

Collision Probabilities for Continuous-Time Systems Without Sampling

Kristoffer M. Frey^{*†}, Ted J. Steiner[†], and Jonathan P. How^{*}

^{*}Department of Aeronautics and Astronautics, MIT

[†]The Charles Stark Draper Laboratory, Inc.

Cambridge, MA 02139

Email: kfrey@mit.edu

Abstract—Demand for high-performance, robust, and safe autonomous systems has grown substantially in recent years. These objectives motivate the desire for efficient risk estimation that can be embedded in core decision-making tasks such as motion planning. On one hand, Monte-Carlo (MC) and other sampling-based techniques provide accurate solutions for a wide variety of motion models but are cumbersome in the context of continuous optimization. On the other hand, “direct” approximations aim to compute (or upper-bound) the failure probability as a smooth function of the decision variables, and thus are convenient for optimization. However, existing direct approaches fundamentally assume discrete-time dynamics and can perform unpredictably when applied to continuous-time systems ubiquitous in the real world, often manifesting as severe conservatism. State-of-the-art attempts to address this within a conventional discrete-time framework require additional Gaussianity approximations that ultimately produce inconsistency of their own. In this paper we take a fundamentally different approach, deriving a risk approximation framework directly in continuous time and producing a lightweight estimate that actually converges as the underlying discretization is refined. Our approximation is shown to significantly outperform state-of-the-art techniques in replicating the MC estimate while maintaining the functional and computational benefits of a direct method. This enables robust, risk-aware, continuous motion-planning for a broad class of nonlinear, partially-observable systems.

I. INTRODUCTION

Robotic motion planning is a decision-making problem that must balance optimality and safety. In the real world, these decisions are complicated by the presence of uncertainty due to imperfect sensing, partial observability, and stochastic dynamics. This uncertainty is often difficult or impossible to explicitly bound, and safety cannot be guaranteed against all realizations of noise and disturbance.

This motivates the replacement of deterministic safety constraints with *risk constraints* [12] that seek to compute or bound the probability of failure. Unfortunately, exact evaluation of this probabilistic risk is challenging and computationally intractable for generic nonlinear systems. While Monte-Carlo (MC) estimation techniques [5, 3, 9] provide a general and powerful workaround, they are still computationally-demanding and difficult to embed within a continuous motion planner. A number of “directly-computable” risk approximations have been proposed [12, 2, 20, 19], but all fundamentally assume the system evolves in discrete time. Many systems we care to control in practice evolve continuously, and while ap-




Fig. 1. Consider a planar system maneuvering from the green initial distribution to the red goal position in the presence of a rectangular obstacle. Shaded ellipses represent the state distribution at a set of discrete timesteps, $\mathcal{T} = \{t_0, t_1, \dots, t_k\}$, and reflect uncertainty under closed-loop execution. A common technique for computing the total collision probability leverages Boole’s inequality to simply sum the violation probabilities at each step (magenta). However, the resulting estimate will be sensitive to the choice of \mathcal{T} – too coarse, and it will underestimate; too fine, and it will “double-count” probability mass corresponding to trajectories that remain in collision across multiple timesteps. This leads to over-conservatism, artificial infeasibility, and ultimately brittle planning. In this paper we introduce a *continuous-time* approximation (cyan) that instead aggregates probability mass *entering* collision over each *interval*. As \mathcal{T} is refined, our result actually improves, eventually converging to the true failure probability.

plication of discrete-time methods is possible in these settings, the ensuing risk estimates will be highly sensitive to the chosen time discretization. As recognized by Ariu et al. [1], Janson et al. [9], and others, they may be either too lax (allowing “corner-cutting”), too conservative (leading to severe sub-optimality or artificial infeasibility), or both simultaneously.

This paper addresses this problem at its source, taking a rigorous look at the evolution of failure probability directly in continuous time. Our approximation is general, applying to partially-observable¹ stochastic systems $\mathbf{x}(t)$ in $X = \mathbb{R}^{n_x}$ with nonlinear Itô dynamics

$$d\mathbf{x}(t) = \mathbf{f}(t, \mathbf{x}(t), \mathbf{u}(t)) dt + G(t, \mathbf{x}(t), \mathbf{u}(t)) d\mathbf{w}(t) \quad (1)$$

where $\mathbf{u}(t)$ and $\mathbf{w}(t)$ represent a vector-valued control process

¹ Note that under partial-observability, the process $\mathbf{x}(t)$ can only be considered Markov if it is “augmented” with the internal state of the estimator/controller. For the purposes of this paper we leave this implicit – because constraints are assumed to involve only on the state of the physical plant, augmentation will contribute no analytic or computational cost.

and Brownian motion, respectively. Under partial observability, $\mathbf{x}(t)$ is not directly available to the controller, and feedback must be accomplished via a parallel observation process represented by the filtration \mathcal{Y}_t (the information available at time t). The special case of fully-observable systems is easily captured here as well. In order to simplify presentation, our main results impose some additional restrictions between the dynamics (1) and the constraints that define the *safe set* $X_{\text{safe}} \subset X$. Even so, they accommodate a wide variety of nonlinear systems and constraints, including the ubiquitous case of second-order systems with position constraints. Moreover, relaxation of some of these requirements will be straightforward, though beyond the scope of this paper.

The types of planning problems we address take the form

$$\min_{\mathbf{u}(\cdot)} \quad \mathbb{E} \int_0^T l(t, \mathbf{u}(t), \mathbf{x}(t)) dt \quad (2)$$

$$\text{subject to: } P\left(\bigvee_{t \in [0, T]} \mathbf{x}(t) \notin X_{\text{safe}}\right) \leq \Delta \quad (3)$$

where the \bigvee symbol is a logical OR, implying existence of a satisfying event among a (possibly uncountable) collection, and $\Delta \in (0, 1)$ represents a given risk tolerance over the finite horizon $[0, T]$. The constraint (3) upper-bounds the probability of failure at *any* time in the planning horizon. This *risk* is challenging to evaluate and optimize against because it couples states across the planning horizon as a whole. As pointed out by Ono et al. [19] and others, joint constraints of this nature can be approached via Lagrangian relaxation; that is, converting the *risk-constrained* problem to a *risk-minimizing* one with objective

$$\mathbb{E} \int_0^T l(t, \mathbf{u}(t), \mathbf{x}(t)) dt + \lambda \left[P\left(\bigvee_{t \in [0, T]} \mathbf{x}(t) \notin X_{\text{safe}}\right) - \Delta \right] \quad (4)$$

for some $\lambda \geq 0$. However, the augmented objective in (4) does not possess the time-additive *Bellman* structure of (2), precluding the application of many optimal control techniques such as dynamic programming [22, 24, 19]. Like [2, 19] and other direct approximations, the approximation presented in this paper restores this time-additive structure while achieving superior accuracy for continuous-time systems.

