

Perception-Aware Chance-Constrained Model Predictive Control for Uncertain Environments

Bonzanini, Angelo Domenico; Mesbah, Ali; Di Cairano, Stefano

TR2021-055 June 04, 2021

Abstract

We consider a known system that operates in an unknown environment, which is discovered by sensing and affects the known system through constraints. However, sensing quality is typically dependent on system operation. Thus, the control decisions should account for both the impact of control on sensing and the impact of sensing on control. Since the information acquired from sensing is of statistical nature, we develop a perception-aware chance-constrained model predictive control (PAC-MPC) strategy that leverages uncertainty propagation models to relate control and sensing decisions to the environment knowledge. We propose conditions for recursive feasibility and provide an overview of the stability properties in such a statistical framework. The performance of the proposed PAC-MPC is demonstrated on a case study inspired by an automated driving application.

American Control Conference (ACC) 2021

Perception-Aware Chance-Constrained Model Predictive Control for Uncertain Environments

Angelo D. Bonzanini, Ali Mesbah, and Stefano Di Cairano

Abstract—We consider a known system that operates in an unknown environment, which is discovered by sensing and affects the known system through constraints. However, sensing quality is typically dependent on system operation. Thus, the control decisions should account for both the impact of control on sensing and the impact of sensing on control. Since the information acquired from sensing is of statistical nature, we develop a perception-aware chance-constrained model predictive control (PAC-MPC) strategy that leverages uncertainty propagation models to relate control and sensing decisions to the environment knowledge. We propose conditions for recursive feasibility and provide an overview of the stability properties in such a statistical framework. The performance of the proposed PAC-MPC is demonstrated on a case study inspired by an automated driving application.

I. INTRODUCTION

Model predictive control (MPC) [1] is especially effective when accurate prediction models can be derived [2] such as in automotive, aerospace and robotics. However, even when accurate prediction models are available, significant uncertainty may still be present in the environment, e.g., other vehicles in autonomous driving, obstacles for drones, and human workers on a factory floor. Sensing devices, such as LIDAR, radar, and cameras, can be used to estimate the environment state, but such knowledge is still subject to uncertainty. In addition, the entire sensing process may depend on the decisions made by the controller. For instance, sensing quality may depend on the system state, e.g., due to its heading angle and distance from the target. Furthermore, the information received from sensors may arrive in vast amounts, which make it impossible to process it in its entirety in real-time. Thus, decisions have to be made on what information to process, i.e., where to focus the “attention” of the sensors, and how much to process it.

The reduction of the environment uncertainty through the sensing process is often partially dependent on the decisions of the controller, either directly due to sensing decisions, or indirectly due to how sensing is affected by the state trajectory. Thus, there is an interdependence between sensing and control. That is, the system should sense better where it is commanded to go, and it should be commanded to go where it can sense better. For example, if the system must be steered to an area in which the environment is highly uncertain, the controller may command to focus the sensing

on that area to reduce the uncertainty before approaching. Similarly, if the sensing quality depends on the system state, the controller may choose to modify the system trajectory to improve the sensing process.

Although most of the literature has focused on *active perception* and control as separate objectives, some research started to consider the two objectives simultaneously. Recently, active perception approaches within the control strategy have been proposed, whereby the controller seeks to allocate the sensing resources such that it strikes a balance between the control objective and reducing the uncertainty in the environment. In [3], [4], the control actions are selected for the exploration of an unknown environment while also maximizing localization accuracy, and in [5], distance-dependent measurement models are considered. In [6], perception-aware model predictive control (PAMPC) is introduced, whereby the goal of reference tracking is balanced with the goal of improving perception, where the latter amounts to maximizing the visibility of a point of interest. Learning-based controllers have also been considered to determine the approximate control inputs, while estimating the uncertainty of the learned controller [7].

In this paper, we consider a known system in an uncertain environment, which affects the system through the constraints. The sensing process that estimates the environment state depends on the system states and inputs and is corrupted by measurement noise, which results in a stochastic environment estimate. We propose a perception-aware chance-constrained MPC (PAC-MPC) that optimizes the control objective and guarantees constraint satisfaction in probability by accounting for the uncertainty in the estimate of the environment and for the impact of the control decisions on it. By incorporating active perception, PAC-MPC exploits sensing more effectively to improve control performance. This, in turn, reduces the uncertainty and enables less conservative control.

Notation: \mathbb{R} , \mathbb{R}_{0+} , \mathbb{R}_+ , \mathbb{Z} , \mathbb{Z}_{0+} , and \mathbb{Z}_+ are the sets of real, nonnegative real, positive real, integer, nonnegative integer, and positive integer numbers, respectively. Intervals are denoted by $\mathbb{Z}_{[a,b)} = \{z \in \mathbb{Z} : a \leq z < b\}$, where \mathbb{Z} can be substituted for any other set. For vectors x , y , we denote the i -th component by $[x]_i$, and the stacking by $(x, y) = [x^\top \ y^\top]^\top$. $\|\cdot\|_p$ denotes the p -norm and $\|x\|_Q^2 = x^\top Qx$. The Cholesky decomposition of Q is $Q^{1/2}$, while the trace is $\text{tr}(Q)$. $\mathbb{P}[A]$ is the probability of event A . For a random vector x , $\mathbb{E}[x] = \mu^x$ is the expectation, Σ^x the covariance matrix, and $x \sim \mathcal{N}(\mu^x, \Sigma^x)$ denotes it is normally distributed. A function $\alpha : \mathbb{R}_{0+} \rightarrow \mathbb{R}_{0+}$ is of class

A.D. Bonzanini and A. Mesbah are with the Department of Chemical and Biomolecular Engineering at the University of California, Berkeley. {adbonzanini, mesbah}@berkeley.edu

S. Di Cairano is with Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA. dicairano@ieee.org

A.D. Bonzanini was an intern at MERL during this work.




