

# Friction-Adaptive Stochastic Nonlinear Model Predictive Control for Autonomous Vehicles

Vaskov, Sean; Quirynen, Rien; Menner, Marcel; Berntorp, Karl

TR2023-064 June 01, 2023

## Abstract

This paper addresses the trajectory-tracking problem under uncertain road-surface conditions for autonomous vehicles. We propose a stochastic nonlinear model predictive controller (SNMPC) that learns a tire-road friction model online using standard automotive-grade sensors. Learning the entire tire-road friction model in real time requires driving in the nonlinear, potentially unstable regime of the vehicle dynamics, using a prediction model that may not have fully converged. To handle this, we formulate the tire-friction model learning in a Bayesian framework, and propose two estimators that learn different aspects of the tire-road friction. The estimators output the estimate of the tire-friction model as well as the uncertainty of the estimate, which expresses the confidence in the model for different driving regimes. The SNMPC exploits the uncertainty estimate in its prediction model to take proper action when the uncertainty is large. We validate the approach in an extensive Monte-Carlo study using real vehicle parameters and in CarSim. The results when comparing to various MPC approaches indicate a substantial reduction in constraint violations, as well as a reduction in closed-loop cost. We also demonstrate the real-time feasibility in automotive-grade processors using a dSPACE MicroAutoBox-II rapid prototyping unit, showing a worst-case computation time of roughly 40ms.

*Vehicle System Dynamics 2023*

© 2023 MERL. This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.



ARTICLE DRAFT

## Friction-Adaptive Stochastic Nonlinear Model Predictive Control for Autonomous Vehicles

Sean Vaskov<sup>a,b</sup>, , Rien Quirynen<sup>a</sup>, Marcel Menner<sup>a</sup>, and Karl Berntorp<sup>a</sup>

<sup>a</sup>Mitsubishi Electric Research Laboratories, Cambridge, MA, 02139, USA.; <sup>b</sup>University of Michigan, Ann Arbor, MI, 48109, USA.

### ARTICLE HISTORY

Compiled January 15, 2023

### ABSTRACT

This paper addresses the trajectory-tracking problem under uncertain road-surface conditions for autonomous vehicles. We propose a stochastic nonlinear model predictive controller (SNMPC) that learns a tire-road friction model online using standard automotive-grade sensors. Learning the entire tire-road friction model in real time requires driving in the nonlinear, potentially unstable regime of the vehicle dynamics, using a prediction model that may not have fully converged. To handle this, we formulate the tire-friction model learning in a Bayesian framework, and propose two estimators that learn different aspects of the tire-road friction. The estimators output the estimate of the tire-friction model as well as the uncertainty of the estimate, which expresses the confidence in the model for different driving regimes. The SNMPC exploits the uncertainty estimate in its prediction model to take proper action when the uncertainty is large. We validate the approach in an extensive Monte-Carlo study using real vehicle parameters and in CarSim. The results when comparing to various MPC approaches indicate a substantial reduction in constraint violations, as well as a reduction in closed-loop cost. We also demonstrate the real-time feasibility in automotive-grade processors using a dSPACE MicroAutoBox-II rapid prototyping unit, showing a worst-case computation time of roughly 40ms.

### KEYWORDS

Vehicle dynamics, stochastic model predictive control, particle filtering, machine learning.

## 1. INTRODUCTION

The strong push recently in the automotive industry for introducing new technologies related to automated driving (AD) and advanced driving-assistance systems (ADAS), has led to predictive information being readily available to the vehicle controllers. For instance, onboard sensors detect the environment in the vicinity of the vehicle, motion planners generate future desired driving profiles [1,2], and motion-prediction modules provides risk assessments for different control actions [3,4]. Consequently, model predictive control (MPC) is suitable for vehicle control as it naturally integrates predictive information in its problem formulation [5,6], and MPC has successfully been applied to vehicle trajectory tracking and stability control [7–11].

The key benefit of MPC for vehicle-control applications is its ability to naturally incorporate constraint handling as well as prediction models in its optimal control


problem (OCP). However, because MPC relies on a vehicle model, it is important that the model represents sufficiently well the actual vehicle. Specifically, since the main actuation of the vehicle is done by altering the forces generated by the tire–road contact, it is imperative to have a sensible guess about the road surface the car operates on, as otherwise safe operation of the vehicle can be endangered [12,13].

In this paper, we develop an MPC that adapts in real time to different road-surface conditions. The interaction between tire and road is highly nonlinear, and the function describing the nonlinear relation varies heavily based on the road surface and other tire properties [14,15]. Fig. 1 shows examples of the function relating the lateral tire force to the tire side slip angle, for different surfaces. The force-slip relation is approximately linear for small slip values, which are typical when driving in normal conditions. However, when driving close to the adhesion limits, which may happen in emergency maneuvers, on unpaved road, or on wet and icy roads, the nonlinear part of the tire-force function may be excited, and hence the full tire curve shape must be considered. Thus, when driving over different surfaces the time-varying tire-force curve needs to be identified in real time, rapidly, and using noisy onboard sensing.

A complicating factor in identifying nonlinear functions is that the entire range of the function should be excited to generate data for identification. In case of the tire force function, obtaining data for the nonlinear part is challenging as it requires to drive the vehicle to the limits of its performance envelope and, usually, this is not done unless an emergency maneuver is needed or a particularly challenging surface is encountered. Furthermore, driving in the nonlinear region of the tire force function before a reliable model of the same is obtained is challenging and possibly dangerous. If the controller approaches the nonlinear region of the force curve without having a reliable model, this can cause control errors possibly leading to vehicle instability. In addition, a difficulty when learning the tire-friction function using automotive-grade sensors is that the amount of sensors is limited, and they are relatively low grade [16]. Moreover, not only do the sensors only provide indirect measurements of the friction, they do not even measure the vehicle state, which is nonlinearly dependent on the tire friction and must therefore be known for learning the tire friction.

To address these issues, in this paper, we couple online tire-friction estimation with a stochastic nonlinear MPC (SNMPC). The proposed SNMPC models the tire-road friction as a disturbance affecting the prediction model and is flexible in the sense that as long as a tire-friction estimator can output the tire-road friction model by the first two moments (i.e., mean and covariance), the SNMPC can readily incorporate the estimates. As an example of this, we integrate the SNMPC with two Bayesian estimators capable of outputting the first two moments; (i), a computationally efficient tire-stiffness estimator modeling the tire-road friction as a linear function of the slip; and (ii), a Gaussian-process (GP) based estimator with unknown and time-varying mean and covariance function [17], leading to a GP state-space model (GP-SSM). In particular, we leverage two recently proposed methods for real-time joint state estimation and learning of the state-transition function in dynamical systems [18,19] based on particle filtering [20] for jointly estimating online the state and associated tire-friction model. These two methods are similar in the sense that they estimate the distribution and associated weighted mean of the tire friction using only sensors available in production cars. Since the methods are based on particle filtering, they have asymptotic convergence guarantees. Due to the estimators being Bayesian, they behave in a manner suitable for safety-critical vehicle control, as they predict larger uncertainty for parts of the system that have not been sufficiently excited.

The resulting control system leverages the information from the estimator in its



**Figure 1. Upper:** Examples of lateral friction as a function of slip angle  $\alpha$  for asphalt, loose snow, and ice. **Lower:** A schematic of the single-track vehicle model and related notation.

prediction model. The respective estimator learns the tire-friction function using a model-focused approach where it updates the tire-friction model as data are gathered, which leads to efficient learning where only a few data points are needed to achieve reliable estimates, resulting in an efficient and safe closed-loop behavior suitable for real-time control. The SNMPC is based on an efficient block-sparse quadratic programming (QP) solver [21] for use within nonlinear optimal control, and the SNMPC with probabilistic chance constraints uses a tailored Jacobian approximation along with an adjoint-based sequential QP (SQP) optimization method [22]. The enabler for integrating the friction estimators with the SNMPC formulation is the insight that although the filters estimate different aspects of the tire-friction relation, they can both output the mean and covariance of the respective friction quantities. This allows a simple and intuitive connection to the SNMPC formulation.

We have previously reported on methods for friction-adaptive MPC in [8,23,24]. The current paper differs in that the work in [8] uses tire-stiffness estimates to choose from a fixed library of predetermined tire-friction functions and does not utilize SNMPC. In [23], we developed a friction estimator targeting the linear region of the tire-friction curve [18], and in [24] we showed how SNMPC can be coupled with GP-based particle filtering for controlling the vehicle while estimating the full tire-friction curve. This paper extends our preliminary work [23,24] in three aspects.

- We generalize the estimation formulation and show that the SNMPC can be connected to any tire-friction estimator, as long as such estimator can output a mean estimate and associated covariance of the tire-friction model.
- We provide an extensive comparison of the closed-loop performance using two different estimators and discuss pros and cons for each of the methods.
- We show that the proposed control system is real-time feasible for current automotive micro-controllers, by implementing it on a dSPACE MicroAutoBox-II rapid prototyping unit, where the worst-case computation time for the proposed vehicle control system is about 40ms, irrespective of the estimator being used.

### 1.1. Outline

We outline the estimation and control model, together with the problem formulation in Sec. 2. Sec. 3 details an overview of the tire-friction estimators employed in the paper. Sec. 4 gives the SNMPC problem formulation and implementation details based on the control model, and we make the connection between the SNMPC and the tire-friction estimator and outline the control system in Sec. 5, which is evaluated in a Monte-Carlo study in Sec. 6. Finally, the paper concludes with Sec. 7.