This paper takes a fresh look at the time-evolution of the failure probability in (3), specifically in continuous-time. Related work is outlined in Section II. Section III leverages the language of first-exit times to produce a time-additive framework for continuous-time risk estimation. The classic theory is rich and well-explored, but to the best of the authors’ knowledge no techniques yet exist to enable computation in the context of generic nonlinear systems and nonlinear constraints. To address this, we propose a piecewise-continuous approximation in Sections IV and V that provably converges as the time discretization is refined – this allows us to losslessly extend classic results for constant-coefficient systems. As a second contribution, we identify a lightweight method to account for “safety-thus-far” that avoids attempting to approximate an explicit posterior and ensures conservatism

without being excessively so. Finally, we identify a large class of systems for which the requisite numerical quadratures can be computed exactly at significantly-reduced dimension, which is critical to ensure computational feasibility. The resulting risk approximation is empirically demonstrated in Section VI to well-approximate MC estimates while retaining the computational simplicity and general applicability of direct methods, paving the way for risk-aware, continuous motion planning onboard real-world systems.

II. RELATED WORK

Risk-aware motion planning is far from a new problem and has received much attention over the years. For example, when probabilistic constraints $P(\mathbf{x}(t) \notin X_{\text{safe}}) \leq \Delta$ are enforced *independently* for each t , suitable extensions of classic algorithms such as RRT [14, 4, 23] and differential dynamic programming (DDP) [24] have been proposed. Alternatively, for problems with discrete time and action spaces, risk-constrained search methods such as RAO* [21, 8] have shown promise.

This paper considers problems continuous in time and action, where safety is most naturally expressed by the joint constraint (3). The field of robust control provides some solutions against *worst-case* (i.e., bounded) disturbances, for example Majumdar and Tedrake [16] and Lopez et al. [13]. However, these approaches are generally restricted to fully-observable systems. Furthermore, their performance may be severely conservative in the average case, motivating a quantification of *probabilistic* risk to allow explicit trade-off between safety and expedience during planning.

A. Sampling-Based Methods

MC techniques provide a general and powerful method for estimating failure probabilities, at the cost of having to run a potentially large number of simulation rollouts. As closed-loop partially-observable systems require simulation of plant, estimator, and controller, this can represent non-negligible amounts of computation. Aside from the issue of sample-complexity, which has been partially addressed by Calafiore and Campi [5] and Janson et al. [9], sample-based estimates are discontinuous and therefore fundamentally cumbersome to incorporate into a continuous optimization framework. For example, [5] proposes enforcement of a deterministic constraint for each sample, [9] resorts to iterative obstacle inflation in RRT, and Blackmore et al. [3] apply Mixed-Integer techniques. This limitation of MC motivates the search for efficient, “direct” risk approximations amenable to online, *continuous* optimization over motion plans.

B. Direct Risk Estimates (in Discrete Time)

A number of direct techniques have been developed in a discrete-time setting, where (3) simplifies to

$$P\left(\bigvee_{i=0}^k \mathbf{x}(t_i) \notin X_{\text{safe}}\right) \leq \Delta. \quad (5)$$

Early work by Li et al. [12] recognized that, under Linear-Time-Varying (LTV) dynamics, dispersions will be distributed

as an $n_x(k+1)$ -dimensioned normal distribution. The joint probability in (5) can then be evaluated as a high-dimensional integral via quadrature methods or sampling, both of which have complexity exponential in dimension. To avoid this unfavorable scaling with k , Blackmore and Ono [2] use Boole’s inequality (a.k.a. the union bound) to decompose the probability in (5) over time as

$$P\left(\bigvee_{i=0}^k \mathbf{x}(t_i) \notin X_{\text{safe}}\right) \leq \sum_{i=0}^k P(\mathbf{x}(t_i) \notin X_{\text{safe}}). \quad (6)$$

This decoupling over timesteps is highly convenient. As pointed out by Ono et al. [19], the right-hand side of (6) is time-additive, allowing the use of dynamic programming to minimize (4). Alternatively, Ono and Williams [18] introduce the *risk-allocation* formulation

$$P(\mathbf{x}(t_i) \notin X_{\text{safe}}) \leq \Delta_i \quad \forall i \leq k \quad \text{and} \quad \sum_{i=0}^k \Delta_i = \Delta \quad (7)$$

which explicitly allocates a risk “budget” between timesteps. This framework has inspired a series of other works including [25, 15, 6], all fundamentally dependent on decomposition (6).

Despite its popularity, the use of Boole’s inequality to decompose joint constraints over time can be severely conservative, as illustrated in Fig. 1. By ignoring correlations between states at adjacent timesteps, the sum in (6) will “double-count” violation probabilities, particularly as the time discretization is refined. This problem is addressed by Patil et al. [20], who recognize that the state distributions should be *conditioned* on the safety of prior timesteps. Because capturing this conditioning exactly is challenging, they approximate the projected and truncated distribution corresponding to each constraint as a single-dimensional Gaussian, allowing closed-form update of the state distribution parameters. However, Gaussianity here is inexact and the resulting estimate may not remain statistically consistent or result in a conservative (upper-bounding) risk estimate.

When addressing continuous-time systems, a standard numeric approach to risk-aware optimization approximates the original system as discrete-time under a given time partition $\mathcal{T} = \{t_0 = 0, t_1, \dots, t_k = T\}$ and then substitutes constraint (5) for (3). This allows well-studied discrete-time techniques to be applied, but, as mentioned earlier, the results will be highly sensitive to the choice of discretization. Depending on the complexity and movement speed of the robot, these effects can lead to unsafe or overly-cautious behavior, or both. This motivates the fundamental contribution of this paper – a direct risk approximation derived fundamentally in continuous time.

C. Stochastic Processes in Continuous-Time

In an effort to explicitly address the continuous-time constraint (3) and avoid “corner-cutting,” Ariu et al. [1] apply Boole’s inequality over *intervals* rather than instantaneous states as

$$P\left(\bigvee_{s \in [0, T]} \mathbf{x}(s) \notin X_{\text{safe}}\right) \leq \sum_{i=0}^{k-1} P\left(\bigvee_{s \in [t_i, t_{i+1})} \mathbf{x}(s) \notin X_{\text{safe}}\right). \quad (8)$$

Note that continuity allows the endpoint $\mathbf{x}(T)$ to be dropped. As with (6), this “interval Boole’s” neglects correlations between events. In particular, their method assumes the state process itself follows a Brownian motion, ultimately resulting in a *doubly-conservative* estimate compared to its discrete-time counterpart. In contrast, the approximation proposed in this paper also operates on intervals, but it addresses the nonlinear dynamics directly and avoids the use of Boole’s inequality in this way, producing a much tighter risk bound.

This paper leans heavily on the classical notion of first-passage times, a well-studied topic in the field of continuous stochastic processes. Indeed, the piecewise approximation presented in Section IV of this paper is similar in spirit to that of Jin and Wang [10], although we consider the case where the process (rather than purely the boundary) is nonlinear.

