $\|\cdot\|_\infty$ is the infinity norm. In Fig. 1, six different robot trajectories r_1-r_6 are displayed.

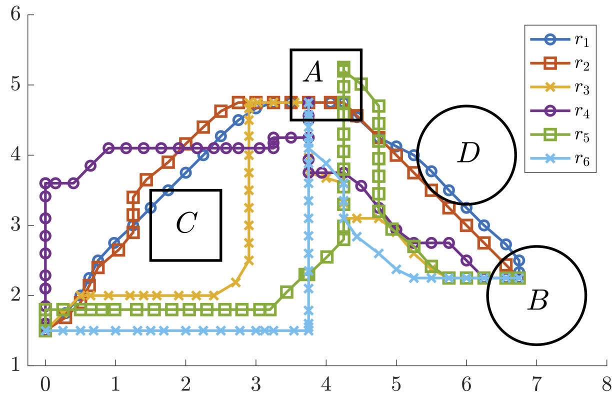

Fig. 1: Shown are six robot trajectories r_1-r_6 along with four regions A, B, C, and D. The specification given in (2) is violated by r_1 and satisfied by r_2-r_6. Furthermore, r_2 only marginally satisfies ϕ, while r_3-r_6 satisfy ϕ robustly.

We omit the exact timings associated with r_1-r_6 in Fig. 1 for readability. However, it can be seen that the signal x_1 that corresponds to r_1 violates ϕ as the region D is entered, while x_2-x_6 satisfy ϕ. In other words, we have $\beta^\phi(x_1,0)=\bot$ and $\beta^\phi(x_j,0)=\top$ for all $j\in\{2,\ldots,6\}$.

2) Robustness: One may additionally be interested in the robustness by which x satisfies ϕ at time t. For this purpose, the *robustness degree* has been introduced in [2, Definition 7]. Define the set of signals that satisfy ϕ at time t as

$$\mathcal{L}^\phi(t):=\{x\in\mathfrak{F}(T,\mathbb{R}^n)|\beta^\phi(x,t)=\top\}.$$

Let us next define the signal metric

$$\kappa(x,x^*):=\sup_{t\in T} d\big(x(t),x^*(t)\big)$$

where $d:\mathbb{R}^n\times\mathbb{R}^n\to\overline{\mathbb{R}}$ is a metric, e.g., the Euclidean distance between two points in $\mathbb{R}^n$. The distance of x to the set $\mathcal{L}^\phi(t)$ is then defined via the metric κ as

$$\text{dist}^\phi(x,t)=\bar{\kappa}\big(x,\text{cl}(\mathcal{L}^\phi(t))\big):=\inf_{x^*\in\text{cl}(\mathcal{L}^\phi(t))}\kappa(x,x^*),$$

where $\text{cl}(\mathcal{L}^\phi(t))$ denotes the closure of $\mathcal{L}^\phi(t)$.

Definition 1 (Robustness Degree): Given an STL formula ϕ and a signal $x\in\mathfrak{F}(T,\mathbb{R}^n)$, the robustness degree at time t is defined as [2, Definition 7]:

$$\mathcal{RD}^\phi(x,t):=\begin{cases}\text{dist}^{\neg\phi}(x,t) \text{ if } x\in\mathcal{L}^\phi(t)\\ -\text{dist}^\phi(x,t) \text{ if } x\notin\mathcal{L}^\phi(t).\end{cases}$$

Intuitively, the robustness degree tells us how much the signal x can be perturbed by additive noise before changing the Boolean truth value of the specification ϕ. If $|\mathcal{RD}^\phi(x,t)|\neq 0$ and $x\in\mathcal{L}^\phi(t)$, all signals $x^*\in\mathfrak{F}(T,\mathbb{R}^n)$ with $\kappa(x,x^*)<|\mathcal{RD}^\phi(x,t)|$ satisfy $x^*\in\mathcal{L}^\phi(t)$.

The robustness degree is a *robust neighborhood*. A robust neighborhood of x is a tube of diameter $\epsilon\geq 0$ around x so that for all x^* in this tube we have $\beta^\phi(x,t)=\beta^\phi(x^*,t)$. Specifically, for $\epsilon\geq 0$ and $x:T\to\mathbb{R}^n$ with $x\in\mathcal{L}^\phi(t)$, a set $\{x'\in\mathfrak{F}(T,\mathbb{R}^n)|\kappa(x,x')<\epsilon\}$ is a robust neighborhood if $x^*\in\{x'\in\mathfrak{F}(T,\mathbb{R}^n)|\kappa(x,x')<\epsilon\}$ implies $x^*\in\mathcal{L}^\phi(t)$.

3) Robust Semantics: In general, it is difficult to calculate the robustness degree $\mathcal{RD}^\phi(x,t)$ as the set $\mathcal{L}^\phi(t)$ is hard to calculate. The authors in [2] introduce the *robust semantics* $\rho^\phi:\mathfrak{F}(T,\mathbb{R}^n)\times T\to\mathbb{R}$ as an alternative way of finding a robust neighborhood.

