

Bayesian Learning of Tire Friction with Automotive-Grade Sensors by Gaussian-Process State-Space Models

Berntorp, K.; Hiroaki, K.

TR2019-145 December 11, 2019

Abstract

The friction dependence between tire and road is highly nonlinear and varies heavily between different surfaces. Knowledge of the tire friction is important for real-time vehicle control, but difficult to estimate with automotive-grade sensors. Based on recent advances in particle filtering and Markov chain Monte-Carlo methods, we propose a batch method for identifying the tire friction as a function of the wheel slip. The unknown function mapping the wheel slip to tire friction is modeled as a Gaussian process (GP) that is included in a dynamic vehicle model relating the GP to the vehicle state. The method is able to efficiently learn the tire friction using only wheel-speed, steering-wheel angle, and inertial automotive-grade sensors. We illustrate the efficacy of the method using several experimental data sets obtained on a snow-covered road.

IEEE Conference on Decision and Control (CDC)

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.

Bayesian Learning of Tire Friction with Automotive-Grade Sensors by Gaussian-Process State-Space Models

Karl Berntorp¹ and Kitano Hiroaki²

Abstract—The friction dependence between tire and road is highly nonlinear and varies heavily between different surfaces. Knowledge of the tire friction is important for real-time vehicle control, but difficult to estimate with automotive-grade sensors. Based on recent advances in particle filtering and Markov chain Monte-Carlo methods, we propose a batch method for identifying the tire friction as a function of the wheel slip. The unknown function mapping the wheel slip to tire friction is modeled as a Gaussian process (GP) that is included in a dynamic vehicle model relating the GP to the vehicle state. The method is able to efficiently learn the tire friction using only wheel-speed, steering-wheel angle, and inertial automotive-grade sensors. We illustrate the efficacy of the method using several experimental data sets obtained on a snow-covered road.

I. INTRODUCTION

Various tire models describing the tire friction as a function of wheel slip have been reported in literature [1]–[4]. Fig. 1 shows the typical tire-friction curves generated by the Pacejka (Magic formula) tire model [1]. Knowing the tire friction over a range of slip values extending into the saturated region of the tire-friction function is important for AD and ADAS [5]. The vehicle states involved in the tire-friction estimation are not directly measured in production vehicles, so indirect methods are needed. The identification approaches in literature typically estimate parameters of specific models. Examples are batch nonlinear optimization methods for identifying the parameters of the Brush model [3], unscented Kalman filter (UKF) for estimating the Pacejka parameters [6], recursive least-squares for estimating the cornering stiffness [7], [8], and nonlinear observer for estimating the peak friction coefficient using the Brush tire model [9]. 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 [10]. 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. Also, it is worth pointing out that few approaches so far, if any, target the estimation of the full tire-friction curve using only production-grade sensors.

We recently proposed a Bayesian approach for identifying the friction-slip dependence [11], where we formulated a nonparametric approach modeling the unknown function describing the tire friction as a Gaussian process [12], which

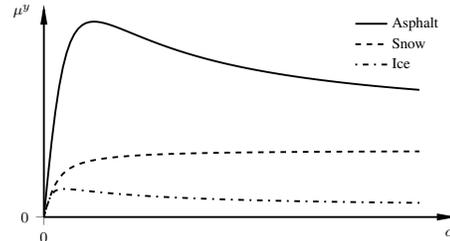


Fig. 1. Illustrations of lateral tire friction μ^y as a function of slip angle α for surfaces corresponding to asphalt, loose snow, and ice. The peak of the respective curve is called the (peak) friction coefficient.

combined with particle filtering (PF) [13] and Markov chain Monte-Carlo (MCMC) methods [14] results in a method for estimating the posterior density function (PDF) of the tire friction, given the measurement data. In [11], the method was evaluated in simulation using multiple step-steer maneuvers.

This paper extends the method in [11] by discussing and implementing ways to improve computational efficiency, and we verify the method using several experimental data sets obtained on a snow-covered road, relying only on production-grade sensors. Since the method is nonparametric, it is not subject to specific modeling constraints that various tire models impose. Still, the method is insensitive to overfitting to the data. As a consequence, the method can be used for extracting parameters for specific tire models, valid for the full range of possible data even when the data used for learning is very limited. We provide an example of this, where we extract the Pacejka tire parameters [1] from the identified tire friction function.

Notation: With $p(\mathbf{x}_{0:t}|\mathbf{y}_{0:t})$, we mean the posterior density function of the state trajectory $\mathbf{x}_{0:t}$ from time index 0 to time index t given the measurement sequence $\mathbf{y}_{0:t} := \{\mathbf{y}_0, \dots, \mathbf{y}_t\}$. We define $\mathbf{f}_t := \mathbf{f}(\mathbf{x}_t)$ for a function \mathbf{f} . Throughout, for a vector \mathbf{x} , $\mathbf{x} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ indicates that \mathbf{x} is Gaussian distributed with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$ and x_n denotes the n th component of \mathbf{x} . Matrices are indicated in capital bold font as \mathbf{X} , and the element on row i and column j is denoted with X_{ij} . The notation $\mathbf{f} \sim \mathcal{GP}(\mathbf{0}, \boldsymbol{\kappa}_{\theta, \mathbf{f}}(\mathbf{x}, \mathbf{x}'))$ means that the function $\mathbf{f}(\mathbf{x})$ is a realization from a GP prior with a given covariance function $\boldsymbol{\kappa}_{\theta, \mathbf{f}}(\mathbf{x}, \mathbf{x}')$ subject to hyperparameters $\boldsymbol{\theta}$, and $\mathcal{IW}(\nu, \boldsymbol{\Lambda})$ is the inverse-Wishart distribution with degree of freedom ν and scale matrix $\boldsymbol{\Lambda}$. Similarly, $\mathcal{MN}(\mathbf{M}, \mathbf{Q}, \mathbf{V})$ is the Matrix-Normal distribution with mean \mathbf{M} , right covariance \mathbf{Q} , and left precision \mathbf{V} .

