Automatica 131 (2021) 109754

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Brief paper

Trajectory planning under environmental uncertainty with finite-sample safety guarantees☆

Vasileios Lefkopoulos [a], Maryam Kamgarpour [a,b,*]

[a] Automatic Control Laboratory, ETH Zurich, Zürich 8092, Switzerland
[b] Electrical and Computer Engineering, University of British Columbia, Vancouver, V6T 1Z4, Canada

ARTICLE INFO

Article history:
Received 14 August 2020
Received in revised form 16 March 2021
Accepted 6 May 2021
Available online 20 June 2021

Keywords:
Autonomous vehicles
Trajectory planning
Stochastic optimal control
Chance constraints

ABSTRACT

We tackle the problem of trajectory planning in an environment comprised of a set of obstacles with uncertain time-varying locations. The uncertainties are modeled using widely accepted Gaussian distributions, resulting in a chance-constrained program. Contrary to previous approaches however, we do not assume perfect knowledge of the moments of the distribution, and instead estimate them through finite samples available from either sensors or past data. We derive tight concentration bounds on the error of these estimates to sufficiently tighten the chance-constraint program. As such, we provide provable guarantees on satisfaction of the chance-constraints corresponding to the nominal yet unknown moments. We illustrate our results with two autonomous vehicle trajectory planning case studies.

1. Introduction

A major challenge in real-world employment of autonomous systems such as robots and autonomous cars is handling uncertainties in the environment. For example, from the perspective of the autonomous car, the driving behavior of the nearby cars is to a large extent unpredictable. In a robotic search-and-rescue, the location of goals and obstacles are a priori uncertain. In several applications of autonomous vehicles, large number of past data on the environment conditions (e.g. driving behavior of cars or environment maps) exist either through camera or LIDAR measurements (Feng, Rosenbaum, & Dietmayer, 2018; Kendall & Gal, 2017), leading to a requirement for an approach that accounts for the potential magnitude of the uncertainty based on these samples and ensures safety while allowing for real-time computation (Michelmore, et al., 2019). Gaussian distributions are often used to approximate sensor noise and are increasingly used in autonomous driving applications to model perturbations from nominal prototype maneuvers when predicting the behavior of other vehicles (Carvalho, Gao, Lefevre, & Borrelli, 2014). The approach proposed in this work addresses this issue of generating safe trajectories while working with an arbitrary finite number of samples, under a Gaussian assumption.

Past work in stochastic trajectory planning considers uncertainties stemming from two sources: (1) the autonomous agent's dynamic and measurement model (Blackmore, Ono, & Williams, 2011; Ono, Blackmore, & Williams, 2010; Ono, Pavone, Kuwata, & Balaram, 2015; Raman, Donzé, Sadigh, Murray, & Seshia, 2015; Vitus, 2012); (2) the obstacle safe sets in the environment (Carvalho et al., 2014; Jha, Raman, Sadigh, & Seshia, 2018; Sessa, Frick, Wood, & Kamgarpour, 2018). Blackmore et al. (2011) consider polyhedral *deterministic* obstacles and a linear plant model perturbed by Gaussian noise with *known* moments. Accordingly, they reformulate the chance constraints and employ a specialized branch-and-bound technique to solve the resulting disjunctive program of obstacle avoidance. This approach is not directly extendable to stochastic obstacles as it would result in a nonlinear variant of the risk allocation problem. Jha et al. (2018) consider polyhedral *uncertain* obstacles whose stochasticity is Gaussian with *known* moments and a *deterministic* vehicle model. They reformulate the chance constraints assuming a Gaussian distribution and then encode the obstacle avoidance problem as a mixed-integer scheme. The above methods work under the assumption of perfect knowledge of the Gaussian moments, whereas in reality one can only estimate an uncertainty model based on sensor measurements or past data. Sessa et al. (2018) consider nonlinear dependence of the dynamics on the noise. They do not make assumptions about the noise distribution, but instead consider

☆ The work of M. Kamgarpour is gratefully supported by the Swiss National Science Foundation, under the grant SNSF 200021_172782, and by the NSERC, Canada Discovery Grant RGPAS-2020-00110. The material in this paper was partially presented at the 18th European Control Conference, June 25–28, 2019, Naples, Italy. This paper was recommended for publication in revised form by Associate Editor Andrey V. Savkin under the direction of Editor Ian R. Petersen.

* Corresponding author at: Electrical and Computer Engineering, University of British Columbia, Vancouver, V6T 1Z4, Canada.
E-mail addresses: vlefkopo@gmail.com (V. Lefkopoulos), maryamk@ece.ubc.ca (M. Kamgarpour).

https://doi.org/10.1016/j.automatica.2021.109754

that a sufficient number of its samples are available. As such, using the scenario approach, they provide probabilistic guarantees of the designed disturbance feedback control. While their work produces control policies that can adhere to safety under general disturbance models, a significant number of samples are required for the approach to produce any safety guarantees. Furthermore, the approach is too computationally intense to be implemented in real time.

In this work we attempt to reach a compromise between the purely data-driven approach of Sessa et al. (2018) and the known distribution assumption of Jha et al. (2018) for trajectory generation in stochastic environments. In particular, we take on the fairly accepted model of a Gaussian distributed uncertainty but consider the realistic case in which the moments of the distribution are *unknown* and are estimated through finite samples. We derive concentration bounds on the estimation error of the Gaussian moments, propose a reformulation of the chance-constrained problem using only the moment estimates and prove its solution's feasibility (in a probabilistic sense) for the original problem. We also extend the proposed method to consider *dynamic* uncertain obstacles in a receding horizon control scheme. Early results of the work in this paper were presented in a conference paper in Lefkopoulos and Kamgarpour (2019). Our work here completes the early studies by extending it in the following ways:

(1) Providing improved moment concentration bounds, as presented in Lemma 5 and thus reducing the conservativeness in Lefkopoulos and Kamgarpour (2019).
(2) Extending the open-loop approach in Lefkopoulos and Kamgarpour (2019) to a closed-loop control scheme, as presented in Section 4.
(3) Providing two case studies to illustrate the above contributions.

The rest of this paper is organized as follows. Section 2 states the problem and reformulates it as a mixed integer second order cone program. Our result on incorporation of the moment uncertainties through concentration bounds is presented in Section 3. In Section 4 a receding horizon scheme utilizing our approach is proposed. Section 5 demonstrates the approach with two case studies. Finally, we conclude in Section 6.