### 1.2. Notation

The notation  $\mathcal{R}(a)$  means the 2D rotation matrix of an angle  $a$ . Vectors are shown in bold,  $\mathbf{x}$ , we denote the stacking of two vectors  $\mathbf{a}$ ,  $\mathbf{b}$  by  $[\mathbf{a}, \mathbf{b}]$ , and constraints between vectors are intended componentwise. Throughout,  $\mathbf{x} \sim \mathcal{N}(\hat{\mathbf{x}}, \mathbf{\Sigma})$  indicates that  $\mathbf{x} \in \mathbb{R}^{n_x}$  is Gaussian distributed with mean  $\hat{\mathbf{x}}$  and covariance  $\text{cov}(\mathbf{x}) = \mathbf{\Sigma}$ . Matrices are indicated in capital bold font as  $\mathbf{X}$ . With  $p(\mathbf{x}_{0:k}|\mathbf{y}_{0:k})$ , we mean the posterior density function of the state trajectory  $\mathbf{x}_{0:k}$  from time step 0 to time step  $k$  given the measurement sequence  $\mathbf{y}_{0:k} := \{\mathbf{y}_0, \dots, \mathbf{y}_k\}$ , and  $\mathbf{x}_{0:k}^i$  is the  $i^{\text{th}}$  realization of  $\mathbf{x}_{0:k}$ . The notation  $\mathbf{f}(\mathbf{x}) \sim \mathcal{GP}(\hat{\mathbf{f}}(\mathbf{x}), \mathbf{\Sigma}(\mathbf{x}))$  means that the function  $\mathbf{f}(\mathbf{x})$  is a realization from a GP with a mean function  $\hat{\mathbf{f}}(\cdot)$  and covariance  $\mathbf{\Sigma}(\cdot)$ . For a continuous-time signal  $\mathbf{x}(t)$  sampled with period  $T_s$ ,  $x_k$  denotes the  $k^{\text{th}}$  sample, i.e.,  $\mathbf{x}_k = \mathbf{x}(kT_s)$  and  $\mathbf{x}_{k+h|k}$  is the value of  $\mathbf{x}$  predicted  $h$  steps ahead from  $k$ , i.e., the predicted value of  $\mathbf{x}((k+h)T_s)$  based on  $\mathbf{x}(kT_s)$ . The quadratic form of the squared weighted norm of a vector  $\mathbf{x}$  and matrix  $Q$  is given by  $\|\mathbf{x}\|_Q^2 = \mathbf{x}^\top Q \mathbf{x}$ .

## 2. Modeling and Problem Formulation

The vehicle model is composed of a chassis model describing the motion of the rigid body due to the forces generated at the tires and a tire model describing what forces the tires generate depending on the chassis and wheel velocities. For the chassis, we consider the standard single-track model, where the left and right track of the car are lumped into a single centered track, as Fig. 1 shows. Hence, only a single front and a single rear tire are considered, and roll and pitch dynamics are ignored, resulting in two translational and one rotational degrees of freedom. While for performance driving it may be advantageous to use a double-track chassis model, which includes lateral and longitudinal load transfer [12], in [12,13] the single-track model was shown to be sufficiently accurate for regular driving conditions, including when tire forces are in the nonlinear region, because in such conditions the roll and pitch angles remain relatively small. Similarly, the single-track model seems sufficient in most evasive maneuvers, because the focus of such maneuvers is on preserving safety rather than achieving optimality, and hence a high-precision model is often unnecessary. On the other hand, the single-track model results in a reduced computing load, which is always desirable in automotive applications [5], particularly for evasive maneuvers.

Taking the longitudinal and lateral velocities in the vehicle frame,  $v^X$ ,  $v^Y$ , and the

yaw rate,  $\dot{\psi}$ , as states, the single-track model is described by

$$\dot{v}^X - v^Y \dot{\psi} = \frac{1}{m}(F_f^x \cos(\delta_f) + F_r^x - F_f^y \sin(\delta_f)), \quad (1a)$$

$$\dot{v}^Y + v^X \dot{\psi} = \frac{1}{m}(F_f^y \cos(\delta_f) + F_r^y + F_f^x \sin(\delta_f)), \quad (1b)$$

$$I_{zz} \ddot{\psi} = l_f F_f^y \cos(\delta_f) - l_r F_r^y + l_f F_f^x \sin(\delta_f), \quad (1c)$$

where  $F_i^x = F_i^z \mu_i^x(\lambda_i)$ ,  $F_i^y = F_i^z \mu_i^y(\alpha_i)$  are the total longitudinal/lateral forces in the tire frame for the lumped left and right tires where the friction functions  $\mu_i^x$ ,  $\mu_i^y$  are provided by the estimator,  $\lambda_i$  is the longitudinal wheel slip,  $\alpha_i$  is the slip angle, and subscripts  $i = f, r$  denote front and rear, respectively. The normal forces resting on the lumped front/rear wheels  $F_i^z$ ,  $i \in \{f, r\}$ , are  $F_f^z = mgl_r/l$ ,  $F_r^z = mgl_f/l$ , where  $g$  is the gravity acceleration,  $l_f$ ,  $l_r$  are the distances of front and rear axles from the center of gravity, and  $l = l_f + l_r$  is the vehicle wheel base. Also,  $m$  is the vehicle mass,  $I_{zz}$  is the vehicle inertia about the vertical axis, and  $\delta_f$  is the front wheel (road) steering angle. The vehicle position in global coordinates  $\mathbf{p} = (p^X, p^Y)$  is obtained from the kinematic equation

$$\begin{bmatrix} \dot{p}^X \\ \dot{p}^Y \end{bmatrix} = \mathcal{R}(\psi) \begin{bmatrix} v^X \\ v^Y \end{bmatrix}. \quad (2)$$

The slip angles  $\alpha_i$  and slip ratios  $\lambda_i$  are defined as in [25],

$$\alpha_i = -\arctan\left(\frac{v_i^y}{v_i^x}\right), \quad \lambda_i = \frac{R_w \omega_i - v_i^x}{\max(R_w \omega_i, v_i^x)}, \quad (3)$$

where  $R_w$  is the wheel radius, and  $v_i^x$  and  $v_i^y$  are the longitudinal and lateral wheel velocities for wheel  $i$ . Given the velocity vector at the center of mass,  $v = [v^X \ v^Y]^\top$ , the velocity vectors at the wheels are

$$\begin{bmatrix} v_i^x \\ v_i^y \end{bmatrix} = \mathcal{R}(\delta_i)^\top \begin{bmatrix} v^X \\ v^Y + c_i \dot{\psi} \end{bmatrix}, \quad c_f = l_f, \quad c_r = -l_r. \quad (4)$$

## 2.1. Problem Formulation

We consider a scenario where we measure the vehicle position  $(p^X, p^Y)$  and the yaw (heading)  $\psi$  (e.g., by GPS), the individual wheel speeds  $\omega_{i,j}$ ,  $i \in \{f, r\}$ , where  $j \in \{l, r\}$  denotes left and right tires, by wheel encoders, which are used to determine also the vehicle speed, approximately equal to  $v^X$ . We measure the vehicle longitudinal,  $a^X = \dot{v}^X - v^Y \dot{\psi}$ , and lateral,  $a^Y = \dot{v}^Y + v^X \dot{\psi}$ , acceleration and the yaw rate  $\dot{\psi}$  by an automotive-grade inertial measurement unit and the front road wheel steering angle  $\delta_f$  by a relative encoder. The control inputs are the front and rear wheel speeds  $\omega_f$ ,  $\omega_r$  and the tire-wheel angle rate of change  $\dot{\delta}_f$ .

We target for implementation of our proposed control system on automotive micro-controllers, which, due to harsh operating conditions, hard real-time requirements, and cost considerations, are significantly less capable than desktop computers [5] in terms of memory and speed. Hence, we aim at limiting the amount of computations and data storage needed by our approach.

Under the above conditions, the objective of this paper is to design a control strategy that makes the vehicle motion follow a time-dependent reference trajectory  $(p_{\text{ref}}^{\mathbf{X}}(\cdot), p_{\text{ref}}^{\mathbf{Y}}(\cdot), \psi_{\text{ref}}(\cdot), v_{\text{ref}}^{\mathbf{X}}(\cdot))$ , possibly generated in real time with an adequate preview, while operating over different road surfaces and environmental conditions. Since some of the vehicle states, e.g.,  $v^{\mathbf{Y}}$ , are not directly measured and the tire-friction model is unknown (or at the very least uncertain) and changes over time, the control strategy needs to estimate the vehicle state and the tire-friction model in real time with the controller activated.

### 3. Bayesian Tire-Friction Learning

In this section, we discuss the estimation framework used in conjunction with the SNMPC to formulate our control strategy. Our SNMPC formulation assumes that the uncertainty of the tire-friction model estimate can be well represented by a mean estimate and associated covariance. We present two methods for real-time estimation of a tire-friction model that both satisfy the requirements of the SNMPC. Both estimators are targeted for embedded automotive-grade hardware and sensors. In fact, they use the same sensor setup and the only difference being the complexity of the tire-friction model it estimates. A difficulty when addressing the tire-friction estimation problem using automotive-grade sensors is that the amount of sensors is limited, and they are relatively low grade [16]. Moreover, not only do the sensors only provide indirect measurements of the friction, they do not even measure the vehicle state, which is nonlinearly dependent on the tire friction and must therefore be known for learning the tire friction. Subsequently, we employ a recently developed method for *jointly* estimating the tire-friction function  $\boldsymbol{\mu}$  and the vehicle state  $\boldsymbol{x}$  only using sensors available in production cars, namely wheel-speed sensors and inexpensive accelerometers and gyroscopes. We briefly outline the formulation of the respective method and refer to [18,19,26] for a more complete description.

**Remark 1.** In this paper, we focus on the estimation of the lateral tire friction at front and rear wheels (i.e.,  $n_{\boldsymbol{\mu}} = 2$ ). The extension to the longitudinal case is analogous. Hence, in the estimation, we focus on the lateral vehicle dynamics, because usually these are the most critical and challenging for vehicle stability and ADAS.

#### 3.1. Estimation Model

The estimation model is based on the single-track equations in (1) and models the tire-friction components as static functions of the slip quantities, which in the lateral case results in


$$\mu_i^y = f_i^y(\alpha_i(\boldsymbol{x}^e, \boldsymbol{u}^e)), \quad i \in \{f, r\}, \quad (5)$$

where the state and input vector for the estimator are defined as  $\boldsymbol{x}^e = [v^{\mathbf{Y}}, \dot{\psi}]^{\top}$ ,  $\boldsymbol{u}^e = [\delta, v^{\mathbf{X}}]^{\top}$ . For brevity, we define the vector  $\boldsymbol{\alpha} = [\alpha_f, \alpha_r]^{\top}$ .

