Online Joint State Inference and Learning of Partially Unknown State-Space Models

Anton Kullberg, Isaac Skog and Gustaf Hendeby

The self-archived postprint version of this journal article is available at Linköping University Institutional Repository (DiVA):

<http://urn.kb.se/resolve?urn=urn:nbn:se:liu:diva-178181>

N.B.: When citing this work, cite the original publication.

Kullberg, A., Skog, I., Hendeby, G., (2021), Online Joint State Inference and Learning of Partially Unknown State-Space Models, *IEEE Transactions on Signal Processing*, 69, 4149-4161.

<https://doi.org/10.1109/TSP.2021.3095709>

Original publication available at:

<https://doi.org/10.1109/TSP.2021.3095709>

Copyright: Institute of Electrical and Electronics Engineers

<http://www.ieee.org/index.html>

©2021 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

Online Joint State Inference and Learning of Partially Unknown State-Space Models

Anton Kullberg, Isaac Skog, *Senior Member, IEEE*, Gustaf Hendeby, *Senior Member, IEEE*

Abstract—A computationally efficient method for online joint state inference and dynamical model learning is presented. The dynamical model combines an a priori known, physically derived, state-space model with a radial basis function expansion representing unknown system dynamics and inherits properties from both physical and data-driven modeling. The method uses an extended Kalman filter approach to jointly estimate the state of the system and learn the unknown system dynamics, via the parameters of the basis function expansion. The key contribution is a computational complexity reduction compared to a similar approach with globally supported basis functions. By using compactly supported radial basis functions and an approximate Kalman gain, the computational complexity is considerably reduced and is essentially determined by the support of the basis functions. The approximation works well when the system dynamics exhibit limited correlation between points well separated in the state-space domain. The method is exemplified via two intelligent vehicle applications where it is shown to: (i) have competitive system dynamics estimation performance compared to the globally supported basis function method, and (ii) be real-time applicable to problems with a large-scale state-space.

Index Terms—Kalman filters, System identification, Computational complexity

I. INTRODUCTION

STATE estimation is a fundamental problem in many areas, such as robotics, target tracking, economics, etc. For a general nonlinear state-space model, typical techniques used in state estimation include Bayesian filters such as the *extended Kalman filter* (EKF), *unscented Kalman filter* (UKF) or the *particle filter* (PF) [1]. Fundamental to these techniques is the specification of a *state-space model* (SSM), which describes the dynamics of the underlying system and how the state of the system relates to the observations. These models can be broadly classified into three different categories: white-box (physically derived), gray-box (semi-physical) and black-box (non-physical) models. Black-box models generally do not assume any prior knowledge of the system. One particular issue with black-box models is that their generalization properties outside training regions is not known. Thus, black-box

models may not be suited for operation critical applications. On the other hand, white-box modeling use domain-specific knowledge to derive a model based on the physical properties of the system. For instance, in the case of vehicle tracking, constant velocity or coordinated turn models are commonly used as first-order physical models to describe the vehicle dynamics [2]. These first-order physical models summarize any unknown behavior as white process noise, which leads to a state estimation bias [3]. One way to overcome the problems with white-box and black-box models is to use a gray-box modeling approach and combine the first-order physical model with a generic black-box model and try to learn (identify) a more refined model on-the-fly. Such a model is nearly as flexible as black-box models, yet can still guarantee minimum performance outside training regions when used for state estimation.

One way of constructing such a gray-box model is to augment a physically derived SSM with a black-box model structure. The black-box model augmentation can be chosen in a variety of ways; for instance, as a *Gaussian process* (GP) [4], [5], as a basis function (BF) expansion [6], or as a neural network. Henceforth, we shall refer to these models as SSM+GP, SSM+BF, etc. Another, highly related gray-box model, is the *Gaussian process state-space model* (GPSSM), which was recently developed for joint state inference and system identification [7]–[10]. In the GPSSM, both the state dynamics and observation model are viewed as separate GPs and are learned jointly. Further, prior system information can be included in the state-space model as a known term [11]–[13].