Definition 2 (STL Robust Semantics): For a signal $x:T\to\mathbb{R}^n$, the robust semantics $\rho^\phi(x,t)$ of an STL formula ϕ are inductively defined as

$$\begin{aligned}
\rho^\top(x,t)&:=\infty,\\
\rho^\mu(x,t)&:=\begin{cases}\text{dist}^{\neg\mu}(x,t) & \text{if } x\in\mathcal{L}^\mu(t)\\ -\text{dist}^\mu(x,t) & \text{otherwise,}\end{cases}\\
\rho^{\neg\phi}(x,t)&:=-\rho^\phi(x,t),\\
\rho^{\phi'\wedge\phi''}(x,t)&:=\min(\rho^{\phi'}(x,t),\rho^{\phi''}(x,t)),\\
\rho^{\phi' U_I\phi''}(x,t)&:=\sup_{t''\in(t\oplus I)\cap T}\Big(\min\big(\rho^{\phi''}(x,t''),\\
&\qquad\inf_{t'\in(t,t'')\cap T}\rho^{\phi'}(x,t')\big)\Big),\\
\rho^{\phi'\underline{U}_I\phi''}(x,t)&:=\sup_{t''\in(t\ominus I)\cap T}\Big(\min\big(\rho^{\phi''}(x,t''),\\
&\qquad\inf_{t'\in(t,t'')\cap T}\rho^{\phi'}(x,t')\big)\Big).
\end{aligned}$$

Importantly, it was shown in [2, Theorem 28] that

$$-\text{dist}^\phi(x,t)\leq\rho^\phi(x,t)\leq\text{dist}^{\neg\phi}(x,t). \tag{3}$$

so that $|\rho^\phi(x,t)|\leq|\mathcal{RD}^\phi(x,t)|$. The robust semantics $\rho^\phi(x,t)$ hence provide a more tractable under-approximation of the robustness degree $\mathcal{RD}^\phi(x,t)$. The robust semantics are sound in the sense that $\beta^\phi(x,t)=\top$ if $\rho^\phi(x,t)>0$ and $\beta^\phi(x,t)=\bot$ if $\rho^\phi(x,t)<0$ [2, Proposition 30].

Example 2: For Example 1 and the trajectories shown in Fig. 1, we obtain $\rho^\phi(x_1,0)=-0.15$, $\rho^\phi(x_2,0)=0.01$, and $\rho^\phi(x_j,0)=0.25$ for all $j\in\{3,\ldots,6\}$ when choosing $d(\cdot)$ as the Euclidean distance. The reason for x_1 having negative robustness lies in r_1 intersecting with the region D. Marginal robustness of x_2 is explained as r_2 only marginally avoids the region D while all other trajectories avoid D robustly.

B. Random Variables and Stochastic Processes

Instead of interpreting an STL specifications ϕ over deterministic signals, we will interpret ϕ over stochastic processes. Consider therefore the *probability space* $(\Omega,\mathcal{F},P)$ where Ω is the sample space, $\mathcal{F}$ is a σ-algebra of Ω, and $P:\mathcal{F}\to[0,1]$ is a probability measure. More intuitively, an element in Ω is an *outcome* of an experiment, while an element in $\mathcal{F}$ is an *event* that consists of one or more outcomes whose probabilities can be measured by the probability measure P.

1) Random Variables: Let Z denote a real-valued *random vector*, i.e., a measurable function $Z:\Omega\to\mathbb{R}^n$.[1] When $n=1$, we say Z is a *random variable*. We refer to $Z(\omega)$ as a realization of the random vector Z where $\omega\in\Omega$. Since Z is a measurable function, a distribution can be assigned to Z

[1]More precisely, we have $Z:\Omega\times\mathcal{F}\to\mathbb{R}^n\times\mathcal{B}^n$ where $\mathcal{B}^n$ is the Borel σ-algebra of $\mathbb{R}^n$, i.e., Z maps a *measurable space* to yet another measurable space. For convenience, this more involved notation is, however, omitted.

and a cumulative distribution function (CDF) $F_Z(z)$ can be defined for Z (see [35, Appendix B]).

Given a random vector Z, we can derive other random variables that we call *derived random variables*. Assume for instance a measurable function $g : \mathbb{R}^n \to \mathbb{R}$ and notice that $G : \Omega \to \mathbb{R}$ with $G(\omega) := g(Z(\omega))$ becomes yet another random variable since function composition preserves measureability. See [36] for a more detailed discussion.

2) Stochastic Processes: A *stochastic process* is a function of the variables $\omega \in \Omega$ and $t \in T$ where T is the time domain. Recall that the time domain is discrete, i.e., $T := \mathbb{N}$, so that we consider discrete-time stochastic processes. This assumption is made for simplicity. The presented results carry over, with some modifications, to the continuous-time case that we defer to another paper. A stochastic process is now a function $X : T \times \Omega \to \mathbb{R}^n$ where $X(t, \cdot)$ is a random vector for each fixed $t \in T$. A stochastic process can be viewed as a collection of random vectors $\{X(t, \cdot) | t \in T\}$ that are defined on a common probability space $(\Omega, \mathcal{F}, P)$ and that are indexed by T. For a fixed $\omega \in \Omega$, the function $X(\cdot, \omega)$ is a *realization* of the stochastic process. Another equivalent definition is that a stochastic process is a collection of deterministic functions of time

$$\{X(\cdot, \omega) | \omega \in \Omega\} \tag{4}$$

that are indexed by Ω. While the former definition is intuitive, the latter allows to define a *random function* mapping from the sample space Ω into the space of functions $\mathfrak{F}(T, \mathbb{R}^n)$.