The two estimators we employ only differ in the way they model (5). In the first estimator, the tire-friction model is linear, such that the friction reads as

$$\mu_i^y = C_i^y \alpha_i, i \in \{f, r\}. \quad (6)$$





**Figure 2.** Linear approximation of the tire friction at slip angles of 0, 4, 8 deg for an asphalt tire model. The solid line is the true tire-friction curve, while dashed lines are the linear approximations for the respective slip angle. The stiffness (slope of the line going through the origin) decreases as the slip angle increases.

The stiffness vector  $\mathbf{C} = [C_f^y, C_r^y]^\top$  in (6) is for each component decomposed into a nominal and unknown part,

$$C_i^y = C_{i,n}^y + \Delta C_{i,k}, \quad (7)$$

where  $C_{i,n}^y$  is the nominal value of the tire stiffness, for example, a priori determined on a nominal surface, and  $\Delta C_{i,k}$  is a time-varying, unknown part, which for each time step  $k$  is Gaussian distributed according to  $\Delta \mathbf{C} \sim \mathcal{N}(\hat{\mathbf{C}}, \mathbf{\Sigma})$ . In the second estimator, the friction vector is modeled as a realization from a GP with mean function  $\hat{\boldsymbol{\mu}}$  and covariance function  $\mathbf{\Sigma}$ ,

$$\boldsymbol{\mu}^y(\boldsymbol{\alpha}(\mathbf{x}^e, \mathbf{u}^e)) \sim \mathcal{GP}(\hat{\boldsymbol{\mu}}(\boldsymbol{\alpha}(\mathbf{x}^e, \mathbf{u}^e)), \mathbf{\Sigma}(\boldsymbol{\alpha}(\mathbf{x}^e, \mathbf{u}^e))). \quad (8)$$

The difference is that (6) is a linear model, that is, it well represents the true tire-friction curve during normal driving conditions but does not capture aggressive maneuvering. On the other hand, (8) targets the full friction curve and therefore has more potential in providing high performance also for aggressive maneuvering, but is more complex and therefore needs more excitation of the driving data.

**Remark 2.** The estimator targeting the tire stiffness uses a model of the tire friction that is linear in the slip angle according to (6). This implies that the tire-stiffness estimator will underestimate the available tire friction as larger slip angles are attained (i.e., as more aggressive maneuvering is performed). Fig. 2 provides an illustration of why this occurs for an asphalt tire model. When the vehicle operates at nonzero slip angles, the estimated tire-friction model can be thought of as a line between the origin and the true tire friction. The slope of this line (the estimated stiffness) decreases as the slip angle increases and the tire-friction curve flattens.

Irrespective of using (6) or (8), with a time discretization using the sampling period  $T_s$ , the vehicle dynamics model used in the estimator can be written as

$$\mathbf{x}_{k+1}^e = \mathbf{f}^e(\mathbf{x}_k^e, \mathbf{u}_k^e) + \mathbf{g}^e(\mathbf{x}_k^e, \mathbf{u}_k^e) \boldsymbol{\mu}(\boldsymbol{\alpha}_k(\mathbf{x}_k^e, \mathbf{u}_k^e)), \quad (9)$$

which is obtained from inserting (6) or (8) into (1). In this paper, a simple forward

Euler method is used to discretize the dynamics for estimation purposes, but any discretization scheme can be applied. The lateral velocity and yaw rate are estimated as part of the state  $\mathbf{x}^e$  according to (1). The longitudinal velocity can be determined from the wheel speeds  $\{\omega_i\}_{i=f,r}$  and the longitudinal acceleration measurement  $a^X$ .

Our measurement model is based on a setup commonly available in production cars, namely the acceleration  $a^Y$ , and yaw-rate measurements  $\dot{\psi}$ , forming the measurement vector  $\mathbf{y} = [a^Y, \dot{\psi}]^\top$ . To relate  $\mathbf{y}_k$  to the estimator state  $\mathbf{x}_k^e$  at each time step  $k$ , note that  $a^Y$  can be extracted from the right-hand side of (1b), and the last measurement is the yaw rate. We model the measurement noise  $\mathbf{e}_k$  as zero-mean Gaussian distributed noise with covariance  $\mathbf{R}$  according to  $\mathbf{e}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R})$ . This results in the measurement model

$$\mathbf{y}_k = \mathbf{h}(\mathbf{x}_k^e, \mathbf{u}_k^e) + \mathbf{D}(\mathbf{x}_k^e, \mathbf{u}_k^e) \boldsymbol{\mu}(\boldsymbol{\alpha}_k(\mathbf{x}_k^e, \mathbf{u}_k^e)) + \mathbf{e}_k. \quad (10)$$

The measurement covariance  $\mathbf{R}$  is assumed known a priori. This is reasonable, since the measurement noise can oftentimes be determined from prior experiments and data sheets. Hence, the estimation problem consists of estimating the vehicle state  $\mathbf{x}_k^e$  and friction statistics  $\{\hat{\boldsymbol{\mu}}, \boldsymbol{\Sigma}\}$ , where  $\{\hat{\boldsymbol{\mu}}, \boldsymbol{\Sigma}\}$  represent the initial slope (the tire stiffness) or the full friction curve, depending on the approach taken.

### 3.2. Joint State and Friction-Function Learning

The vehicle state  $\mathbf{x}_k^e$  needs to be estimated concurrently with the friction function  $\boldsymbol{\mu}$  at each time step  $k$ . This is easier to do reliably for moderate slip angles corresponding to normal driving as data gathering in the nonlinear regime is potentially difficult. Furthermore, the sensors are automotive grade and therefore have significant noise. Hence, to properly account for the inherent uncertainty and to be able to integrate properly with SNMPC, we approach the estimation problem in a Bayesian framework. The estimation problem is highly non-Gaussian, which is further amplified when we model the tire-friction model using a GP. We therefore leverage a particle-filter based estimator that estimates the vehicle state  $\mathbf{x}^e$  and friction function  $\boldsymbol{\mu}$ . The particle filter approximates the joint posterior density

$$p(\boldsymbol{\mu}_k, \mathbf{x}_{0:k}^e | \mathbf{y}_{0:k}) \quad (11)$$

at each time step  $k$ , that is, the posterior density function of the friction function  $\boldsymbol{\mu}_k$  and state trajectory  $\mathbf{x}_{0:k}^e = \{\mathbf{x}_0^e, \dots, \mathbf{x}_k^e\}$  from time step 0 to  $k$ , given the measurement history  $\mathbf{y}_{0:k} = \{\mathbf{y}_0, \dots, \mathbf{y}_k\}$ . We use the standard decomposition of (11) into

$$p(\boldsymbol{\mu}_k, \mathbf{x}_{0:k}^e | \mathbf{y}_{0:k}) = p(\boldsymbol{\mu}_k | \mathbf{x}_{0:k}^e, \mathbf{y}_{0:k}) p(\mathbf{x}_{0:k}^e | \mathbf{y}_{0:k}). \quad (12)$$

The two densities on the right-hand side of (12) can be estimated recursively. We estimate the state trajectory density  $p(\mathbf{x}_{0:k}^e | \mathbf{y}_{0:k})$  by a set of  $N_{\text{PF}}$  weighted state trajectories as

$$p(\mathbf{x}_{0:k}^e | \mathbf{y}_{0:k}) \approx \sum_{i=1}^{N_{\text{PF}}} q_k^i \delta_{\mathbf{x}_{0:k}^e, i}(\mathbf{x}_{0:k}^e), \quad (13)$$

where  $q_k^i$  is the weight of the  $i^{\text{th}}$  state trajectory  $\mathbf{x}_{0:k}^{e,i}$  and  $\delta(\cdot)$  is the Dirac delta. Given the state trajectory, we can compute the sufficient statistics necessary to approximate  $p(\boldsymbol{\mu}_k | \mathbf{x}_{0:k}^{e,i}, \mathbf{y}_{0:k})$  for each particle. Because we determine the state trajectory from the particle filter, the computations leading up to the estimation of  $p(\boldsymbol{\mu}_k | \mathbf{x}_{0:k}^{e,i}, \mathbf{y}_{0:k})$  are analytic.

To determine the covariance and mean function, we can marginalize out the state trajectory from  $p(\boldsymbol{\mu}_k | \mathbf{x}_{0:k}^{e,i}, \mathbf{y}_{0:k})$  according to

$$p(\boldsymbol{\mu}_k | \mathbf{y}_{0:k}) = \int p(\boldsymbol{\mu}_k | \mathbf{x}_{0:k}^e, \mathbf{y}_{0:k}) p(\mathbf{x}_{0:k}^e | \mathbf{y}_{0:k}) d\mathbf{x}_{0:k}^e \approx \sum_{i=1}^{N_{\text{PF}}} q_k^i p(\boldsymbol{\mu}_k | \mathbf{x}_{0:k}^{e,i}, \mathbf{y}_{0:k}), \quad (14)$$

from where the mean function  $\hat{\boldsymbol{\mu}}_k(\boldsymbol{\alpha}_k(\mathbf{x}_k^e, \mathbf{u}_k^e))$  and covariance function  $\boldsymbol{\Sigma}_k(\boldsymbol{\alpha}_k(\mathbf{x}_k^e, \mathbf{u}_k^e))$  needed in the SNMPC problem (21) can be extracted.

### 3.3. Computationally Efficient Particle-Filter Implementation

To enable real-time implementation of the particle filter using automotive-grade sensors, we rely on the marginalization concept, implying that parts of the estimation problem is analytic. For the case of tire-stiffness estimation, by modeling the unknown friction slope in (6) as a Gaussian random variable according to  $\Delta \mathbf{C} \sim \mathcal{N}(\hat{\mathbf{C}}, \boldsymbol{\Sigma})$ , we can estimate the state trajectory with a particle filter, and the sufficient statistics of the tire-stiffness parameters are subsequently updated analytically by relying on suitable *conjugate priors* [18].