II. MODELING AND PROBLEM FORMULATION

In this section we briefly summarize the different models used by the proposed learning method. For details about

¹Mitsubishi Electric Research Laboratories (MERL), 02139 Cambridge, MA, USA. Email: karl.o.berntorp@ieee.org. ² Mitsubishi Electric Corp., Adv. Technology R&D Center, Amagasaki, 661-8661, Japan.

the used model, see [11]. We use a single-track chassis model that includes the lateral velocity v^Y and the yaw rate $\dot{\psi}$. Consequently, in this paper we focus on the lateral dynamics, but the method extends straightforwardly to also handle longitudinal dynamics. Hence, the state vector is $\mathbf{x} = [v^Y \ \dot{\psi}]^T \in \mathbb{R}^2$. A single-track model is sufficiently accurate for purposes where the tire friction reaches the nonlinear region but the maneuvers are not aggressive enough to result in large roll angles [5]. The presented framework can be extended to handle a double-track model, but it increases computation time and modeling complexity. The single-track model lumps together the left and right wheel on each axle, and roll and pitch dynamics are neglected. Thus, the model

The tire friction components μ_i^y , $i \in \{f, r\}$ are modeled as static functions of the slip quantities,

$$\mu_i^y = f_i^y(\alpha_i(\mathbf{x})), \quad (1)$$

α is the slip angle. The lateral velocity is estimated in the proposed method, whereas the longitudinal velocity is determined from the measured wheel-speeds $\{\omega_i\}_{i=1}^4$. For brevity, we define the vector $\boldsymbol{\alpha} = [\alpha_f \ \alpha_r]^T$. We write (1) as

$$\boldsymbol{\mu} = [f_f^y \ f_r^y]^T, \quad (2)$$

and model the friction vector as a realization from a zero-mean Gaussian process prior

$$\boldsymbol{\mu}(\boldsymbol{\alpha}) \sim \mathcal{GP}(\mathbf{0}, \boldsymbol{\kappa}_{\theta, \boldsymbol{\mu}}(\boldsymbol{\alpha}, \boldsymbol{\alpha}')), \quad (3)$$

where the covariance function $\boldsymbol{\kappa}_{\theta, \boldsymbol{\mu}}(\boldsymbol{\alpha}, \boldsymbol{\alpha}')$ is chosen in advance. In this work the hyperparameters $\boldsymbol{\theta}$ are determined a priori but can also be included in the learning process [15].

A. Estimation Model

After discretization with sampling period T_s and using $\mathbf{u} = [\delta \ v^X]^T$ as the known input vector, the complete vehicle model can be written as

$$\mathbf{x}_{t+1} = \mathbf{a}(\mathbf{x}_t, \mathbf{u}_t) + \mathbf{G}(\mathbf{x}_t, \mathbf{u}_t)\boldsymbol{\mu}(\boldsymbol{\alpha}_t), \quad (4)$$

where \mathbf{a} and \mathbf{G} are the (known) parts of the vehicle model, and $\boldsymbol{\mu}$ is the unknown function we want to estimate.

Our measurement model is based on a setup commonly available in production cars, namely the lateral acceleration a_m^Y and the yaw rate $\dot{\psi}_m$, forming the measurement vector $\mathbf{y} = [a_m^Y \ \dot{\psi}_m]^T$. The yaw-rate measurement is directly related to the yaw rate state, but the lateral acceleration depends on the vehicle model. We model the measurement noise \mathbf{e}_t as zero-mean Gaussian distributed noise with covariance \mathbf{R} according to $\mathbf{e}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{R})$. The measurement model can be written as

$$\mathbf{y}_t = \mathbf{h}(\mathbf{x}_t, \mathbf{u}_t) + \mathbf{D}(\mathbf{x}_t, \mathbf{u}_t)\boldsymbol{\mu}(\boldsymbol{\alpha}_t) + \mathbf{e}_t. \quad (5)$$

The measurement model (5) is decomposed into known parts of the dynamics, \mathbf{h} and \mathbf{D} , and an unknown part, $\boldsymbol{\mu}$. 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.

The estimation model consisting of (4) and (5) is a GP-SSM where the tire friction is a GP. The reason for modeling

the tire friction as a GP is its ability to model the inherent uncertainty stemming from the measurement data, not only the uncertainty from the stochastic noise term \mathbf{e}_t , which affects the estimation quality, but also that the measurement data may contain few measurements in certain regions of the state space.

B. Problem Formulation

We want to estimate the nonlinear function $\boldsymbol{\mu}$ describing the tire friction. We approach this problem as follows. Given the system model (4), (5), and a Gaussian process prior (3) on the tire friction, we want to infer the posterior distribution of $\boldsymbol{\mu}(\boldsymbol{\alpha})$ given a set of measurement data $\mathbf{y}_{0:T}$,

$$p(\boldsymbol{\mu}|\mathbf{y}_{0:T}). \quad (6)$$

Since the tire-friction estimate will depend on the vehicle state, we solve for (6) by approximating the joint posterior $p(\boldsymbol{\mu}, \mathbf{x}_{0:T}|\mathbf{y}_{0:T})$ and performing the marginalization step

$$\begin{aligned} p(\boldsymbol{\mu}|\mathbf{y}_{0:T}) &= \int p(\boldsymbol{\mu}, \mathbf{x}_{0:T}|\mathbf{y}_{0:T}) d\mathbf{x}_{0:T} \\ &= \int p(\boldsymbol{\mu}|\mathbf{x}_{0:T}, \mathbf{y}_{0:T})p(\mathbf{x}_{0:T}|\mathbf{y}_{0:T}) d\mathbf{x}_{0:T}. \end{aligned} \quad (7)$$

III. REDUCED-RANK GP-SSMS AND PARTICLE FILTERING

In this section we briefly review background material on GP-SSMs and PF necessary for understanding the proposed learning method described in Sec. IV.