C. Risk Measures

A *risk measure* is a function $R : \mathfrak{F}(\Omega, \mathbb{R}) \to \mathbb{R}$ that maps from the set of real-valued random variables to the real numbers. In particular, we refer to the input of a risk measure R as the *cost random variable* since typically a cost is associated with the input of R. Risk measures hence allow for a risk assessment in terms of such cost random variables. Commonly used risk measures are the expected value, the variance, or the conditional value-at-risk [24]. A particular property of R that we need in this paper is monotonicity. For two cost random variables $Z, Z' \in \mathfrak{F}(\Omega, \mathbb{R})$, the risk measure R is monotone if $Z(\omega) \leq Z'(\omega)$ for all $\omega \in \Omega$ implies that $R(Z) \leq R(Z')$.

Remark 2: In [35, Appendix C], we summarize other desireable properties of R such as translation invariance, positive homogeneity, subadditivity, commotone additivity, and law invariance. We also provide a summary of existing risk measures. We emphasize that our presented method is compatible with any of these risk measures as long as they are monotone.

III. Risk of STL Specifications

While an STL formula ϕ as defined in Section II-A is defined over deterministic signals x, we will interpret ϕ over a stochastic process X as defined in Section II-B.

For a particular realization $X(\cdot, \omega)$ of the stochastic process X, we can evaluate whether or not $X(\cdot, \omega)$ satisfies ϕ. For the stochastic process X, however, it is not clear how to interpret the satisfaction of ϕ by X. In fact, some realizations of X may satisfy ϕ while some other realizations of X may violate ϕ. To bridge this gap, we use risk measures as introduced in Section II-C to argue about the risk of the stochastic process X not satisfying the specification ϕ.

Before going into the main parts of this paper, we remark that all important symbols that have been or will be introduced are summarized in Table I.

A. Measurability of STL Semantics and Robustness Degree

Note that the semantics and the robust semantics as well as the robustness degree become stochastic entities when evaluated over a stochastic process X, i.e., the functions $\beta^\phi(X, t)$, $\rho^\phi(X, t)$, and $RD^\phi(X, t)$ become stochastic entities. We first provide conditions under which $\beta^\phi(X, t)$ and $\rho^\phi(X, t)$ become (derived) random variables, which boils down to showing that $\beta^\phi(X(\cdot, \omega), t)$ and $\rho^\phi(X(\cdot, \omega), t)$ are measurable in ω for a fixed $t \in T$.

Theorem 1: Let X be a discrete-time stochastic process and let ϕ be an STL specification. Then $\beta^\phi(X(\cdot, \omega), t)$ and $\rho^\phi(X(\cdot, \omega), t)$ are measurable in ω for a fixed $t \in T$ so that $\beta^\phi(X, t)$ and $\rho^\phi(X, t)$ are random variables.

By Theorem 1, the probabilities $P(\beta^\phi(X, t) \in B)$ and $P(\rho^\phi(X, t) \in B)$[2] are well defined for measurable sets B from the corresponding measurable space.

We next show measurability of the distance function $\text{dist}^\phi(X(\cdot, \omega), t)$ and the robustness degree $\mathcal{RD}^\phi(X(\cdot, \omega), t)$.

Theorem 2: Let X be a discrete-time stochastic process and let ϕ be an STL specification. Then $\text{dist}^\phi(X(\cdot, \omega), t)$ and $\mathcal{RD}^\phi(X(\cdot, \omega), t)$ are measurable in ω for a fixed $t \in T$ so that $\text{dist}^\phi(X, t)$ and $\mathcal{RD}^\phi(X, t)$ are random variables.

B. The STL Robustness Risk

Towards defining the risk of not satisfying a specification ϕ, note that the expression $R(\beta^\phi(X, t) = \perp)$ is not well defined as opposed to $P(\beta^\phi(X, t) = \perp)$ that indicates the probability of not satisfying ϕ. The reason for this is that the function R takes a real-valued cost random variable as its input. We can instead evaluate $R(-\beta^\phi(X, t))$, but not much information will be gained due to the binary encoding of the STL semantics $\beta^\phi(X, t)$.

1) The risk of not satisfying ϕ robustly: Instead, we will define the risk of the stochastic process X not satisfying ϕ robustly by considering $\text{dist}^{\neg\phi}(X, t)$. As shown, $\text{dist}^{\neg\phi}(X, t)$ is a random variable indicating the distance between realizations of the stochastic process and the set of signals $\mathcal{L}^{\neg\phi}(t)$ that violate ϕ. We refer to the following definition as the *STL robustness risk* for brevity.