The use of a GP model is attractive because of its non-parametric nature. However, since it uses all of the accumulated data during learning and inference, it is in practice still necessary to parameterize it to be able to use it efficiently. Particularly, the computational complexity of the standard GP is $\mathcal{O}(N^3)$ and $\mathcal{O}(N^2)$ during learning and inference, respectively, where N is the number of data points. Further, the storage complexity is $\mathcal{O}(N^2)$. One way of reducing the complexity is to split the GP input into smaller, assumed independent, regions and fit a separate GP per region, such as in [14], [15]. In those cases, it becomes important to handle the boundaries of the regions and it is not completely clear how to choose the shape and size of the regions. Another alternative is to parametrize the GP with inducing points, see e.g., [16], [17]. This is a way of summarising the accumulated function observations in a set of discrete points, thereby bounding the computational complexity during learning and inference. The computational complexity then drops to $\mathcal{O}(NM^2)$ and

A. Kullberg, I. Skog and G. Hendeby are with the Department of Electrical Engineering, Division of Automatic Control, Linköping University, Linköping, SE-58183, Sweden (E-mail: {anton.kullberg, isaac.skog, gustaf.hendeby}@liu.se).

This paper has supplementary downloadable material available at <https://ieeexplore.ieee.org>, provided by the author. The material includes a derivation of the motion model used in the second numerical example. This material is 8 MB in size.

This work was partially supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

$\mathcal{O}(M^2)$ for learning and inference, respectively, where M is the number of inducing points. Further, the storage complexity is reduced to $\mathcal{O}(MN + M^2)$. Hence, to see any major benefit, the number of inducing points, M should be much smaller than N [16]. Essentially all of [7]–[10] use inducing point approaches to reduce the computational burden of the GP.

A related approach is the SSM+BF model, which instead parametrizes the unknown function using a BF expansion. In the limit, *i.e.*, when the number of basis functions goes to infinity, these two function representations are equivalent, under certain conditions, and share many properties. Particularly, [11]–[13] use a regularized BF expansion that can be viewed as a SSM+GP model. The benefit of using a BF expansion is that it offers an intuitive appeal with regard to function approximation and is easy to include in the state-space model. However, in its standard formulation, it still suffers from the same computational issues as the SSM+GP methods. Recently, [13] extended the GPSSM inference framework in [11] to an online setting and used the method to estimate lateral tire-friction in real-time. However, [13] was focused on limited state-spaces, such that the function expansion requires only a limited amount of basis functions (< 100). The common limitation of the standard SSM+BF and SSM+GP models, is that all of the parameters of the GP or BF expansion are required for each prediction/update step of the filtering algorithms that are applied to them. This effectively limits the methods to problems where the number of inducing points/basis functions can be kept less than a few thousand. Hence, they are not suitable when the bandwidth or the support of the function to be learned grows.

Therefore, in this paper, a scalable online method for joint state inference and the system identification of SSM+BF gray-box models is presented. By using basis functions with a limited support the method is able to handle models with large-scale state-spaces. The connection between the suggested SSM+BF model and the SSM+GP model is highlighted. Further, the computational complexity of performing state inference and model learning with the suggested SSM+BF model using an extended Kalman filter is analyzed. The performance of the suggested SSM+BF model and inference framework is demonstrated using two intelligent vehicle applications.

II. PROBLEM FORMULATION

A general gray-box SSM with partially unknown state dynamics is given by

$$\mathbf{x}_{k+1} = \mathbf{f}_k(\mathbf{x}_k, \mathbf{u}_k, \mathbf{w}_k, \mathbf{g}_k(\mathbf{x}_k, \mathbf{u}_k)) \quad (1a)$$

$$\mathbf{y}_k = \mathbf{h}_k(\mathbf{x}_k, \mathbf{u}_k, \mathbf{e}_k). \quad (1b)$$

Here, $\mathbf{x}_k$ and $\mathbf{y}_k$ denotes the unknown system state and observations at time k , respectively. Further, $\mathbf{u}_k$ is a known input to the system. Moreover, $\mathbf{w}_k$ and $\mathbf{e}_k$ denote the process and observation noise, respectively. The process and observation noises are assumed to be mutually independent white noise processes with covariance matrices $\mathbf{Q}_k$ and $\mathbf{R}_k$, respectively. The functions $\mathbf{f}_k$ and $\mathbf{h}_k$, which specify the state dynamics

and the observation model, are chosen a priori, and are here assumed to be differentiable. Lastly, $\mathbf{g}_k$ is an unknown function, dependent on the state $\mathbf{x}_k$ and the known input $\mathbf{u}_k$, which is to be inferred from data. Further, depending on how the function $\mathbf{f}_k$ is chosen, $\mathbf{g}_k$ can be interpreted in different ways. For instance, if $\mathbf{f}_k$ is chosen as a *constant velocity* (CV) model [2], $\mathbf{g}_k$ represents the system acceleration, see Section V-A. Moreover, note that $\mathbf{g}_k$ could be additive or multiplicative and could affect but a subset of $\mathbf{x}_k$.