A. Reduced-Rank GP-SSMs

We rely on GP priors for learning the tire friction function, where the covariance function $\boldsymbol{\kappa}(\mathbf{x}, \mathbf{x}')$ encodes the prior assumptions. A bottleneck in some of the proposed GP-SSM methods is the computational load. In this paper we use the computationally efficient reduced-rank GP-SSM framework presented in [15], [16]. For a thorough derivation and convergence proofs, see [16]. Following the notation in [16], isotropic covariance functions (that only depend on the Euclidean norm $\|\mathbf{x} - \mathbf{x}'\|$) can be approximated in terms of Laplace operators on the form:

$$\boldsymbol{\kappa}_{\theta}(\mathbf{x}, \mathbf{x}') \approx \sum_{j_1, \dots, j_d=1}^m \mathcal{S}_{\theta}(\lambda^{j_1, \dots, j_d}) \phi^{j_1, \dots, j_d}(\mathbf{x}) \phi^{j_1, \dots, j_d}(\mathbf{x}'), \quad (8)$$

where we for simplicity assume m basis functions for each state dimension. In (8), \mathcal{S}_{θ} is the spectral density of $\boldsymbol{\kappa}_{\theta}$ and

$$\lambda^{j_1, \dots, j_d} = \sum_{n=1}^d \left(\frac{\pi j_n}{2L_n} \right)^2, \quad (9a)$$

$$\phi^{j_1, \dots, j_d} = \prod_{n=1}^d \frac{1}{\sqrt{L_n}} \sin \left(\frac{\pi j_n (x_n + L_n)}{2L_n} \right), \quad (9b)$$

are the Laplace operator eigenvalues and eigenfunctions, respectively, defined on the intervals $[-L_n, L_n]$. For brevity, we will in the rest of the paper denote j_1, \dots, j_d with \mathbf{j} . Note that according to (8), (9), only the spectral density depends on the hyperparameters $\boldsymbol{\theta}$. Furthermore, (8) can be

interpreted as an optimal parametric expansion with respect to the covariance function in the GP prior [15].

From the approximation (8) using Laplace operators, [16] provides a relation between basis function expansions of a function f and GPs based on the Karhunen-Loeve expansion. With the basis functions chosen as (9b),

$$f(\mathbf{x}) \sim \mathcal{GP}(0, \kappa(\mathbf{x}, \mathbf{x}')) \Leftrightarrow f(\mathbf{x}) \approx \sum_j \gamma^j \phi^j(\mathbf{x}), \quad (10)$$

with

$$\gamma^j \sim \mathcal{N}(0, S(\lambda^j)). \quad (11)$$

For a state-space model $\mathbf{x}_{t+1} = \mathbf{f}(\mathbf{x}_t) + \mathbf{w}_t$, (10) implies the reduced-rank GP-SSM

$$\mathbf{x}_{t+1} = \begin{bmatrix} \gamma_1^1 & \cdots & \gamma_1^m \\ \vdots & & \vdots \\ \gamma_d^1 & \cdots & \gamma_d^m \end{bmatrix} \begin{bmatrix} \phi^1(\mathbf{x}_t) \\ \vdots \\ \phi^m(\mathbf{x}_t) \end{bmatrix} + \mathbf{w}_t, \quad (12)$$

where γ_n^j are the weights to be learned, m is the total number of basis functions (i.e., m^d in (8)), and \mathbf{w}_t is zero-mean Gaussian distributed noise with covariance \mathbf{Q} . In Sec. IV, (12) in combination with PF forms the basis for learning the tire friction.

Remark 1: In this paper we have the state dimension $d = 2$. However, when more complicated vehicle models are used, the state dimension can be much larger. This will render exponential growth of m with the state dimension. There are ways to alleviate this, some of which are discussed in [15]. This is not pursued further here because of the small state dimension used for the results in this paper.

B. Sequential Monte Carlo and Particle Filtering

Sequential Monte-Carlo (SMC) methods, such as PFs, constitute a class of techniques that estimate the posterior distribution in SSMs, and SMCs have recently emerged as a useful tool in learning of SSMs (e.g., [15]). PFs approximate the posterior density $p(\mathbf{x}_t | \mathbf{y}_{0:t})$ by a set of N weighted state trajectories as

$$p(\mathbf{x}_t | \mathbf{y}_{0:t}) \approx \sum_{i=1}^N q_t^i \delta_{\mathbf{x}_t^i}(\mathbf{x}_t), \quad (13)$$

where q_t^i is the importance weight of the i th particle \mathbf{x}_t^i and $\delta(\cdot)$ is the Dirac delta mass. The PF recursively estimates (13) by repeated application of Bayes' rule, where the states are sampled according to a proposal density $\pi(\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{y}_t)$, which in the simplest case is the dynamical model. This yields the state samples at each time step as

$$\mathbf{x}_t^i \sim p(\mathbf{x}_t | \mathbf{x}_{t-1}^i), \quad i \in \{1, \dots, N\}. \quad (14)$$

The importance weights are updated using the likelihood as

$$q_t^i \propto q_{t-1}^i p(\mathbf{y}_t | \mathbf{x}_t^i). \quad (15)$$

The PF algorithm iterates between (14) and (15), combined with a resampling step that removes particles with low weights and replaces them with more likely particles.

In this paper, we adapt a conditional PF with ancestor sampling (CPF-AS) [14] to generate the state trajectories needed

to learn the function μ describing the tire friction. CPF-AS generates the state trajectories by a procedure similar to the standard PF, except for that the PF is conditioned on one prespecified *reference trajectory* $\mathbf{x}'_{0:T}$, which is retained throughout the procedure. When used within an MCMC procedure [17], it can be shown that after a burn-in period, the state trajectories generated by CPF-AS are samples drawn from the smoothing distribution $p(\mathbf{x}_{0:T} | \mathbf{y}_{0:T})$ for any finite $N > 1$ [14], [18], that is, the second distribution on the right-hand side of (7).