Definition 3 (STL Robustness Risk): Given an STL formula ϕ and a stochastic process $X : T \times \Omega \to \mathbb{R}^n$, the risk of X not satisfying ϕ robustly at time t is defined as

$$R(-\text{dist}^{\neg\phi}(X, t)).$$

Fig. 2 illustrates the idea underlying Definition 3 and shows $\text{dist}^{\neg\phi}(X(\cdot, \omega_i), t)$ for realizations $X(\cdot, \omega_i)$ of the

[2]We use the shorthand notations $P(\beta^\phi(X, t) \in B)$ and $P(\rho^\phi(X, t) \in B)$ instead of the more complex notations $P(\{\omega \in \Omega | \beta^\phi(X(\cdot, \omega), t) \in B\})$ and $P(\{\omega \in \Omega | \rho^\phi(X(\cdot, \omega), t) \in B\})$, respectively.

Symbol	Meaning
x	Deterministic signal $x : T \to \mathbb{R}^n$
$\mathfrak{F}(T, \mathbb{R}^n)$	Set of all measurable functions mapping from T to $\mathbb{R}^n$
$\beta^\phi(x,t)$	Boolean semantics $\beta^\phi : \mathfrak{F}(T, \mathbb{R}^n) \times T \to \mathbb{B}$ of an STL formula ϕ
$\mathcal{L}^\phi(t)$	Set of deterministic signals x that satisfy ϕ at time t
$\text{dist}^\phi(x,t)$	Distance of the signal x to the set $\mathcal{L}^\phi(t)$
$\mathcal{RD}^\phi(x,t)$	Robustness degree $\mathcal{RD}^\phi : \mathfrak{F}(T, \mathbb{R}^n) \times T \to \mathbb{R}$ of an STL formula ϕ
$\rho^\phi(x,t)$	Robust semantics $\rho^\phi : \mathfrak{F}(T, \mathbb{R}^n) \times T \to \mathbb{R}$ of an STL formula ϕ
X	Stochastic Process $X : T \times \Omega \to \mathbb{R}^n$
$R(-\text{dist}^{\neg\phi}(X,t))$	The STL robustness risk, i.e., the risk of the stochastic process X not satisfying the STL formula ϕ robustly at time t
$R(-\rho^\phi(X,t))$	The approximate STL robustness risk

TABLE I: Summary of robustness and risk notions for signal temporal logic.

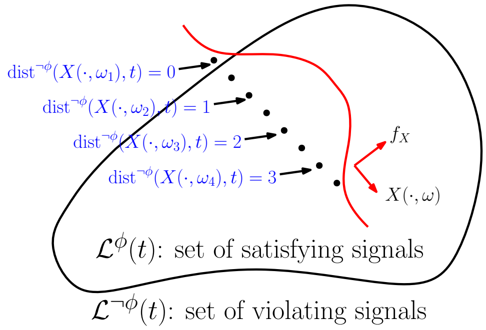

Fig. 2: Illustration of $\text{dist}^{\neg\phi}(X(\cdot,\omega_i),t)$, i.e., the distance between realizations $X(\cdot,\omega_i)$ and the set $\mathcal{L}^{\neg\phi}(t)$. Note that the figure shows the function space $\mathfrak{F}(T,\mathbb{R}^n)$ and not $\mathbb{R}^n$.

stochastic process X where $\omega_i \in \Omega$, i.e., $\text{dist}^{\neg\phi}(X(\cdot,\omega_i),t)$ is the distance between the realization $X(\cdot,\omega_i)$ and the set $\mathcal{L}^{\neg\phi}(t)$. Positive values of $\text{dist}^{\neg\phi}(X(\cdot,\omega_i),t)$ indicate that the realization $X(\cdot,\omega_i)$ satisfies ϕ at time t, while the value zero indicates that the realization $X(\cdot,\omega_i)$ either marginally satisfies ϕ at time t or does not satisfy ϕ at time t. Furthermore, large positive values of $\text{dist}^{\neg\phi}(X(\cdot,\omega_i),t)$ indicate robust satisfaction and are hence desirable. This is the reason why $-\text{dist}^{\neg\phi}(X,t)$ is considered in Definition 3 as the cost random variable. To complement Fig. 2, note that the red curve sketches a possible distribution of X and hence the probability by which a realization occurs. Note that the robustness degree of the corresponding realizations in Fig. 2 would be $\mathcal{RD}^\phi(X(\cdot,\omega_1),t) < 0$, $\mathcal{RD}^\phi(X(\cdot,\omega_2),t) = 1$, $\mathcal{RD}^\phi(X(\cdot,\omega_3),t) = 2$, and $\mathcal{RD}^\phi(X(\cdot,\omega_4),t) = 3$.

Example 3: Consider the value-at-risk at level $\beta := 0.95$ (see [35, Appendix C]), which is also known as the $1-\beta$ risk quantile. Assume that we obtained $VaR_\beta(-\text{dist}^{\neg\phi}(X,t)) = -1$ for a given stochastic process X and a given STL formula ϕ. The interpretation is now that with a probability of 0.05 the robustness $\text{dist}^{\neg\phi}(X,t)$ is smaller than (or equal to) $|VaR_\beta(-\text{dist}^{\neg\phi}(X,t))| = 1$. Or in other words, with a probability of 0.95, the robustness $\text{dist}^{\neg\phi}(X,t)$ is greater than $|VaR_\beta(-\text{dist}^{\neg\phi}(X,t))| = 1$.