Notation

We denote the subset from a to b of an ordered set $\mathbb{X}$ by $\mathbb{X}[a, b]$. We denote a conjunction by $\wedge$ and a disjunction by $\vee$. We denote the concatenation of two row vectors $a \in \mathbb{R}^{1\times n}$, $b \in \mathbb{R}^{1\times m}$ as $[a, b] \in \mathbb{R}^{1\times(n+m)}$. By $\mathcal{N}(\mu, \Sigma)$ we denote the n-dimensional multivariate Gaussian distribution with mean $\mu \in \mathbb{R}^n$ and covariance $\Sigma \in \mathbb{R}^{n\times n}$. By $\Psi^{-1}(\cdot)$ we denote the inverse cumulative distribution function of $\mathcal{N}(0, 1)$. A set of random variables d_i are i.i.d. if they are independent and identically distributed.

2. Problem statement

We consider the system evolving according to the time-varying and linear dynamics:

$$x_{t+1} = A_t x_t + B_t u_t \,, \tag{1}$$

where $x_t \in \mathbb{R}^{n_x}$ is the state, $u_t \in \mathbb{R}^{n_u}$ is the input and $A_t \in \mathbb{R}^{n_x\times n_x}$, $B_t \in \mathbb{R}^{n_x\times n_u}$ are the system dynamics' matrices at time $t \in \mathbb{N}$. Given a horizon length $N \in \mathbb{N}_{>0}$, we denote the finite-length input sequence as $\boldsymbol{u} := (u_0, \ldots, u_{N-1}) \in \mathbb{R}^{Nn_u}$. Given an initial state x_0 and $\boldsymbol{u}$, the evolution of the state trajectory is denoted as $\boldsymbol{x} := (x_1, \ldots, x_N) \in \mathbb{R}^{Nn_x}$. The control inputs are constrained to a convex set $\mathcal{U}$, $u_t \in \mathcal{U}$ for all $t \in \mathbb{N}[0, N-1]$.[1]

For trajectory planning in uncertain environments, the state needs to avoid a set of obstacles with *uncertain* locations. We consider obstacles modeled by polyhedrons. Let $N_o \in \mathbb{N}$ be the number of obstacles (indexed by j), with F_j number of faces (indexed by i). The j-th obstacle's interior $\mathcal{O}_j^t \in \mathbb{R}^{n_x}$ at time t can be expressed as the conjunction of the linear constraints of its faces:

$$\mathcal{O}_j^t := \left\{ x \in \mathbb{R}^{n_x} : \wedge_{i=1}^{F_j} {a_{ij}^t}^\top x + b_{ij}^t \le 0 \right\} , \tag{2}$$

where $a_{ij}^t \in \mathbb{R}^{n_x}$ and $b_{ij}^t \in \mathbb{R}$ define the j-th's obstacle i-th face at time t, and are concatenated in $d_{ij}^t := [{a_{ij}^t}^\top, b_{ij}^t]^\top \in \mathbb{R}^{n_x+1}$. The complement of the obstacle set is defined via the disjunction of the constraints:

$$\overline{\mathcal{O}_j^t} := \left\{ x \in \mathbb{R}^{n_x} : \vee_{i=1}^{F_j} {a_{ij}^t}^\top x + b_{ij}^t > 0 \right\} . \tag{3}$$

The safe set over the planning horizon is $\boldsymbol{\mathcal{X}} \subset \mathbb{R}^{Nn_x}$:

$$\boldsymbol{\mathcal{X}} := \left\{ \boldsymbol{x} \in \mathbb{R}^{Nn_x} : \wedge_{t=1}^{N} x_t \in \mathcal{X}_t \right\} , \tag{4}$$

$$\mathcal{X}_t := \left\{ x \in \mathbb{R}^{n_x} : \wedge_{j=1}^{N_o} \vee_{i=1}^{F_j} {a_{ij}^t}^\top x + b_{ij}^t > 0 \right\} . \tag{5}$$

The uncertainty of the obstacles' locations is captured by considering $d_{ij}^t \sim \mathcal{D}_{ij}^t$, where $\mathcal{D}_{ij}^t$ is a probability distribution on $\mathbb{R}^{n_x+1}$, making the set $\boldsymbol{\mathcal{X}}$ stochastic. We formulate the safety requirement as a *joint chance constraint* $\Pr(\boldsymbol{x} \in \boldsymbol{\mathcal{X}}) \ge 1-\epsilon$, where ϵ is a prescribed safety margin.

We assume the objective of the autonomous vehicle is captured through a convex cost function $J(x_0, \cdot) : \mathbb{R}^{n_u} \to \mathbb{R}$. Thus, the chance-constraint safe trajectory planning is:

$$\min_{\boldsymbol{u}} \quad J(x_0, \boldsymbol{u}) \tag{6a}$$

$$\text{s.t.} \quad \boldsymbol{x}, \boldsymbol{u} \text{ satisfy (1) with initial state } x_0 \tag{6b}$$

$$\boldsymbol{u} \in \boldsymbol{\mathcal{U}} \tag{6c}$$

$$\Pr(\boldsymbol{x} \in \boldsymbol{\mathcal{X}}) \ge 1 - \epsilon \tag{6d}$$

where $\boldsymbol{\mathcal{U}} := \mathcal{U} \times \ldots \times \mathcal{U} \subseteq \mathbb{R}^{Nn_u}$.

2.1. Single chance constraints formulation

Similarly to Jha et al. (2018) and Lefkopoulos and Kamgarpour (2019), we reformulate the non-convex disjunction of (3) using the Big-M method (Schouwenaars, Moor, Feron, & How, 2001) and Boole's inequality (Casella & Berger, 2001). Hence, a sufficient condition for satisfying constraint (6d) is:

$$\wedge_{t=1}^{N} \wedge_{j=1}^{N_o} \wedge_{i=1}^{F_j} \Pr({a_{ij}^t}^\top x_t + b_{ij}^t + M z_{ij}^t > 0) \ge 1 - \epsilon_{ij}^t \,, \tag{7}$$

$$\sum_{i=1}^{F_j} z_{ij}^t < F_j, \quad \forall t \in \mathbb{N}[1, N], \quad \forall j \in \mathbb{N}[1, N_o] \,, \tag{8}$$

$$\sum_{t=1}^{N} \sum_{j=1}^{N_o} \sum_{i=1}^{F_j} \epsilon_{ij}^t \le \epsilon \,, \tag{9}$$

for a sufficiently large constant $M \in \mathbb{R}_{>0}$, where $z_{ij}^t \in \mathbb{B}$, $\epsilon_{ij}^t \in (0, 0.5)$ are the binary and risk variables associated with each obstacle face.

[1] The assumption of the set $\mathcal{U}$ being time-invariant is only for simplicity in notation.

2.2. Risk allocation

A simple although generally conservative way to allocate the risk ϵ_{ij}^t is to do so uniformly (Blackmore, Hui Li, & Williams, 2006; Nemirovski & Shapiro, 2007):

$$\epsilon_{\text{uni}} := \epsilon_{ij}^t = \frac{\epsilon}{N \sum_{j=1}^{N_o} F_j}. \tag{10}$$

In literature, different approaches have been proposed to improve over (10), such as optimizing over the decision variables and the risk allocation separately (Jha et al., 2018; Vitus & Tomlin, 2011) or even simultaneously (Blackmore & Ono, 2009; Blackmore et al., 2011; Ono et al., 2010); with all approaches resulting in an increase of computation time. Our first result below shows that, for the problem at hand, one can significantly improve upon the uniform risk allocation without the need for additional optimization.