On the other hand, for the GP-based friction estimator, a bottleneck in some of the proposed GP-SSM methods is the computational load. In this paper, we use a computationally efficient reduced-rank GP-SSM framework presented in [27,28], where the GP is approximated by a basis function expansion using the Laplace operator eigenfunctions

$$\phi^j(x) = \frac{1}{\sqrt{L}} \sin\left(\frac{\pi j(x+L)}{2L}\right), \quad (15)$$

defined on the interval  $[-L, L]$ , such that

$$\mu_i^y \approx \sum_j \gamma_i^j \phi^j(\alpha_i), \quad (16)$$

where the weights  $\gamma_i^j$  are Gaussian random variables with unknown mean and covariance, whose prior depends on the spectral density that is a function of the eigenvalues in (15). For convenience, we express the prior on the coefficients  $\gamma_i^j$  at time step  $k=0$  as a zero-mean matrix-normal ( $\mathcal{MN}$ ) distribution over  $\mathbf{A}$ ,

$$\mathbf{A} \sim \mathcal{MN}(\mathbf{0}, \mathbf{Q}, \mathbf{V}), \quad (17)$$

with right covariance  $\mathbf{Q}$  and left covariance  $\mathbf{V}$ . We can write the basis-function ex-

pansion in matrix form as

$$\boldsymbol{\mu} = \underbrace{\begin{bmatrix} \gamma_f^1 & \cdots & \gamma_f^m & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \cdots & \mathbf{0} & \gamma_r^1 & \cdots & \gamma_r^m \end{bmatrix}}_{\mathbf{A} = \begin{bmatrix} \mathbf{A}_f & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_r \end{bmatrix}} \underbrace{\begin{bmatrix} \phi^1(\alpha_f) \\ \vdots \\ \phi^m(\alpha_f) \\ \phi^1(\alpha_r) \\ \vdots \\ \phi^m(\alpha_r) \end{bmatrix}}_{\boldsymbol{\varphi}(\boldsymbol{\alpha})}, \quad (18)$$

where  $\gamma_i^j$  are the weights to be learned and  $m$  is the total number of basis functions. This formulation leads to a computationally efficient marginalized particle-filter implementation where, similar to the linear case, we can update the sufficient statistics analytically.

**Remark 3.** According to well-established tire models (e.g., [29], see Fig. 1), it is known that the tire-friction function is antisymmetric. At the same time, the basis-function expansion (16) is a weighted sum of sinusoids for each dimension. Hence, we can trivially enforce antisymmetry by restricting the sum in (16) to even values for  $j$  (i.e.,  $j = 2, 4, \dots, m$ ). This substantially reduces the number of parameters to estimate and at the same time ensures that the estimated friction function (5) passes through the origin.

**Remark 4.** Because of the approximation of the GP as a basis-function expansion, the infinite-dimensional function learning problem is reduced to a finite-dimensional parameter learning problem. Hence, in the end, both methods, while algorithmically slightly different, conceptually both estimate parameters.

#### 4. STOCHASTIC NONLINEAR MODEL PREDICTIVE CONTROL FOR REAL-TIME VEHICLE CONTROL

Based on (1)–(3) and using a fourth-order Runge-Kutta method with sampling period  $T_s$ , the complete nonlinear vehicle model used for control can be written in the form

$$\mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k, \mathbf{u}_k, \boldsymbol{\mu}(\mathbf{x}_k, \mathbf{u}_k)), \quad (19)$$

where  $\mathbf{x}_k \in \mathbb{R}^{n_x}$  denotes the state,  $\mathbf{u}_k \in \mathbb{R}^{n_u}$  the control inputs, and  $\boldsymbol{\mu}(\mathbf{x}_k, \mathbf{u}_k)$  is the tire-friction function, which is the disturbance (the process noise) that we estimate at each time step  $k$  using the methods described in Sec. 3. The state and control vector in the SNMPC formulation are

$$\begin{aligned} \mathbf{x} &:= [p^{\mathbf{X}}, p^{\mathbf{Y}}, \psi, v^{\mathbf{X}}, v^{\mathbf{Y}}, \dot{\psi}, \delta_f]^{\top}, \\ \mathbf{u} &:= [\dot{\delta}_f, \omega_f, \omega_r]^{\top}, \end{aligned} \quad (20)$$

such that  $n_x = 7$  and  $n_u = 3$ .

At each sampling time, based on the given state estimate  $\hat{\mathbf{x}}_t$  and covariance  $\mathbf{P}_t$  at

the current time step  $t$ , the SNMPC solves the following tracking-type OCP

$$\min_{\mathbf{x}_k, \mathbf{u}_k, \mathbf{P}_k} \sum_{k=0}^{N_{\text{MPC}}-1} l_k(\mathbf{x}_k, \mathbf{u}_k) + l_{N_{\text{MPC}}}(\mathbf{x}_{N_{\text{MPC}}}) \quad (21a)$$

$$\text{s.t. } \mathbf{x}_0 = \hat{\mathbf{x}}_t, \quad \mathbf{P}_0 = \mathbf{P}_t, \quad (21b)$$

$$\forall k \in \{0, \dots, N_{\text{MPC}} - 1\} :$$

$$\mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k, \mathbf{u}_k, \hat{\boldsymbol{\mu}}_t(\mathbf{x}_k, \mathbf{u}_k)), \quad (21c)$$

$$\mathbf{P}_{k+1} = \mathbf{F}_k \mathbf{P}_k \mathbf{F}_k^\top + \mathbf{G}_k \boldsymbol{\Sigma}_t(\mathbf{x}_k, \mathbf{u}_k) \mathbf{G}_k^\top, \quad (21d)$$

$$Pr(\mathbf{c}(\mathbf{x}_k, \mathbf{u}_k) \leq 0) \geq 1 - \epsilon, \quad (21e)$$

where the control action is in the feedforward-feedback form  $\mathbf{u}_k = \mathbf{u}_{\text{ref},k} + \mathbf{K}(\mathbf{x}_k - \mathbf{x}_{\text{ref},k}) + \Delta \mathbf{u}_k$  due to a prestabilizing controller, where  $\mathbf{u}_{\text{ref},k}$  and  $\mathbf{x}_{\text{ref},k}$  are provided by a reference generator, for example, a motion planner [1,2]. The matrices  $\mathbf{F}_k(\cdot)$ ,  $\mathbf{G}_k(\cdot)$  are the Jacobian matrices of (19) linearized with respect to  $\mathbf{x}$  and  $\boldsymbol{\mu}$ , respectively, which are functions of the state  $\mathbf{x}_k$ , control  $\mathbf{u}_k$  and mean friction estimate  $\hat{\boldsymbol{\mu}}_t$ . The state covariance propagation in (21d) corresponds to the extended Kalman filter (EKF) equations to define the state covariance matrix  $\mathbf{P}_k$  for  $k = 1, \dots, N_{\text{MPC}}$ .

In linearization-based SNMPC, the model disturbance is Gaussian assumed and state independent, for example, see [22]. When we model the friction as a Gaussian random variable, as we do for the estimator targeting the initial slope of the tire-friction curve, the estimator and SNMPC can be straightforwardly connected. However, when the tire-friction function is a GP and therefore state dependent, we need to adapt the covariance propagation (21d) to account for this state dependence by integrating the estimated tire-friction function with the SNMPC (see Sec. 5). This amounts to modifying the calculation of the Jacobian matrices  $\mathbf{F}_k(\cdot)$  and  $\mathbf{G}_k(\cdot)$ .

#### 4.1. Objective Function and Inequality Constraints

In our tracking-type OCP formulation, we consider the stage cost and terminal cost in (21a) to be the least-squares functions

$$\begin{aligned} l_k(\cdot) &= \frac{1}{2} \|\mathbf{x}_k - \mathbf{x}_{\text{ref},k}\|_Q^2 + \frac{1}{2} \|\mathbf{u}_k - \mathbf{u}_{\text{ref},k}\|_R^2 + r_s s_k, \\ l_{N_{\text{MPC}}}(\cdot) &= \frac{1}{2} \|\mathbf{x}_{N_{\text{MPC}}} - \mathbf{x}_{\text{ref},N_{\text{MPC}}}\|_{Q_{N_{\text{MPC}}}}^2, \end{aligned} \quad (22)$$

which includes a term for both state and control reference tracking, and an L1 term for penalizing the slack variable  $s_k \geq 0$ . The constraints  $\mathbf{c}(\mathbf{x}_k, \mathbf{u}_k) \leq 0$  in the OCP (21) consist of geometric and physical limitations on the system. In practice, it is important to reformulate these requirements as soft constraints, based on the slack variable  $s_k$ . Otherwise, the problem may become infeasible due to unknown disturbances and modeling errors, and the controller will cease to operate. The constraints in (21e) therefore

include soft bounds on

$$p_{\min}^{\mathbf{Y}} - s_k \leq p_k^{\mathbf{Y}} \leq p_{\max}^{\mathbf{Y}} + s_k, \quad (23a)$$

$$-\delta_{f,\max} - s_k \leq \delta_{f,k} \leq \delta_{f,\max} + s_k, \quad (23b)$$

$$-\dot{\delta}_{f,\max} - s_k \leq \dot{\delta}_{f,k} \leq \dot{\delta}_{f,\max} + s_k, \quad (23c)$$

$$\omega_{\min} - s_k \leq \omega_{i,k} \leq \omega_{\max} + s_k, \quad i \in \{f, r\}, \quad (23d)$$

$$-\alpha_{\max} - s_k \leq \alpha_{i,k} \leq \alpha_{\max} + s_k, \quad i \in \{f, r\}, \quad (23e)$$

where the slack variable is restricted to be nonnegative  $s_k \geq 0$ . Eq. (23a) bounds the lateral position and is used to ensure that the vehicle stays on the road. Eqs. (23b)-(23e) bound the wheel angle, inputs, and slip angles, where the slip angles are defined in (3). The input constraints in (23c)-(23d) are imposed as soft constraints, due to the state feedback formulation for the control action  $\mathbf{u}_k = \mathbf{u}_{\text{ref},k} + K(\mathbf{x}_k - \mathbf{x}_{\text{ref},k}) + \Delta \mathbf{u}_k$  in (21). An affine scaling of the slack variable in the soft constraints (23) can be used to account for the different dimensions of each of the variables. Using one slack variable  $s_k \geq 0$  in each of the soft constraints (23) corresponds to penalizing the maximum violation of the constraints at each time step. The weight value  $r_s > 0$  in (22) for the L1 slack penalty needs to be chosen sufficiently large [30], such that all slack variables are zero, i.e.,  $s_k = 0$  if a feasible solution exists. Alternatively, a separate slack variable and separate weight value could be used for each of the constraints. However, this would increase the total number of optimization variables and therefore increase the computational complexity of solving the resulting OCP in general.