Remark 3: An alternative to Definition 3 would be to consider $R(-\mathcal{RD}^\phi(X,t))$ instead of $R(-\text{dist}^{\neg\phi}(X,t))$. We, however, refrain from such a definition since the meaning of $\mathcal{RD}^\phi(X(\cdot,\omega_i),t)$ for realizations $w_i \in \Omega$ with $\mathcal{RD}^\phi(X(\cdot,\omega_i),t) < 0$ is not what we aim for here. In particular, when $\mathcal{RD}^\phi(X(\cdot,\omega_i),t) < 0$ we have that $\mathcal{RD}^\phi(X(\cdot,\omega_i),t) = -\text{dist}^\phi(X,t)$. In this case, $|\mathcal{RD}^\phi(X(\cdot,\omega_i),t)|$ indicates the robustness by which $X(\cdot,\omega_i)$ satisfies $\neg\phi$, while we are interested in the opposite.

Unfortunately, the risk of not satisfying ϕ robustly, i.e., $R(-\text{dist}^{\neg\phi}(X,t))$, can in most of the cases not be computed. Instead, we will focus on $R(-\rho^\phi(X,t))$ as an approximate risk of not satisfying ϕ robustly, i.e., the *approximate STL robustness risk*. We next show that this approximation has the desirable property of being an over-approximation.

2) Approximating the risk of not satisfying ϕ robustly: A desirable property, which we will show to hold, is that $R(-\rho^\phi(X,t))$ over-approximates $R(-\text{dist}^{\neg\phi}(X,t))$ so that $R(-\rho^\phi(X,t))$ is more risk-aware than $R(-\text{dist}^{\neg\phi}(X,t))$, i.e., that it holds that $R(-\text{dist}^{\neg\phi}(X,t)) \leq R(-\rho^\phi(X,t))$.

Theorem 3: Let R be a monotone risk measure. Then it holds that $R(-\text{dist}^{\neg\phi}(X,t)) \leq R(-\rho^\phi(X,t))$.

This indeed enables us to use $R(-\rho^\phi(X,t))$ instead of $R(-\text{dist}^{\neg\phi}(X,t))$. For two stochastic processes X_1 and X_2, note that $R(-\rho^\phi(X_1,t)) \leq R(-\rho^\phi(X_2,t))$ means that X_1 has less risk than X_2 with respect to the specification ϕ.

Oftentimes, one may be interested in associating a monetary cost with $\text{dist}^{\neg\phi}(X,t)$ that reflects the severity of an event with low robustness. One may hence want to assign high costs to low robustness and low costs to high robustness. Let us define an increasing cost function $C : \mathbb{R} \to \mathbb{R}$ that reflects this preference.

Corollary 1: Let R be a monotone risk measure and C be an increasing cost function. Then it holds that $R(C(-\text{dist}^{\neg\phi}(X,t))) \leq R(C(-\rho^\phi(X,t)))$.

IV. Data-Driven Estimation of the STL Robustness Risk

There are two main challenges in computing the approximate STL robustness risk. First, note that calculation of $R(-\rho^\phi(X,t))$ requires knowledge of the CDF of $\rho^\phi(X,t)$ no matter what the choice of the risk measure R will be. However, the CDF of $\rho^\phi(X,t)$ is not known (only the CDF of X is known) and deriving the CDF of $\rho^\phi(X,t)$ is often not possible. Second, calculating $R(-\rho^\phi(X,t))$ may involve solving high dimensional integrals.

In this paper, we assume a data-driven setting where realizations X^i of the stochastic process X are observed, e.g., obtained from experiments, and where not even the CDF of X is known. We present a data-driven sample average approximation $\overline{R}(-\rho^\phi(X^i,t))$ of $R(-\rho^\phi(X,t))$ for the value-at-risk and show that this approximation has the favorable property of being an upper bound to $R(-\rho^\phi(X,t))$, i.e., that $R(-\rho^\phi(X,t)) \leq \overline{R}(-\rho^\phi(X^i,t))$ with high probability.

A. Sample Average Approximation of the Value-at-Risk (VaR)

Let us obtain a sample average approximation of the value-at-risk (VaR) and define for convenience the random variable

$$Z := -\rho^\phi(X,t).$$

For further convenience, let us define the tuple

$$\mathcal{Z} := (Z^1, \ldots, Z^N)$$

where $Z^i := -\rho^\phi(X^i,t)$ and where $X^1, \ldots, X^N$ are N independent copies of X. Consequently, all Z^i are independent and identically distributed.