Lemma 1. *Using the risk allocation:*

$$\epsilon_{ij}^t = \frac{\epsilon}{NN_o} > \epsilon_{uni}\,, \tag{11}$$

and the binary variable constraints

$$\sum_{i=1}^{F_j} z_{ij}^t = F_j - 1\,, \tag{12}$$

any feasible point for problem $\mathcal{P}$:={(7), (11), (12)} is also feasible for problem $\mathcal{P}'$:={(8), (7), (9)} and vice versa.

Proof. First, let us show that any feasible point for $\mathcal{P}$ is feasible for $\mathcal{P}'$. Notice that every binary variable z_{ij}^t that satisfies (12) also satisfies (8). For a given obstacle j, let i^* denote the unique binary term above such that $z_{i^*j}^t = 0$. It follows that for all $i \neq i^*$, the corresponding constraints $a_{ij}^t x_t + b_{ij}^t + M z_{ij}^t$ hold regardless of the values the random variables take due to the choice of M in the big-M method. Consequently, (11) is sufficient to ensure (7). It follows that any feasible point for $\mathcal{P}$ is also feasible for $\mathcal{P}'$. Next, let us show that any feasible point for $\mathcal{P}'$ is feasible for $\mathcal{P}$. Suppose there exists an index i such that $z'^t_{ij} = 0$ in problem $\mathcal{P}'$ with its corresponding x'_t, whereas this term is $z_{ij}^t = 1$ in $\mathcal{P}$. From the choice of M in big-M method, it follows that x'_t also satisfies constraint i in $\mathcal{P}$ for any values of the random variables (since $z_{ij}^t = 1$ means that x_t can be set arbitrarily.) Hence, the risk corresponding to these constraints can be set to zero.

2.3. Mixed-integer second-order cone formulation

Consistent with several existing trajectory planning works (Blackmore et al., 2011; Carvalho et al., 2014; Jha et al., 2018), we assume Gaussian uncertainties to allow the analytic reformulation of the chance constraints.

Assumption 2. Each d_{ij}^t is a Gaussian random variable, $d_{ij}^t \sim \mathcal{N}(\mu_{ij}^t, \Sigma_{ij}^t)$.

Under Assumption 2, each chance constraint of (7) is equivalent to the following second-order cone constraint (Calafiore & El Ghaoui, 2006, Theorem 2.1):

$$\Psi^{-1}(1-\epsilon_{ij}^t)\left\|(\Sigma_{ij}^t)^{1/2}\tilde{x}_t\right\|_2 \leq \mu_{ij}^{t\top}\tilde{x}_t + M z_{ij}^t\,, \tag{13}$$

where $\tilde{x} := [x^\top, 1]^\top \in \mathbb{R}^{n_x+1}$.

Finally, as a direct result of Sections 2.1–2.3, we can conservatively satisfy the joint chance constraint (6d) through the implication (13) $\wedge$(12) $\Rightarrow$ (6d).

3. Moments Robust Approach

To ensure satisfaction of the chance constraints (7) when the moments μ_{ij}^t and Σ_{ij}^t are estimated from data, the estimates' uncertainties need to be accounted for in the constraint formulation. In order to do so and to guarantee constraint satisfaction with high confidence, we derive tight moment concentration bounds in Section 3.1. Using these results we reformulate the original optimization problem into a tractable chance-constrained one based on data in Section 3.2, with our work's main result, Theorem 8, providing the safety guarantees of its solution. For the sake of brevity the indices t, i, and j are omitted from the random variable d_{ij}^t and its moments.

3.1. Moment concentration inequalities

Let us consider having access to N_s i.i.d. samples $d_1, \ldots, d_{N_s}$ of d. We form the sample mean and covariance estimates:

$$\hat{\mu} = \frac{1}{N_s}\sum_{i=1}^{N_s} d_i\,, \tag{14a}$$

$$\hat{\Sigma} = \frac{1}{N_s - 1}\sum_{i=1}^{N_s}(d_i - \hat{\mu})(d_i - \hat{\mu})^\top\,. \tag{14b}$$

Assumption 3. $\hat{\Sigma}$ is positive definite, i.e. $\hat{\Sigma} \succ 0$.[2]

The concentration bound r_1 of Lemma 4 corresponding to the μ is identical to the one derived in Lefkopoulos and Kamgarpour (2019, Lemma 1), and hence its proof is omitted. The new concentration bound r_2 of Lemma 5 corresponding to Σ improves the one in Lefkopoulos and Kamgarpour (2019, Lemma 2).

Lemma 4 (*From Lefkopoulos & Kamgarpour, 2019*)**.** *Under Assumptions 2–3 and using the sample estimates (14a) and (14b), with probability $1-\beta$, $\beta \in (0,1)$:*

$$\left\|\mu - \hat{\mu}\right\|_2 \leq r_1 := \sqrt{\frac{T^2_{n,N_s-1}(1-\beta)}{N_s \lambda_{\min}(\hat{\Sigma}^{-1})}}\,, \tag{15}$$

where $T^2_{a,b}(p)$ denotes the p-th quantile of the Hotelling's T-squared distribution (Hotelling, 1931) with parameters a and b.

Lemma 5. *Under Assumptions 2–3 and using the sample estimates (14a) and (14b), define the constant:*

$$r_2 := \max\left\{\left|1 - \frac{N_s - 1}{\chi^2_{N_s-1,1-\beta/2}}\right|, \left|1 - \frac{N_s - 1}{\chi^2_{N_s-1,\beta/2}}\right|\right\}, \tag{16}$$

where $\chi^2_{k,p}$ is the p-th quantile of the chi-squared distribution with k degrees of freedom. Then, with probability $1-\beta$, $\beta \in (0,1)$:

$$\left|x^\top(\Sigma - \hat{\Sigma})x\right| \leq x^\top \hat{\Sigma} x r_2. \tag{17}$$

Proof. It is known (Mardia, Kent, & Bibby, 1979) that $(N_s-1)\hat{\Sigma}$ follows a Wishart distribution, i.e.:

$$(N_s - 1)\hat{\Sigma} \sim W_n(\Sigma, N_s - 1)\,, \tag{18}$$

where $W_n(\Sigma, N_s - 1)$ denotes the Wishart distribution with positive definite scale matrix Σ and $N_s - 1$ degrees of freedom associated with an n-variate Gaussian distribution. If $A \in \mathbb{R}^{n\times n}$

[2] Assumption 3 is not reasonable for a very small number of samples $N_s < n$, since the sample covariance matrix is a sum of N_s rank-1 matrices. In the trajectory generation application, however, it is reasonable to assume that N_s is sufficiently larger than n as will be discussed in the examples.