III. HOW RISK EVOLVES IN TIME

Consider a càdlàg², Markov process $\mathbf{z}(t) \in \mathbb{R}^m$ with associated probability space (Ω, \mathcal{F}, P) . Here, Ω represents the *outcome* space, \mathcal{F} is a *sigma-algebra* over Ω (a collection of events $E \subseteq \Omega$) such that $\mathbf{z}(t)$ is measurable for all t , and P is a probability measure over \mathcal{F} . Note that for now we do *not* assume $\mathbf{z}(t)$ has continuous sample paths. We begin by exploring the “nearly” time-additive evolution of the *exit cumulant*

$$F_{\mathbf{z}}(t) \triangleq P\left(\bigvee_{s \in [0, t]} \mathbf{z}(s) \notin D\right) \quad (9)$$

for some closed set $D \subset \mathbb{R}^m$.

For clarity in the following discussion, we adopt the language of *passage times*, also known as *hitting* or *exit times* [11, 26]. Define a parameterized family of exit times with respect to process $\mathbf{z}(t)$ as

$$z_{\tau_t} \triangleq \inf \{s \in [t, T] \mid \mathbf{z}(s) \notin D\} \quad (10)$$

where the infimum of the empty set is assigned to ∞ . That is, z_{τ_t} refers to the first time (after t) that $\mathbf{z}(s)$ “exits” D . For both $F_{\mathbf{z}}$ and z_{τ_t} we will drop the explicit \mathbf{z} specification when the corresponding process is unambiguous.

The following useful properties can be identified.

Lemma 1. *For any $0 \leq t_1 < t_2 \leq T$, the following hold:*

- $\tau_{t_1} \geq t_2 \iff \mathbf{z}(s) \in D \quad \forall s \in [t_1, t_2)$.
- $\tau_0 \in [t_1, t_2) \iff \tau_0 \geq t_1, \tau_{t_1} < t_2$.
- For each outcome $\omega \in \Omega$, either
 - 1) $\tau_0 = \infty$ (“no exit”), or
 - 2) $\mathbf{z}(\tau_0) \in \partial D \cup D^c$ (“smooth exit” or a “jump out”).
 where ∂D and D^c are the boundary and complement of D , respectively.

It is straightforward to verify that the total exit probability can be written as $F(T) = P(\tau_0 \leq T)$, and thus τ_0 offers a means by which to analyze the time-evolution of $F(t)$. In some cases $F(t)$ is known to be time-differentiable [10], and thus computing this derivative (called the *first-passage density*)

²a.k.a. right-continuous with left limits.

would seem to be a natural goal. However, this computation is challenging for generic nonlinear processes, and instead we will settle for an interval-based “integration” scheme that reflects how $F(T)$ can be approximated in practice. In later sections, through both analysis and experiment we demonstrate that this approximation indeed converges to the true value as the time discretization is refined.

Proceeding, assume a given partition $\mathcal{T} = \{t_0 = 0, t_1, t_2, \dots, t_k = T\}$ of the fixed horizon $[0, T]$. A crucial advantage of the language of the *first*-exit time is that it provides a natural disjointness between events, and in particular

$$F(T) = \sum_{i=0}^{k-1} P(\tau_0 \in [t_i, t_{i+1})) + P(\tau_0 = T). \quad (11)$$

Note that, in contrast to (8), the relation (11) holds with equality. Proceeding from here and adopting the shorthand $\mathbf{z}_i \triangleq \mathbf{z}(t_i)$, $\tau_i \triangleq \tau_{t_i}$, and so on, the above can be written

$$\sum_{i=0}^{k-1} \left(P(\tau_0 \in [t_i, t_{i+1}), \mathbf{z}_i \in D) + P(\tau_0 = t_i, \mathbf{z}_i \notin D) \right) + P(\tau_0 = T, \mathbf{z}(T) \in D^c) \quad (12)$$

$$= \sum_{i=0}^{k-1} P(\tau_0 \in [t_i, t_{i+1}), \mathbf{z}_i \in D) + \sum_{i=0}^k P(\tau_0 = t_i, \mathbf{z}_i \in D^c) \quad (13)$$

where in (12) we split the probability over the event that $\mathbf{z}_i \in D$. Note that the right-hand summation in (13) involves the probabilities that discontinuous sample paths “jump out” of D at each of the partition points t_i . Examining the first set of terms, Lemma 1 allows

$$P(\tau_0 \in [t_i, t_{i+1}), \mathbf{z}_i \in D) = P(\tau_i < t_{i+1} \mid \tau_0 \geq t_i, \mathbf{z}_i \in D) P(\tau_0 \geq t_i, \mathbf{z}_i \in D). \quad (14)$$

The conditioning on $\tau_0 \geq t_i$ and $\mathbf{z}_i \in D$ in (14) implies that the process has not exited “yet,” imposing a specific posterior over \mathbf{z}_i that we will call the *anthropic* distribution³

$$\bar{\pi}_t(d\mathbf{z}) \triangleq P(\mathbf{z}(t) \in d\mathbf{z} \mid \tau_0 \geq t, \mathbf{z}(t) \in D). \quad (15)$$

Before proceeding further, define the function

$$\Phi_{\mathbf{z}}(t_i, t_{i+1}; \mu_i) \triangleq P(\tau_i < t_{i+1} \mid \mathbf{z}_i \sim \mu_i) \quad (16)$$

for any (not necessarily normalized) measure μ_i over \mathbb{R}^m . (16) captures the probability of an exit (not necessarily the first) in the interval $[t_i, t_{i+1})$, given that \mathbf{z}_i is distributed according to μ_i . As before, the \mathbf{z} subscript will be left implicit where possible. Because the process $\mathbf{z}(t)$ is assumed Markov and applying (14), sum (13) can be re-written in terms of Φ and

³ In cosmology, the *anthropic principle* remarks that life can only observe universes that themselves allow for the existence of life.

$\bar{\pi}_t$ as

$$F(T) = \sum_{i=0}^{k-1} \Phi(t_i, t_{i+1}; \bar{\pi}_i) P(\tau_0 \geq t_i, \mathbf{z}_i \in D) + \sum_{i=0}^k P(\tau_0 = t_i, \mathbf{z}_i \in D^c). \quad (17)$$

A key challenge in evaluation of (17) is computation of $\bar{\pi}_i$. For one thing, its support is clearly limited to D , which is sufficient to ensure non-Gaussianity. This motivates us to avoid attempting to approximate or bound $\bar{\pi}_t$ directly and instead decompose it via Baye’s rule, producing

$$\bar{\pi}_t(d\mathbf{z}) = \frac{P(\tau_0 \geq t, \mathbf{z}(t) \in D \mid \mathbf{z}(t) = \mathbf{z}) P(\mathbf{z}(t) \in d\mathbf{z})}{P(\tau_0 \geq t, \mathbf{z}(t) \in D)} \triangleq \frac{\bar{\Psi}_t(\mathbf{z}) \bar{b}_t(d\mathbf{z})}{P(\tau_0 \geq t, \mathbf{z}(t) \in D)}. \quad (18)$$

where \bar{b}_t is the *a priori* distribution of $\mathbf{z}(t)$, and we refer to $\bar{\Psi}_t : \mathbb{R}^m \mapsto [0, 1]$ as the *anthropic likelihood*. It is straightforward to show

$$\Phi(t_i, t_{i+1}; \bar{\pi}_i) P(\tau_0 \geq t_i, \mathbf{z}_i \in D) = \Phi(t_i, t_{i+1}; \bar{\Psi}_i \bar{b}_i), \quad (19)$$

and therefore (17) can be written

$$F(T) = \sum_{i=0}^{k-1} \Phi(t_i, t_{i+1}; \bar{\Psi}_i \bar{b}_i) + \sum_{i=0}^k P(\tau_0 = t_i, \mathbf{z}_i \in D^c). \quad (20)$$

As discussed in Section V-B, and in contrast to the case of $\bar{\pi}_t$, identifying conservative approximations for $\bar{\Psi}_t$ will be both straightforward and effective.