For a risk level of $\beta \in (0,1)$, the VaR of Z is given by

$$VaR_\beta(Z) := \inf\{\alpha \in \mathbb{R} | F_Z(\alpha) \geq \beta\}$$

where we recall that F_Z is the CDF of Z. To estimate $F_Z(\alpha)$, define the empirical CDF

$$\widehat{F}(\alpha, \mathcal{Z}) := \frac{1}{N}\sum_{i=1}^{N} \mathbb{I}(Z^i \leq \alpha)$$

where we recall that $\mathbb{I}$ denotes the indicator function. Let $\delta \in (0,1)$ be a probability threshold. Inspired by [37] and [38], we calculate an upper bound of $VaR_\beta(Z)$ as

$$\overline{VaR}_\beta(\mathcal{Z},\delta) := \inf\left\{\alpha \in \overline{\mathbb{R}} | \widehat{F}(\alpha,\mathcal{Z}) - \sqrt{\frac{\ln(2/\delta)}{2N}} \geq \beta\right\} \tag{5}$$

and a lower bound as

$$\underline{VaR}_\beta(\mathcal{Z},\delta) := \inf\left\{\alpha \in \overline{\mathbb{R}} | \widehat{F}(\alpha,\mathcal{Z}) + \sqrt{\frac{\ln(2/\delta)}{2N}} \geq \beta\right\}$$

where we recall that $\inf \emptyset = \infty$, for $\emptyset$ being the empty set, due to the extended definition of the infimum operator. We next show that $\overline{VaR}_\beta(\mathcal{Z},\delta)$ is an upper bound of $VaR_\beta(Z)$ with a probability of at least $1-\delta$.

Theorem 4: Let $\delta \in (0,1)$ be a probability threshold and $\beta \in (0,1)$ be a risk level. Let $\overline{VaR}_\beta(\mathcal{Z},\delta)$ and $\underline{VaR}_\beta(\mathcal{Z},\delta)$ be based on $\mathcal{Z}$ that are N independent copies of Z. With probability of at least $1-\delta$, it holds that

$$\underline{VaR}_\beta(\mathcal{Z},\delta) \leq VaR_\beta(Z) \leq \overline{VaR}_\beta(\mathcal{Z},\delta).$$

Theorem 4 provides an upper bound $\overline{VaR}_\beta(\mathcal{Z},\delta)$ for $VaR_\beta(Z)$ as desired, while the lower bound $\underline{VaR}_\beta(\mathcal{Z},\delta)$ indicates how conservative our upper bound may have been.

As a result of Theorems 3 and 4 and the fact that the VaR is a monotone risk measure, we have now a procedure to find a tight upper bound $\overline{VaR}_\beta(\mathcal{Z},\delta)$ of $VaR_\beta(Z)$ with high probability. To summarize, we observe N realizations $X^1, \ldots, X^N$ of the stochastic process X. We then select $\delta, \beta \in (0,1)$ and are guaranteed that with probability of at least $1-\delta$

$$VaR_\beta(-\text{dist}^{\neg\phi}(X,t)) \leq VaR_\beta(Z) \leq \overline{VaR}_\beta(\mathcal{Z},\delta).$$

Remark 4: Upper and lower bounds for other risk measures than the value-at-risk can often be derived. For the expected value $E(-\rho^\phi(X,t))$, concentration inequalities for the sample average approximation of $E(-\rho^\phi(X,t))$ can be obtained by applying Hoeffding's inequality when $\rho^\phi(X,t)$ is bounded. For the conditional value-at-risk $CVaR(-\rho^\phi(X,t))$, concentration inequalities are presented in [39]–[42]. We plan to address this in future work.

V. Case Study

We continue with the case study presented in Example 1. Now, however, the environment is uncertain as the regions C and D in which humans operate are not exactly known. Let therefore c and d be Gaussian random vectors as

$$c \sim \mathcal{N}\Big(\begin{bmatrix}2\\3\end{bmatrix}, \begin{bmatrix}0.125 & 0\\ 0 & 0.125\end{bmatrix}\Big),$$
$$d \sim \mathcal{N}\Big(\begin{bmatrix}6\\4\end{bmatrix}, \begin{bmatrix}0.125 & 0\\ 0 & 0.125\end{bmatrix}\Big),$$

where $\mathcal{N}$ denotes a multivariable Gaussian distribution with according mean vector and covariance matrix.

Consequently, the signals x_1-x_6 become stochastic processes denoted by X_1-X_6. Our goal is now to calculate

$$\overline{VaR}_\beta(-\rho^\phi(X_j^1,t), \ldots, -\rho^\phi(X_j^N,t), A, \delta)$$

for each $j \in \{1, \ldots, 6\}$ to compare the risk between the six robot trajectories r_1-r_6. We set $\delta := 0.001$ and $N := 6500$. For different β, the resulting $\overline{VaR}_\beta$ are shown in the following table.

j \ β	0.9	0.925	0.95	0.975
1	0.336	0.363	0.395	0.539
2	0.162	0.187	0.220	0.336
3	-0.17	-0.152	-0.121	-0.008
4	-0.249	-0.249	-0.249	-0.19
5	-0.25	-0.25	-0.25	-0.149
6	-0.249	-0.249	-0.249	-0.249

Across all β, the table indicates that trajectories r_1 and r_2 are not favorable in terms of the induced STL robustness risk. Trajectory r_3 is better compared to trajectories r_1 and r_2, but worse than r_4-r_6 in terms of the robustness risk of ϕ. For trajectories r_4-r_6, note that a β of 0.9, 0.925, and 0.95 provides the information that the trajectories have roughly the same robustness risk. However, once the risk level β is increased to 0.975, it becomes clear that r_6 is preferable over r_4 that is again preferable over r_5. This matches with what one would expect by closer inspection of Fig. 1.