Data from a single system allows the identification (learning) of the unknown function $\mathbf{g}_k$. However, data from multiple systems, whose dynamics depend on $\mathbf{g}_k$, is often available. For instance, ships sailing in certain regions are all influenced by the same currents, neighboring buildings are all affected by the same winds, and road vehicles traversing the same road network are all subject to the same road conditions. Hence, a particularly interesting application of the model (1) is when collaborative learning of $\mathbf{g}_k$ is possible. It is also worth mentioning that, within this paper, explicit references to several systems have been omitted for notational brevity.

A. Basis Function Expansion

As previously mentioned, the function $\mathbf{g}_k$ can be modeled in a variety of ways. Here, it is modeled as a basis function expansion, *i.e.*,

$$[\mathbf{g}_k]_j = (\phi_k^j)^\top \boldsymbol{\theta}_k^j \quad j = 1, \dots, J \quad (2)$$

where $[\mathbf{g}_k]_j$ denotes the j :th component of $\mathbf{g}_k$. Further,

$$\begin{aligned} \phi_k^j &= \left(\phi_1^j \quad \dots \quad \phi_{n_\theta^j}^j \right)^\top, \\ \boldsymbol{\theta}_k^j &= \left(\theta_{k,1}^j \quad \dots \quad \theta_{k,n_\theta^j}^j \right)^\top, \end{aligned}$$

where $\phi_i^j = \phi_i^j(\mathbf{x}_k, \mathbf{u}_k)$ is the i :th basis function corresponding to component j of $\mathbf{g}_k$ and $\theta_{k,i}^j$ is the weight of this basis function at time k . Moreover, the dependence of $\mathbf{g}_k$ on $\mathbf{x}_k$ and $\mathbf{u}_k$ is solely through the basis functions ϕ_i^j . Lastly, n_θ^j denotes the number of basis functions used to represent component j of $\mathbf{g}_k$.

The basis function expansion (2) is closely related to the GP representation of $\mathbf{g}_k$. With a Gaussian prior on the weights, $\boldsymbol{\theta}_0^j \sim \mathcal{N}(0, \sigma_j^2 I)$, $[\mathbf{g}_k]_j$ will correspond to a GP with kernel

$$\kappa^j(\mathbf{x}, \mathbf{x}') = \sigma_j^2 \sum_{i=1}^{n_\theta^j} \phi_i^j(\mathbf{x}) \phi_i^j(\mathbf{x}'), \quad (3)$$

if the basis functions ϕ_i^j are such that the matrix

$$\mathbf{K}^j = \sigma_j^2 \begin{bmatrix} \phi_1^j(\mathbf{x}) \phi_1^j(\mathbf{x}') & \dots & \phi_1^j(\mathbf{x}) \phi_{n_\theta^j}^j(\mathbf{x}') \\ \vdots & \ddots & \vdots \\ \phi_{n_\theta^j}^j(\mathbf{x}) \phi_1^j(\mathbf{x}') & \dots & \phi_{n_\theta^j}^j(\mathbf{x}) \phi_{n_\theta^j}^j(\mathbf{x}') \end{bmatrix} \quad (4)$$

is at least positive semi-definite. Here, $\mathbf{x}$ and $\mathbf{x}'$ are two data points, see, *e.g.*, [18]. Particularly, if the basis function is chosen as a Gaussian, *i.e.*,

$$\phi_i^j(\mathbf{x}) = \exp \left(-\frac{(\mathbf{x} - \boldsymbol{\xi}_i^j)^2}{2l^2} \right), \quad (5)$$

with center $\xi_i^j \in \Xi$, where Ξ is the set of all ξ_i^j , and scaling parameter l . Then, if Ξ is a regular, equally-spaced grid, the basis function expansion will correspond to a GP with a standard squared exponential kernel as the number of basis functions $n_\theta^j \rightarrow \infty$ [18].

