European Control Conference 2022

Upper Bounds for Continuous-Time End-to-End Risks in Stochastic Robot Navigation

Apurva Patil Takashi Tanaka

July 15, 2022

Outline

Outline

Background

- End-to-end risks in stochastic robot navigation
  - Characterize safety of the planned trajectories.
  - Plan risk optimal trajectories.

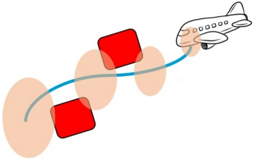

Background

- End-to-end risks in stochastic robot navigation
  - Characterize safety of the planned trajectories.
  - Plan risk optimal trajectories.

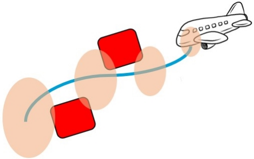

- Continuous-time end-to-end risk

$$\mathcal{R} = P\left(\bigcup_{t \in [0,T]} \boldsymbol{x}^{sys}(t) \in \mathcal{X}_{obs} \right).$$

Background

- End-to-end risks in stochastic robot navigation
  - Characterize safety of the planned trajectories.
  - Plan risk optimal trajectories.

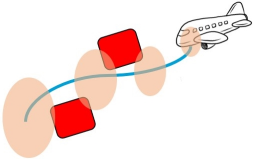

- Continuous-time end-to-end risk

$$\mathcal{R} = P\left(\bigcup_{t\in[0,T]} \boldsymbol{x}^{sys}(t) \in \mathcal{X}_{obs}\right).$$

- $\mathcal{R}$ is challenging to compute
  - $\boldsymbol{x}^{sys}(t)$ across $[0, T]$ are correlated.
  - We derive two upper bounds using **properties of Brownian motion**, and **Boole** and **Hunter's inequalities** .

Outline

Existing Approaches

- Monte Carlo methods [1]
  - Computationally expensive and cumbersome to embed in planning algorithms.

[1] (Janson, Schmerling, and Pavone 2018), (Blackmore et al. 2010)

[2] (Patil and Tanaka 2021)

[3] (Shah, Pahlajani, and Tanner 2011), (Chern et al. 2021)

[4] (Ariu et al. 2017)

Existing Approaches

- Monte Carlo methods [1]
  - Computationally expensive and cumbersome to embed in planning algorithms.
- Discrete-time approximations [2]
  - $\mathcal{R} \approx P\left(\bigcup_{i=0}^{N} \boldsymbol{x}^{sys}(t_i) \in \mathcal{X}_{obs}\right) \leq \sum_{i=0}^{N} P\left(\boldsymbol{x}^{sys}(t_i) \in \mathcal{X}_{obs}\right).$
  - Sensitive to the chosen time discretization.

[1] (Janson, Schmerling, and Pavone 2018), (Blackmore et al. 2010)

[2] (Patil and Tanaka 2021)

[3] (Shah, Pahlajani, and Tanner 2011), (Chern et al. 2021)

[4] (Ariu et al. 2017)

Existing Approaches

- Monte Carlo methods [1]
  - Computationally expensive and cumbersome to embed in planning algorithms.
- Discrete-time approximations [2]
  - $\mathcal{R} \approx P\left(\bigcup_{i=0}^{N} \boldsymbol{x}^{sys}(t_i) \in \mathcal{X}_{obs}\right) \leq \sum_{i=0}^{N} P\left(\boldsymbol{x}^{sys}(t_i) \in \mathcal{X}_{obs}\right).$
  - Sensitive to the chosen time discretization.
- Continuous-time methods
  - PDE-based methods [3]: closed-form solution not tractable.
  - Reflection-principle-based method [4]

1 (Janson, Schmerling, and Pavone 2018), (Blackmore et al. 2010)

2 (Patil and Tanaka 2021)

3 (Shah, Pahlajani, and Tanner 2011), (Chern et al. 2021)

4 (Ariu et al. 2017)

Outline

Problem Formulation

- Planned trajectory
  - A finite sequences of positions $\{x_j^{plan} \in \mathcal{X}_{free}\}_{j=0,1,\ldots,N}$ and control inputs $\{v_j^{plan} \in \mathbb{R}^n\}_{j=0,1,\ldots,N-1}$.
  - Let $\mathcal{T} = (0 = t_0 < \ldots < t_N = T)$ s.t. $v_j^{plan}\Delta t_j = x_{j+1}^{plan} - x_j^{plan}$.