#### 4.1.1. Stability Constraints for the Tire-Stiffness SNMPC

As the tire-stiffness estimator targets the linear region of the tire-friction curve, for aggressive maneuvering we need to ensure that the SNMPC does not operate excessively outside the validity region of the estimator. To this end, when the SNMPC is connected to the tire-stiffness estimator, we include the additional constraints

$$|\dot{\psi}_k v_k^x| \leq 0.85\mu_c g, \quad \left| \frac{v_k^y}{v_k^x} \right| \leq \tan^{-1}(0.02\mu_c g), \quad (24)$$

where  $\mu_c$  is the friction coefficient. The constraints in (24) prevent the vehicle from entering regions of high lateral acceleration and side slip, and can be found in [25, Chapter 8]. We refer to (24) as *stability constraints*. The stability constraints depend on the (peak) friction coefficient  $\mu_c$ . Experimental studies suggest that using a monotonic relationship is sufficient to differentiate between asphalt and snow [14,31]. In this work, we use a linear relationship to approximate the friction coefficient as a function of the tire-stiffness estimate,

$$\mu_c \approx \min \left( \frac{a(C_{f,n}^y + \Delta C_f^y + C_{r,n}^y + \Delta C_r^y)}{2}, 1 \right), \quad (25)$$

where  $a$  is a constant that in this work is adjusted from Pacejka models for asphalt and snow. The central idea of (25) is that the bounds on the acceleration and sideslip should tighten as the friction coefficient, and consequently the cornering stiffness, decreases. For surfaces such as wet asphalt, which may have a high cornering stiffness but lower

road friction, (25) is conservative because the stiffness estimator underestimates the true stiffness as the tires saturate (as in Fig. 2, see Remark 2).

#### 4.2. Probabilistic Chance Constraints

To enforce the probabilistic chance constraints in (21e), we reformulate them as deterministic constraints as in [32], where the  $j^{\text{th}}$  constraint is written as

$$c_j(\mathbf{x}_k, \mathbf{u}_k) + \nu \sqrt{\frac{\partial c_j}{\partial \mathbf{x}_k} \mathbf{P}_k \frac{\partial c_j}{\partial \mathbf{x}_k}^T} \leq 0, \quad (26)$$

where  $\nu$  is referred to as the back-off coefficient and depends on the desired probability threshold  $\epsilon$  and assumptions about the resulting state distribution. The backoff coefficient for Cantelli's inequality,  $\nu = \sqrt{\frac{1-\epsilon}{\epsilon}}$ , holds regardless of the underlying distribution but is conservative. In this work, we approximately assume normal-distributed state trajectories and set

$$\nu = \sqrt{2} \operatorname{erf}^{-1}(1 - 2\epsilon), \quad (27)$$

where  $\operatorname{erf}^{-1}(\cdot)$  is the inverse error function.

#### 4.3. Online SNMPC and Software Implementation

The resulting SNMPC implementation, proposed recently in [22], uses the QP solver PRESAS [21], which applies block-structured factorization techniques with low-rank updates to preconditioning of an iterative solver within a primal active-set algorithm with tailored initialization methods. This results in a simple, efficient and reliable QP solver suitable for embedded control hardware. The algorithm uses a tailored Jacobian approximation along with an adjoint-based SQP method that allows for the numerical elimination of the covariance matrices from the SQP subproblem, which reduces the computation time when compared to standard SQP methods applied to SNMPC [22]. The nonlinear function and derivative evaluations, for the preparation of each SQP subproblem, are performed using algorithmic differentiation (AD) and C code generation in CasADi [33]. In addition, a standard line search method is used [34] to improve the closed-loop convergence of the SQP-based SNMPC controller.

We solve the nonlinear OCP (21) only approximately at each control time step by a tailored implementation of the real-time iteration (RTI) scheme [35] using the adjoint-based SQP method [22]. RTI performs one SQP iteration per time step of control, and uses a continuation-based warm starting of the state and control trajectories from one time step to the next [36]. RTI may not approximate well the OCP if the problem is linearized far from a local minimum. However, as we envision an autonomous-driving system [37], the MPC tracks a trajectory generated by a motion planner to be a suitable reference for linearization, e.g., kinematically feasible and constraint aware. Thus, RTI seems a suitable approach for this application.

**Remark 5.** Due to time delays resulting from vehicle network communication and computation delays, it is beneficial to introduce time-delay compensation. Denote the time delay as  $T_d$ . Time-delay compensation can be done by defining the estimates at

time  $t$  entering the OCP (21) to equal the predicted estimates at time  $t+T_d$ , computed from the estimates at time  $t$  and the input signals up until time  $t$  [8,37].

## 5. Adaptation of SNMPC to the Learned Tire-Friction Function

In determining the mean  $\hat{\boldsymbol{\mu}}_t$  and covariance  $\boldsymbol{\Sigma}_t$  of the tire-friction in (21), we use that each particle retains its own estimate  $\boldsymbol{\mu}_k^i$  together with the weight  $q_k^i$ . Then, we can estimate the mean and covariance by choosing the  $i$ th particle that fulfills  $i = \arg \max_{i \in [1, N]} q_k^i$ , or we can choose to estimate the mean and covariance as the weighted mean among the particles.

When integrating the tire-stiffness estimator with the SNMPC, the connection is already there because the tire-stiffness estimator directly outputs the mean and covariance of the tire-stiffness estimate, which is already in the form that the SNMPC admits. However, for the GP case we need to make some adjustments. Using the single-track model and the basis-function expansion of the friction,  $\boldsymbol{\mu} \approx \mathbf{A}\boldsymbol{\varphi}(\boldsymbol{\alpha}(\mathbf{x}))$ , we can write the vehicle model (19) used by the SNMPC in the same format as (9),

$$\mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k, \mathbf{u}_k) + \mathbf{g}(\mathbf{x}_k, \mathbf{u}_k)\hat{\boldsymbol{\mu}}_k(\boldsymbol{\alpha}(\mathbf{x}_k)). \quad (28)$$

Note that  $\delta_f$  is part of the state in the MPC formulation (20) and the slip angle  $\boldsymbol{\alpha}(\cdot)$  is therefore a function of only  $\mathbf{x}$ .

In linearization-based SNMPC [22], the parametric uncertainty is typically modeled according to a Gaussian random variable,  $\boldsymbol{\mu} \sim \mathcal{N}(\hat{\boldsymbol{\mu}}, \boldsymbol{\Sigma})$ , where the uncertainty  $\boldsymbol{\mu}$  is state independent. This leads to Jacobian matrices  $\mathbf{F}_k = \frac{\partial \mathbf{f}}{\partial \mathbf{x}}(\mathbf{x}_k, \mathbf{u}_k, \hat{\boldsymbol{\mu}}_k)$  and  $\mathbf{G}_k = \frac{\partial \mathbf{f}}{\partial \boldsymbol{\mu}}(\mathbf{x}_k, \mathbf{u}_k, \hat{\boldsymbol{\mu}}_k)$  in (21d). However, in our case, the friction uncertainty is a function that is modeled according to a GP at each time step  $k$ ,  $\boldsymbol{\mu}(\boldsymbol{\alpha}(\mathbf{x})) \sim \mathcal{GP}(\hat{\boldsymbol{\mu}}_k(\boldsymbol{\alpha}(\mathbf{x})), \boldsymbol{\Sigma}(\boldsymbol{\alpha}(\mathbf{x})))$ . Hence, we want to determine the equivalent to (21d) for the state-dependent uncertainty  $\boldsymbol{\mu}(\boldsymbol{\alpha}(\mathbf{x}))$ , by finding expressions for the involved Jacobians  $\mathbf{F}_k$ ,  $\mathbf{G}_k$ . Starting from (28), by using the chain rule and  $\boldsymbol{\mu} \approx \mathbf{A}\boldsymbol{\varphi}(\boldsymbol{\alpha}(\mathbf{x}))$ ,

$$\begin{aligned} \mathbf{F}_k &= \frac{\partial \mathbf{f}}{\partial \mathbf{x}}(\mathbf{x}_k, \mathbf{u}_k) + \frac{\partial \mathbf{g}}{\partial \mathbf{x}}(\mathbf{x}_k, \mathbf{u}_k)\hat{\boldsymbol{\mu}}_t(\boldsymbol{\alpha}(\mathbf{x}_k)) \\ &\quad + \mathbf{g}(\mathbf{x}_k, \mathbf{u}_k)\frac{\partial \hat{\boldsymbol{\mu}}_t(\boldsymbol{\alpha}(\mathbf{x}_k))}{\partial \mathbf{x}} \\ &= \frac{\partial \mathbf{f}}{\partial \mathbf{x}}(\mathbf{x}_k, \mathbf{u}_k) + \frac{\partial \mathbf{g}}{\partial \mathbf{x}}(\mathbf{x}_k, \mathbf{u}_k)\hat{\mathbf{A}}_t\boldsymbol{\varphi}(\boldsymbol{\alpha}(\mathbf{x}_k)) \\ &\quad + \mathbf{g}(\mathbf{x}_k, \mathbf{u}_k)\hat{\mathbf{A}}_t\frac{\partial \boldsymbol{\varphi}}{\partial \boldsymbol{\alpha}}(\boldsymbol{\alpha}(\mathbf{x}_k))\frac{\partial \boldsymbol{\alpha}}{\partial \mathbf{x}}(\mathbf{x}_k). \end{aligned} \quad (29)$$