The basis function expansion can, in other words, be seen as a parametrization of a GP. The main difference between the approaches is what is being “learned.” In an inducing point approach, function values are learned whereas in a basis function expansion, basis function weights are learned. Regardless of the semantic differences, both of the approaches suffer from the same problems, which primarily concern a restrictive computational complexity. Henceforth, we shall focus on the basis function expansion (2) and begin with casting the problem of learning (identifying) the function $\mathbf{g}_k$ into a state estimation problem. Thereafter, we will present a method to reduce the computational complexity.

To that end, form the augmented state vector and process noise vector

$$\mathbf{x}_k^e = (\mathbf{x}_k^\top \quad \boldsymbol{\theta}_k^\top)^\top \quad \mathbf{w}_k^e = (\mathbf{w}_k^\top \quad \dot{\mathbf{w}}_k^\top)^\top, \quad (6)$$

where

$$\boldsymbol{\theta}_k = ((\boldsymbol{\theta}_k^1)^\top \quad \dots \quad (\boldsymbol{\theta}_k^J)^\top)^\top.$$

Then, the model described by (1) and (2) can be written as

$$\mathbf{x}_{k+1}^e = \mathbf{f}_k^e(\mathbf{x}_k^e, \mathbf{u}_k, \mathbf{w}_k^e) \quad (7a)$$

$$\mathbf{f}_k^e \triangleq \begin{pmatrix} \mathbf{f}_k(\mathbf{x}_k, \mathbf{u}_k, \mathbf{w}_k, \Phi_k \boldsymbol{\theta}_k) \\ \boldsymbol{\theta}_k + \dot{\mathbf{w}}_k \end{pmatrix}, \quad (7b)$$

where

$$\Phi_k = \text{diag}((\phi_k^1)^\top \quad \dots \quad (\phi_k^J)^\top),$$

and $\dot{\mathbf{w}}_k$ is a white process noise with covariance matrix $\boldsymbol{\Sigma}_k$, that allows $\mathbf{g}_k$ to change over time and is assumed to be independent of $\mathbf{w}_k$ and $\mathbf{e}_k$. Hence, the problem of learning the function $\mathbf{g}_k$ has been cast into a state estimation problem, and both the states $\mathbf{x}_k$ and function $\mathbf{g}_k$ can be inferred simultaneously.

B. Computational Complexity

Focusing on the model defined by (1) and (7), the computational complexity of standard state estimation algorithms is in the order of $\mathcal{O}(d^3)$ or more [19], where d is the dimension of the state vector. Hence, nothing has yet been gained in terms of computational complexity, thus leaving the problem of handling models in which the function $\mathbf{g}_k$ varies quickly or has a large support, unsolved. To simplify the forthcoming discussion about how the computation complexity can be reduced, we will focus on estimating the states of (7) using an EKF. This because (i) the EKF is one of the most commonly used nonlinear filtering approaches, and (ii) the EKF is a computationally efficient approximation of the general recursive Bayesian filter, which fits well into the focus of computational efficiency here. Noteworthy is that the linearizations in the EKF may cause problems with highly nonlinear systems and particularly if the function $\mathbf{g}_k$ being learned is highly nonlinear in the states $\mathbf{x}_k$. In such cases, it may be necessary to substitute the proposed EKF for a

marginalized PF, as in [11], [13], or some type of linear regression Kalman filter, such as the UKF [20], or the S^2KF [21]. Even though the complexity reduction method presented next is not directly applicable to these cases, large parts of it can be used.

The EKF propagates the first two moments in each iteration and hence, only the mean vector $\hat{\mathbf{x}}^e$ and covariance matrix $\mathbf{P}$, need to be stored, defined as

$$\hat{\mathbf{x}}_k^e = \begin{bmatrix} \hat{\mathbf{x}}_k \\ \hat{\boldsymbol{\theta}}_k \end{bmatrix} = \mathbb{E} \begin{bmatrix} \mathbf{x}_k \\ \boldsymbol{\theta}_k \end{bmatrix} \quad (8a)$$

$$\mathbf{P}_k = \mathbb{E}[(\mathbf{x}_k^e - \hat{\mathbf{x}}_k^e)(\mathbf{x}_k^e - \hat{\mathbf{x}}_k^e)^\top]. \quad (8b)$$