A. Analogous Development Under Filtration \mathcal{F}_t

Though not the primary focus of this paper, we note that the preceding development can be applied analogously in the context on a filtration \mathcal{F}_t representing information that becomes available *during* execution as opposed to *a priori*. For example, in the context of $\mathbf{x}(t)$ as defined in (1), we might consider the observation filtration $\mathcal{F}_t = \mathcal{Y}_t$.

Though the information generated by \mathcal{F}_t itself evolves randomly, modification of the above discussion is straightforward and we simply provide some analogous definitions here for clarity. As in (15), we can define the *anthropic belief* $\pi_t(d\mathbf{z}) \triangleq P(\mathbf{z}(t) \in d\mathbf{z} \mid \tau_0 \geq t, \mathbf{z}(t) \in D, \mathcal{F}_t)$, and like (15) it has the structure

$$\pi_t(d\mathbf{z}) = \frac{\Psi_t(\mathbf{z}) b_t(d\mathbf{z})}{P(\tau_0 \geq t, \mathbf{z}(t) \in D \mid \mathcal{F}_t)} \quad (21)$$

where the familiar estimation belief b_t and anthropic likelihood Ψ_t are adapted to \mathcal{F}_t . An analogous identity to (19) can be established, yielding

$$F(T) = \mathbb{E} \left[\sum_{i=0}^{k-1} \Phi(t_i, t_{i+1}; \Psi_i b_i) \right] + \sum_{i=0}^k P(\tau_0 = t_i, \mathbf{z}_i \in D^c). \quad (22)$$

where the expectation is taken over \mathcal{F}_t .

IV. A PIECEWISE-CONTINUOUS APPROXIMATION

The preceding section introduced a framework for computing the failure probability based on a helper function Φ defined in (16). Rather than attempting to evaluate $\Phi_{\mathbf{x}}$ directly for the nonlinear process $\mathbf{x}(t)$, this section introduces an approximating process *in the constraint space*, for which computation is made tractable. This simplified process is piecewise-continuous according to the given time discretization \mathcal{T} , which controls the “accuracy” of the approximation. Though \mathcal{T} in practice will likely be dictated by the computational resources available, we prove that as it is refined the corresponding failure probability estimate converges to $F_{\mathbf{x}}(T)$. That is to say, the approximation is asymptotically lossless.

Assume the feasible set $X_{\text{safe}} \subset X$ is defined by the sublevel sets $g_j(\mathbf{x}) \leq 0$ for a set of m twice-differentiable, real-valued functions $\{g_j\}$. Let $\mathbf{g}(\mathbf{x})$ refer to the stacked vector in \mathbb{R}^m , and let $\mathbf{a}_j(\mathbf{x})$ and $H_j(\mathbf{x})$ refer to the gradient vector and Hessian of each g_j , respectively. Note that $\mathbf{x} \in X_{\text{safe}}$ is equivalent to the statement that $g_j(\mathbf{x}) \leq 0$ for all j . This naturally motivates a consideration of $\mathbf{y}(t) \triangleq \mathbf{g}(\mathbf{x}(t))$ as a process with respect to the non-positive orthant $\mathcal{O}^- \subset \mathbb{R}^m$.

Before we proceed, however, computation and analysis will require some regularity conditions on the control process $\mathbf{u}(t)$, which evolves according to \mathcal{Y}_t and some (potentially non-deterministic) policy. Rather than complicate the discussion by attempting to account for all such possibilities, we make the following simplifying assumption.

Assumption 1 (*$\mathbf{y}(t)$ Locally Independent of Control*). *We require that either*

- 1) *we have state-feedback $\mathbf{u}(t) = \kappa(t, \mathbf{x}(t))$, with κ deterministic and Lipschitz, or*
- 2) *$\mathbf{a}_j^T(\mathbf{x})\mathbf{f}(t, \mathbf{x}, \mathbf{u})$ and $\mathbf{a}_j^T(\mathbf{x})G(t, \mathbf{x}, \mathbf{u})$ are independent of \mathbf{u} for all constraints j .*

Assumption 1 ensures that the dynamics of process $\mathbf{y}(t)$ depend on the control only via the state $\mathbf{x}(t)$, and notably accommodates second-order systems under workspace constraints (i.e., obstacle avoidance). Partial relaxation of this assumption will in many cases require only minor modifications of the analysis and computations presented in this paper.

Proceeding, we consider a piecewise approximation of this process constructed from the time discretization \mathcal{T} . Using the shortcuts $\mathbf{f}_t \triangleq \mathbf{f}(t, \mathbf{x}(t), \mathbf{u}(t))$, $G_t \triangleq G(t, \mathbf{x}(t), \mathbf{u}(t))$, and so on, define the process $\hat{\mathbf{y}}(s)$ such that

$$\hat{\mathbf{y}}(s) \triangleq \mathbf{g}(\mathbf{x}(t)) + \int_t^s \mathbf{h}_t ds' + \int_t^s \Sigma_t d\mathbf{w}(s') \quad (23)$$

$$\mathbf{h}_{t,j} \triangleq \mathbf{a}_j^T(\mathbf{x}(t))\mathbf{f}_t + \frac{1}{2}\text{tr}\left(G_t^T H_j(\mathbf{x}(t))G_t\right) \quad (24)$$

$$\Sigma_t \triangleq \begin{bmatrix} \mathbf{a}_1^T(\mathbf{x}(t)) \\ \vdots \\ \mathbf{a}_m^T(\mathbf{x}(t)) \end{bmatrix} G_t \triangleq \begin{bmatrix} \boldsymbol{\sigma}_1(t, \mathbf{x}(t)) \\ \vdots \\ \boldsymbol{\sigma}_m(t, \mathbf{x}(t)) \end{bmatrix} \quad (25)$$

where $t = \max\{t_i \in \mathcal{T} \mid t_i \leq s\}$ corresponds to the start point of the interval $[t_i, t_{i+1})$ containing s . Here \mathbf{h}_s and Σ_s can be identified as the drift and diffusion coefficients of




Fig. 2. An illustration of the continuous process $y_j(t) = g_j(\mathbf{x}(t))$ and our approximation $\hat{y}_j(t)$. The two processes coincide at each discrete timestep $t_i \in \mathcal{T}$, and as \mathcal{T} is refined, Prop. 1 establishes that the two converge pathwise-uniformly. Zero-crossings of $y_j(t)$ imply constraint violations, and therefore exit-times of $\hat{\mathbf{y}}(t)$ from the non-positive orthant approximate those of $\mathbf{x}(t)$ from X_{safe} .