IV. LEARNING THE TIRE FRICTION BY GP-SSMS

The objective is to infer the posterior distribution (6) of the unknown function μ . In the presentation of the method we focus on the lateral dynamics, that is, we learn the lateral tire friction of front and rear wheels. However, the extension to the longitudinal case is analogous.

A. Adapting the Model for Learning

The Bayesian learning method we leverage assumes dynamical systems on the form $\mathbf{x}_{t+1} = \mathbf{f}_t + \mathbf{w}_t$, where the full state-transition function \mathbf{f}_t is to be learned. Hence, we need to adapt the vehicle model (4). Specifically, by manipulation of (4) and using the basis function expansion approach (10), the model can be written on the form [11]

$$\zeta_{t+1} = \underbrace{\begin{bmatrix} \gamma_1^1 & \cdots & \gamma_1^m \\ \vdots & & \vdots \\ \gamma_d^1 & \cdots & \gamma_d^m \end{bmatrix}}_{\mathbf{A}} \underbrace{\begin{bmatrix} \phi^1(\alpha_t) \\ \vdots \\ \phi^m(\alpha_t) \end{bmatrix}}_{\varphi(\alpha_t)} + \mathbf{w}_t \quad (16)$$

for some $\zeta_t = [\zeta_{1,t} \ \zeta_{2,t}]^T$.

B. Tire Friction Learning with GP-SSM

With the reduced-rank GP-SSM (16) in combination with the measurement model (5), we are now ready to formulate our learning approach. Using (16), the problem of estimating the distribution (6) now amounts to infer the distribution of \mathbf{A} and \mathbf{Q} , that is, to estimate the distribution

$$p(\mathbf{A}, \mathbf{Q} | \mathbf{y}_{0:T}), \quad (17)$$

where the components in \mathbf{A} have a prior Gaussian distribution according to (11). To estimate the covariance matrix \mathbf{Q} , we impose the additional assumption that the prior of \mathbf{Q} is inverse-Wishart (\mathcal{IW}) distributed according to

$$\mathbf{Q} \sim \mathcal{IW}(\ell_Q, \mathbf{\Lambda}_Q). \quad (18)$$

The \mathcal{IW} distribution is a distribution over (real) positive definite matrices, and has the degrees of freedom ℓ_Q and positive definite scale matrix $\mathbf{\Lambda}_Q$ as hyperparameters. Letting an unknown covariance matrix have the \mathcal{IW} distribution as prior distribution is common due to its properties and has been used in automotive applications before (e.g., [19], [20]).

Due to that the components of the system matrix \mathbf{A} in (17) have a Gaussian prior (11), \mathbf{A} is Matrix-Normal (\mathcal{MN}) distributed according to

$$\mathbf{A} \sim \mathcal{MN}(\mathbf{0}, \mathbf{Q}, \mathbf{V}). \quad (19)$$

With \mathbf{A} \mathcal{MN} distributed and \mathbf{Q} \mathcal{IW} distributed, the joint prior $p(\mathbf{A}, \mathbf{Q})$ is $\mathcal{MN}\mathcal{IW}$ distributed according to [21]

$$p(\mathbf{A}, \mathbf{Q}) = \mathcal{MN}\mathcal{IW}(\mathbf{A}, \mathbf{Q} | \mathbf{0}, \mathbf{V}, \ell_Q, \mathbf{\Lambda}_Q), \quad (20)$$

where \mathbf{V} has the inverse spectral density of the covariance function as diagonal entries [15],

$$\mathbf{V} = \text{diag}([S^{-1}(\lambda^1) \quad \dots \quad S^{-1}(\lambda^m)]), \quad (21)$$

and where $\text{diag}(\cdot)$ is the diagonal matrix. To estimate (17), we need the two densities $p(\mathbf{A}, \mathbf{Q} | \mathbf{x}_{0:T}, \mathbf{y}_{0:T})$ and $p(\mathbf{x}_{0:T} | \mathbf{y}_{0:T})$ similar to the right-hand side in (7).

1) *Estimating the State Posterior:* The state posterior $p(\mathbf{x}_{0:T} | \mathbf{y}_{0:T})$ depends on the tire friction through (4), which implies

$$p(\mathbf{x}_{0:T} | \mathbf{y}_{0:T}) = \int p(\mathbf{x}_{0:T} | \mathbf{A}, \mathbf{Q}, \mathbf{y}_{0:T}) d\mathbf{A} d\mathbf{Q}. \quad (22)$$

Hence, we need to sample from $p(\mathbf{x}_{0:T} | \mathbf{A}, \mathbf{Q}, \mathbf{y}_{0:T})$. We use CPF-AS, outlined in Algorithm 1, which produces samples that are asymptotically consistent with (22) when encapsulated into an MCMC procedure [14].

Algorithm 1 CPF-AS

Input: $\mathbf{x}_{0:T}(k)$, $\mathbf{u}_{0:T-1}$, N , model $\{\mathbf{a}, \mathbf{G}, \mu, \mathbf{Q}, \mathbf{D}, \mathbf{R}\}$.

Output Trajectory $\mathbf{x}_{0:T}(k+1)$.