Problem Formulation

- Planned trajectory
  - A finite sequences of positions $\{x_j^{plan} \in \mathcal{X}_{free}\}_{j=0,1,\ldots,N}$ and control inputs $\{v_j^{plan} \in \mathbb{R}^n\}_{j=0,1,\ldots,N-1}$.
  - Let $\mathcal{T} = (0 = t_0 < \ldots < t_N = T)$ s.t. $v_j^{plan} \Delta t_j = x_{j+1}^{plan} - x_j^{plan}$.
- Robot dynamics
  - Controlled Itô process

$$d\boldsymbol{x}^{sys}(t) = \boldsymbol{v}^{sys}(t)dt + R^{\frac{1}{2}} d\boldsymbol{w}(t),$$

where $\boldsymbol{v}^{sys}(t) = v_j^{plan} \qquad \forall\ t \in [t_j, t_{j+1})$.

Problem Formulation

- Planned trajectory
  - A finite sequences of positions $\{x_j^{plan} \in \mathcal{X}_{free}\}_{j=0,1,\ldots,N}$ and control inputs $\{v_j^{plan} \in \mathbb{R}^n\}_{j=0,1,\ldots,N-1}$.
  - Let $\mathcal{T} = (0 = t_0 < \ldots < t_N = T)$ s.t. $v_j^{plan}\Delta t_j = x_{j+1}^{plan} - x_j^{plan}$.
- Robot dynamics
  - Controlled Itô process

$$d\boldsymbol{x}^{sys}(t) = \boldsymbol{v}^{sys}(t)dt + R^{\frac{1}{2}}d\boldsymbol{w}(t),$$

where $\boldsymbol{v}^{sys}(t) = v_j^{plan} \quad \forall\, t \in [t_j, t_{j+1})$.

- Continuous-time end-to-end risk
  - If $\mathcal{T}_j = [t_{j-1}, t_j]$,

$$\mathcal{R} = P\left(\bigcup_{t\in[0,T]} \boldsymbol{x}^{sys}(t) \in \mathcal{X}_{obs}\right) = P\left(\bigcup_{j=1}^{N}\bigcup_{t\in\mathcal{T}_j} \boldsymbol{x}^{sys}(t) \in \mathcal{X}_{obs}\right).$$

Outline

Risk in terms of 1-D Brownian Motions

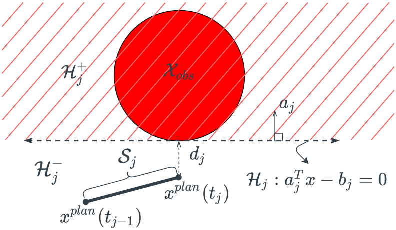

$$\mathcal{R} = P\left(\bigcup_{j=1}^{N} \bigcup_{t\in\mathcal{T}_j} \boldsymbol{x}^{sys}(t) \in \mathcal{X}_{obs}\right)$$

$$\leq P\left(\bigcup_{j=1}^{N} \bigcup_{t\in\mathcal{T}_j} a_j^T \boldsymbol{x}(t) \geq d_j\right)$$

where $\boldsymbol{x}(t) = \boldsymbol{x}^{sys}(t) - x^{plan}(t)$.