Fig. 1. Schematic of system, environment, and control architecture.

\mathcal{K} if it is continuous, strictly increasing, $\alpha(0) = 0$, and of class \mathcal{K}_∞ if $\lim_{c \rightarrow \infty} \alpha(c) = \infty$, also.

II. MODELING AND PROBLEM DEFINITION

Consider the discrete-time linear system

$$x_{k+1} = Ax_k + Bu_k, \quad (1a)$$

$$y_k = Ex_k, \quad (1b)$$

where (A, B) is stabilizable, $x \in \mathbb{R}^{n_x}$ is the state vector, $u \in \mathbb{R}^{n_u}$ is the input vector, and $y \in \mathbb{R}^{n_y}$ is the performance output vector. System (1) is subject to constraints

$$x \in \mathcal{X} = \{x : H_x x \leq b_x\}, u \in \mathcal{U} = \{u : H_u u \leq b_u\}. \quad (2)$$

While (1) and (2) are perfectly known, the system operates in an uncertain environment, which imposes additional constraints on (1). The environment is represented by a vector $w \in \mathbb{R}^{n_w}$, which we simply call *the environment*. We model w_k at any $k \in \mathbb{Z}_{0+}$ as a Gaussian random vector, $w_k \sim \mathcal{N}(\mu_k^w, \Sigma_k^w)$. Due to w being random, the constraints imposed by the environment on (1) are formulated as individual chance constraints (ICCs) [8]

$$\mathbb{P}[h_s^\top x_k + \eta_s^\top w_k \leq b_s] \geq 1 - \varepsilon_s, s \in \mathbb{Z}_{[1, n_s]}, \quad (3)$$

where ε_s is the allowed violation probability for the s^{th} ICC.

We model the environment dynamics and its measurements as

$$w_{k+1} = A^w w_k + B^w \xi_k, \quad (4a)$$

$$\psi_k = C^w(x_k, u_k)w_k + D^w(x_k, u_k)\zeta_k, \quad (4b)$$

where $\psi_k \in \mathbb{R}^{n_\psi}$ is the measurement of the environment, $\xi_k \sim \mathcal{N}(\mu_k^\xi, \Sigma_k^\xi)$ is the process noise, and $\zeta_k \sim \mathcal{N}(\mu_k^\zeta, \Sigma_k^\zeta)$ is the measurement noise. In (4b), ψ depends on the state and input vectors of (1), which allows to represent a variable sensing quality depending on x, u . In what follows $\xi_k \sim \mathcal{N}(0, I)$, $\zeta_k \sim \mathcal{N}(0, I)$, since standard steps can be applied to (4) to reformulate into this case.

The estimate of w is denoted by a Gaussian random variable $\hat{w} \sim \mathcal{N}(\mu^{\hat{w}}, \Sigma^{\hat{w}})$, where

$$(\mu_{k+1}^{\hat{w}}, \Sigma_{k+1}^{\hat{w}}) = g(\mu_k^{\hat{w}}, \Sigma_k^{\hat{w}}, \psi_k, x_k, u_k), \quad (5)$$

and g is a general estimator, which includes a model of (4) and depends on the state and input of (1) due to (4b).

The problem tackled in this paper is to control (1) such that the performance output y tracks a reference $r \in \mathbb{R}^{n_y}$ while

enforcing constraints (2) and chance constraints (3) based on the estimate \hat{w} of the environment w . Figure 1 shows a schematic of the control architecture, the environment, and the system. In the optimal control problem (OCP), we propagate the uncertainty moments over the prediction horizon by a model \hat{g} of the estimator g , where, ideally, $\hat{g} = g$, although in practice \hat{g} may be an approximation of g . At each time step k , given x_k and the estimate \hat{w}_k of w_k , the PAC-MPC solves the OCP

$$J^*(x_k, \mathbf{r}_k) = \min_{U_k} J(x_k, U_k, \mathbf{r}_k) \quad (6a)$$

$$\text{s.t. } x_{j+1|k} = Ax_{j|k} + Bu_{j|k} \quad (6b)$$

$$\hat{w}_{j|k} \sim \mathcal{N}(\mu_{j|k}^{\hat{w}}, \Sigma_{j|k}^{\hat{w}}) \quad (6c)$$

$$(\mu_{j+1|k}^{\hat{w}}, \Sigma_{j+1|k}^{\hat{w}}) = \hat{g}(\mu_{j|k}^{\hat{w}}, \Sigma_{j|k}^{\hat{w}}, \psi_{j|k}, x_{j|k}, u_{j|k}) \quad (6d)$$

$$(x_{j|k}, u_{j|k}) \in \mathcal{X} \times \mathcal{U} \quad (6e)$$

$$\mathbb{P}[h_s^\top x_{j|k} + \eta_s^\top \hat{w}_{j|k} \leq b_s] \geq 1 - \varepsilon_s, s \in \mathbb{Z}_{[1, n_s]} \quad (6f)$$