- 1: Sample $\mathbf{x}_0^i \sim p(\mathbf{x}_0)$, $\forall i \in \{1, \dots, N-1\}$.
 - 2: Set $\mathbf{x}_0^N = \mathbf{x}_0(k)$.
 - 3: **for** $t \leftarrow 0$ to $T-1$ **do**
 - 4: Compute slip angles $\alpha_t^i \forall i \in \{1, \dots, N\}$.
 - 5: Set $q_t^i \propto \mathcal{N}(\mathbf{y}_t | \mathbf{h}_t^i + \mathbf{D}_t^i \mu(\alpha_t^i), \mathbf{R})$, $\forall i \in \{1, \dots, N\}$.
 - 6: Sample a_t^i with $\mathbb{P}(a_t^i = j) \propto q_t^j$, $\forall i \in \{1, \dots, N\}$.
 - 7: Sample $\mathbf{x}_{t+1}^i \sim \mathcal{N}(\mathbf{x}_{t+1} | \mathbf{a}_t^{a_t^i} + \mathbf{G}_t^{a_t^i} \mu(\alpha_t^{a_t^i}), \mathbf{Q})$, $\forall i \in \{1, \dots, N\}$.
 - 8: Set $\mathbf{x}_{t+1}^N = \mathbf{x}_{t+1}(k)$.
 - 9: Sample $a_t^N \mathbb{P}(a_t^N = j) \propto q_t^j \mathcal{N}(\mathbf{x}_{t+1}^N | \mathbf{a}_t^j + \mathbf{G}_t^j \mu(\alpha_t^j), \mathbf{Q})$.
 - 10: Set $\mathbf{x}_{1:t+1}^i = \{\mathbf{x}_{1:t}^{a_t^i}, \mathbf{x}_{t+1}^i\}$, $\forall i \in \{1, \dots, N\}$.
 - 11: **end for**
 - 12: Draw J with $\mathbb{P}(i = J) \propto q_t^i$.
 - 13: Set $\mathbf{x}_{0:T}(k+1) = \mathbf{x}_{0:T}^J$.
-

2) *Learning the Tire Friction:* To learn the posterior (17), that is, to learn the PDF of the function describing the tire friction and process-noise covariance that accounts for modeling errors, we use Bayes' rule,

$$p(\mathbf{A}, \mathbf{Q} | \mathbf{x}_{0:T}, \mathbf{y}_{0:T}) \propto p(\mathbf{x}_{0:T}, \mathbf{y}_{0:T} | \mathbf{A}, \mathbf{Q}) p(\mathbf{A}, \mathbf{Q}). \quad (23)$$

The likelihood $p(\mathbf{x}_{0:T}, \mathbf{y}_{0:T} | \mathbf{A}, \mathbf{Q})$ can be written as

$$p(\mathbf{x}_{0:T}, \mathbf{y}_{0:T} | \mathbf{A}, \mathbf{Q}) = \underbrace{\prod_{t=0}^{T-1} p(\mathbf{x}_{t+1} | \mathbf{x}_t, \mathbf{A}, \mathbf{Q})}_{p(\mathbf{x}_{0:T} | \mathbf{A}, \mathbf{Q})} \underbrace{\prod_{t=0}^T p(\mathbf{y}_t | \mathbf{x}_t, \mathbf{A}, \mathbf{Q})}_{p(\mathbf{y}_{0:T} | \mathbf{x}_{0:T}, \mathbf{A}, \mathbf{Q})}. \quad (24)$$

Conditioned on \mathbf{A} and \mathbf{Q} , the vehicle model (4) and measurement model (5) are Gaussian, implying that the two terms $p(\mathbf{x}_{0:T} | \mathbf{A}, \mathbf{Q})$ and $p(\mathbf{y}_{0:T} | \mathbf{x}_{0:T}, \mathbf{A}, \mathbf{Q})$ in (24) are Gaussian. The density $p(\mathbf{x}_{0:T}, \mathbf{y}_{0:T} | \mathbf{A}, \mathbf{Q})$ is Gaussian since it is a product of Gaussians. Therefore, we can utilize the concept

of conjugate priors. If a prior distribution belongs to the same family as the posterior distribution, the prior is conjugate to the likelihood. For Gaussian distributed data, an $\mathcal{MN}\mathcal{IW}$ distribution is a conjugate prior [22], which results in closed-form expressions for the update of \mathbf{A} and \mathbf{Q} [15]. Define

$$\mathbf{\Phi} = \sum_{t=0}^T \zeta_t \zeta_t^T, \quad (25a)$$

$$\mathbf{\Psi} = \sum_{t=0}^T \zeta_t \varphi(\alpha_t)^T, \quad (25b)$$

$$\mathbf{\Sigma} = \sum_{t=0}^T \varphi(\alpha_t) \varphi(\alpha_t)^T. \quad (25c)$$

Then it follows that the joint posterior is

$$p(\mathbf{A}, \mathbf{Q} | \mathbf{x}_{0:T}, \mathbf{y}_{0:T}) = p(\mathbf{A} | \mathbf{Q}, \mathbf{x}_{0:T}, \mathbf{y}_{0:T}) p(\mathbf{Q} | \mathbf{x}_{0:T}, \mathbf{y}_{0:T}), \quad (26)$$

where

$$p(\mathbf{Q} | \mathbf{x}_{0:T}, \mathbf{y}_{0:T}) = \mathcal{IW}(\mathbf{Q} | T + \ell_Q, \mathbf{\Lambda}_Q + \mathbf{\Phi} - \mathbf{\Psi}(\mathbf{\Sigma} + \mathbf{V})^{-1} \mathbf{\Psi}^T), \quad (27)$$

$$p(\mathbf{A} | \mathbf{Q}, \mathbf{x}_{0:T}, \mathbf{y}_{0:T}) = \mathcal{MN}(\mathbf{A} | \mathbf{\Psi}(\mathbf{\Sigma} + \mathbf{V})^{-1}, \mathbf{Q}, (\mathbf{\Sigma} + \mathbf{V})^{-1}). \quad (28)$$

Algorithm 2 Proposed method for tire-friction learning

Input: $\mathbf{y}_{0:T}$, $\mathbf{u}_{0:T-1}$, priors (20), (21).