Risk in terms of 1-D Brownian Motions

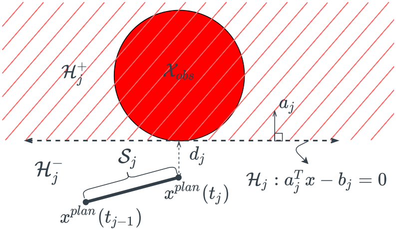

$$\mathcal{R} = P\left(\bigcup_{j=1}^{N} \bigcup_{t \in \mathcal{T}_j} \boldsymbol{x}^{sys}(t) \in \mathcal{X}_{obs}\right)$$

$$\leq P\left(\bigcup_{j=1}^{N} \bigcup_{t \in \mathcal{T}_j} a_j^T \boldsymbol{x}(t) \geq d_j\right)$$

where $\boldsymbol{x}(t) = \boldsymbol{x}^{sys}(t) - x^{plan}(t)$.

$\boldsymbol{w}_j(t) = a_j^T \boldsymbol{x}(t)$ is a one-dimensional Brownian motion for $t \in [0, T]$ that starts in the origin.

$$\mathcal{R} \leq P\left(\bigcup_{j=1}^{N} \max_{t \in \mathcal{T}_j} \boldsymbol{w}_j(t) \geq d_j\right).$$

Outline

First-Order Risk Bound

- Using Boole's inequality

$$\mathcal{R} \leq P\left(\bigcup_{j=1}^{N} \max_{t \in \mathcal{T}_j} \boldsymbol{w}_j(t) \geq d_j\right) \leq \boxed{\sum_{j=1}^{N} \underbrace{P\left(\max_{t \in [t_{j-1}, t_j]} \boldsymbol{w}_j(t) \geq d_j\right)}_{p_j}}.$$

First-Order Risk Bound

- Using Boole's inequality

$$\mathcal{R} \leq P\left(\bigcup_{j=1}^{N} \max_{t \in \mathcal{T}_j} \boldsymbol{w}_j(t) \geq d_j\right) \leq \boxed{\sum_{j=1}^{N} \underbrace{P\left(\max_{t \in [t_{j-1}, t_j]} \boldsymbol{w}_j(t) \geq d_j\right)}_{p_j}}.$$

- This bound possesses the time-additive structure.

First-Order Risk Bound

- Using Boole's inequality

$$\mathcal{R} \leq P\left(\bigcup_{j=1}^{N} \max_{t \in \mathcal{T}_j} \boldsymbol{w}_j(t) \geq d_j\right) \leq \boxed{\sum_{j=1}^{N} \underbrace{P\left(\max_{t \in [t_{j-1}, t_j]} \boldsymbol{w}_j(t) \geq d_j\right)}_{p_j}}.$$

- This bound possesses the time-additive structure.
- p_j is the continuous-time risk associated with the time segment $\mathcal{T}_j = [t_{j-1}, t_j]$.

First-Order Risk Bound

- Using Boole's inequality

$$\mathcal{R} \leq P\left(\bigcup_{j=1}^{N} \max_{t\in\mathcal{T}_j} \boldsymbol{w}_j(t) \geq d_j\right) \leq \boxed{\sum_{j=1}^{N} \underbrace{P\left(\max_{t\in[t_{j-1},t_j]} \boldsymbol{w}_j(t) \geq d_j\right)}_{p_j}}.$$

- This bound possesses the time-additive structure.
- p_j is the continuous-time risk associated with the time segment $\mathcal{T}_j = [t_{j-1}, t_j]$.

How to compute p_j?

Reflection principle of Brownian Motion

If $\boldsymbol{w}(t)$, $t \geq 0$ is a one-dimensional Brownian motion started in the origin and $d > 0$ is a threshold value, then

$$P\left(\sup_{s\in[0,t]} \boldsymbol{w}(s) \geq d\right) = 2P\left(\boldsymbol{w}(t) \geq d\right).$$

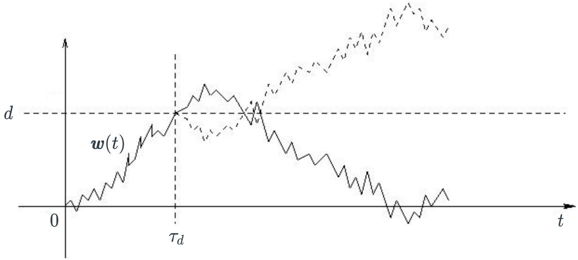

Durrett 2019.