$$x_{0|k} = x_k, \hat{\mu}_{0|k}^{\hat{w}} = \mu_k^{\hat{w}}, \Sigma_{0|k}^{\hat{w}} = \Sigma_k^{\hat{w}}, \quad (6g)$$

where $N \in \mathbb{Z}_+$ is the prediction horizon, $U_k = (u_{0|k}, \dots, u_{N|k})$ is the control sequence and $\mathbf{r}_k = (r_{0|k}, \dots, r_{N|k})$ is the reference trajectory, which is anticipatively known, and (6d) propagates the uncertainty, i.e., the distribution of \hat{w} , along the prediction horizon. Uncertainty propagation via \hat{g} , which includes a dependency on the measurement equation (4b) and as a consequence on the state and input of (1), enables predicting the impact of control actions onto the estimate uncertainty, hence enabling perception-aware control. Therefore, using (6d), the controller makes decisions that may reduce the predicted uncertainty, which impacts the constraints (6f), in order to optimize the cost (6a) and hence the control performance.

Remark 1: Since our focus is on the uncertainty due to environment sensing, here (1) is assumed known. For handling model uncertainty our approach can be merged with the methods in [9]–[11]. ■

In the next sections we propose designs for (6) that yield a tractable formulation and provide conditions for stability and recursive feasibility in the stochastic setting.

III. UNCERTAINTY PROPAGATION AND CHANCE CONSTRAINTS

Next we develop designs for the uncertainty propagation (6d), the chance constraints (6f), and the cost function (6a). In what follows, we use the short-hand notation $C_k^w = C^w(x_k, u_k)$, $D_k^w = D^w(x_k, u_k)$.

A. Model-based Uncertainty Propagation

A first approach for the uncertainty propagation (6d) is based on the availability of a model for (5). Specifically, we consider (5) to be a linear estimator with gain $L_k = L(x_k, u_k)$ that possibly depends on the states and inputs of (1). This includes the Luenberger-type observers, including the stationary Kalman filter, and can be extended immediately to the Kalman filter, where the gain is time-varying yet completely predictable. In this case, the mean

and covariance of \hat{w} are propagated by

$$\mu_{j+1|k}^{\hat{w}} = \Lambda_{j|k} \mu_{j|k}^{\hat{w}} - L_{j|k} \psi_{j|k} \quad (7a)$$

$$\Sigma_{j+1|k}^{\hat{w}} = \Lambda_{j|k} \Sigma_{j|k}^{\hat{w}} \Lambda_{j|k}^\top + Q_{j|k} + R_{j|k}, \quad (7b)$$

where $\Lambda_k := (A^w + L_k C_k^w)$, $Q_k := B^w (B^w)^\top$, and $R_k := L_k D_k^w (L_k D_k^w)^\top$.

Remark 2: In (7a), $\psi_{j|k}$ is the predicted measurement, which can be selected in several ways, $\psi_{j|k} = \psi_{0|k} = \psi_k$ or $\psi_{j|k} = C^w \mu_{j|k}^{\hat{w}}$, where, in the second case, the mean evolves in open-loop yet the estimator still affects the covariance. ■

Remark 3: The gain dependency on x and u allows for modeling the variability of the sensing process as a function of the system state and sensing decisions. For instance, when the measurement noise decreases due to a shorter sensing distance, or when the sensing process focuses more resources on the point of interest, the estimator gain can be higher. ■

B. Learning-based Uncertainty Propagation

Instead of deriving model-based uncertainty propagation, if historical data is available, machine learning (ML) models can be learned to directly propagate the mean and covariance in (6d). Here, we choose Gaussian Process regression (GPR) because of its non-parametric form, which lends itself well to learning arbitrary functions of states and inputs, and yields the uncertainty associated with its predictions [12]. Given M training points from previously collected system states and inputs, as well as measurements and estimates of the environment, let $X_j = (\mu_j^{\hat{w}}, \Sigma_j^{\hat{w}}, x_j, u_j, \psi_j)$ and $Y_j = (\mu_{j+1}^{\hat{w}}, \Sigma_{j+1}^{\hat{w}})$, and define the training dataset as

$$\mathcal{D} = \left\{ \mathbf{X} = [X_0, \dots, X_{M-1}]^\top, \quad \mathbf{Y} = [Y_0, \dots, Y_{M-1}]^\top \right\}.$$

Given a test point X_j , the Gaussian posterior distribution mean and covariance conditioned to \mathcal{D} are

$$[\mu_{j+1}^{\hat{w}}(X_j)]_a = [m(X_j)]_a + k_{X_j \mathbf{X}}^a (k_{\mathbf{X} \mathbf{X}}^a + \sigma_a^2 I)^{-1} ([\mathbf{Y}]_a - [m(\mathbf{X})]_a), \quad (8a)$$

$$[\Sigma_{j+1}^{\hat{w}}(X_j)]_a = k_{X_j X_j}^a - k_{X_j \mathbf{X}}^a (k_{\mathbf{X} \mathbf{X}}^a + \sigma_a^2 I)^{-1} k_{\mathbf{X} X_j}^a, \quad (8b)$$

where $a = \{1, \dots, 2n_w\}$ is the a^{th} dimension of \mathbf{Y} , σ_a^2 is the a^{th} diagonal term of the noise covariance of the training outputs, $m^a(\cdot)$ is the mean function of the GP prior, and $k_{X_1 X_2}^a = k^a(X_1, X_2)$ is the kernel function, such as the squared exponential kernel function [12]. Concatenating the individual predictions in (8), the GP-based predictor is

$$\hat{w}_{j+1|k} \sim \mathcal{N} \left(\mu_{j+1|k}^{\hat{w}}(\mu_{j|k}^{\hat{w}}, \Sigma_{j|k}^{\hat{w}}, x_{j|k}, u_{j|k}, \psi_{j|k}), \Sigma_{j+1|k}^{\hat{w}}(\mu_{j|k}^{\hat{w}}, \Sigma_{j|k}^{\hat{w}}, x_{j|k}, u_{j|k}, \psi_{j|k}) \right).$$