If an EKF is directly designed based upon model (7), the filter has a computational complexity of $\mathcal{O}((n_x^e)^3)$, where $n_x^e = \dim(\mathbf{x}_k^e)$. However, using the fact that the weights $\boldsymbol{\theta}$ are modeled as a random walk process, the computational complexity can be reduced. By using this structure for the time update, the prediction mean is given by

$$\hat{\mathbf{g}}_{k|k} = (\mathbf{I} \otimes \phi_k^\top(\hat{\mathbf{x}}_{k|k})) \hat{\boldsymbol{\theta}}_{k|k} \triangleq \Phi_k(\hat{\mathbf{x}}_{k|k}) \hat{\boldsymbol{\theta}}_{k|k} \quad (9a)$$

$$\hat{\mathbf{x}}_{k+1|k} = \mathbf{f}_k(\hat{\mathbf{x}}_{k|k}, \mathbf{u}_{k+1}, \mathbf{w}_{k+1}, \hat{\mathbf{g}}_{k|k}) \quad (9b)$$

$$\hat{\boldsymbol{\theta}}_{k+1|k} = \hat{\boldsymbol{\theta}}_{k|k}. \quad (9c)$$

Further, the predictive covariance is given by

$$\begin{aligned} \mathbf{P}_{k+1|k} &= \begin{bmatrix} \mathbf{P}_{k+1|k}^{xx} & \mathbf{P}_{k+1|k}^{x\theta} \\ \mathbf{P}_{k+1|k}^{\theta x} & \mathbf{P}_{k+1|k}^{\theta\theta} \end{bmatrix} = \mathbf{F}_k \mathbf{P}_{k|k} \mathbf{F}_k^\top + \mathbf{Q}_k^e \\ &= \begin{bmatrix} \mathbf{F}_x & \mathbf{F}_\theta \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{P}_{k|k}^{xx} & \mathbf{P}_{k|k}^{x\theta} \\ \mathbf{P}_{k|k}^{\theta x} & \mathbf{P}_{k|k}^{\theta\theta} \end{bmatrix} \begin{bmatrix} \mathbf{F}_x^\top & \mathbf{0} \\ \mathbf{F}_\theta^\top & \mathbf{I} \end{bmatrix} \\ &\quad + \begin{bmatrix} \mathbf{Q}_k & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\Sigma}_k \end{bmatrix} \end{aligned} \quad (10)$$

which can be written as

$$\begin{aligned} \mathbf{P}_{k+1|k}^{xx} &= \mathbf{F}_x \mathbf{P}_{k|k}^{xx} \mathbf{F}_x^\top + \mathbf{F}_x \mathbf{P}_{k|k}^{x\theta} \mathbf{F}_\theta^\top \\ &\quad + \mathbf{F}_\theta \mathbf{P}_{k|k}^{\theta x} \mathbf{F}_x^\top + \mathbf{F}_\theta \mathbf{P}_{k|k}^{\theta\theta} \mathbf{F}_\theta^\top + \mathbf{Q}_k \end{aligned} \quad (11a)$$

$$\mathbf{P}_{k+1|k}^{x\theta} = \mathbf{F}_x \mathbf{P}_{k|k}^{x\theta} + \mathbf{F}_\theta \mathbf{P}_{k|k}^{\theta\theta} \quad (11b)$$

$$\mathbf{P}_{k+1|k}^{\theta x} = (\mathbf{P}_{k+1|k}^{x\theta})^\top \quad (11c)$$

$$\mathbf{P}_{k+1|k}^{\theta\theta} = \mathbf{P}_{k|k}^{\theta\theta} + \boldsymbol{\Sigma}_k. \quad (11d)$$

Here,

$$\mathbf{F}_k \triangleq \nabla_{\mathbf{x}^e} \mathbf{f}_k^e = \begin{bmatrix} \nabla_{\mathbf{x}} \mathbf{f}_k & \nabla_{\boldsymbol{\theta}} \mathbf{f}_k \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \triangleq \begin{bmatrix} \mathbf{F}_x & \mathbf{F}_\theta \\ \mathbf{0} & \mathbf{I} \end{bmatrix}. \quad (12)$$