First-Order Risk Bound: Computation of p_j

- Approach proposed by Ariu et al.[5]
  - Compute an upper bound to p_j

$$p_j = P\left(\max_{t\in[t_{j-1},t_j]} \boldsymbol{w}_j(t) \geq d_j\right) \leq P\left(\max_{t\in[0,t_j]} \boldsymbol{w}_j(t) \geq d_j\right).$$

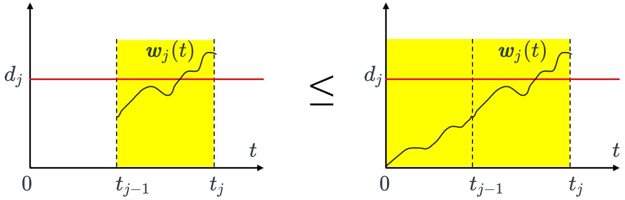

[5] (Ariu et al. 2017)

First-Order Risk Bound: Computation of p_j

- Approach proposed by Ariu et al.[5]
  - Compute an upper bound to p_j

$$p_j = P\left(\max_{t\in[t_{j-1},t_j]} \boldsymbol{w}_j(t) \geq d_j\right) \leq P\left(\max_{t\in[0,t_j]} \boldsymbol{w}_j(t) \geq d_j\right).$$

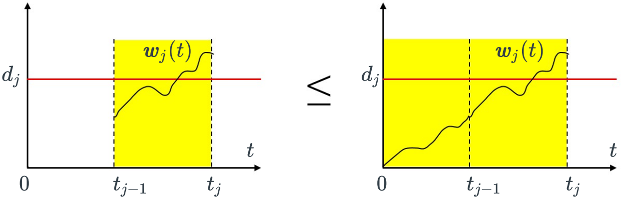

  - Using the reflection principle

$$P\left(\max_{t\in[0,t_j]} \boldsymbol{w}_j(t) \geq d_j\right) = 2P(\boldsymbol{w}_j(t_j) \geq d_j) = 2P\left(a_j^T \boldsymbol{x}_j \geq d_j\right).$$

---

[5] (Ariu et al. 2017)

Markov property of Brownian Motion

Let $\boldsymbol{w}(t)$, $t \geq 0$ be an n-dimensional Brownian motion started in $z \in \mathbb{R}^n$. Let $s \geq 0$, then the process $\widetilde{\boldsymbol{w}}(t) = \boldsymbol{w}(t+s) - \boldsymbol{w}(s)$, $t \geq 0$ is again a Brownian motion started in the origin and it is independent of the process $\boldsymbol{w}(t)$, $0 \leq t \leq s$.

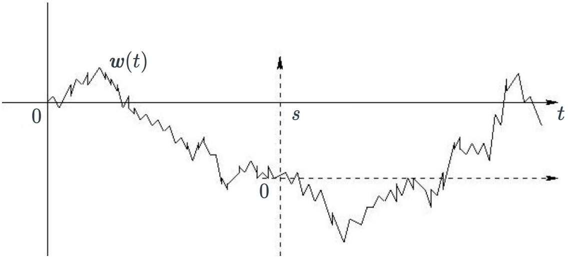

Mörters and Peres 2010.

First-Order Risk Bound: Computation of p_j

- Our approach

Compute p_j exactly using the Markov property and reflection principle of Brownian motion

$$p_j = P\left(\max_{t\in[t_{j-1},t_j]} \boldsymbol{w}_j(t) \geq d_j\right) = P\left(\max_{t\in[t_{j-1},t_j]} \boldsymbol{w}_j(t) \geq d_j,\ \boldsymbol{w}_j(t_{j-1}) \geq d_j\right) + P\left(\max_{t\in[t_{j-1},t_j]} \boldsymbol{w}_j(t) \geq d_j,\ \boldsymbol{w}_j(t_{j-1}) < d_j\right).$$