$\mathbf{y}(s)$ as prescribed by Itô’s formula [26, Thm. 5.5]. Thus, the piecewise-continuous $\hat{\mathbf{y}}(s)$ has constant coefficients over each interval and coincides with $\mathbf{y}(s)$ for every $s = t \in \mathcal{T}$.

Approximating $F_{\hat{\mathbf{y}}}$ will be much easier than working directly with $F_{\mathbf{x}}$, and in the following we show that, under mild conditions, the substitution of $\hat{\mathbf{y}}(s)$ for $\mathbf{y}(s)$ can be considered asymptotically lossless.

A. Convergence Analysis

Consider evaluating $F_{\hat{\mathbf{y}}}(T)$ for the process $\hat{\mathbf{y}}(t)$ – the result will, of course, depend on $\hat{\mathbf{y}}(t)$ and therefore the underlying time discretization \mathcal{T} . We show that as \mathcal{T} is refined, $\hat{\mathbf{y}}(t)$ converges to $\mathbf{y}(t)$ in a specific sense, and that $F_{\hat{\mathbf{y}}}(T)$ converges to $F_{\mathbf{y}}(T) = F_{\mathbf{x}}(T)$. Proofs are left to the supplementary material [7, Appendix A].

Definition 1 (Pathwise-Uniform Convergence). *For a sequence of processes $\mathbf{z}^{(n)}(t)$ and an associated real-valued sequence $\varepsilon_n \rightarrow 0$, we say $\mathbf{z}^{(n)}(t)$ converges ε_n -pathwise-uniformly to process $\mathbf{z}(t)$ if there exists for each $\omega \in \Omega$ an $0 < M_\omega < \infty$ almost surely such that*

$$P\left(\sup_{s \in [0, T]} \left\| \mathbf{z}^{(n)}(s) - \mathbf{z}(s) \right\|_2 > \varepsilon_n M_\omega\right) \rightarrow 0.$$

Note that pathwise-uniform convergence (PUC) does not imply that $\mathbf{z}^{(n)}(\cdot; \omega)$ converges to $\mathbf{z}(\cdot; \omega)$ for any fixed outcome $\omega \in \Omega$. Additionally, it is specific to the given sequence ε_n . Nonetheless, we will see that this condition is sufficient to ensure a non-trivial convergence-of-probabilities result.

Definition 2 (Supremum Process). *For a process $\mathbf{z}(t)$, the corresponding supremum process is $\mathbf{z}^*(t) \triangleq \sup_{s \in [0, t]} \mathbf{z}(s)$.*

Theorem 1. *If a given sequence of \mathbb{R}^m -valued processes $\mathbf{z}^{(n)}(t)$ converges ε_n -pathwise-uniformly to $\mathbf{z}(t)$, and we have that $P(\mathbf{z}_j^*(T) = 0) = 0$ for all j , then*

$$P\left(\bigvee_{s \in [0, T]} \mathbf{z}^{(n)}(s) \notin \mathcal{O}^-\right) \rightarrow P\left(\bigvee_{s \in [0, T]} \mathbf{z}(s) \notin \mathcal{O}^-\right).$$

For technical reasons, we will be required to make the following additional assumption relating the diffusion matrix G and the constraint gradients $\{\mathbf{a}_j\}$.

Assumption 2 (Paired Lipschitz:). *With respect to a given function $g : X \mapsto \mathbb{R}$ and diffusion matrix $G(t, \mathbf{x}, \mathbf{u})$, either*

- 1) G is constant and the function gradient, \mathbf{a} , is Lipschitz,
- 2) \mathbf{a} is constant (i.e., g is linear) and G is Lipschitz, or
- 3) both \mathbf{a} and G are bounded and Lipschitz.

Remark 1. The requirements that Assumption 2 imposes on G in no way imply that disturbances are bounded. Furthermore, when constraints are linear G is only required to be Lipschitz, so cases of multiplicative noise can be handled here as well.

This leads to our main result.

Proposition 1. Let Assumptions 1 and 2 hold for the m constraint functions $\{g_j\}$. Additionally, assume:

- 1) \mathbf{f}_t is pathwise-bounded a.s., and $\mathbb{E} \int_0^T \|\mathbf{f}_t\|_2^2 dt < \infty$.
- 2) $\mathbb{E} \|G_t\|_2^2$ bounded over $[0, T]$, and $\mathbb{E} \int_0^T \|G_t\|_2^2 dt < \infty$.

Then, for a sequence of partitions $\mathcal{T}_n = \{0 = t_0, t_1, \dots, t_{k_n} = T\}$ over the compact interval $[0, T]$ with mesh $\delta_n \rightarrow 0$ and $k_n \delta_n \leq cT$ for fixed $c > 0$, there exists a sequence $\varepsilon_n \rightarrow 0$ such that the sequence of processes $\hat{\mathbf{y}}^{(n)}(s)$ converges ε_n -pathwise-uniformly to $\mathbf{y}(s)$.

From Thm. 1 and the definition of $\mathbf{y}(t)$ the following corollary immediately follows.

Corollary 1. If the assumptions of Prop. 1 hold and $P(\mathbf{y}^*(T) = 0) = 0$ for all j then $F_{\hat{\mathbf{y}}^{(n)}}(T) \rightarrow F_{\mathbf{x}}(T)$.

Remark 2. The condition $P(\mathbf{y}_j^*(T) = 0) = 0$ is somewhat difficult to verify; while conditions over which the variable $y_j(t)$ admits a density have been somewhat examined, the case of the supremum process is less explored. For the purposes of the paper we consider this, as well as the regularity conditions required by Prop. 1, to be fairly mild.

Corollary 1 makes clear that the substitution of $\hat{\mathbf{y}}(s)$ for $\mathbf{y}(s)$ is (asymptotically) loss-less for the purposes of computing the risk probability $F_{\mathbf{x}}(T)$. It is worth noting that this result does not anywhere assume that the dynamics or observation model are Gaussian or LTV, although practical computation of b_t may still require this assumption.

V. COMPUTATION

The results of Section IV indicate that we can perform computation with the simplified process $\hat{\mathbf{y}}(t)$ and the result will approximate the true $F_{\mathbf{x}}(T)$ in an asymptotic sense. While $\hat{\mathbf{y}}(t)$ is only piecewise-continuous and thus ‘‘jump out’’ probabilities in (22) may be non-zero, our empirical results indicate that they can be safely ignored. Identifying a rigorous proof is left as a topic for future work. Additionally, for the sake of presentation we make the simplifying assumption that $P(\mathbf{x}_0 \notin X_{\text{safe}}) = 0$. In light of these considerations, we focus on computing

$$\hat{F}(T) \triangleq \sum_{i=0}^{k-1} \Phi_{\hat{\mathbf{y}}}(t_i, t_{i+1}; \bar{\Psi}_i \bar{b}_i). \quad (26)$$

In order to compute (26), we require both a means of computing the quantity $\Phi_{\hat{\mathbf{y}}}(t_i, t_{i+1}; \mu_i)$ for arbitrary measure μ_i and a means of approximating $\bar{\Psi}_i \bar{b}_i$. The former is the subject of V-A and the latter of V-B.