Output K MCMC samples from $p(\mathbf{A}, \mathbf{Q}, \mathbf{x}_{0:T} | \mathbf{y}_{0:T})$

- 1: Sample initial guess $\mathbf{x}_{0:T}(0)$, $\mathbf{Q}(0)$, $\mathbf{A}(0)$.
 - 2: **for** $k \leftarrow 0$ to $K-1$ **do**
 - 3: Sample $\mathbf{x}_{0:T}(k+1)$ given $\mathbf{Q}(k)$, $\mathbf{A}(k)$ using Algorithm 1.
 - 4: Compute slip angle $\alpha_{0:T}$.
 - 5: Compute $\varphi(\alpha_{0:T})$ in (16) using (9b).
 - 6: Compute $\zeta_{0:T}$ in (16).
 - 7: Compute $\mathbf{\Phi}$, $\mathbf{\Psi}$, $\mathbf{\Sigma}$ using (25).
 - 8: Sample $\mathbf{Q}(k+1)$ given $\mathbf{A}(k)$, $\mathbf{x}_{0:T}(k+1)$ using (27).
 - 9: Sample $\mathbf{A}(k+1)$ given $\mathbf{x}_{0:T}(k+1)$, $\mathbf{Q}(k+1)$ using (28).
 - 10: Set $\mu = \mathbf{A}(k+1)\varphi$.
 - 11: **end for**
-

V. IMPLEMENTATION ASPECTS

Automotive-grade inertial sensors usually have bias that affects the estimation performance. Since the proposed method uses batches of measurement data, it is possible to remove the constant part of the bias beforehand. If the data set is longer than a few minutes, such that the time-varying aspects of the bias dominate, bias estimation can be directly incorporated into the framework. By modeling the bias dynamics as a random walk process, bias estimation can be implemented in a computationally efficient manner by using a Rao-Blackwellized particle filter (RBPF) in Algorithm 1, where the bias states are estimated using Kalman filters constrained on the particle trajectories [23].

Prior knowledge of the tire friction can be used to initialize the algorithm and therefore possibly improve convergence speed. The tire friction has a linear dependence on the wheel slip for small slip angles (the tire stiffness) and there is a decrease in the slope until, usually but not always, the friction

reaches its peak value (the peak friction) [3]. To leverage this, we can split up the friction function into two parts,

$$\boldsymbol{\mu}(\boldsymbol{\alpha}_t) = \tilde{\boldsymbol{\mu}}(\boldsymbol{\alpha}_t) + \Delta\boldsymbol{\mu}(\boldsymbol{\alpha}_t), \quad (29)$$

where $\tilde{\boldsymbol{\mu}}(\boldsymbol{\alpha}_t)$ is the prior information. Hence, the estimation problem of determining the tire friction $\boldsymbol{\mu}(\boldsymbol{\alpha}_t)$ amounts to estimating the variation $\Delta\boldsymbol{\mu}(\boldsymbol{\alpha}_t)$ around the initial guess $\tilde{\boldsymbol{\mu}}(\boldsymbol{\alpha}_t)$, which reduces initial uncertainty.

We rely on inertial sensors to measure the vehicle state. The yaw rate is directly measured, but the accelerometer only implicitly measures the velocity. Therefore, the tire friction enters in both process model and measurement model. Prior knowledge about the tire friction can be used in the measurement model, by modeling the measurement equation as a known function for a range of the slip values. For instance, for slip angles in the range $-\beta \leq \alpha_j \leq \beta$, $j \in \{f, r\}$ for a small β (e.g., $\beta = 1\text{deg}$), we can model the measurement equation as linear in the slip. Similar to the initialization (29), it can have significant effect on the convergence speed.

VI. EXPERIMENTAL RESULTS

We evaluate Algorithm 2 on three data sets obtained on snow, all roughly 250s long. The data has been recorded from test drives using a mid-size SUV on multiple loops of a track. Due to the proximity in time between the test drives, the road conditions are similar in all of the data sets.

Algorithm 2 is executed for $K = 500$ iterations, which is a factor of 20 less than used in [11] and the number of burn-in samples is $K_{\text{bi}} = 100$. We use the initialization routines described in Sec. V for speeding up convergence. We use a sampling period of $T_s = 0.04\text{s}$ for time discretization of the vehicle dynamics, resulting in (4), and for obtaining the measurements (5). The chosen sampling period corresponds to each data set having roughly 6000 measurements. We split up each of the three data sets into two sets; one set is used for the actual learning and the other is used for validation. Hence, the lengths of the data sets used for learning are about 120s. We use 10 basis functions each for the front and rear tire, which gives $m = 100$ basis functions in total. The underlying particle filter uses $N = 500$ particles.

Fig. 2 shows the estimated mean of the lateral tire-friction curve for the front (first row) and rear (second row) tire in black solid for the three data sets. We do not have ground truth for the tire friction. However, the shape of the curve and the peak friction are similar to what has been reported in literature for snow [3]. For instance, the estimated peak friction $\mu_{f,\text{max}}$ for the six plots are between $0.25 \lesssim \mu_{f,\text{max}} \lesssim 0.38$, which is consistent with [3], where peak friction coefficients between $0.25 \lesssim \mu_{f,\text{max}} \lesssim 0.45$ were reported for a high-precision test rig, depending on tire pressure, studded/regular winter tire, and other factors. Although not being ground truth, it still indicates that the estimates are reasonable. Fig. 2 also displays the front and rear tire-friction curves of a fitted Pacejka tire model [1] (black dash-dotted) to the learned function for all three data sets. We have identified the Pacejka model by minimizing the norm of the deviation

of the Pacejka tire model to the respective estimate, scaled with the covariance $\text{cov}(\mathbf{A}\boldsymbol{\varphi}(\alpha_i))$,

$$\min_{\mu_i, B_i, C_i, E_i} \int \frac{1}{\text{cov}(\mathbf{A}\boldsymbol{\varphi}(\alpha_i))} \|\hat{\mathbf{A}}\boldsymbol{\varphi}(\alpha_i) - F_i(\alpha_i)\| d\alpha_i, \quad (30)$$

for $i \in \{f, r\}$, where μ , B , C , and E are the peak, stiffness, shape, and curvature factor, respectively. Due to the weighting with the uncertainty in (30), the estimates where most data are located are given a higher confidence.