Remark 4: An advantage of GPR is that the worst-case quality of the model can be systematically quantified through its covariance. Thus, model accuracy can be ensured up to a pre-defined confidence interval. ■

C. Chance Constraints Formulation

Since $\hat{w}_{j|k} \sim \mathcal{N}(\mu_{j|k}^{\hat{w}}, \Sigma_{j|k}^{\hat{w}})$, $\hat{w}_{j|k} = \mu_{j|k}^{\hat{w}} + (\Sigma_{j|k}^{\hat{w}})^{1/2} \xi_{j|k}$, where $\xi_{j|k} \sim \mathcal{N}(0, I)$, and (3) is formulated as

$$\mathbb{P} \left[h_s^\top x_{j|k} + \eta_s^\top \mu_{j|k}^{\hat{w}} + \left(\eta_s^\top \Sigma_{j|k}^{\hat{w}} \eta_s \right)^{1/2} \xi_{j|k} \leq b_s \right] \geq 1 - \varepsilon_s.$$

Since constraints are linear and $\xi_{j|k}$ is Gaussian ICCs are formulated as the deterministic constraints

$$h_s^\top x_k + \eta_s^\top \mu_k^{\hat{w}} + \alpha_s (\eta_s^\top \Sigma_k^{\hat{w}} \eta_s)^{1/2} \leq b_s, \quad (9)$$

with $\alpha_s = F_N^{-1}(1 - \varepsilon_s)$, where $F_N^{-1}(\cdot)$ is the standard normal inverse cumulative distribution function (CDF). Since we consider ICCs, the probability of all constraints being satisfied is $\pi_{\text{sat}} \geq \prod_{s=1}^{n_s} (1 - \varepsilon_s)$.

D. Cost Function

The cost function

$$J(x_k, U_k, \mathbf{r}_k) = \sum_{j=0}^{N-1} \ell(x_{j|k}, u_{j|k}, r_{j|k}) + F(x_{N|k}, r_{N|k}) \quad (10)$$

balances the control objective with the acquisition of information on the environment. In (10), $\ell(x, u, r) = \|x - r^x\|_{Q_c}^2 + \|u - r^u\|_{R_c}^2 + S_c (\text{tr}[\Sigma^{\hat{w}}] - \text{tr}[\Sigma_{\text{ss}}^{\hat{w}}])^2$ and $F(x, r) = \|x - r^x\|_{P_c}^2$, where Q_c , R_c , and P_c are positive (semi)definite weight matrices, and the state and input reference trajectories, $r_{j|k}^x$ and $r_{j|k}^u$, respectively, are generated from the performance output reference $r_{j|k}$ by the standard parametrization $(r_{j|k}^x, r_{j|k}^u) = [T_x, T_u]^\top r_{j|k}$ [13]. The term $(\text{tr}[\Sigma^{\hat{w}}] - \text{tr}[\Sigma_{\text{ss}}^{\hat{w}}])^2$ aims at reducing the uncertainty in the estimate of the environment by driving the covariance to its steady-state. $S_c \in \mathbb{R}_{0+}$ determines the trade-off between the control objective and the uncertainty reduction.

IV. PAC-MPC AND ITS PROPERTIES

At time k , the PAC-MPC designed with components described based on Section III solves

$$J^*(x_k, \mathbf{r}_k) = \min_{U_k} J(x_k, U_k, \mathbf{r}_k) \quad (11a)$$

$$\text{s.t. } x_{j+1|k} = A x_{j|k} + B u_{j|k} \quad (11b)$$

$$(\mu_{j+1|k}^{\hat{w}}, \Sigma_{j+1|k}^{\hat{w}}) = \hat{g}(\hat{w}_{j|k}, \Sigma_{j|k}^{\hat{w}}, \psi_{j|k}, x_{j|k}, u_{j|k}) \quad (11c)$$

$$\psi_{j|k} = h_w(x_{j|k}, u_{j|k}, \psi_k), \quad (11d)$$

$$(x_{j|k}, u_{j|k}) \in \mathcal{X} \times \mathcal{U}, \quad (11e)$$

$$h_s^\top x_{j|k} + \eta_s^\top \mu_{j|k}^{\hat{w}} + \alpha_s (\eta_s^\top \Sigma_{j|k}^{\hat{w}} \eta_s)^{1/2} \leq b_s \quad (11f)$$

$$(x_{N|k}, r_{N|k}) \in \mathcal{Z}_f(\gamma_{N|k}) \quad (11g)$$

$$x_{0|k} = x_k, \quad \psi_{0|k} = \psi_k, \quad s \in \mathbb{Z}_{[1, n_s]}, \quad (11h)$$

where (11c) is the uncertainty propagation function, which can take the form (7) for model-based PAC-MPC, or (8) for learning-based PAC-MPC, (11d) is the function that is used to predict the measurement, e.g., $h_w(\cdot) = \psi_{0|k}$ or $h_w(\cdot) = C^w \mu_{0|k}^{\hat{w}}$, $\gamma \in \mathbb{R}^{n_s}$ is a short hand notation for the effect of \hat{w} onto the constraints, i.e., $[\gamma]_s = \eta_s^\top \mu_{j|k}^{\hat{w}} + \alpha_s (\eta_s^\top \Sigma_{j|k}^{\hat{w}} \eta_s)^{1/2}$, so that (11f) becomes $h_s x + [\gamma]_s \leq b_s$, and (11g) is the terminal constraint, which can be made trivial by setting