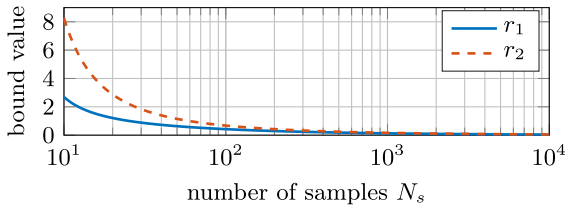

Fig. 1. Concentration bounds r_1 (blue) and r_2 (red) for a varying number of samples N_s, for $n = 3$ and $\beta = 10^{-3}$. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

follows a Wishart distribution $W_n(\Sigma, m)$ and $x \in \mathbb{R}^n$ is a nonzero vector, then (Rao, 1965, p. 535):

$$x^\top A x \sim x^\top \Sigma x \chi^2_m \,, \tag{19}$$

where χ^2_m is the chi-squared distribution with m degrees of freedom. Combining (18) and (19) we obtain:

$$x^\top \hat{\Sigma} x \sim \frac{1}{N_s - 1} x^\top \Sigma x \chi^2_{N_s-1} \,. \tag{20}$$

Hence, in a similar fashion to Krishnamoorthy (2006, p. 133), we can construct the $1-\beta$ confidence interval for (20):

$$x^\top \Sigma x \in \left[\frac{(N_s-1)x^\top \hat{\Sigma} x}{\chi^2_{N_s-1,1-\beta/2}}, \frac{(N_s-1)x^\top \hat{\Sigma} x}{\chi^2_{N_s-1,\beta/2}} \right], \tag{21}$$

holds with probability $1-\beta$. By subtracting $x^\top \hat{\Sigma} x$ from (21) it follows that:

$$x^\top(\Sigma - \hat{\Sigma})x \in \left[\frac{(N_s-1)x^\top \hat{\Sigma} x}{\chi^2_{N_s-1,1-\beta/2}} - x^\top \hat{\Sigma} x, \frac{(N_s-1)x^\top \hat{\Sigma} x}{\chi^2_{N_s-1,\beta/2}} - x^\top \hat{\Sigma} x \right]$$

$$= \left[x^\top \hat{\Sigma} x \left(\frac{(N_s-1)}{\chi^2_{N_s-1,1-\beta/2}} - 1 \right), x^\top \hat{\Sigma} x \left(\frac{(N_s-1)}{\chi^2_{N_s-1,\beta/2}} - 1 \right) \right]$$

$$= x^\top \hat{\Sigma} x \left[\frac{(N_s-1)}{\chi^2_{N_s-1,1-\beta/2}} - 1, \frac{(N_s-1)}{\chi^2_{N_s-1,\beta/2}} - 1 \right],$$

from which the statement of the lemma readily follows.

Remark 6. The concentration bounds r_1 and r_2 of (15) and (16) asymptotically converge to zero as the number of samples N_s grows, as illustrated in Fig. 1.

3.2. Robustifying the chance constraints

Using Lemmas 4 and 5 on the concentration of the sample moments around their true values, we conservatively reformulate a chance constraint as follows.

Lemma 7. *Under Assumptions 2–3 and using the sample estimates (14a) and (14b), the chance constraint:*

$$\Pr(d^\top \tilde{x} + Mz > 0) \geq 1 - \epsilon \,, \tag{22}$$

holds with a probability of at least $1-2\beta$, provided that:

$$\Psi^{-1}(1-\epsilon)\sqrt{1+r_2} \left\| \hat{\Sigma}^{1/2} \tilde{x} \right\|_2 + r_1 \left\| \tilde{x} \right\|_2 \leq \hat{\mu}^\top \tilde{x} + Mz \,. \tag{23}$$

Proof. In Section 2.3 we established that the chance constraint (22) is equivalent to:

$$\Psi^{-1}(1-\epsilon) \left\| \Sigma^{1/2} \tilde{x} \right\|_2 - \mu^\top \tilde{x} \leq Mz \,. \tag{24}$$

According to Lemmas 4 and 5 the bounds (15) and (17) each hold with probability $1-\beta$ and thus hold jointly[3] with probability $(1-\beta)^2$. Thus, with probability $(1-\beta)^2 > 1-2\beta$:

$$\begin{aligned}
&\Psi^{-1}(1-\epsilon) \left\| \Sigma^{1/2} \tilde{x} \right\|_2 - \mu^\top \tilde{x} \\
&= \Psi^{-1}(1-\epsilon)\sqrt{\tilde{x}^\top \hat{\Sigma} \tilde{x} + \tilde{x}^\top (\Sigma - \hat{\Sigma}) \tilde{x}} - (\mu - \hat{\mu})^\top \tilde{x} - \hat{\mu}^\top \tilde{x} \\
&\leq \Psi^{-1}(1-\epsilon)\sqrt{\tilde{x}^\top \hat{\Sigma} \tilde{x} + \tilde{x}^\top \hat{\Sigma} \tilde{x} r_2} + \left\| \mu - \hat{\mu} \right\|_2 \left\| \tilde{x} \right\|_2 - \hat{\mu}^\top \tilde{x} \\
&\leq \Psi^{-1}(1-\epsilon)\sqrt{1+r_2} \left\| \hat{\Sigma}^{1/2} \tilde{x} \right\|_2 + r_1 \left\| \tilde{x} \right\|_2 - \hat{\mu}^\top \tilde{x}
\end{aligned}$$

where the first inequality is an application of the Cauchy–Schwarz inequality and the second inequality follows from Lemmas 4 and 5.

We are now ready to address Problem (6a) given sample data on the position of obstacles.

$$\min_{\boldsymbol{u}, z^t_{ij}} \quad J(x_0, \boldsymbol{u}) \tag{25a}$$

$$\text{s.t} \quad \boldsymbol{x}, \boldsymbol{u} \text{ satisfy (1) with initial state } x_0 \tag{25b}$$

$$\boldsymbol{u} \in \mathcal{U} \tag{25c}$$

$$\Psi^{-1}(1-\epsilon^t_{ij})\sqrt{1+r_2} \left\| (\hat{\Sigma}^t_{ij})^{1/2} \tilde{x}_t \right\|_2 + r^t_{1,ij} \left\| \tilde{x}_t \right\|_2 \leq \hat{\mu}^{t\top}_{ij} \tilde{x}_t + M z^t_{ij} \tag{25d}$$

$$\sum_{i=1}^{F_j} z^t_{ij} = F_j - 1 \tag{25e}$$

where constraint (25d) must hold for all $t \in \mathbb{N}[1, N], j \in \mathbb{N}[1, N_o], i \in \mathbb{N}[1, F_j]$ and constraint (25e) must hold for all $t \in \mathbb{N}[1, N], j \in \mathbb{N}[1, N_o]$.