The lowest row in Fig. 2 shows the excitation level of the underlying data for the range of slip angles. The slip angles shown are the estimates from the last of the K iterations of Algorithm 2. We stress that the data in the lowest plots are *not* used for learning but are outputs from Algorithm 2. For learning, we only employ the onboard automotive-grade wheel-speed sensors for computing the forward velocity, the steering angle, the yaw rate, and lateral acceleration measurements. Due to the zero-mean prior of the function coefficients in (11), the estimates (black solid) do not suffer from overfitting issues outside of the available data range. Instead, they smoothly converge to the zero-mean prior.

To validate the accuracy of the learned tire friction functions, Fig. 3 shows the measured yaw rate (red dashed, upper) and lateral acceleration (red dashed, lower), together with the predicted quantities (black solid) when simulating the system using the steering and longitudinal velocity inputs for a portion of the first of the three data sets. Note that it is a pure simulation of the vehicle model that has been used to generate the trajectories, with the average estimate of the tire friction (i.e., the black solid lines in the two upper left-most plots in Fig. 2) to predict the forces. The resulting tire models give accurate prediction capabilities when comparing to the measured quantities.

VII. CONCLUSION

We presented an experimental evaluation of a novel method for learning the nonlinear function describing the dependence between wheel slip and tire friction. The method is fully Bayesian and is based on recent developments in particle MCMC and GP-SSMs. A key feature is that the method only uses inertial, steering-wheel, and wheel-speed sensors, which are typically installed in production vehicles.

From the experimental results on three data sets obtained on snow, the proposed method seems capable of learning the nonlinear tire-friction curve in a Bayesian framework. We do not have access to ground truth, but the estimates are consistent with estimates reported in literature.

REFERENCES

- [1] H. B. Pacejka, *Tire and Vehicle Dynamics*, 2nd ed. Oxford, United Kingdom: Butterworth-Heinemann, 2006.
- [2] U. Kiencke and L. Nielsen, *Automotive Control Systems—For Engine, Driveline and Vehicle*, 2nd ed. Berlin Heidelberg: Springer-Verlag, 2005.
- [3] J. Svendenius, “Tire modeling and friction estimation,” Ph.D. dissertation, Dept. Automatic Control, Lund University, Sweden, Apr. 2007.
- [4] C. Canudas-de Wit, P. Tsiotras, E. Velenis, M. Basset, and G. Gissinger, “Dynamic friction models for road/tire longitudinal interaction,” *Veh. Syst. Dyn.*, vol. 39, no. 3, pp. 189–226, 2003.

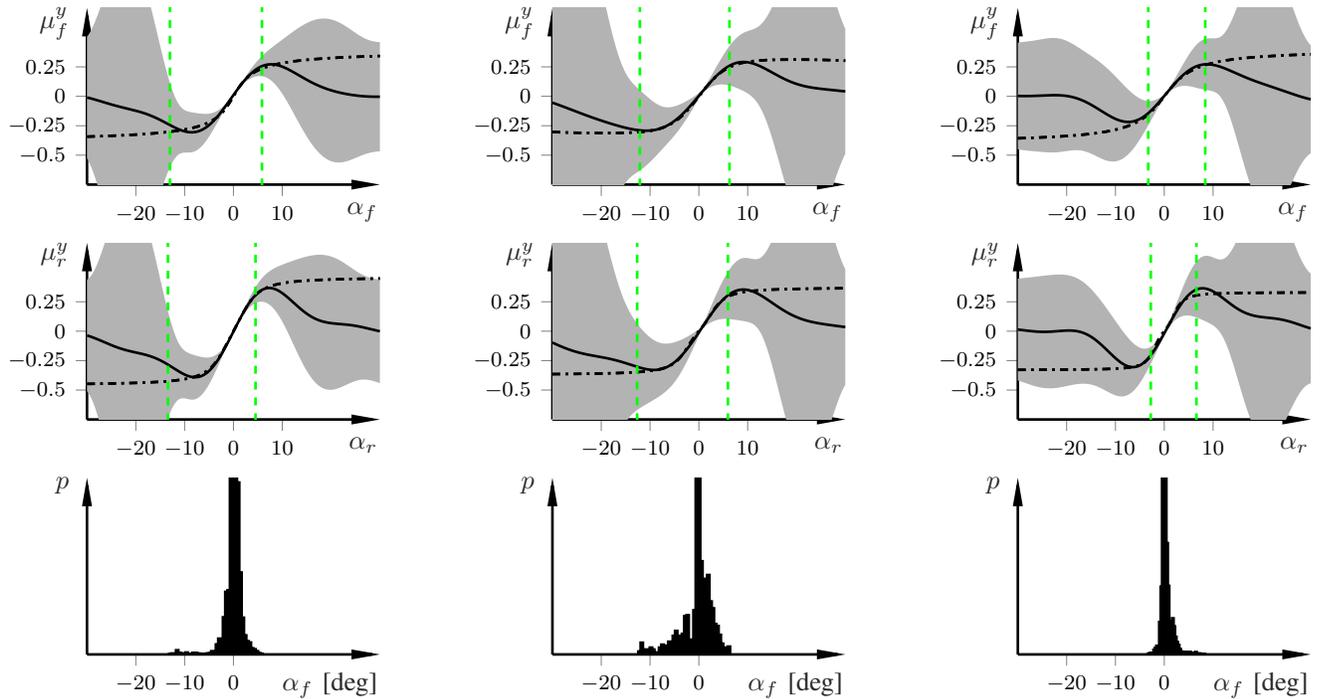


Fig. 2. Bayesian learning of the tire friction with inertial sensing of the front tire (first row), rear tire (second row), and the estimated distribution of the underlying data (third row), for the three data sets used in the evaluation. The black line is the posterior mean of $p(\mu|y_{0:T})$ generated from Algorithm 2 and the estimated 3σ boundaries are indicated by the shadowed areas. The dash-dotted lines are the respective Pacejka tire model fitted according to (30). The bars in the third row show the distribution of the estimated mean slip angle of the front wheel obtained from the estimation algorithm, and the green vertical dashed lines in the two upper rows indicate the range of the underlying data (c.f. third row for front tire).