$\mathcal{Z}_f(\gamma_{N|k}) = \mathbb{R}^{n_x+n_y}$ for all $\gamma \in \mathbb{R}^{n_s}$. The optimal solution of (11) is denoted by $U_k^* = (u_{0|k}^*, \dots, u_{N-1|k}^*)$. Then, the PAC-MPC law is

$$u_k = \kappa(x_k, \mu_k^{\hat{w}}, \Sigma_k^{\hat{w}}, \psi_k) = u_{0|k}^*. \quad (12)$$

Next, we present preliminary results to achieve recursive feasibility and stability for (12) based on (11) in this stochastic setting. For the remainder of this section we consider model-based PAC-MPC where (11c) is implemented by (7).

A. Recursive Feasibility and Stability

The terminal set $\mathcal{Z}_f(\gamma)$ in (11g) is designed to ensure recursive feasibility. Let $u = Kx$ be a stabilizing control gain for (1), and consider the control law

$$u = K(x - r_x) + r_u = Kx + T_c r \quad (13)$$

where $T_c = T_u - KT_x$. Consider (1), (13), resulting in $x_{k+1} = A_{cl}x_k + B_{cl}r_k$, and auxiliary constant dynamics $r_{k+1} = r_k$, $\gamma_{k+1} = \gamma_k$, the re-formulation of ICCs as $H_c x_k + \gamma_k \leq b_c$, and the admissible references as $H_r r_k \leq b_r$. Let $z = (x, r, \gamma)$, under mild assumptions, we can compute [14] a set $\mathcal{O} \subseteq \mathcal{Z}_0 = \{z : H_z z \leq b_z\}$ which is positive invariant, i.e., if $z \in \mathcal{O}$ then $A_{cl}^z z \in \mathcal{O}$. Let γ_k be such that $[\gamma_k]_s = \eta_s^\top \mu_k^{\hat{w}} + \alpha_s (\eta_s^\top \Sigma_k^w \eta_s)^{1/2}$, and $r_{k+1} = r_k$. If $(x_k, r_k, \gamma_k) \in \mathcal{O}$, then $(A_{cl}x_k + B_{cl}r_k, r_{k+1}, \gamma_{k+1}) \in \mathcal{O}$ for every $\gamma_{k+1} \leq \gamma_k$. Thus, if $z \in \mathcal{O}$, r constant, and γ does not increase, the constraints are satisfied for all future steps.

The conditions for recursive feasibility are summarized by the following assumptions.

Assumption 1: $\mu_{k+1}^{\hat{w}} = \mu_{1|k}^{\hat{w}}$, $\Sigma_{k+1}^{\hat{w}} = \Sigma_{1|k}^{\hat{w}}$.

Assumption 2: $\gamma_{N|k+1} \leq \gamma_{N|k}$ is admissible.

Assumption 3: $\psi_{k+1} = \psi_{1|k}$.

Although Assumptions 1-3 are challenging to satisfy in general, below we discuss minor modifications to the PAC-MPC that guarantee recursive feasibility even when these assumptions are relaxed.

Theorem 1: Let $\mathcal{Z}_f(\gamma) = \{(x, r) : (x, r, \gamma) \in \mathcal{O}\}$ and assume that the OCP (11) with (11c) implemented by (7) is feasible at time k . If, in addition, assumptions 1-3 are satisfied, then, constraints (2), (3) are satisfied at time $k+1$ and (11) is feasible at time $k+1$ with probability π_{sat} .

Proof (sketch): If (11) is feasible at time k , by the chance constraints under the state assumptions, there is a probability at least π_{sat} that the constraints are satisfied at time $k+1$. Let the optimal solution of the OCP (11) at time k be $U_k^* = (u_{0|k}^* \dots u_{N-1|k}^*)$ and $X_k^* = (x_{0|k}^* \dots x_{N|k}^*)$, and define $\Gamma_k^* = (\gamma_{0|k}^* \dots \gamma_{N|k}^*)$. Construct the candidate solution $\tilde{U}_{k+1} = (u_{1|k}^*, \dots, u_{N|k}^*, Kx_{N|k}^* + T_c r_{N|k}^*)$, and $\tilde{X}_{k+1} = (x_{1|k}^*, \dots, x_{N|k}^*, A_{cl}x_{N|k}^* + B_{cl}r_{N|k}^*)$. Since $\mu_{k+1}^{\hat{w}} = \mu_{1|k}^{\hat{w}}$, $\Sigma_{k+1}^{\hat{w}} = \Sigma_{1|k}^{\hat{w}}$, $\psi_{k+1} = \psi_{1|k}$, and by (11d) we also have $\tilde{\Gamma}_{k+1} = (\gamma_{1|k}^*, \dots, \gamma_{N|k}^*, \gamma_{N|k+1}^*)$. Since U_k^* , X_k^* , Γ_k^* satisfied the constraints, \tilde{U}_{k+1} , \tilde{X}_{k+1} , $\tilde{\Gamma}_{k+1}$ satisfy (11e), (11f). Finally, $(x_{N|k}, r_{N|k}, \gamma_{N|k}) \in \mathcal{O}$ implies $(x_{N|k+1}, r_{N|k+1}, \gamma_{N|k+1}) = (A_{cl}x_{N|k}^* + B_{cl}r_{N|k}^*, r_{N|k}, \gamma_{N|k}) \in \mathcal{O}$ by invariance. Since $\gamma_{N|k+1} \leq \gamma_{N|k}$, then $(x_{N|k+1}, r_{N|k+1}) \in \mathcal{Z}_f(\gamma_{N|k+1})$ is satisfied. ■