Theorem 8. *Under Assumptions 2–3, using the sample estimates (14a) and (14b), and the risk allocation (11), a solution to Problem (25a) is a feasible solution to Problem (6a) with a probability of at least $1-2\beta NN_o$. Furthermore, the solution of Problem (25a) asymptotically converges to the solution of the exact moment problem, with constraint (13) instead of (25d), as the number of available samples N_s converges to infinity.*

Proof. According to Lemma 7 each constraint of (25d) implies the corresponding constraint (7) with probability $1-2\beta$. By requiring that this implication holds jointly for all $k = NN_o$ non-vacuous constraints, and noting that $(1-2\beta)^k > 1-2\beta k$,[4] we can conclude that a solution to Problem (25a) is a feasible solution to Problem (6a) with a probability of at least $1-2\beta NN_o$. We do not consider the vacuous constraints by applying the same reasoning as the proof of Lemma 1 but for β instead of ϵ^t_{ij}. The asymptotic convergence of the solution to the case of exact moment knowledge follows from the fact that $\hat{\mu} \to \mu$, $\hat{\Sigma} \to \Sigma$ as $N_s \to \infty$ and Remark 6.

4. Receding horizon implementation

In a receding horizon framework, we consider the dynamics of an obstacle and measurements of its state being made available at each planning stage. We consider a time-varying model so as to incorporate the possibility of nonlinear dynamics through linearization around a nominal trajectory:

$$\chi_{t+1} = E_t \chi_t + F_t + w_t \,, \tag{26a}$$

[3] By Cochran's theorem, for Gaussian distributions the sample mean $\hat{\mu}$ and the sample covariance $\hat{\Sigma}$ are independent (Krishnamoorthy, 2006).

[4] By considering the function $f(\beta) = (1-2\beta)^k + 2k\beta - 1$ and verifying that $f(0) = 0$ and $f'(\beta) > 0$ for $k \geq 1$ and $0 < \beta < 1/2$.

$$y_t = H_t \chi_t + v_t \,, \tag{26b}$$

where $\chi_t \in \mathbb{R}^{n_\chi}$ is the obstacle's state, $\chi_0 \sim \mathcal{N}(\mu_{\chi_0}, \Sigma_{\chi_0})$ is the initial state, $y_t \in \mathbb{R}^{n_y}$ is the measurement and $w_t \sim \mathcal{N}(0, \Sigma_{w_t})$, $v_t \sim \mathcal{N}(0, \Sigma_{v_t})$ are white noise signals. The obstacle is represented as a polyhedron with G faces, with each face described by:

$$\{x \,|\, {a_i^t}^\top x + {c_i^t}^\top \chi_t + d_i = 0\} \,, \tag{27}$$

where $t \in \mathbb{N}[1, N]$, $i \in \mathbb{N}[1, G]$, the constants a_i^t, c_i^t and d_i are known based on the obstacle's shape. The constraint that x_t must be outside of the obstacle is written as:

$$\overline{\mathcal{O}^t} := \left\{ x_t \in \mathbb{R}^{n_x} : \vee_{i=1}^{G} {a_i^t}^\top x_t + {c_i^t}^\top \chi_t + d_i > 0 \right\} \,, \tag{28}$$

which is of type (3), with a_{ij}^t now being deterministic.

4.1. Estimation and prediction of obstacle's motion

We use a Kalman filter to estimate the position of the obstacle (Kalman, 1960). Let $\hat{\chi}_{t|t-1}$ and $\hat{\chi}_{t|t}$ be the a priori and a posteriori state estimates at time t, with $\hat{\Sigma}_{t|t-1}$ and $\hat{\Sigma}_{t|t}$ the corresponding error covariances. The prediction step of the Kalman filter is:

$$\hat{\chi}_{t|t-1} = E_{t-1}\hat{\chi}_{t-1|t-1} + F_{t-1} \,, \tag{29a}$$

$$\hat{\Sigma}_{t|t-1} = E_{t-1}\hat{\Sigma}_{t-1|t-1}{E_{t-1}}^\top + \Sigma_{w_{t-1}} \,, \tag{29b}$$

with the initialization $\hat{\chi}_{0|0} = \mu_{\chi_0}$ and $\hat{\Sigma}_{0|0} = \Sigma_{\chi_0}$. The update step of the Kalman filter is:

$$K_t = \hat{\Sigma}_{t|t-1} + H_t^\top \left(\Sigma_{v_t} + H_t \hat{\Sigma}_{t|t-1} H_t^\top \right)^{-1} \,, \tag{30a}$$

$$\hat{\chi}_{t|t} = \hat{\chi}_{t|t-1} + K_t \left(y_t - H_t \hat{\chi}_{t|t-1} \right) \,, \tag{30b}$$

$$\hat{\Sigma}_{t|t} = (I - K_t H_t)\, \hat{\Sigma}_{t|t-1} \,. \tag{30c}$$

In order to plan safe trajectories we need to predict the future positions of the obstacle with quantifiable confidence. Although this prediction could be done by propagating estimates $\hat{\chi}_{\tau|\tau}$, $\hat{\Sigma}_{\tau|\tau}$ using (29), in practice the moments of w_t will not be known exactly and thus we cannot quantify the probability of constraint violation. To this end, we utilize past sample data to get a confidence bound on the predicted trajectory of the obstacle.

Specifically, we assume that we have N_s samples of w_t available for the next N steps of the planning horizon, i.e. samples $w_t^{(1)}$, $\ldots$, $w_t^{(N_s)}$ for all $t \in \mathbb{N}[1, N]$, alternatively denoted as $\{w_t^{(i)}\}_{i=1}^{N_s}$. The availability of these samples is justified either through the use of a generative model (as done for example in Lefkopoulos & Kamgarpour, 2019, Section IV) or by considering that the obstacle maneuvers are usually extracted through a clustering algorithm based on previously collected sample data (Carvalho, 2016; Lefèvre, Vasquez, & Laugier, 2014). In both cases a set of trajectory data $\{\chi_t^{(i)}\}_{i=1}^{N_s}$ is available, from which the model uncertainty $\{w_t^{(i)}\}_{i=1}^{N_s}$ can be calculated by comparison with the nominal response of linear model (26a). We use these samples directly to forward simulate the process model (26a) N steps and to estimate the uncertainty moments required in the Kalman filter equations above. Consequently, we can form the sample estimates (14a) and (14b) of the obstacle's predicted state. The results of Section 3 can then be used to robustify against the uncertainty of the dynamic obstacle's future trajectory.