VI. Conclusion

We defined the risk of a stochastic process not satisfying a signal temporal logic specification robustly which we referred to as the "STL robustness risk". We also presented an approximation of the STL robustness risk that is an upper bound of the STL robustness risk when the used risk measure is monotone. For the case of the value-at-risk, we presented a data-driven method to estimate the approximate STL robustness risk.

Acknowledgments

The authors would like to thank Alëna Rodionova and Matthew Cleaveland for proofreading parts of this paper.

References

[1] O. Maler and D. Nickovic, "Monitoring temporal properties of continuous signals," in *Proc. Int. Conf. FORMATS FTRTFT*, Grenoble, France, September 2004, pp. 152–166.

[2] G. E. Fainekos and G. J. Pappas, "Robustness of temporal logic specifications for continuous-time signals," *Theoret. Comp. Science*, vol. 410, no. 42, pp. 4262–4291, 2009.

[3] A. Donzé and O. Maler, "Robust satisfaction of temporal logic over real-valued signals," in *Proc. Int. Conf. FORMATS*, Klosterneuburg, Austria, September 2010, pp. 92–106.

[4] N. Mehdipour, C. Vasile, and C. Belta, "Arithmetic-geometric mean robustness for control from signal temporal logic specifications," in *Proc. Am. Control Conf.*, Philadelphia, PA, July 2019, pp. 1690–1695.

[5] I. Haghighi, N. Mehdipour, E. Bartocci, and C. Belta, "Control from signal temporal logic specifications with smooth cumulative quantitative semantics," in *Proc. Conf. Decis. Control*, Nice, France, December 2019, pp. 4361–4366.

[6] T. Akazaki and I. Hasuo, "Time robustness in MTL and expressivity in hybrid system falsification," in *Proc. Int. Conf. Comp. Aided Verif.*, San Francisco, CA, July 2015, pp. 356–374.

[7] P. Varnai and D. V. Dimarogonas, "On robustness metrics for learning STL tasks," in *Proc. Am. Control Conf.*, Denver, CO, July 2020, pp. 5394–5399.

[8] C. Madsen, P. Vaidyanathan, S. Sadraddini, C.-I. Vasile, N. A. DeLateur, R. Weiss, D. Densmore, and C. Belta, "Metrics for signal temporal logic formulae," in *Proc. Conf. Decis. Control*, Nice, France, December 2018, pp. 1542–1547.

[9] A. Rodionova, E. Bartocci, D. Nickovic, and R. Grosu, "Temporal logic as filtering," in *Proc. Int. Conf. Hybrid Syst.: Comp. Control*, Vienna, Austria, April 2016, pp. 11–20.

[10] M. Tiger and F. Heintz, "Incremental reasoning in probabilistic signal temporal logic," *Int. Journal Approx. Reasoning*, vol. 119, pp. 325–352, 2020.

[11] J. Li, P. Nuzzo, A. Sangiovanni-Vincentelli, Y. Xi, and D. Li, "Stochastic contracts for cyber-physical system design under probabilistic requirements," in *Proc. Int. Conf. Form. Meth. Mod. Syst. Design Formal*, Vienna, Austria, September 2017, pp. 5–14.

[12] P. Kyriakis, J. V. Deshmukh, and P. Bogdan, "Specification mining and robust design under uncertainty: A stochastic temporal logic approach," *ACM Trans. Embed. Comp. Syst.*, vol. 18, no. 5s, pp. 1–21, 2019.

[13] D. Sadigh and A. Kapoor, "Safe control under uncertainty with probabilistic signal temporal logic," in *Proc. Robot.: Science Syst.*, AnnArbor, Michigan, June 2016.

[14] S. Jha, V. Raman, D. Sadigh, and S. A. Seshia, "Safe autonomy under perception uncertainty using chance-constrained temporal logic," *Journal Autom. Reasoning*, vol. 60, no. 1, pp. 43–62, 2018.

[15] L. Lindemann, G. J. Pappas, and D. V. Dimarogonas, "Control barrier functions for nonholonomic systems under risk signal temporal logic specifications," in *Proc. Conf. Decis. Control*, Jeju Island, Republic of Korea, December 2020, pp. 1422–1428.

[16] S. Safaoui, L. Lindemann, D. V. Dimarogonas, I. Shames, and T. H. Summers, "Control design for risk-based signal temporal logic specifications," *IEEE Control Syst. Lett.*, vol. 4, no. 4, pp. 1000–1005, 2020.

[17] S. S. Farahani, R. Majumdar, V. S. Prabhu, and S. Soudjani, "Shrinking horizon model predictive control with signal temporal logic constraints under stochastic disturbances," *IEEE Trans. Autom. Control*, vol. 64, no. 8, pp. 3324–3331, 2018.

[18] S. Bharadwaj, R. Dimitrova, and U. Topcu, "Synthesis of surveillance strategies via belief abstraction," in *Proc. Conf. Decis. Control*, Miami, FL, Dec. 2018, pp. 4159–4166.

[19] C.-I. Vasile, K. Leahy, E. Cristofalo, A. Jones, M. Schwager, and C. Belta, "Control in belief space with temporal logic specifications," in *Proc. Conf. Decis. Control*, Las Vegas, NV, Dec. 2016, pp. 7419–7424.