Assumptions 1 and 2 are satisfied as follows. Let $\psi_{k+1} = \psi_{1|k}$, i.e., Assumption 3 is satisfied. Then, $\mu_{k+1}^{\hat{w}} = \mu_{1|k}^{\hat{w}}$, $\Sigma_{k+1}^{\hat{w}} = \Sigma_{1|k}^{\hat{w}}$ are satisfied when (7a), (7b) are the uncertainty propagation equations used in (11c). The condition $\gamma_{N|k+1} \leq \gamma_{N|k}$ (Assumption 2) holds if the estimator guarantees that $\gamma_{k+1} \leq \gamma_k$, componentwise, i.e., if the uncertainty on the constraints does not increase, since $\psi_{k+1} = \psi_{1|k}$. This can be enforced by the estimator design, and also by the PAC-MPC control over the uncertainty reduction by sensing.

The remaining condition is to ensure Assumption 3, i.e., that we can correctly predict the next measurement, which is challenging to satisfy in general. An incorrect prediction of the measurement ψ_{k+1} affects the prediction of the mean $\mu_{1|k+1}^{\hat{w}}$, but not the prediction of the covariance $\Sigma_{1|k+1}^{\hat{w}}$. While in practice an error in prediction of the mean may not render (11) infeasible, and its effect on the constraint tightening may be compensated for by the covariance estimate, we propose a slight modification to the PAC-MPC that accounts for the measurement prediction error.

Let $\Delta\psi_{j|k} = \psi_{k+j} - \psi_{j|k} \sim p_{\Delta\psi(k,j)}$ be a random vector that captures the prediction error in the measurement, and let $\tilde{p}_{\Delta\psi(k,j)} = \mathcal{N}(\tilde{\mu}_{j|k}^{\Delta\psi}, \tilde{\Sigma}_{j|k}^{\Delta\psi})$ be our model for the distribution $p_{\Delta\psi(k,j)}$. We modify (7) as

$$\begin{aligned} \mu_{j+1|k}^{\hat{w}} &= \Lambda_{j|k} \mu_{j|k}^{\hat{w}} - L_{j|k} \psi_{j|k} - L_{j|k} \tilde{\mu}_{j|k}^{\Delta\psi} \\ \Sigma_{j+1|k}^{\hat{w}} &= \Lambda_{j|k} \Sigma_{j|k}^{\hat{w}} \Lambda_{j|k}^\top + Q_{j|k} + R_{j|k} + L_{j|k} \tilde{\Sigma}_{j|k}^{\Delta\psi} L_{j|k}^\top. \end{aligned} \quad (14)$$

By decomposing $\psi_{k+j} = \Delta\psi_{j|k} + \psi_{j|k}$, we obtain the uncertainty propagation (14) which includes the additional uncertainty due to $\Delta\psi_{j|k}$. Thus, if the measurement prediction error is correctly modeled, i.e., $\tilde{p}_{\Delta\psi(k,j)} = p_{\Delta\psi(k,j)}$, each chance constraint is satisfied with the desired probability $1 - \varepsilon_s$. If instead $\tilde{p}_{\Delta\psi(k,j)} \neq p_{\Delta\psi(k,j)}$, the chance constraint satisfaction probabilities may be different from the desired ones, but will be larger with increasing uncertainty in the measurement prediction error. Then, $\tilde{p}_{\Delta\psi(k,j)}$ can be used as design parameter that trades off conservativeness in the trajectory and probability of constraint satisfaction.

Remark 5: Since the cost is deterministic and the uncertainty does not affect the dynamics (1) existing stability results can be applied [1], [9] for the stability of the PAC-MPC. Therefore, under the assumptions in Theorem 1, given the terminal controller $u = Kx$, and if $S_c = 0$, the PAC-MPC (12) is asymptotically stable to the reference r_k . The case $S_c > 0$ requires co-design of controller and estimator (5), and is subject of ongoing research. ■

V. CASE STUDY: AUTOMATED VEHICLE CONTROL

We consider a linearized model of the lateral vehicle dynamics with respect to the centerlane, see, e.g., [15], discretized with sampling period $T_s = 0.050$ s, which results in (1) where $x = (e_1, e_2, \dot{e}_1, \dot{e}_2)$, $u = \delta$, and $y = ([x]_1, [x]_2)$, e_1 is the distance from the lane center, e_2 is the orientation error with respect to the road, and δ is the wheel steering angle. The values of A , B , E in (1) are from real vehicle data [15]. We include two additional inputs $[u]_2$, $[u]_3$, which




Fig. 2. Regulation to straight driving. Lateral position trajectories by standard MPC and model-based PAC-MPC with distance-dependent measurement (15). PAC-MPC overshoots the reference to approach the top road boundary, which improves sensing.

determine the amount of sensing on the left and right road boundary, respectively, so that $u = (\delta, [u]_2, [u]_3)$.