4.2. Closed loop

Given a planning horizon N, let $u_{t-1|\tau}$ and $x_{t|\tau}$, with $t \in \mathbb{N}[\tau + 1, \tau + N]$, denote the optimization problem solutions at time τ. The Chance-Constrained Receding Horizon (CCRH) is presented in Algorithm 1.

Algorithm 1 Chance-Constrained Receding Horizon

1: Given samples $\{w_t^{(i)}\}_{i=1}^{N_s}$ for $t \in \mathbb{N}[0, 2N-1]$ from past data or generative model
2: **for** $\tau = 0$ to $N-1$ **do**
3: **if** $\tau = 0$ **then**
4: initialize $\hat{\chi}_{0|0}$, $\hat{\Sigma}_{0|0}$ to μ_{χ_0}, Σ_{χ_0}
5: **else**
6: obtain measurement y_τ from Eq. (26b)
7: calculate $\hat{\chi}_{\tau|\tau}$, $\hat{\Sigma}_{\tau|\tau}$ from Eq. (30b)
8: propagate samples $\{\chi_\tau^{(i)}\}_{i=1}^{N_s}$ using $\{w_t^{(i)}\}_{i=1}^{N_s}$ and Eq. (26a) to get $\{\chi_t^{(i)}\}_{i=1}^{N_s}$ for $t \in \mathbb{N}[\tau+1, \tau+N]$
9: calculate sample moments $\hat{\chi}_{t|\tau}$, $\hat{\Sigma}_{t|\tau}$ using Eq. (14a), Eq. (14b) for $t \in \mathbb{N}[\tau+1, \tau+N]$
10: measure x_τ and solve Problem (25) using Eq. (11)
11: apply first input $u_{\tau|\tau}$

Assumption 9. Each of the N optimization problems of Algorithm 1 are feasible.

Theorem 10. *The sequence of states $x_{1|0}, x_{2|1}, \ldots, x_{N|N-1}$ that is obtained as a result of Algorithm 1 satisfies:*

$$\Pr(\wedge_{t=1}^{N} x_{t|t-1} \in \mathcal{X}_t) \geq 1 - \epsilon \,, \tag{31}$$

with a probability according to Theorem 8.

Proof. In Algorithm 1 every problem is solved using the risk allocation (11), which implies that:

$$\Pr(x_{\tau+t|\tau} \in \mathcal{X}_{\tau+t}) \geq 1 - \frac{\epsilon}{N}, \quad \forall t \in \mathbb{N}[1, N] \,, \tag{32}$$

holds for any $\tau \in \mathbb{N}[0, N-1]$. Summing up the probabilities of the complement events of (32) we get:

$$\sum_{\tau=0}^{N-1} \Pr(x_{\tau+1|\tau} \notin \mathcal{X}_{\tau+1}) \leq \epsilon \,. \tag{33}$$

Using Boole's inequality and (33) we obtain:

$$\Pr(\vee_{\tau=0}^{N-1} x_{\tau+1|\tau} \notin \mathcal{X}_{\tau+1}) \leq \epsilon \,, \tag{34}$$

from the complement of which the statement of the theorem follows.

5. Simulations

We will present two simulations to demonstrate the applicability and effectiveness of our approach. All computations were carried out on an Intel i5 CPU at 2.50 GHz with 8 GB of memory using YALMIP (Lofberg, 2004) and CPLEX (International Business Machines Corporation (IBM), 2017).

5.1. Open-loop case study

We consider a robot with the objective to reach a desired target position while avoiding two *uncertain* walls that block its path. The robot motion model is $\dot{x} = u$, where the state $x \in \mathbb{R}^2$ is its position (x_1, x_2) and is constrained to the set $\mathcal{X} := \{x \in \mathbb{R}^2 : 0 \leq x_i \leq 9, i = \mathbb{N}[1, 2]\}$. The velocities (u_1, u_2) are the control inputs and are constrained to the set $\mathcal{U} := \{u \in \mathbb{R}^2 : \|u\|_\infty \leq 1\,\text{m/s}\}$. The objective is to minimize the distance between the robot's position and the target position $x_d = [8, 7]^\top$. The initial state is $x_0 = [1, 1]^\top$.

The coefficients describing the two walls $d_1^t, d_2^t \in \mathbb{R}^3$ are uncertain and at time t are distributed as:

$$d_1^t \sim \mathcal{N}([-1, 0, 2]^\top, 0.001 I_3) \,, \tag{35a}$$

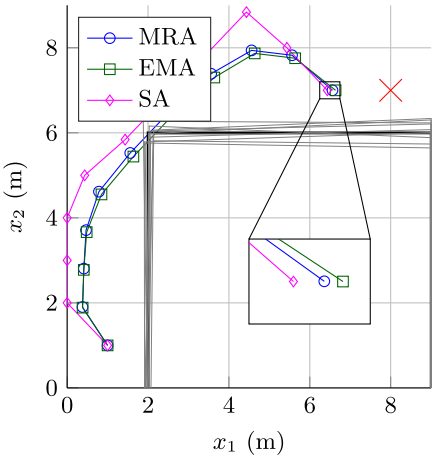

Fig. 2. One simulation of the robot's trajectories for the MRA (blue circles), EMA (green squares) and SA (magenta diamonds). The expected positions of the two uncertain walls are displayed (black) along with 10 exemplary wall samples (gray) and the terminal target position (red cross). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

$$d_2^t \sim \mathcal{N}([0, 1, 6]^\top, 0.001 I_3)\,, \tag{35b}$$

for all $t \in \mathbb{N}[1, N]$. In order to deal with the uncertain walls, we enforce the joint chance constraint:

$$\Pr(\wedge_{t=1}^{N} \vee_{i=1}^{2} [x_{1,t}, x_{2,t}, 1] d_i^t > 0) \geq 1 - \epsilon\,. \tag{36}$$

We consider the moments of the distributions (35a) and (35b) to be *unknown*, and thus draw N_s i.i.d. samples from each one and reformulate the resulting problem as per the Moments Robust Approach (MRA) of Section 3. For comparison reasons we also solve the problem assuming perfect knowledge of the moments, henceforth called "Exact Moments Approach" (EMA), and using the Scenario Approach (SA).

The dynamics are discretized with sampling time $T_s = 1\,\text{s}$ for a planning horizon of $N = 10$ (i.e. 10 s). The joint chance constraint (36) is imposed with $\epsilon = 0.05$ and $\beta = 10^{-3}$ for the whole horizon. The number of samples used is $N_s^{\text{MRA}} = N_s^{\text{SA}} = 1259$.[5] The resulting optimization problem is a MISOCP with 40 continuous variables, 20 binary variables and 132 constraints in the case of the MRA/EMA and a MIQP[6] with 40 continuous variables, 20 binary variables and 25 292 constraints in the case of the SA. The problem is solved with all methods 100 different times, with different realizations of the disturbance. The risk distribution iterative algorithm of Jha et al. (2018) is used for the MRA and the EMA to improve the risk allocation with a small increase in computation time.