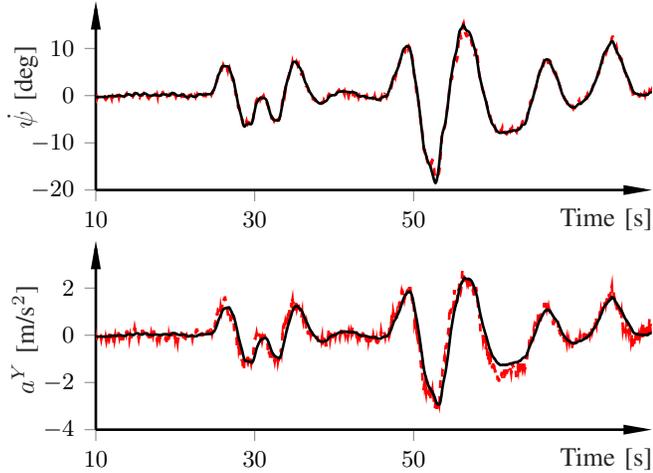


Fig. 3. Predicted (black solid) and ground truth (red dashed) yaw rate and lateral acceleration, for the first of the three data sets (left column in Fig. 2).

[5] R. Quirynen, K. Berntorp, and S. Di Cairano, "Embedded optimization algorithms for steering in autonomous vehicles based on nonlinear model predictive control," in *Amer. Control Conf.*, Milwaukee, WI, June 2018.

[6] B. Goldfain, P. Drews, C. You, M. Barulic, O. Velev, P. Tsiotras, and J. M. Rehg, "Autorally: An open platform for aggressive autonomous driving," *arXiv preprint arXiv:1806.00678*, 2018.

[7] C. Lundquist and T. B. Schön, "Recursive identification of cornering stiffness parameters for an enhanced single track model," in *15th IFAC Symp. System Identification*, Saint-Malo, France, July 2009.

[8] S. Lee, K. Nakano, and M. Ohori, "On-board identification of tyre cornering stiffness using dual Kalman filter and GPS," *Veh. Syst. Dyn.*, vol. 53, no. 4, pp. 437–448, 2015.

[9] C. Ahn, H. Peng, and H. E. Tseng, "Robust estimation of road friction coefficient using lateral and longitudinal vehicle dynamics," *Veh. Syst. Dyn.*, vol. 50, no. 6, pp. 961–985, 2012.

[10] F. Gustafsson, "Automotive safety systems," *IEEE Signal Process. Mag.*, vol. 26, no. 4, pp. 32–47, July 2009.

[11] K. Berntorp, "Bayesian tire-friction learning by Gaussian-process state-space models," in *Eur. Control Conf.*, Naples, Italy, June 2019.

[12] C. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA, USA: MIT Press, 2006.

[13] A. Doucet and A. M. Johansen, "A tutorial on particle filtering and smoothing: Fifteen years later," in *Handbook of Nonlinear Filtering*, D. Crisan and B. Rozovsky, Eds. Oxford University Press, 2009.

[14] F. Lindsten, M. I. Jordan, and T. B. Schön, "Particle Gibbs with ancestor sampling," *J. Machine Learning Res.*, vol. 15, no. 1, pp. 2145–2184, 2014.

[15] A. Svensson and T. B. Schön, "A flexible state-space model for learning nonlinear dynamical systems," *Automatica*, vol. 80, pp. 189–199, 2017.

[16] A. Solin and S. Särkkä, "Hilbert space methods for reduced-rank Gaussian process regression," *arXiv preprint arXiv:1401.5508*, 2014.

[17] C. Robert and G. Casella, *Monte Carlo statistical methods*. Springer, 2013.

[18] C. Andrieu, A. Doucet, and R. Holenstein, "Particle Markov chain Monte Carlo methods," *J. Royal Statistical Society: Series B (Statistical Methodology)*, vol. 72, no. 3, pp. 269–342, 2010.

[19] E. Özkan, V. Šmídl, S. Saha, C. Lundquist, and F. Gustafsson, "Marginalized adaptive particle filtering for nonlinear models with unknown time-varying noise parameters," *Automatica*, vol. 49, no. 6, pp. 1566–1575, 2013.

[20] K. Berntorp and S. Di Cairano, "Tire-stiffness and vehicle-state estimation based on noise-adaptive particle filtering," *IEEE Trans. Control Syst. Technol.*, 2018, in press.

[21] A. Wills, T. B. Schön, F. Lindsten, and B. Ninness, "Estimation of linear systems using a Gibbs sampler," in *IFAC Symp. System Identification*, Brussels, Belgium, 2012.

[22] A. P. Dawid, "Some matrix-variate distribution theory: notational considerations and a Bayesian application," *Biometrika*, vol. 68, no. 1, pp. 265–274, 1981.

[23] T. B. Schön, F. Gustafsson, and P.-J. Nordlund, "Marginalized particle filters for mixed linear nonlinear state-space models," *IEEE Trans. Signal Process.*, vol. 53, pp. 2279–2289, 2005.