The environment is modeled as $w \in \mathbb{R}^2$, for representing the lateral coordinate of the road boundaries, with $A^w = I_2$ and $B^w = 0$, with $w_0 = [3.5, -0.5]^\top$. We consider deterministic input constraints, $[-1.5 \ 0 \ 0]^\top \leq u_k \leq [1.5 \ 1 \ 1]^\top$, and ICCs imposing to remain within the road boundaries, $\mathbb{P}[h_s^\top x_k + \eta_s^\top w_k \leq 0] \geq 0.95$, $s = \{1, 2\}$, where $h_1 = [1 \ 0 \ 0 \ 0]^\top$, $\eta_1 = [-1 \ 0]^\top$, $h_2 = [-1 \ 0 \ 0 \ 0]^\top$, $\eta_2 = [0 \ 1]^\top$.

We consider two measurement models with $C^w = I_2$ for sensing the environment: (i) distance-dependent measurements, where the noise increases with the distance [16],

$$[D_k^w]_{i,i} = \text{diag}(P([w_k]_i)) \|[x_k]_i - [w_k]_i\| / L_L |i|^\beta, \quad (15)$$

$\text{diag}(P(\cdot))$ depends on the boundary position w_k , L_L is a length-scale constant, and β is deterioration rate by distance; and (ii), input-dependent measurements, where $[u]_2$ and $[u]_3$ give direct control over the uncertainty,

$$D_k^w = P(w_k)(I_2 - \beta \cdot \text{diag}([u]_2, [u]_3)), \quad (16)$$

so that, $[u]_2$ and $[u]_3$ are the amount of sensor processing for each side of the road.

A. Regulation to straight driving

First, we regulate the vehicle starting at $e_1 = 1$ m to straight driving with $r_k = [2, 0]^\top$ (m, rad), under the distance-dependent measurement (15). The road clearance (width minus vehicle size) is 4 m, with boundaries at $[y]_1 = 3.5$ m and $[y]_1 = -0.5$ m, initially unknown to the controller. We set $S_c = 100$ and limit D^w to only depend on the $[y_k]_1 - [w_k]_1$, i.e., only the left boundary is considered.

Fig. 2 shows the closed-loop trajectory of PAC-MPC with uncertainty propagation (7). Since in (15) D^w does not depend on $[u]_2$, $[u]_3$, $[u]_2 = [u]_3 = 0$, constantly.

Due to $S_c > 0$, the model-based PAC-MPC balances the control objective, i.e., tracking r , and the acquisition of environment information, i.e., driving the covariance to steady-state. Hence, as compared to a standard MPC ($S_c = 0$), PAC-MPC overshoots $[r]_1$ to better sense the road boundary at $[y]_1 = 3.5$ m, which makes the covariance converge faster.




Fig. 3. Double lane change maneuver. Standard MPC vs model-based PAC-MPC with the input-dependent measurement model (16). Trajectories of lateral position (top) and sensing inputs (bottom). The sensing inputs for the standard MPC are fixed to zero, while model-based PAC-MPC manipulates the sensing inputs.

B. Double Lane Change

Next we consider a double lane change maneuver, with the input-dependent measurement model (16) and $S_c = 0$, so that sensing decisions are driven only by the constraints. The road is as in the previous test, the vehicle starts at $e_1 = 1$ and the lane centerlines are at 0 m and at 3 m. In prediction, the controller knows the future reference trajectory.

Fig. 3 shows the comparison of: (i) MPC with fixed sensing, $[u]_2 = [u]_3 = 0$; and (ii) PAC-MPC with model-based uncertainty propagation (7) that can allocate the sensing resources. Table I reports the absolute average errors (AAEs).

Since the MPC cannot optimize sensing to reduce the environment uncertainty, the constraint tightening in the ICCs (9) forces the vehicle to follow a more conservative trajectory, where the MPC cannot reach the desired setpoints, since these are too close to the estimated road boundaries.

In contrast, PAC-MPC can track the setpoints with practically no steady-state offset by exploiting $[u]_2$ and $[u]_3$ (Fig. 3 (bottom)) to reduce D^w . This reduces the covariance of the environment measurement $\Sigma_{j|k}^w$ through the term $R_{j|k}$ in (7), which, in turn, reduces the constraint tightening in (9). Fig. 3 (bottom) shows that when the vehicle operates close to $[y]_1 = -0.5$ m, $[u]_3$ is increased to reduce the right boundary measurement uncertainty, whereas when the vehicle operates close to $[y]_1 = 3.5$ m, $[u]_2$ is increased to

TABLE I

AVERAGE ABSOLUTE ERRORS FOR THE CLOSED-LOOP TRAJECTORIES

Strategy	MPC	PAC (model)	PAC (model)	PAC (learning)
ρ	N/A	N/A	$\rho = 0.5$	$\rho = 0.5$
AAE (m)	0.64	0.31	0.45	0.37

reduce the left boundary measurement uncertainty. Thus, the AAE of the model-based PAC-MPC is 51% lower than that of the MPC (Table I). At steady-state, the AAE of the model-based PAMPC is only approximately 0.08 m. The controller reduces the measurement uncertainty despite $S_c = 0$ because it is worth incurring an input cost for $[u]_2$ and $[u]_3$ for getting closer to the reference trajectory.