[20] M. Lahijanian, S. B. Andersson, and C. Belta, "Formal verification and synthesis for discrete-time stochastic systems," *IEEE Trans. Autom. Control*, vol. 60, no. 8, pp. 2031–2045, 2015.

[21] M. Guo and M. M. Zavlanos, "Probabilistic motion planning under temporal tasks and soft constraints," *IEEE Trans. Autom. Control*, vol. 63, no. 12, pp. 4051–4066, 2018.

[22] E. Bartocci, L. Bortolussi, L. Nenzi, and G. Sanguinetti, "On the robustness of temporal properties for stochastic models," in *Proc. Int. Workshop Hybrid Syst. Biology*, Taormina, Italy, Sept. 2013, pp. 3–19.

[23] ——, "System design of stochastic models using robustness of temporal properties," *Theoret. Comp. Science*, vol. 587, pp. 3–25, 2015.

[24] R. T. Rockafellar, S. Uryasev *et al.*, "Optimization of conditional value-at-risk," *Journal of risk*, vol. 2, pp. 21–42, 2000.

[25] R. T. Rockafellar and S. Uryasev, "Conditional value-at-risk for general loss distributions," *Journal of banking & finance*, vol. 26, no. 7, pp. 1443–1471, 2002.

[26] A. Majumdar and M. Pavone, "How should a robot assess risk? towards an axiomatic theory of risk in robotics," in *Robotics Research*. Springer, 2020, pp. 75–84.

[27] S. Samuelson and I. Yang, "Safety-aware optimal control of stochastic systems using conditional value-at-risk," in *Proc. Am. Control Conf.*, Milwaukee, WI, June 2018, pp. 6285–6290.

[28] E. Hyeon, Y. Kim, and A. G. Stefanopoulou, "Fast risk-sensitive model predictive control for systems with time-series forecasting uncertainties," in *Proc. Conf. Decis Control*, Jeju Island, Republic of Korea, Dec. 2020, pp. 2515–2520.

[29] S. G. McGill, G. Rosman, T. Ort, A. Pierson, I. Gilitschenski, B. Araki, L. Fletcher, S. Karaman, D. Rus, and J. J. Leonard, "Probabilistic risk metrics for navigating occluded intersections," *IEEE Robot. Autom. Lett.*, vol. 4, no. 4, pp. 4322–4329, 2019.

[30] 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, 2018.

[31] D. S. Kalogerias, L. F. Chamon, G. J. Pappas, and A. Ribeiro, "Better safe than sorry: Risk-aware nonlinear bayesian estimation," in *Proc. Int. Conf. Acoustics, Speech Signal Proc.*, Barcelona, Spain, May 2020, pp. 5480–5484.

[32] M. P. Chapman, J. Lacotte, A. Tamar, D. Lee, K. M. Smith, V. Cheng, J. F. Fisac, S. Jha, M. Pavone, and C. J. Tomlin, "A risk-sensitive finite-time reachability approach for safety of stochastic dynamic systems," in *Proc. Am. Control Conf.*, Philadelphia, PA, June 2019, pp. 2958–2963.

[33] M. Ahmadi, M. Ono, M. D. Ingham, R. M. Murray, and A. D. Ames, "Risk-averse planning under uncertainty," in *Proc. Am. Control Conf.*, Denver, CO, July 2020, pp. 3305–3312.

[34] M. Schuurmans and P. Patrinos, "Learning-based distributionally robust model predictive control of markovian switching systems with guaranteed stability and recursive feasibility," in *Proc. Conf. Decis. Control*, Jeju Island, Republic of Korea, December 2020, pp. 4287–4292.

[35] L. Lindemann, N. Matni, and G. J. Pappas, "STL robustness risk over discrete-time stochastic processes," *arXiv preprint arXiv:2104.01503*, 2021.

[36] R. Durrett, *Probability: theory and examples*. Cambridge university press, 2019, vol. 49.

[37] B. Szorenyi, R. Busa-Fekete, P. Weng, and E. Hüllermeier, "Qualitative multi-armed bandits: A quantile-based approach," in *Int. Conf. Machine Learning*, Lille, France, July 2015, pp. 1660–1668.

[38] P. Massart, "The tight constant in the dvoretzky-kiefer-wolfowitz inequality," *The annals of Probability*, pp. 1269–1283, 1990.

[39] S. P. Bhat and P. LA, "Concentration of risk measures: A wasserstein distance approach," *Proc. Adv. Neural Inform. Process. Syst.*, vol. 32, pp. 11 762–11 771, Dec. 2019.

[40] D. B. Brown, "Large deviations bounds for estimating conditional value-at-risk," *Operations Research Letters*, vol. 35, no. 6, pp. 722–730, 2007.

[41] P. Thomas and E. Learned-Miller, "Concentration inequalities for conditional value at risk," in *Proc. Int. Conf. Mach. Learning*, Long Beach, CA, June 2019, pp. 6225–6233.

[42] Z. Mhammedi, B. Guedj, and R. C. Williamson, "Pac-bayesian bound for the conditional value at risk," *Proc. Adv. Neural Inform. Process. Syst.*, vol. 33, pp. 17 919–17 930, Dec. 2020.