A. Bounding $\Phi_{\hat{\mathbf{y}}}$

The constant-coefficient approximation $\hat{\mathbf{y}}(t)$ allows us to leverage classic results in the study of first-exit times. First, recalling (16) we apply the union bound to decouple individual constraints, producing a sum over element-wise diffusions \hat{y}_j

$$\Phi_{\hat{\mathbf{y}}}(t_i, t_{i+1}; \mu_i) \leq \sum_{j=1}^m \Phi_{\hat{y}_j}(t_i, t_{i+1}; \mu_i). \quad (27)$$

Letting ∂X_{safe} refer to the boundary of X_{safe} , a classic result (see Karatzas and Shreve [11, Ch. 3, Eq. 5.12]) provides the following lemma.

Lemma 2. Under Assumption 1, for any (not-necessarily normalized) measure μ_i over X_{safe} , and such that $\mu_i(\partial X_{\text{safe}}) = 0$,

$$\begin{aligned} \Phi_{\hat{y}_j}(t_i, t_{i+1}; \mu_i) &= P\left(\sup_{s \in [t_i, t_{i+1}]} \hat{y}_j(s) > 0 \mid \hat{y}_j(t_i) \sim g_j(\mu_i)\right) \\ &= \int_{X_{\text{safe}}} \psi\left(g_j(\mathbf{x}), h_j(t_i, \mathbf{x}), \|\sigma_j(t_i, \mathbf{x})\|_2, \Delta_i\right) d\mu_i(\mathbf{x}) \end{aligned} \quad (28)$$

where $\Delta_i = t_{i+1} - t_i$ and we define $\psi(z, h, \sigma, \Delta t) \triangleq$

$$1 - \phi\left(\frac{-h\Delta t - z}{\sigma\sqrt{\Delta t}}\right) + \exp\left\{-\frac{2hz}{\sigma^2}\right\} \phi\left(\frac{-h\Delta t + z}{\sigma\sqrt{\Delta t}}\right) \quad (29)$$

and ϕ refers to the CDF of the standard normal distribution.

Note that the condition $\mu_i(\partial X_{\text{safe}}) = 0$ is satisfied for all measures which admit a density (a mild assumption).

Like the instantaneous probability $P(\mathbf{x}(t_i) \notin X_{\text{safe}})$, computing (28) in general requires evaluating a multi-dimensional integral over the state space $X = \mathbb{R}^{n_x}$. Numerical quadrature techniques have exponential complexity in dimension (the curse of dimensionality). Fortunately, depending on the structure of the dynamics and constraints, significant dimensionality reductions can often be found. In particular, consider Gaussian-distributed second-order systems with workspace constraints (such as obstacle avoidance). If system noise is injected only in the velocity and acceleration dynamics (corresponding to $\sigma_j = 0$ for all j), then the integral in (28) reduces to workspace dimension n_p (commonly 2 or 3). This ensures that $\Phi_{\hat{\mathbf{y}}}$ can be evaluated efficiently for a large class of dynamics. Details are provided in the supplementary material [7, Appendix B].

B. Approximating Anthropoc Belief

Given we can compute (27), we still require some means of estimating the anthropic state distribution via $\bar{\pi}_t$ or $\bar{\Psi}_t \bar{b}_t$. Though inexact, Gaussian approximations for \bar{b}_i are applicable and widely-used whenever the state dynamics are sufficiently smooth and, as can be seen in our experimental results, accurately model the dispersion of even nonlinear MC trajectories. However, modeling the anthropic information represented by $\bar{\Psi}_i$ or $\bar{\pi}_i$ is not so straightforward.

1) *Approximating $\bar{\pi}_i$ as a Gaussian*: In existing work, Patil et al. [20] propose a Gaussian approximation for the anthropic distribution $\bar{\pi}_i$ (or analogously, anthropic belief π_i) at each timestep t_i . This is attractive because it implies $\bar{\pi}_i$ can be propagated in parallel with \bar{b}_i . Likewise, we can consider approximating $\bar{\pi}_i$ by conditioning on safety (via a recursive Gaussian approximation) at each prior timestep

$$\hat{\bar{\pi}}_i(d\mathbf{x}) \approx P(\mathbf{x}_i \in d\mathbf{x} \mid \mathbf{x}_l \in X_{\text{safe}} \forall l \leq i). \quad (30)$$

We refer the interested reader to their paper for implementation details.

Again considering (26), we can apply identity (19) and directly produce an upper-bound in terms of $\bar{\pi}_i$

$$\hat{F}(T) \leq \sum_{i=0}^{k-1} \Phi_{\hat{\mathbf{y}}}(t_i, t_{i+1}; \bar{\pi}_i). \quad (31)$$

Together with our Gaussian approximation $\hat{\bar{\pi}}_i$ from (30), we refer to the resulting estimate as `ival_gauss`.

While the independence and Gaussianity assumptions involved here are convenient, they are also heuristic. This can lead to inconsistent $\bar{\pi}_t$ estimates and ultimately unpredictable risk estimation, as will be seen in Section VI.

2) *A Simple Alternative*: Rather than attempting to estimate $\bar{\pi}_t$, we propose a simple, yet surprisingly powerful alternative based on the anthropic likelihood $\bar{\Psi}_t$. A clear upper-bound follows directly from the definition as

$$\bar{\Psi}_t(\mathbf{x}) = P(\mathbf{x}\tau_0 \geq t, \mathbf{x}(t) \in X_{\text{safe}} \mid \mathbf{x}(t) = \mathbf{x}) \leq \mathbb{1}_{X_{\text{safe}}}(\mathbf{x}). \quad (32)$$

Applying (32) to (26) yields a conservative estimate

$$\hat{F}(T) \leq \sum_{i=0}^{k-1} \Phi_{\hat{\mathbf{y}}}(t_i, t_{i+1}; \mathbb{1}_{X_{\text{safe}}} \bar{b}_i) \quad (33)$$

which accumulates “new” exits over each interval – note the contrast with the naive “interval Boole’s” approximation given by (8). Because it restricts exit flow to probability mass which is “safe” at the start of each interval, we refer to approximation (33) as `ival_safe`.

In principle, (33) shares one of the same weaknesses of the discrete-time Boole’s approximation in (6) – it is not bounded above, and indeed may diverge as \mathcal{T} is refined. However, it avoids “locally” double-counting by counting only probability mass *leaving* X_{safe} over each interval. This interpretation is illustrated in Fig. 1, and our empirical results suggest `ival_safe` does very well in practice. Furthermore, compared to [20] or (31) it requires no “extra” belief propagation at all, making this approximation particularly lightweight.