Inspecting (9)–(11), the main computational burden in the time update is related to (11), and more specifically, the update of $\mathbf{P}^{xx}$ in (11a), where the complexity of the product $\mathbf{F}_\theta \mathbf{P}^{\theta\theta} \mathbf{F}_\theta^\top$ is in the order of $\mathcal{O}(n_x n_\theta^2 + n_x^2 n_\theta) < \mathcal{O}((n_x^e)^3)$. Further, after a measurement, the corrected mean is given by

$$\hat{\mathbf{x}}_{k|k} = \hat{\mathbf{x}}_{k|k-1} + \mathbf{K}_k^x (\mathbf{y}_k - \mathbf{H}_k^x \hat{\mathbf{x}}_{k|k-1}) \quad (13a)$$

$$\hat{\boldsymbol{\theta}}_{k|k} = \hat{\boldsymbol{\theta}}_{k|k-1} + \mathbf{K}_k^\theta (\mathbf{y}_k - \mathbf{H}_k^x \hat{\mathbf{x}}_{k|k-1}), \quad (13b)$$

and the corresponding covariance update, in the less commonly used Joseph’s form, is given by

$$\mathbf{P}_{k|k} = (\mathbf{I} - \mathbf{K}_k \mathbf{H}_k) \mathbf{P}_{k|k-1} (\mathbf{I} - \mathbf{K}_k \mathbf{H}_k)^\top + \mathbf{K}_k \mathbf{R} \mathbf{K}_k^\top. \quad (14)$$

Here,

$$\mathbf{H}_k \triangleq \nabla_{\mathbf{x}_{k|k}^e} \mathbf{h}_k(\hat{\mathbf{x}}_{k|k}, \mathbf{u}_k, \mathbf{e}_k) = [\mathbf{H}_k^x \quad \mathbf{0}] \quad (15a)$$

$$\mathbf{S}_k \triangleq \mathbf{R}_k + \mathbf{H}_k^x \mathbf{P}_{k|k-1}^{xx} \mathbf{H}_k^{x\top} \quad (15b)$$

$$\mathbf{K}_k^x \triangleq \mathbf{P}_{k|k-1}^{xx} \mathbf{H}_k^{x\top} \mathbf{S}_k^{-1} \quad (15c)$$

$$\mathbf{K}_k^\theta \triangleq \mathbf{P}_{k|k-1}^{\theta x} \mathbf{H}_k^{x\top} \mathbf{S}_k^{-1}, \quad (15d)$$

where (15a) follows from the independence between $\mathbf{h}_k$ and θ_k . The covariance update in (14) can also be written in its separate components as

$$\mathbf{P}_{k|k}^{xx} = (\mathbf{I} - \mathbf{K}_k^x \mathbf{H}_k^x) \mathbf{P}_{k|k-1}^{xx} (\mathbf{I} - \mathbf{K}_k^x \mathbf{H}_k^x)^\top + \mathbf{K}_k^x \mathbf{R}_k \mathbf{K}_k^{x\top} \quad (16a)$$

$$\begin{aligned} \mathbf{P}_{k|k}^{x\theta} &= (\mathbf{I} - \mathbf{K}_k^x \mathbf{H}_k^x) \mathbf{P}_{k|k-1}^{xx} (-\mathbf{K}_k^\theta \mathbf{H}_k^x)^\top \\ &\quad + (\mathbf{I} - \mathbf{K}_k^x \mathbf{H}_k^x) \mathbf{P}_{k|k-1}^{x\theta} + \mathbf{K}_k^x \mathbf{R}_k \mathbf{K}_k^{\theta\top} \end{aligned} \quad (16b)$$

$$\mathbf{P}_{k|k}^{\theta x} = (\mathbf{P}_{k|k}^{x\theta})^\top \quad (16c)$$

$$\begin{aligned} \mathbf{P}_{k|k}^{\theta\theta} &= \mathbf{P}_{k|k-1}^{\theta\theta} + \mathbf{K}_k^\theta \mathbf{H}_k^x \mathbf{P}_{k|k-1}^{xx} \mathbf{H}_k^{x\top} \mathbf{K}_k^{\theta\top} \\ &\quad - \mathbf{K}_k^\theta \mathbf{H}_k^x \mathbf{P}_{k|k-1}^{x\theta} - \mathbf{P}_{k|k-1}^{\theta x} \mathbf{H}_k^{x\top} \mathbf{K}_k^{\theta\top} + \mathbf{K}_k^\theta \mathbf{R}_k \mathbf{K}_k^{\theta\top}. \end{aligned} \quad (16d)$$

Note that (14) and (16) are guaranteed to be positive semi-definite regardless of the choice of gain $\mathbf{K}_k$, which shall be of importance later.