First-Order Risk Bound: Computation of p_j

- Our approach

Compute p_j exactly using the Markov property and reflection principle of Brownian motion

$$p_j = P\left(\max_{t\in[t_{j-1},t_j]} \boldsymbol{w}_j(t) \geq d_j\right) = P\left(\max_{t\in[t_{j-1},t_j]} \boldsymbol{w}_j(t) \geq d_j,\ \boldsymbol{w}_j(t_{j-1}) \geq d_j\right) + P\left(\max_{t\in[t_{j-1},t_j]} \boldsymbol{w}_j(t) \geq d_j,\ \boldsymbol{w}_j(t_{j-1}) < d_j\right).$$

If $\boldsymbol{z}_j^s = \boldsymbol{w}_j(t_{j-1}),\ \boldsymbol{z}_j^e := \boldsymbol{w}_j(t_j),\ \boldsymbol{z}_j := \begin{bmatrix}\boldsymbol{z}_j^s & \boldsymbol{z}_j^e\end{bmatrix}^T$

$$p_j = \int_{z_j^s=d_j}^{\infty} \mu_{\boldsymbol{z}_j^s}(z_j^s)dz_j^s + 2\int_{z_j^s=-\infty}^{d_j}\int_{z_j^e=d_j}^{\infty} \mu_{\boldsymbol{z}_j}(z_j)dz_j^e dz_j^s.$$

Outline

Second-Order Risk Bound

- Using Hunter's inequality

$$\mathcal{R} \leq P\left(\bigcup_{j=1}^{N} \underbrace{\max_{t\in\mathcal{T}_j} \boldsymbol{w}_j(t) \geq d_j}_{\mathcal{E}_j}\right) \leq \boxed{\sum_{j=1}^{N} p_j - \sum_{j=1}^{N-1} p_{j,j+1}.}$$

- $p_{j,j+1} = P(\mathcal{E}_j \cap \mathcal{E}_{j+1})$ is the joint risk associated with the time segments $\mathcal{T}_j$ and $\mathcal{T}_{j+1}$.

Second-Order Risk Bound

- Using Hunter's inequality

$$\mathcal{R} \leq P\left(\bigcup_{j=1}^{N} \underbrace{\max_{t\in\mathcal{T}_j} \boldsymbol{w}_j(t) \geq d_j}_{\mathcal{E}_j}\right) \leq \boxed{\sum_{j=1}^{N} p_j - \sum_{j=1}^{N-1} p_{j,j+1}.}$$

- $p_{j,j+1} = P(\mathcal{E}_j \cap \mathcal{E}_{j+1})$ is the joint risk associated with the time segments $\mathcal{T}_j$ and $\mathcal{T}_{j+1}$.
- This bound possesses the time-additive structure.

Second-Order Risk Bound

- Using Hunter's inequality

$$\mathcal{R} \leq P\left(\bigcup_{j=1}^{N} \underbrace{\max_{t\in\mathcal{T}_j} \boldsymbol{w}_j(t) \geq d_j}_{\mathcal{E}_j}\right) \leq \boxed{\sum_{j=1}^{N} p_j - \sum_{j=1}^{N-1} p_{j,j+1}}.$$

- $p_{j,j+1} = P(\mathcal{E}_j \cap \mathcal{E}_{j+1})$ is the joint risk associated with the time segments $\mathcal{T}_j$ and $\mathcal{T}_{j+1}$.
- This bound possesses the time-additive structure.

How to compute $p_{j,j+1}$?