Hence, for ease of notation, if we momentarily define  $\mathbf{g}_k := \mathbf{g}(\mathbf{x}_k, \mathbf{u}_k)$ ,  $\boldsymbol{\varphi}_k := \boldsymbol{\varphi}(\boldsymbol{\alpha}(\mathbf{x}_k))$ , we can approximate the covariance propagation as

$$\begin{aligned} \mathbb{E}[\mathbf{x}_{k+1}\mathbf{x}_{k+1}^\top] &\approx \mathbb{E}[(\mathbf{F}_k\mathbf{x}_k + \mathbf{g}_k\hat{\mathbf{A}}_t\boldsymbol{\varphi}_k)(\mathbf{F}_k\mathbf{x}_k + \mathbf{g}_k\hat{\mathbf{A}}_t\boldsymbol{\varphi}_k)^\top] \\ &= \mathbf{F}_k\mathbf{P}_k\mathbf{F}_k^\top + \mathbb{E}(\mathbf{g}_k\hat{\mathbf{A}}_t\boldsymbol{\varphi}_k\boldsymbol{\varphi}_k^\top\hat{\mathbf{A}}_t^\top\mathbf{g}_k^\top) \\ &= \mathbf{F}_k\mathbf{P}_k\mathbf{F}_k^\top + \underbrace{\mathbf{g}_k}_{\mathbf{G}_k}\underbrace{\text{cov}(\hat{\mathbf{A}}_t\boldsymbol{\varphi}_k)}_{\boldsymbol{\Sigma}_k}\mathbf{g}_k^\top. \end{aligned} \quad (30)$$

Eq. (30) leads to the SNMPC formulation, where the integration of the tire-friction



estimates is done by using the estimated mean in the state prediction (21c) and the estimated covariance function through the Jacobians in the covariance propagation (21d). Algorithm 1 summarizes our proposed friction-adaptive SNMPC strategy.

---

**Algorithm 1** Proposed Real-Time SNMPC with Friction Adaptation

---

- 1: **for** each time step  $k$  **do**
  - 2:     Estimate current state vector  $\hat{\mathbf{x}}_k$ , tire-friction estimate  $\hat{\boldsymbol{\mu}}_k(\boldsymbol{\alpha}(\mathbf{x}_k^e, \mathbf{u}_k^e))$  and covariance function  $\boldsymbol{\Sigma}_k(\boldsymbol{\alpha}(\mathbf{x}_k^e, \mathbf{u}_k^e))$ , using either of [18] or [19].
  - 3:     Perform SQP iteration(s) to approximately solve SNMPC problem (21), using state estimate in (21b), mean friction function in (21c), and covariance propagation in (21d) based on (29)-(30) for tire-friction estimation.
  - 4: **end for**
- 

## 6. Friction-Adaptive Simulation Results

In this section, we validate the proposed friction-adaptive SNMPC in simulation and in a real-time computing environment. In the simulation studies, we consider a sequence of nine double lane-change maneuvers similar to the standardized ISO 3888-2 [38] double lane-change maneuver commonly used in vehicle stability tests, with the middle three on snow and the rest on dry asphalt. To investigate the learning behavior of the controllers, the surface change occurs during a straight portion, where the friction curve is unobservable. The reference is generated with Bezier polynomials and the position, heading, and longitudinal velocity are given to the controllers to track. For simplicity, we set the velocity reference to be constant and evaluate the proposed control strategy for two different reference velocities. First, we evaluate different combinations of controllers and estimators in Matlab simulations using the same vehicle model for both generating synthetic data and the controller, but with different friction models. Second, we evaluate the controllers in simulation using CarSim [39]. Third, we verify the real-time feasibility by executing the proposed method in a dSPACE MicroAutoBox-II equipped with an IBM PPC 900MHz processor and 16 MB main memory, which closely resembles the capabilities of current and near-future embedded microcontrollers for automotive applications.

The vehicle parameters are from a mid-size SUV, and the tire parameters for the different surfaces are taken from [40]. The simulation model is the nonlinear single-track model (1) with a Pacejka tire model [29] with the friction ellipse for modeling combined slip,

$$\begin{aligned}
 F_i^x &= \mu_i^x F_i^z \sin(D_i^x \arctan(B_i^x(1 - E_i^x)\lambda_i + E_i^x \arctan(B_i^x \lambda_i))), \\
 F_i^y &= \eta_i \mu_i^y F_i^z \sin(D_i^y \arctan(B_i^y(1 - E_i^y)\alpha_i + E_i^y \arctan(B_i^y \alpha_i))), \\
 \eta_i &= \sqrt{1 - \left(\frac{F_i^x}{\mu_i^x F_i^z}\right)^2},
 \end{aligned} \tag{31}$$

where  $\mu_i^j$ ,  $B_i^j$ ,  $D_i^j$  and  $E_i^j$ , for  $i \in \{f, r\}, j \in \{x, y\}$ , are the friction coefficients and stiffness, shape, and curvature factors, with values from [40]. The noise levels in the measurement model (10) are taken from those of a low-cost inertial measurement unit typical for automotive applications. The estimators use  $N_{\text{PF}} = 100$  particles

and the tire-friction estimator uses  $m = 10$  basis functions, five in each direction by exploiting antisymmetry. The initial estimates and the different tuning parameters in the estimator are fairly generic, and the same as in [26].

We evaluate the following seven controllers:

- **FRICION-SNMPC**: the proposed adaptive SNMPC method (Algorithm 1) with tire-friction estimator from [19].
- **STIFFNESS-SNMPC**: the proposed adaptive SNMPC method (Algorithm 1) with tire-stiffness estimator from [18].
- **FRICION-NMPC**: Nominal adaptive NMPC with the tire-friction estimator in [19], i.e., it does not use the covariance propagation and chance constraints.
- **STIFFNESS-NMPC**: Nominal adaptive NMPC with the tire-stiffness estimator in [18], i.e., it does not use the covariance propagation and chance constraints.
- **SNOW-NMPC**: Nominal NMPC that uses fixed friction coefficient values in the Pacejka tire model (31) corresponding to a snow surface.
- **ASPHALT-NMPC**: Nominal NMPC that uses fixed friction coefficient values in the Pacejka tire model (31) corresponding to an asphalt surface.
- **ORACLE**: Nominal NMPC with the exact nonlinear tire-friction Pacejka model. This controller acts as ground truth and cannot be implemented in practice.

We include the ORACLE controller to provide a lower bound on cost and constraint violations for the simulations. Its performance cannot be achieved in practice because it is given the exact tire force curve used by the simulation model; in reality, there will be model mismatch due to inaccuracies in both the tire force and single-track vehicle models. All of the controllers perform 1 SQP iteration per planning instance. The estimators are executed at 100Hz and the different MPCs at 20Hz, with a prediction horizon of 2s. The constraint satisfaction probability for the SNMPCs (FRICION-SNMPC and STIFFNESS-SNMPC) are set to 95%, i.e.,  $\epsilon = 0.05$ . The least-squares cost (22) prioritizes the lateral position and wheel speed inputs. The metrics we use to evaluate the controllers are cost and score, and are computed as

$$\text{Cost} = \sum_k l(\mathbf{x}_k, \mathbf{u}_k), \quad (32)$$


$$\text{Score} = \sum_k ((y_k - y_{\max})_+ + (y_{\min} - y_k)_+) T_s, \quad (33)$$

summed over the simulation time, where  $(\cdot)_+ = \max(\cdot, 0)$  and  $l(\cdot)$  corresponds to the stage cost in the MPC objective function.

**Remark 6.** In all of the simulations for the considered maneuver, ASPHALT-NMPC diverged after the transition to snow and is therefore omitted from the results section.

### 6.1. Closed-loop Estimation Results

Fig. 3 shows the front tire-friction estimates using the proposed GP-based estimator in Algorithm 1 in closed loop for one realization of the considered maneuver, with the resulting  $\alpha_f$ . The estimates are initialized for the parameters on snow. As more measurements are gathered, the estimates converge. The changes from asphalt to snow and vice versa are clearly indicated by the estimator reacting to the abrupt changes. At the switch from asphalt to snow ( $t \approx 13$ s), the short transient behavior of the estimator can be seen from the slightly larger slip angles around just after the first



**Figure 3.** Estimated tire-friction function of the front wheel for one realization in closed-loop simulation with multiple double lane-change maneuvers and surface changes for  $v_{\text{ref}} = 22\text{m/s}$ . The trace of  $\alpha_f$  is indicated in black and its projection to the  $t$ - $\alpha_f$  plane is in black dashed.

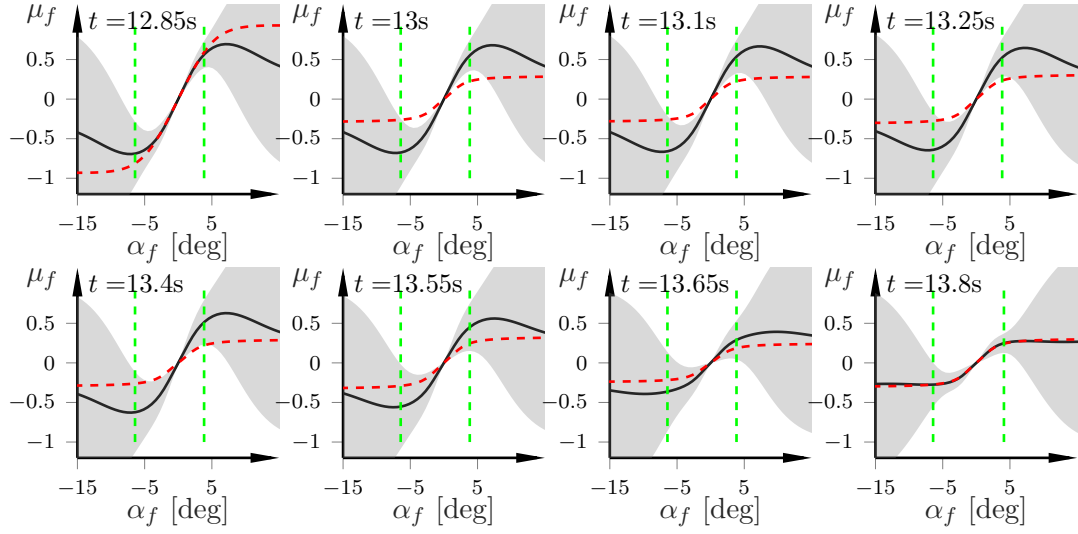
lane change is initiated on the new surface. However, the increased slip angles result in the estimator learning about the new nonlinearity, and the controller quickly adapts.