The solutions to one instance of the problem are presented in Fig. 2. Under moment uncertainty (MRA) the robot chooses a slightly wider maneuver compared to perfect moment knowledge (EMA), in order to avoid the uncertain position of the walls, and has a final position slightly farther away from the target position. The SA on the other hand performs a wider maneuver and ends up even farther away from the desired target. As expected, robustifying against moment uncertainty using the MRA produced a slightly worse solution, to account for the increased wall uncertainty compared to the EMA.

We evaluated the empirical violation probability of each method through Monte Carlo simulations using 10^5 new realizations of the walls' unknown positions. The MRA had a violation probability of approximately 1 % to 2 %, the EMA approximately 4.8 %, and the SA was the most conservative with a probability of 0 % to 1 %.

5.2. Closed-loop case study

We examine a closed-loop trajectory planning scenario. The controlled "ego car" is driving with an initial forward velocity of 19.44 m/s on the left lane of a two-lane highway when an uncontrolled "adversary car" driving on the right lane starts merging in front of the ego car with an unchanging orientation, under the assumption of a slow maneuver.

We model the ego car dynamics as a double integrator $\dot{x} = Ax + Bu$, where the state $x \in \mathbb{R}^4$ contains the two-dimensional position (x_1, x_2) of the car and the corresponding velocities (x_3, x_4). The accelerations (u_1, u_2) are the control inputs and are constrained to the set $\mathcal{U} := \{u \in \mathbb{R}^2 : |u_1| \leq 10\,\text{m/s}^2, |u_2| \leq 2\,\text{m/s}^2\}$. The objective is to minimize the deviation of the car's position from the center of the left lane and its velocity from 19.44 m/s, while also minimizing the use of the control inputs. The dynamics are discretized with sampling time $T_s = 0.2\,\text{s}$. The initial state is $x_0 = [0\,\text{m}, 2\,\text{m}, 19.44\,\text{m/s}, 0\,\text{m/s}]^\top$.

The adversary car follows the dynamics $\dot{\chi} = A\chi + B\nu$, where the state $\chi \in \mathbb{R}^4$ contains the position (χ_1, χ_2) of the car, which is measured, and the corresponding velocities (χ_3, χ_4). The adversary car longitudinal model is detected as one of constant velocity, corresponding to $\nu_1 = 0$. The lateral model is assumed to be a second-order response converging to the center of the left lane, corresponding to $\nu_2 = -[k_1, k_2][\chi_2, \chi_4]^\top$ where the gains are provided by a Linear–Quadratic Regulator (LQR). The dynamics are discretized with sampling time $T_s = 0.2\,\text{s}$, and process/measurement additive noise is introduced, resulting in the discrete form (26). The initial state χ_0, process noise w_t and measurement noise v_t are distributed according to:

$$\chi_0 \sim \mathcal{N}([0, -2, 19.44, 0]^\top, \text{diag}([0, 0, 1.23, 0.08]))\,, \tag{37a}$$

$$w_t \sim \mathcal{N}(\mathbb{0}, \text{diag}([0, 0, 0.04, 0.001]))\,, \tag{37b}$$

$$v_t \sim \mathcal{N}(\mathbb{0}, \text{diag}([1, 0.04]))\,, \tag{37c}$$

In order to deal with the uncertain pose χ_t of the adversary car, we enforce the joint chance constraint:

$$\Pr(\wedge_{t=1}^{N} \vee_{i=1}^{4} [x_{1,t}, x_{2,t}, 1] d_i^t > 0) \geq 1 - \epsilon\,. \tag{38}$$

We consider the moments of the distributions (37a), (37b) and (37c) to be *unknown*, and thus draw N_s i.i.d. samples from each one and reformulate the resulting problem as per the CCRH approach of Section 4.

A planning horizon of $N = 25$ (i.e. 5 s) is chosen. Both cars have length 4 m and width 2 m, which are taken into account for the relevant inequality constraints. We also enforce lane constraints. The joint chance constraint (38) is imposed with $\epsilon = 0.05$ and $\beta = 10^{-3}$ for the whole horizon. The number of samples drawn from the distributions of x_0 and is $N_s = 1000$ for every time step; obtainable for example from real-world data of lane change maneuvers of this kind (or a similar generative model). The resulting optimization problem is a MISOCP with 150 continuous variables, 100 binary variables and 379 constraints.

The simulation is presented in Fig. 3. The ego car brakes and steers towards the left, in order to avoid the merging adversary car whose position is uncertain. As a result of the initially high estimate uncertainty of the current and future positions of the

[5] The number of samples N_s^{SA} was chosen such that the SA would provide the safety and certainty guarantees prescribed (Sessa et al., 2018). The number of samples N_s^{MRA} was chosen to be the same for comparison reasons.

[6] Mixed-Integer Quadratic Program.

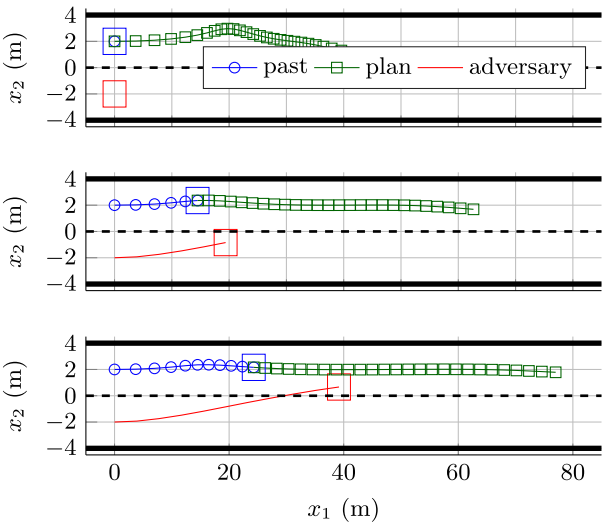

Fig. 3. Simulation of the ego car's trajectory. Pictured are time steps $t = 0$ s, 1 s and 2 s. At every time step the ego car's past closed-loop trajectory is displayed (blue circles), along with the future open-loop plan (green squares) and the adversary car's past trajectory (red). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

adversary car, the first open-loop solution calculated (in the first frame of Fig. 3) is to react sharply. In contrast, the closed-loop solution exhibits a less severe and more comfortable maneuver. The more aggressive maneuver of the initial open-loop solution incurs a higher cost of 3785.2, whereas the closed-loop solution has a lower cost of 1766.3. The improved performance of the closed-loop trajectory is due to the receding horizon implementation's ability to react to the adversary car's latest movement. As more measurements are acquired and passed through the Kalman filter's update steps, the estimates of the adversary's car position improve.