Finally, we compare model-based PAC-MPC (Section III-A) with learning-based PAC-MPC (Section III-B), which uses a GPR learned model (8) to propagate the uncertainty mean and covariance in prediction. We also introduce a “sensing budget” constraint, $[u]_2 + [u]_3 \leq \rho$ which could arise due to limitations in computation or energy.

Fig. 4 shows the closed-loop trajectories of the model-based PAC-MPC and learning-based PAC-MPC with $\rho = 0.5$. The performance of the model-based PAC-MPC has deteriorated compared to Fig. 3, shown by a larger tracking error (Table I) due to the sensing budget constraint ($\rho = 0.5$). Even though the learning-based PAC-MPC is subject to the same constraint, it is able to track the desired reference trajectory with minimal offset, and performs, on average, 18% better than the model-based PAC-MPC, see Table I. Further simulations revealed that, in general, learning-based PAC-MPC performed equally or better than model-based PAC-MPC. This can be attributed to the GPR, which was trained using simulation data. Due to the finite number of samples, the empirical distribution of the environment measurement uncertainties may not be exactly Gaussian. Consequently, it may be better captured through the GPR (8). Also, learning-based PAC-MPC may outperform model-based PAC-MPC because it intrinsically accounts for the measurement prediction error, therefore adjusting the tightening in (14), without assuming any distribution for $\psi_{j|k}$ or treating the distribution as a tuning parameter.

VI. CONCLUSIONS

We presented a perception-aware chance-constrained MPC (PAC-MPC) for a system operating in an unknown environment that affects the system through constraints, and where the environment is discovered through sensing, which depends on how the system is operated. Due to using stochastic uncertainty propagation, which can be obtained from models or data, and PAC-MPC enforces chance constraints and leverages stochastic MPC approaches. We provided preliminary conditions for recursive feasibility and stability, which are being extended by ongoing research.

ACKNOWLEDGMENTS

The authors would like to thank Dr. Karl Berntorp and Dr. Rien Quirynen for valuable discussions on this work.

REFERENCES

- [1] J. B. Rawlings and D. Q. Mayne, *Model predictive control: Theory and design*. Nob Hill Pub., 2009.
- [2] S. Di Cairano and I. V. Kolmanovsky, “Real-time optimization and model predictive control for aerospace and automotive applications,” in *Proc. Amer. Control Conf.*, 2018, pp. 2392–2409.




Fig. 4. Double lane change maneuver. Model-based PAC-MPC vs learning-based PAC-MPC using the input-dependent measurement model (16) with a sensing budget constraint where $\rho = 0.5$. Trajectories of lateral position (top) and sensing inputs (bottom).

- [3] F. Bourgault, A. A. Makarenko, S. B. Williams, B. Grocholsky, and H. F. Durrant-Whyte, “Information based adaptive robotic exploration,” *IEEE International Conference on Intelligent Robots and Systems*, vol. 1, pp. 540–545, 2002.
- [4] G. Costante, C. Forster, J. Delmerico, P. Valigi, and D. Scaramuzza, “Perception-aware Path Planning,” *arXiv preprint*, pp. 1–16, 2016. [Online]. Available: <http://arxiv.org/abs/1605.04151>
- [5] N. E. Du Toit and J. W. Burdick, “Robot motion planning in dynamic, uncertain environments,” *IEEE Transactions on Robotics*, vol. 28, no. 1, pp. 101–115, 2011.
- [6] D. Falanga, P. Foehn, P. Lu, and D. Scaramuzza, “Pampc: Perception-aware model predictive control for quadrotors,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2018, pp. 1–8.
- [7] K. Lee, G. N. An, V. Zakharov, and E. A. Theodorou, “Perceptual Attention-based Predictive Control,” *arXiv preprint*, 2019. [Online]. Available: <http://arxiv.org/abs/1904.11898>
- [8] M. Farina, L. Giulioni, and R. Scattolini, “Stochastic linear model predictive control with chance constraints—a review,” *Jour. Process Control*, vol. 44, pp. 53–67, 2016.
- [9] B. Kouvaritakis and M. Cannon, *Model Predictive Control: Classical, Robust, and Stochastic*. Springer, 2015.
- [10] A. Mesbah, “Stochastic model predictive control: An overview and perspectives for future research,” *IEEE Control Systems*, vol. 36, no. 6, pp. 30–44, 2016.
- [11] A. Mesbah, I. V. Kolmanovsky, and S. Di Cairano, “Stochastic model predictive control,” in *Handbook of Model Predictive Control*. Springer, 2019, pp. 75–97.
- [12] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006, vol. 2, no. 3.
- [13] D. Limón, I. Alvarado, T. Alamo, and E. F. Camacho, “Mpc for tracking piecewise constant references for constrained linear systems,” *Automatica*, vol. 44, no. 9, pp. 2382–2387, 2008.
- [14] E. Garone, S. Di Cairano, and I. Kolmanovsky, “Reference and command governors for systems with constraints: A survey on theory and applications,” *Automatica*, vol. 75, pp. 306–328, 2017.
- [15] S. Di Cairano, U. Kalabić, and K. Berntorp, “Vehicle tracking control on piecewise-clothoidal trajectories by mpc with guaranteed error bounds,” in *Proc. 55th IEEE Conf. Decision and Control*, 2016, pp. 709–714.
- [16] B. Huang, L. Xie, and Z. Yang, “Tdoa-based source localization with distance-dependent noises,” *IEEE Transactions on Wireless Communications*, vol. 14, no. 1, pp. 468–480, 2014.