To give a more quantitative measure of how the estimator performs, Fig. 4 shows the results for the same realization based on different snapshots, when an abrupt surface change from dry asphalt to snow is about to occur. For the region of the slip angle corresponding to where data have been acquired, the estimates follow closely to the true friction model, and it is clear that where small amounts of data have been gathered, for example, for slip angles close to or beyond the peak, the uncertainty increases accordingly. The intuition is that such increase of uncertainty can be utilized by the SNMPC when determining the appropriate control commands. The estimator is quick to adapt to the surface change and converges within roughly 1s. Note, however, that even though it takes about 1s for the estimator to fully converge, the change in surface is reflected in the uncertainty estimates faster than that.


Similarly, Fig. 5 displays the stiffness estimates for the tire-stiffness estimator in Algorithm 1. In this closed-loop simulation, for  $v_{\text{ref}} = 22\text{m/s}$ , the nonlinear region of the tire-friction curve is excited. From the discussion in Remark 2, this leads to an underestimation of the stiffness. However, the true value is still covered in the 95% confidence intervals. From previous papers, when driving in the linear region, the estimates are highly accurate [18].

## 6.2. Closed-loop Control Results


To illustrate the reference path of the maneuver with surface changes, Fig. 6 shows the trajectories of FRICTION-SNMPC, SNOW-NMPC, and ORACLE when setting the velocity reference to  $v_{\text{ref}} = 22\text{m/s}$ . The results demonstrate that the reference trajectory is tracked well by the adaptive FRICTION-SNMPC controller and the ORACLE controller. The SNOW-NMPC controller performs conservatively in the asphalt sections, and also has difficulty tracking the reference in the snow portion; due to the aggressive nature of the maneuver and inaccuracies between the MPC and simulation model.



**Figure 4.** The tire-friction estimates (black solid) and estimated  $2\sigma$  confidence (gray area) at different time steps for one realization in closed-loop simulation with multiple double lane-change maneuvers and surface changes (see Fig. 3 for whole realization). The true tire-friction function is in red dashed, and the estimated range of the  $\alpha$  values that have been excited are indicated by the green dashed vertical lines. An abrupt surface change from asphalt to snow occurs right before  $t = 13s$ , and the estimator converges in less than 1s.



**Figure 5.** Estimated tire stiffness for one realization in closed-loop simulation with multiple double lane-change maneuvers and surface changes. The red line is the true stiffness, the slope of the nonlinear tire-friction curve at  $\alpha_i = 0$ . The black solid and gray area are the mean and  $2\sigma$  confidence interval from the stiffness estimator, respectively. The stiffness is underestimated when the tire friction saturates, see Remark 2. The true stiffness corresponds to the initial slopes in Fig. 4.



**Figure 6.** Resulting path for the various approaches during the lane-change maneuvers. Red and green dashed lines are the constraints and reference, respectively. The gray vertical dashed lines indicate the surface changes from asphalt-snow-asphalt. Color coding: FRICITION-SNMPC (black), ORACLE (blue), and SNOW-NMPC (red).

After the estimator transients, FRICITION-SNMPC seems to track the path slightly better. Note that ORACLE, which is an ideal controller that cannot be implemented in practice, does not seem to clearly outperform the proposed adaptive controller.

To numerically quantify the performance and show the merits of the proposed methods, we have executed two sets of 200 Monte-Carlo runs for the considered maneuver, with the velocity reference  $v_{\text{ref}} = 19\text{m/s}$  and  $v_{\text{ref}} = 22\text{m/s}$ , respectively. For each Monte-Carlo simulation, the Pacejka parameters for each road surface are randomly perturbed, with samples drawn from a uniform distribution up to  $\pm 10\%$  on asphalt and  $\pm 20\%$  on snow. In addition, the measurements are subject to zero-mean Gaussian noise with covariance values chosen based on specifications from automotive-grade sensors. Table 1 and 2 show the results. Clearly, when comparing the friction-adaptive controllers, the stochastic formulations give substantially less constraint violations than their nominal counterparts; for example, the combination of SNMPC with GP-based tire-friction estimation (FRICITION-SNMPC) leads to constraint satisfaction at all times. Notably, the nominal NMPC with tire-friction estimation (FRICITION-NMPC) destabilizes for large velocities (Table 2), but this does not happen for the nominal counterpart using tire-stiffness estimation. Insight into this is given by inspection of Figs. 4 and 5. The tire-stiffness estimator gives an underestimation of the actual stiffness when exciting the nonlinear region of the tire-friction curve, see Remark 2. Hence, the effect of having slightly biased estimates, is that the controller acts conservatively and therefore implicitly injects robustness into the approach.

In summary, the tire-stiffness estimator gives a good compromise in the trade-off between cost and score, whereas leveraging the potential of full tire-friction curve estimation in an SNMPC formulation gives the best performance. However, care needs to be taken in the design of the MPC used in combination with estimating the full tire-friction curve, as such an estimator demands more from the employed controller.


Finally, in our evaluation, we show the stability constraints (24) for STIFFNESS-SNMPC in the considered realization. As the surface shifts from asphalt to snow, the constraints tighten to maintain performance of the control strategy. These additional constraints turned out to be important for the adequate functioning when estimating a linear tire-friction curve.

**Table 1.** Results for 200 Monte-Carlo runs with  $v_{\text{ref}} = 19\text{m/s}$ .


Method	Mean Cost	Max Cost	Mean Score	Max Score
FRICITION-SNMPC	0.938	1.008	0	0
STIFFNESS-SNMPC	1.254	2.058	0.008	0.097
FRICITION-NMPC	0.917	0.951	0	0
STIFFNESS-NMPC	1.740	4.929	0.018	0.071
SNOW-NMPC	3.332	3.439	0.035	0.114
ORACLE	0.710	0.725	0	0

**Table 2.** Results for 200 Monte-Carlo runs with  $v_{\text{ref}} = 22\text{m/s}$ .


Method	Mean Cost	Max Cost	Mean Score	Max Score
FRICITION-SNMPC	2.187	2.868	0	0
STIFFNESS-SNMPC	2.716	3.098	0.002	0.055
FRICITION-NMPC	557	72,626	22	2,973
STIFFNESS-NMPC	2.833	5.90	0.006	0.0677
SNOW-NMPC	4.391	5.05	0.077	0.111
ORACLE	1.682	1.709	0	0



**Figure 7.** Stability constraints enforced for the controllers using the stiffness estimator (STIFFNESS-SNMPC in magenta and STIFFNESS-NMPC in blue), where the middle portion is on snow. Red dashed lines are the constraint boundaries, where the road friction is calculated with (25) using the estimator output. The constraints tighten during the snow portion. Although the differences are minor in this particular realization, STIFFNESS-NMPC has more constraint violations.



**Figure 8.** CarSim results of lateral position for different controllers during the lane-change maneuvers. Black and blue dashed lines are the constraints and reference, respectively. The gray area indicates snow surface. The black, red, and green lines indicate the trajectories for FRICION-SNMPC, SNOW-NMPC, and ASPHALT-NMPC.



**Figure 9.** Front wheel steering angle commands and closed-loop cost for the CarSim validation results of the vehicle maneuver with surface changes as illustrated in Fig. 8.

### 6.3. CarSim Validation Results

In this section, we further validate the method using the high-fidelity vehicle simulator CarSim. From Tables 1 and 2, it is clear that FRICION-SNMPC is the method that performs best according to the metrics used. Consequently, we choose to evaluate this control system in CarSim and compare with nominal controllers. For reproducibility, we use modules that are standard in CarSim for all components such as suspension, springs, tires, and so on, and the vehicle parameters are from a mid-size SUV. We set the longitudinal velocity reference to be constant,  $v_{\text{ref}}^X = 18\text{m/s}$ .

The performance of the estimator has been validated in several papers [19,26]. Here, we focus on the closed-loop control performance. Fig. 8 shows a comparison of the lateral position tracking performance between the three different controllers. The controller denoted by ASPHALT-NMPC shows significant overshoot and poor tracking performance once the surface switches to snow (indicated by the gray area). The results

**Table 3.** Timing results for Algorithm 1 using the GP-based tire-friction estimator for two sets of combinations of control horizon  $N_{\text{MPC}}$ , prediction length  $T$ , and number of particles  $N_{\text{PF}}$  on a dSPACE MicroAutoBox-II.

	$N_{\text{MPC}} = 15$ $T = 0.75\text{s}$ $N_{\text{PF}} = 200$		$N_{\text{MPC}} = 20$ $T = 1\text{s}$ $N_{\text{PF}} = 100$	
	Mean	Max	Mean	Max
QP iterations	3.65	25	6.57	25
SNMPC	16.4ms	29ms	23.1ms	37.8ms
Tire-friction estimator	10.4ms	10.5ms	5.5ms	5.6ms
Total turnaround time	26.7ms	39.4ms	28.6ms	43.3ms

demonstrate that the reference trajectory is tracked well by the FRICTION-SNMPC and SNOW-NMPC controllers, and that there are some transients in FRICTION-SNMPC due to the convergence time of the tire-friction estimator. The SNOW-NMPC controller performs conservatively as expected.

Fig. 9 displays the wheel steering angle and closed-loop cost for the three controllers. Note that the MPC cost function consists of more terms than the lateral position tracking error, and focusing solely on the lateral position can lead to wrong conclusions about performance. Fig. 9 shows that the surface change at about 21s destabilizes the vehicle when using ASPHALT-NMPC. Also, we can see that the adaptive FRICTION-SNMPC improves performance compared to the SNOW-NMPC throughout the maneuver, reducing the closed-loop cost from over 0.6 to about 0.1.