VI. RESULTS

To evaluate the accuracy of our approximation, we simulate a second-order Dubin’s car system in the plane. System details are provided in the supplementary material [7, Appendix C]. Noise in dynamics and observations render the system partially-observable, and an LQG-style feedback control policy stabilizes around a nominal trajectory. We evaluate four

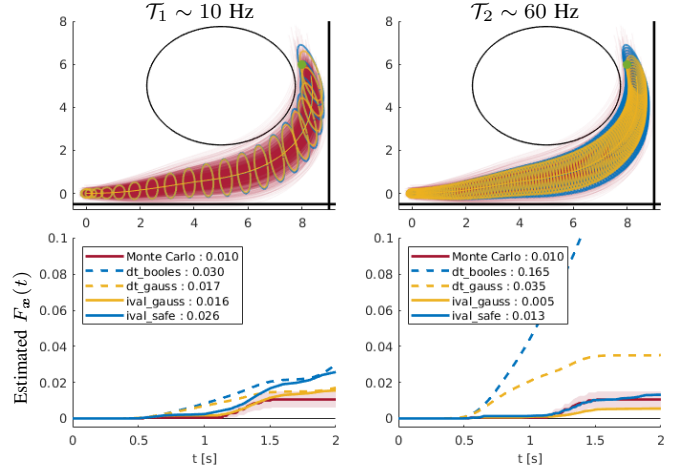


Fig. 3. We simulate a noisy Dubin’s car system navigating in a narrow passageway, computing collision probabilities under coarse (left) and fine (right) time discretizations, \mathcal{T}_1 and \mathcal{T}_2 . MC trajectories (red) simulate a fixed 60Hz feedback control rate, and a corresponding *a priori* state distribution \bar{b}_i (blue) is propagated via an LTV-Gaussian assumption. The anthropic estimates $\{\hat{\bar{\pi}}_i\}_j$ (yellow) are computed according to each \mathcal{T}_j . The underlying plots confirm the divergence of `dt_booles` under fine discretizations, and illustrate the susceptibility of `dt_gauss` and `ival_gauss` to inconsistency in the Gaussian $\hat{\bar{\pi}}_i$ estimates. In contrast, our method `ival_safe` avoids estimating $\bar{\pi}$ entirely, producing a lightweight and consistent risk estimate.

principle approximations. Discrete-time methods `dt_booles` and `dt_gauss` refer to the naive Boole’s approximation (6) and the conditioned variant proposed by [20], respectively. Our methods `ival_gauss` and `ival_safe` are described in the preceding section.

First, we successfully use our `ival_safe` estimate to optimize a risk-constrained ($\Delta = 2\%$) plan in a tight environment, shown in Fig. 3. We then computed risk estimates for this plan under multiple time discretizations. As predicted, `dt_booles` diverges severely as \mathcal{T} is refined while `ival_safe` converges correctly. Interestingly, `dt_gauss` and `ival_gauss` converge to incorrect values, likely due to inconsistency in the underlying Gaussian $\hat{\bar{\pi}}_t$ estimate.

Next, we perform a larger statistical evaluation over random, programmatically-generated environments. We consider two different regimes: *nominally*-safe trajectories, which ensure only that the nominal trajectory is collision-free, and *risk*-constrained trajectories optimized using `ival_safe` with $\Delta = 10\%$. This segregation allows us to identify potentially distinct statistical performance in scenarios representing the full range of risk values $[0, 1]$ and over the smaller range of “planning-relevant” values in $[0, \Delta]$. These environments are non-convex, and the locally-optimal trajectories were generated via MATLAB’s `fmincon` routine [17]. To ensure that the batch of experiments is not dominated by “uninteresting” cases, we specifically reject scenarios where the *unconstrained* optimal has negligible failure probability (i.e., when the obstacles are not relevant). Computation times for our MATLAB implementation of various methods are shown in Table I – as predicted, our method is significantly faster than MC while representing some additional complexity compared to discrete-




Fig. 4. Sample Dubin’s car trajectories in randomly-generated environments over a $T = 2.5$ [s] horizon under 60 Hz LQG feedback control towards the (green) goal. The two scenarios on the left demonstrate *nominally-safe* plans, which can experience significant risk upon stochastic execution. On the right are risk-constrained plans (generated using our `ival_safe` approximation with $\Delta = 10\%$). MC trajectories are shown in red, and dispersions are well-approximated by the *a priori* distributions $\{b_i\}$ shown in blue. The Gaussian anthropic belief estimates $\{\hat{\pi}_i\}$ are shown in yellow. The lower plots show $F_x(t)$ estimates as compared to MC (estimated standard error is indicated by the shaded region). Note that `dt_booles` is prone to extreme over-estimation, while the two $\{\hat{\pi}_i\}$ -dependent methods `dt_gauss` and `ival_gauss` suffer from the inconsistency of the underlying Gaussianity assumption. In contrast, our proposed method `ival_safe` achieves the best approximation performance and avoids explicit $\hat{\pi}_i$ approximation.

TABLE I
COMPUTE TIMES FOR DUBIN’S CAR MATLAB IMPLEMENTATION ON CONSUMER LAPTOP. QD REFERS TO QUADRATURE DIMENSION.

Monte Carlo ($N = 1000$)	LTV-Gaussian \bar{b}_i	Gaussian $\hat{\pi}_i$	$\sum_{i=0}^k P(x(t_i) \notin X_{\text{safe}})$ QD: $(n_p - 1) = 1$	$\sum_{i=0}^{k-1} \Phi(t_i, t_{i+1}; \mathbb{1}_{X_{\text{safe}}} \bar{b}_i)$ QD: $n_p = 2$
39.1 [s]	0.034 [s]	0.054 [s]	0.034 [s]	0.316 [s]

time methods. A more detailed discussion is presented in the supplementary material [7, Appendix B].

Some representative scenarios are shown in Fig. 4, and statistics are reported in Table II over 100 nominally-safe and 50 risk-constrained scenarios. For reference, the average MC-evaluated risk in each of the two test batches is listed in the top of the column. The **Bias** column lists the mean (signed) difference between the estimate and a 1000-sample MC estimate, and $P(\text{Conservative})$ reports the percentage of cases where the estimate was greater than (or within 5%) of the MC “truth.” The proposed method `ival_safe` outperforms or matches all others in bias, root-mean-squared error (RMSE), and median relative error (MRE), while remaining conservative (i.e., safe) in a significant majority of trials. Both `dt_gauss` and `ival_gauss` outperform `dt_booles` but suffer from inconsistency of the Gaussian $\hat{\pi}_i$ estimate.

VII. CONCLUSION

This paper addresses the challenging problem of efficiently and accurately estimating failure probabilities to enable risk-aware, continuous motion planning. By developing a rigorous framework directly in continuous-time and leaning heavily on the classic study of first-exit times, it becomes straightforward to identify a lightweight approximation (`ival_safe`) that dramatically outperforms existing methods. Furthermore, our approximation restores a convenient Bellman structure required for optimal control, enabling practical application for a wide variety of nonlinear systems.