6. Conclusion

We tackled the problem of trajectory planning with uncertain polyhedral obstacles using chance-constrained optimization. We reformulated the problem into a deterministic and tractable form based on the uncertainty's moments. We estimated said moments from collected samples, resulting in noisy estimates. We derived tight and asymptotic concentration bounds on said estimates that were used to formulate a robust tractable optimization problem whose solution is feasible with regards to the original problem up to a prescribed confidence bound. We extended this framework to a receding horizon, including obstacles with partially known dynamics. We illustrated both open-loop and closed-loop results through two case studies. Our method produced safe trajectories using only finite sample moment estimates, outperforming previous solutions, with the closed-loop improving over the open-loop results.

Acknowledgments

The authors would like to thank I. Usmanova for helping with Lemma 5.

References

Blackmore, L., Hui Li, & Williams, B. (2006). A probabilistic approach to optimal robust path planning with obstacles. In *Proc. of the 2006 American control conference*.

Blackmore, Lars, & Ono, Masahiro (2009). Convex chance constrained predictive control without sampling. In *AIAA guidance, navigation, and control conference*.

Blackmore, L., Ono, M., & Williams, B. C. (2011). Chance-constrained optimal path planning with obstacles. *IEEE Transactions on Robotics and Automation*, *27*(6), 1080–1094.

Calafiore, G. C., & El Ghaoui, L. (2006). On distributionally robust chance-constrained linear programs. *Journal of Optimization Theory and Applications*, *130*(1), 1–22.

Carvalho, Ashwin M. (2016). *Predictive control under uncertainty for safe autonomous driving: Integrating data-driven forecasts with control design* (Ph.D. thesis), UC Berkeley.

Carvalho, Ashwin, Gao, Yiqi, Lefevre, Stéphanie, & Borrelli, Francesco (2014). Stochastic predictive control of autonomous vehicles in uncertain environments. In *12th international symposium on advanced vehicle control*.

Casella, George, & Berger, Roger L. (2001). *Statistical inference*. Cengage Learning.

Feng, D., Rosenbaum, L., & Dietmayer, K. (2018). Towards safe autonomous driving: Capture uncertainty in the deep neural network for lidar 3D vehicle detection. In *2018 21st international conference on intelligent transportation systems* (pp. 3266–3273).

Hotelling, Harold (1931). The generalization of student's ratio. *The Annals of Mathematical Statistics*, *2*(3), 360–378.

International Business Machines Corporation (IBM) (2017). IBM ILOG CPLEX optimization studio.

Jha, Susmit, Raman, Vasumathi, Sadigh, Dorsa, & Seshia, Sanjit A. (2018). Safe autonomy under perception uncertainty using chance-constrained temporal logic. *Journal of Automated Reasoning*, *60*(1), 43–62.

Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. *Journal of Fluids Engineering*, *82*(1), 35–45.

Kendall, Alex, & Gal, Yarin (2017). What uncertainties do we need in Bayesian deep learning for computer vision? In *Proceedings of the 31st international conference on neural information processing systems* (pp. 5574–5584).

Krishnamoorthy, K. (2006). *Handbook of statistical distributions with applications*. Chapman and Hall/CRC.

Lefèvre, Stéphanie, Vasquez, Dizan, & Laugier, Christian (2014). A survey on motion prediction and risk assessment for intelligent vehicles. *ROBOMECH Journal*, *1*(1), 1.

Lefkopoulos, V., & Kamgarpour, M. (2019). Using uncertainty data in chance-constrained trajectory planning. In *2019 European control conference* (pp. 2264–2269).

Lofberg, J. (2004). YALMIP : a toolbox for modeling and optimization in MATLAB. In *Proc. of the 2004 IEEE international conference on robotics and automation* (pp. 284–289).

Mardia, K., Kent, J., & Bibby, J. (1979). *Multivariate analysis*. Academic Press.

Michelmore, Rhiannon, Wicker, Matthew, Laurenti, Luca, Cardelli, Luca, Gal, Yarin, & Kwiatkowska, Marta (2019). Uncertainty quantification with statistical guarantees in end-to-end autonomous driving control. arXiv:1909.09884v1 [cs.LG].

Nemirovski, Arkadi, & Shapiro, Alexander (2007). Convex approximations of chance constrained programs. *SIAM Journal on Optimization*, *17*(4), 969–996.

Ono, M., Blackmore, L., & Williams, B. C. (2010). Chance constrained finite horizon optimal control with nonconvex constraints. In *Proceedings of the 2010 American control conference* (pp. 1145–1152).

Ono, Masahiro, Pavone, Marco, Kuwata, Yoshiaki, & Balaram, J. (2015). Chance-constrained dynamic programming with application to risk-aware robotic space exploration. *Autonomous Robots*, *39*(4), 555–571.

Raman, Vasumathi, Donzé, Alexandre, Sadigh, Dorsa, Murray, Richard M., & Seshia, Sanjit A. (2015). Reactive synthesis from signal temporal logic specifications. In *Proc. of the 18th international conference on hybrid systems: computation and control* (pp. 239–248).

Rao, C. R. (1965). *Linear statistical inference and its applications*. John Wiley and Sons, Inc.

Schouwenaars, Tom, Moor, Bart De, Feron, Eric, & How, Jonathan P. (2001). Mixed integer programming for multi-vehicle path planning. In *2001 European control conference* (pp. 2603–2608).

Sessa, Pier Giuseppe, Frick, Damian, Wood, Tony A., & Kamgarpour, Maryam (2018). From uncertainty data to robust policies for temporal logic planning. In *Proc. of the 21st international conference on hybrid systems: computation and control* (pp. 157–166).

Vitus, Michael P. (2012). *Stochastic control via chance constrained optimization and its application to unmanned aerial vehicles* (Ph.D. thesis, Stanford University.

Vitus, M. P., & Tomlin, C. J. (2011). On feedback design and risk allocation in chance constrained control. In *Proc. of the 50th IEEE conference on decision and control and european control conference* (pp. 734–739).

Vasileios Lefkopoulos holds an M.Sc. degree in Robotics, Systems and Control from ETH Zurich, Switzerland, and a Diploma degree in Electrical and Computer Engineering from the Aristotle University of Thessaloniki, Greece. His research interests include model predictive control and optimal control theory with applications in trajectory planning and control of autonomous systems.

Maryam Kamgarpour is with the faculty of Electrical and Computer Engineering at the University of British Columbia, Vancouver, Canada. She holds a Doctor of Philosophy in Engineering from the University of California, Berkeley and a Bachelor of Applied Science from University of Waterloo, Canada. Her research is on safe decision-making and control under uncertainty, game theory and mechanism design, mixed integer and stochastic optimization and control. Her theoretical research is motivated by control challenges arising in intelligent transportation networks, robotics, power grid systems and healthcare. She is the recipient of NASA High Potential Individual Award, NASA Excellence in Publication Award, and the European Union (ERC) Starting Grant.