Second-Order Risk Bound: Computation of $p_{j,j+1}$

- Compute $p^{LB}_{j,j+1}$: an lower bound to $p_{j,j+1}$

$$\mathcal{R} \leq \sum_{j=1}^{N} p_j - \sum_{j=1}^{N-1} p^{LB}_{j,j+1}.$$

Second-Order Risk Bound: Computation of $p_{j,j+1}$

- Compute $p^{LB}_{j,j+1}$: an lower bound to $p_{j,j+1}$

$$\mathcal{R} \leq \sum_{j=1}^{N} p_j - \sum_{j=1}^{N-1} p^{LB}_{j,j+1}.$$

- If $t_{j-1} = \hat{t}^0_j < \hat{t}^1_j < \ldots < \hat{t}^{r_j}_j = t_j$ is a discretization of the time segment $\mathcal{T}_j$, and $\boldsymbol{z}^i_j := \boldsymbol{w}_j(\hat{t}^i_j) = a^T_j \boldsymbol{x}(\hat{t}^i_j)$, $\mathcal{D}_j := \left(\boldsymbol{z}^0_j < d_j\right) \cap \left(\boldsymbol{z}^1_j < d_j\right) \cap \ldots \cap \left(\boldsymbol{z}^{r_j}_j < d_j\right)$, then $p_{j,j+1}$ is lower bounded by $p^{LB}_{j,j+1}$ given as

$$\boxed{p^{LB}_{j,j+1} = 1 - P(\mathcal{D}_j) - P(\mathcal{D}_{j+1}) + P(\mathcal{D}_j \cap \mathcal{D}_{j+1}).}$$

Outline

Simulation Results

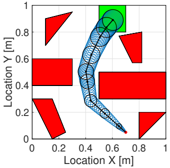

(a) 95%

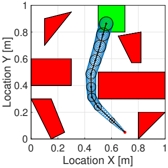

(b) 75%

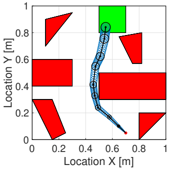

(c) 50%

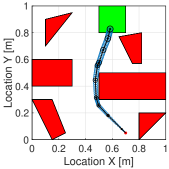

(d) 25%

Trajectories planned with the instantaneous safety

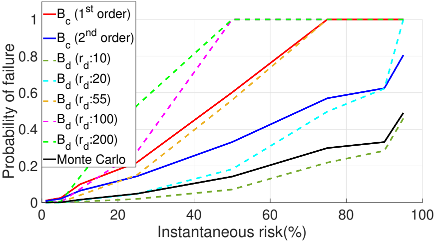

B_c: continuous-time risk bounds
B_d: discrete-time risk bounds
r_d: rate of time discretization

Pedram et al. 2021.

Simulation Results

Comparison of different risk estimates over 100 trajectories

Risk Estimates	Avg. Time	Bias	RMSE	%Cons.
Monte Carlo	101.50 s	0	0	-
Discrete-time				
r_d : 5	0.14 s	-0.14	0.18	28%
r_d : 10	0.26 s	-0.002	0.16	59%
r_d : 20	0.52 s	0.31	0.57	82%
r_d : 55	1.53 s	1.50	2.33	100%
r_d : 100	2.87 s	2.98	4.53	100%
Continuous-time				
Ariu et al. [6]	1.39 s	0.97	1.33	100%
Our 1^{st} order	1.47 s	0.66	0.90	100%
Our 2^{nd} order	2.23 s	0.28	0.36	100%

6 (Ariu et al. 2017)

Outline

Summary

- We derived two upper bounds for the continuous-time end-to-end risk using properties of Brownian motion.
- These bounds possess the time-additive structure, making them useful for risk-aware motion planning.
- Numerical validation demonstrates that our bounds outperform the state-of-the-art discrete-time bound and are cheaper in computation than the Monte Carlo method.

Check out the paper for more details and results.

[7] Patil et al. "Chance-Constrained Stochastic Optimal Control via Path Integral and Finite Difference Methods." arXiv preprint arXiv:2205.00628 (2022).

Summary

- We derived two upper bounds for the continuous-time end-to-end risk using properties of Brownian motion.
- These bounds possess the time-additive structure, making them useful for risk-aware motion planning.
- Numerical validation demonstrates that our bounds outperform the state-of-the-art discrete-time bound and are cheaper in computation than the Monte Carlo method.

Check out the paper for more details and results.

- **Future work**: Risk-constrained optimal control and risk analysis of systems with generalized nonlinear stochastic dynamics via an HJB-PDE [7].

[7] Patil et al. "Chance-Constrained Stochastic Optimal Control via Path Integral and Finite Difference Methods." arXiv preprint arXiv:2205.00628 (2022).