TABLE II
 $F_x(T)$ ESTIMATION STATISTICS FOR DUBIN’S CAR SYSTEM.

Batch / Method	Bias	RMSE	MRE	$P(\text{Conservative})$
Nominally-safe	MC : 0.2649			
<code>dt_booles</code>	+2.9780	4.908	948 %	100 %
<code>dt_gauss</code>	+0.3156	0.923	70 %	89 %
<code>ival_gauss</code>	+0.0073	0.257	49 %	46 %
<code>ival_safe</code>	+0.0915	0.173	32 %	92 %
Risk-constrained	MC : 0.0321			
<code>dt_booles</code>	+0.4849	0.957	1162 %	100 %
<code>dt_gauss</code>	+0.0510	0.067	224 %	96 %
<code>ival_gauss</code>	-0.0138	0.025	39 %	30 %
<code>ival_safe</code>	+0.0057	0.019	41 %	86 %

Ultimately, the framework and concepts introduced in this paper (particularly Section III) motivate a number of future investigations. Of particular interest are robust methods of estimating anthropic belief via π_t or Ψ_t , as “survival-thus-far” may represent a useful and yet-uncaptured source of information for aspects of autonomous decision-making not limited to risk estimation. Also, fusing MC and Gaussian-based risk-approximations in a hybrid approach may provide the best of both worlds: high accuracy in the face of nonlinearity and compatibility with continuous optimization.

ACKNOWLEDGMENTS

This work was supported by the Education Office at The Charles Stark Draper Laboratory, Inc. and by ARL DCIST under Cooperative Agreement Number W911NF-17-2-0181.

REFERENCES

- [1] Kaito Ariu, Cheng Fang, Marcio Arantes, Claudio Toledo, and Brian Williams. Chance-constrained path planning with continuous time safety guarantees. In *Workshops at the Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- [2] Lars Blackmore and Masahiro Ono. Convex chance constrained predictive control without sampling. In *AIAA Guidance, Navigation, and Control Conference*, page 5876, 2009.
- [3] Lars Blackmore, Masahiro Ono, Askar Bektassov, and Brian C Williams. A probabilistic particle-control approximation of chance-constrained stochastic predictive control. *IEEE Transactions on Robotics*, 26(3):502–517, 2010.
- [4] Adam Bry and Nicholas Roy. Rapidly-exploring random belief trees for motion planning under uncertainty. In *Proc. IEEE Conf. Robot. Autom. (ICRA)*, pages 723–730. IEEE, 2011.
- [5] Giuseppe C Calafiore and Marco C Campi. The scenario approach to robust control design. *IEEE Trans. Automat. Contr.*, 51(5):742–753, 2006.
- [6] Siyu Dai, Shawn Schaffert, Ashkan Jasour, Andreas Hofmann, and Brian Williams. Chance constrained motion planning for high-dimensional robots. In *Proc. IEEE Conf. Robot. Autom. (ICRA)*, pages 8805–8811. IEEE, 2019.
- [7] Kristoffer M. Frey, Ted J. Steiner, and Jonathan P. How. Collision Probabilities for Continuous-Time Systems Without Sampling [with Appendices], 2020.
- [8] Xin Huang, Ashkan Jasour, Matthew Deyo, Andreas Hofmann, and Brian C Williams. Hybrid risk-aware conditional planning with applications in autonomous vehicles. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 3608–3614. IEEE, 2018.
- [9] Lucas Janson, Edward Schmerling, and Marco Pavone. Monte Carlo motion planning for robot trajectory optimization under uncertainty. In *Robotics Research*, pages 343–361. Springer, 2018.
- [10] Zhiyong Jin and Liqun Wang. First passage time for Brownian motion and piecewise linear boundaries. *Methodology and Computing in Applied Probability*, 19(1):237–253, 2017.
- [11] Ioannis Karatzsas and Steven E Shreve. *Brownian motion and stochastic calculus*, volume 113 of *Graduate Texts in Mathematics*. Springer, 1988.
- [12] Pu Li, Moritz Wendt, and Günter Wozny. A probabilistically constrained model predictive controller. *Automatica*, 38(7):1171–1176, 2002.
- [13] Brett T Lopez, Jonathan P How, and Jean-Jacques E Slotine. Dynamic tube MPC for nonlinear systems. In *American Control Conference (ACC)*, pages 1655–1662. IEEE, 2019.
- [14] Brandon Luders, Mangal Kothari, and Jonathan How. Chance constrained RRT for probabilistic robustness to environmental uncertainty. In *AIAA guidance, navigation, and control conference*, page 8160, 2010.
- [15] Yudong Ma, Sergey Vichik, and Francesco Borrelli. Fast stochastic MPC with optimal risk allocation applied to building control systems. In *Conf. on Dec. and Control (CDC)*, pages 7559–7564. IEEE, 2012.
- [16] Anirudha Majumdar and Russ Tedrake. Robust online motion planning with regions of finite time invariance. In *Algorithmic Foundations of Robotics X*, pages 543–558. Springer, 2013.
- [17] MATLAB Optimization Toolbox. MATLAB optimization toolbox, 2019b. The MathWorks Inc., Natick, MA, USA.
- [18] Masahiro Ono and Brian C Williams. Iterative risk allocation: A new approach to robust model predictive control with a joint chance constraint. In *Conf. on Dec. and Control (CDC)*, pages 3427–3432. IEEE, 2008.
- [19] Masahiro Ono, Marco Pavone, Yoshiaki Kuwata, and J Balaram. Chance-constrained dynamic programming with application to risk-aware robotic space exploration. *Autonomous Robots*, 39(4):555–571, 2015.
- [20] Sachin Patil, Jur Van Den Berg, and Ron Alterovitz. Estimating probability of collision for safe motion planning under Gaussian motion and sensing uncertainty. In *Proc. IEEE Conf. Robot. Autom. (ICRA)*, pages 3238–3244. IEEE, 2012.
- [21] Pedro Santana, Sylvie Thiébaux, and Brian Williams. RAO*: an algorithm for chance-constrained POMDP’s. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- [22] Emanuel Todorov and Weiwei Li. A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic systems. In *Proceedings of the 2005, American Control Conference, 2005.*, pages 300–306. IEEE, 2005.
- [23] Jur Van Den Berg, Pieter Abbeel, and Ken Goldberg. LQG-MP: Optimized path planning for robots with motion uncertainty and imperfect state information. *Int. J. of Robotics Research*, 30(7):895–913, 2011.
- [24] Jur Van Den Berg, Sachin Patil, and Ron Alterovitz. Motion planning under uncertainty using iterative local optimization in belief space. *Int. J. of Robotics Research*, 31(11):1263–1278, 2012.
- [25] Michael P Vitus and Claire J Tomlin. On feedback design and risk allocation in chance constrained control. In *Conf. on Dec. and Control (CDC)*, pages 734–739. IEEE, 2011.
- [26] Jiongmin Yong and Xun Yu Zhou. *Stochastic controls: Hamiltonian systems and HJB equations*, volume 43. Springer Science & Business Media, 1999.