#### 6.4. Real-time Feasibility for Embedded Implementation

To conclude the evaluation, we assess the real-time computational feasibility of the proposed methods. In Tables 3 and 4, we show the closed-loop computation times of the different friction-adaptive SNMPC on a dSPACE MicroAutoBox-II rapid prototyping unit. The table includes the computation time for the same maneuver as before for two different velocity references and two sets of combinations of the control horizon length  $N_{\text{MPC}}$ , time prediction length  $T$ , and the number of particles  $N_{\text{PF}}$  that give good performance for the respective velocity reference. From these results, it is clear that the computation time for both the SNMPC and the estimators scale linearly and the total time remains well below the desired sampling time of 50ms.

## 7. Conclusion

The proposed friction-adaptive control strategy uses a recently developed, computationally efficient SNMPC formulation. In this paper, the proposed SNMPC models the tire-road friction as an external disturbance and is flexible in the sense that it can be combined with any friction estimator, as long as the estimator is capable of outputting the first two moments of the friction estimate. We implemented the control strategy by leveraging two recently developed Bayesian friction estimators that model different aspects of the tire-friction characteristics. The first estimator is based on GPs and models the full nonlinear tire-friction curve, whereas the second estimator restricts the estimation of the tire stiffness. However, note that the proposed control strategy



**Table 4.** Timing results for Algorithm 1 using the tire-stiffness estimator for two sets of combinations of control horizon  $N_{\text{MPC}}$ , prediction length  $T$ , and number of particles  $N_{\text{PF}}$  on a dSPACE MicroAutoBox-II.

	$N_{\text{MPC}} = 15$ $T = 0.75\text{s}$ $N_{\text{PF}} = 1000$		$N_{\text{MPC}} = 20$ $T = 1\text{s}$ $N_{\text{PF}} = 500$	
	Mean	Max	Mean	Max
QP iterations	4.03	20	3.97	25
SNMPC	10.4ms	20.6ms	13ms	32.7ms
Tire-stiffness estimator	12.2ms	14.1ms	6.2ms	7ms
Total turnaround time	22.6ms	34.2ms	19.1ms	39.5ms

is not limited to these two estimators

Our Monte-Carlo studies show that friction-adaptive control based on SNMPC is generally superior to a nominal NMPC in terms of tracking error and constraint violations, and that the tire-friction estimator is superior to the tire-stiffness estimator. However, while the tire-friction estimator has more potential than the estimator based on a linear friction model, our results indicate that the GP-based tire-friction estimator should be combined with SNMPC and not nominal NMPC, while the tire-stiffness estimator seems to work sufficiently well irrespective of choosing SNMC or nominal NMPC. While these results alone are inconclusive, as they were obtained with a specific set of parameters and for a particular maneuver, they show that care should be taken with respect to parameter choices and the controller to integrate with, when using an estimator targeting the full tire-friction curve.

The timing results, which we obtained using a dSPACE MicroAutoBox-II, indicate that the presented method is suitable for implementation on automotive embedded platforms. While the tire-friction estimator is more computationally demanding than the tire-stiffness estimator, both can run online with a sufficiently large number of particles. Furthermore, it is the SNMPC and not the estimator that takes up most of the computational demands. Hence, computational considerations do not seem to be the most important factor in determining which estimator to employ.

## References

- [1] Paden B, Cap M, Yong SZ, et al. A survey of motion planning and control techniques for self-driving urban vehicles. *IEEE Trans Intell Veh.* 2016;1(1):33–55.
- [2] Berntorp K. Path planning and integrated collision avoidance for autonomous vehicles. In: *Amer. Control Conf.*; May; Seattle, WA; 2017.
- [3] Lefèvre S, Vasquez D, Laugier C. A survey on motion prediction and risk assessment for intelligent vehicles. *Robomech J.* 2014;1(1):1.
- [4] Okamoto K, Berntorp K, Di Cairano S. Driver intention-based vehicle treat assessment using random forests and particle filtering. In: *IFAC World Congress*; Jul.; Toulouse, France; 2017.
- [5] Di Cairano S, Kolmanovsky IV. Real-time optimization and model predictive control for aerospace and automotive applications. In: *Amer. Control Conf.*; Jun.; Milwaukee, WI; 2018.
- [6] Hrovat D, Di Cairano S, Tseng HE, et al. The development of model predictive control in

- automotive industry: A survey. In: *Int. Conf. Control Applications*; Dubrovnik, Croatia; 2012.
- [7] Falcone P, Borrelli F, Asgari J, et al. Predictive active steering control for autonomous vehicle systems. *IEEE Trans Control Syst Technol.* 2007;15(3):566–580.
- [8] Berntorp K, Quirynen R, Uno T, et al. Trajectory tracking for autonomous vehicles on varying road surfaces by friction-adaptive nonlinear model predictive control. *Veh Syst Dyn.* 2019;58(5):705–725.
- [9] Carvalho A, Lefèvre S, Schildbach G, et al. Automated driving: The role of forecasts and uncertainty - a control perspective. *Eur J Control.* 2015;24:14–32.
- [10] Di Cairano S, Kalabić U, Berntorp K. Vehicle tracking control on piecewise-clothoidal trajectories by MPC with guaranteed error bounds. In: *Conf. Decision and Control*; Dec.; Las Vegas, NV; 2016.
- [11] Borrelli F, Falcone P, Keviczky T, et al. MPC-based approach to active steering for autonomous vehicle systems. *Int J Veh Auton Syst.* 2005;3(2–4):265–291.
- [12] Berntorp K, Olofsson B, Lundahl K, et al. Models and methodology for optimal trajectory generation in safety-critical road-vehicle manoeuvres. *Veh Syst Dyn.* 2014;52(10):1304–1332.
- [13] Quirynen R, Berntorp K, Di Cairano S. Embedded optimization algorithms for steering in autonomous vehicles based on nonlinear model predictive control. In: *Amer. Control Conf.*; Jun.; Milwaukee, WI; 2018.
- [14] Gustafsson F. Slip-based tire-road friction estimation. *Automatica.* 1997;33:1087–1099.
- [15] Svendenius J. Tire modeling and friction estimation [dissertation]. Dept. Automatic Control, Lund University, Sweden; 2007.
- [16] Gustafsson F. Automotive safety systems. *IEEE Signal Processing magazine.* 2009 July; 26(4):32–47.
- [17] Rasmussen C, Williams C. Gaussian processes for machine learning. Cambridge, MA, USA: MIT Press; 2006.
- [18] Berntorp K, Di Cairano S. Tire-stiffness and vehicle-state estimation based on noise-adaptive particle filtering. *IEEE Trans Control Syst Technol.* 2018;27(3):1100–1114.
- [19] Berntorp K. Online Bayesian inference and learning of Gaussian-process state-space models. *Automatica.* 2021;129:109613.
- [20] Doucet A, Johansen AM. A tutorial on particle filtering and smoothing: Fifteen years later. In: Crisan D, Rozovsky B, editors. *Handbook of nonlinear filtering*. Oxford University Press; 2009.
- [21] Quirynen R, Di Cairano S. PRESAS: Block-structured preconditioning of iterative solvers within a primal active-set method for fast model predictive control. *Optimal Control Applications and Methods.* 2020;.
- [22] Feng X, Di Cairano S, Quirynen R. Inexact adjoint-based SQP algorithm for real-time stochastic nonlinear MPC. In: *IFAC World Congress*; Jul.; Berlin, Germany; 2020.
- [23] Vaskov S, Berntorp K, Quirynen R. Cornering stiffness adaptive, stochastic nonlinear model predictive control for vehicles. In: *Amer. Control Conf.*; May; New Orleans, LA; 2021.
- [24] Vaskov S, Quirynen R, Menner M, et al. Friction-adaptive stochastic predictive control for trajectory tracking of autonomous vehicles. In: *Amer. Control Conf.*; Jun.; Atlanta, GA; 2022. Submitted.
- [25] Rajamani R. *Vehicle dynamics and control*. Springer-Verlag; 2006.
- [26] Berntorp K. Online Bayesian tire-friction learning by Gaussian-process state-space models. In: *IFAC World Congress*; Jul.; Berlin, Germany; 2020.
- [27] Solin A, Särkkä S. Hilbert space methods for reduced-rank Gaussian process regression. *Statistics and Computing.* 2020;30(2):419–446.
- [28] Svensson A, Schön TB. A flexible state-space model for learning nonlinear dynamical systems. *Automatica.* 2017;80:189–199.
- [29] Pacejka HB. *Tire and vehicle dynamics*. 2nd ed. Oxford, United Kingdom: Butterworth-Heinemann; 2006.

- [30] Kerrigan EC, Maciejowski JM. Soft constraints and exact penalty functions in model predictive control. In: Control 2000 Conference, Cambridge; 2000. p. 2319–2327.
- [31] Ahn C, Peng H, Tseng HE. Robust estimation of road friction coefficient using lateral and longitudinal vehicle dynamics. *Veh Syst Dyn.* 2012;50(6):961–985.
- [32] Telen D, Vallerio M, Cabianca L, et al. Approximate robust optimization of nonlinear systems under parametric uncertainty and process noise. *J Proc Control.* 2015;33:140–154.
- [33] Andersson JAE, Gillis J, Horn G, et al. Casadi—a software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation.* 2018;.
- [34] Nocedal J, Wright SJ. Numerical optimization. 2nd ed. Springer; 2006. Springer Series in Operations Research and Financial Engineering.
- [35] Diehl M, Findeisen R, Allgöwer F, et al. Nominal stability of the real-time iteration scheme for nonlinear model predictive control. *IEE Proc-Control Theory Appl.* 2005; 152(3):296–308.
- [36] Gros S, Zanon M, Quirynen R, et al. From linear to nonlinear MPC: bridging the gap via the real-time iteration. *International Journal of Control.* 2020;93(1):62–80.
- [37] Berntorp K, Hoang T, Quirynen R, et al. Control architecture design of autonomous vehicles. In: Conf. Control Technol. and Applications; Aug.; Copenhagen, Denmark; 2018. Invited paper.
- [38] ISO 3888-2: 2002. Passenger cars—test track for a severe lane change manoeuvre—part 2: Obstacle avoidance ; 2002.
- [39] Mechanical Simulation. Carsim [www.carsim.com]; 2021.
- [40] Berntorp K. Particle filtering and optimal control for vehicles and robots [dissertation]. Department of Automatic Control, Lund University, Sweden; 2014.