From (13)–(16), it is clear that the main computational complexity of the measurement update is related to (16d), which is in the order of $\mathcal{O}(n_\theta n_y^2 + n_\theta^2 n_y)$, where n_y is the size of the observation vector $\mathbf{y}$. Specifically, products involving both $\mathbf{K}_k^\theta$ and $\mathbf{K}_k^{\theta\top}$ are the culprits, e.g., the product $\mathbf{K}_k^\theta \mathbf{R}_k \mathbf{K}_k^{\theta\top}$. If it is (reasonably) assumed that $n_\theta \gg n_x$ and $n_\theta \gg n_y$, the computational complexity is thus approximately $\mathcal{O}(n_\theta^2)$. Even though this is a substantial reduction in computational complexity, it is still prohibitive if n_θ is large.

The main computational burden is caused by the basis functions ϕ . Assume that the basis functions have global support, i.e., $\phi(\mathbf{x}) \neq 0, \forall \mathbf{x}$, which could be, for instance, the Gaussian basis function (5). Essentially, this means that regardless of $\mathbf{x}$, all of the basis functions will be non-zero and are required in both the time update and measurement update of the EKF. In particular, the matrices $\mathbf{P}^{\theta x}$, $\mathbf{P}^{x\theta}$ and $\mathbf{P}^{\theta\theta}$ will be dense. The naive solution would be to decrease n_θ to yield a tractable complexity for the problem at hand. However, if the domain of $\mathbf{g}$ is large, then n_θ needs to be large to cover the entire domain. Similarly, if $\mathbf{g}$ varies quickly in regard to the input, then n_θ needs to be large enough to accurately represent the function. This leads to a conundrum as n_θ needs to be small enough to remain with a tractable algorithm, yet also large enough to represent complex functions. Hence, with a globally supported basis function, there is a trade-off between computational complexity and the types of functions $\mathbf{g}$ can represent.

To reduce the complexity of learning $\mathbf{g}$ and estimating the state $\mathbf{x}_k$, we seek a basis function ϕ so that only a few of the basis function weights θ need to be used at each filter step. In the following section, we present a basis function, which fulfills this criterion. We now restrict the basis functions ϕ_i^j to the family of *radial basis functions* (RBFs) [22], as they are

universal approximators [23]. It does simplify the discussion, but we stress that it is possible to use other basis functions and reach similar conclusions.

III. COMPUTATIONAL COMPLEXITY REDUCTION USING COMPACT BASIS FUNCTIONS

The computational complexity is largely caused by the number of basis functions n_θ . In particular, the problem is that, given a set of globally supported basis functions, all of the basis functions need to be evaluated in each time step, since $\mathbf{x}_k$ depends on the entire θ_k (recall $\mathbf{g}_k = \Phi_k \theta_k$). If, instead, $\mathbf{x}_k$, at every time instant, only depended on a subset of θ_k , only a few weights would need to be updated in each step of the filtering algorithm. This restricts the possible choice of basis function to those with *compact support*, i.e., basis functions which are only non-zero on a subset of $\mathcal{X}$, where $\mathbf{x}_k \in \mathcal{X}$. Let $\tilde{n}_\theta$ be the number of non-zero basis functions at any time k . Then, if $\tilde{n}_\theta \ll n_\theta$, the computational complexity can be greatly reduced.

A. Basis Functions with Compact Support

In the family of RBFs, a number of basis functions have compact support [18], [22], i.e.,

$$\exists c, \quad \text{s.t.} \quad \phi_i^j(\mathbf{x}_k) = 0, \quad \text{if} \quad r_{k,i}^j = \|\mathbf{x}_k - \xi_i^j\| > c. \quad (17)$$

These basis functions are commonly referred to as *compactly supported radial basis functions* (CSRBFs). CSRBFs have been studied extensively in many different scientific disciplines, for instance, the geostatistical community [24], [25]. They have also been used in ensemble Kalman filters for atmospheric data assimilation [26], [27]. In the geostatistical community, CSRBFs have mainly been used for data interpolation, where the choice of basis function is of great importance in recovering the “true” function exactly. As alluded to before, in the context of this paper, the choice is important only for the interpretability as a GP and in general, the choice can be made more freely. Further, smoothness and continuous differentiability are convenient and sometimes necessary. For instance, if a standard EKF is used, differentiability is required. Hence, we shall restrict ourselves to Wendland functions [28], which have been shown to be positive definite, up to some dimension d . Thus, with Wendland basis functions and a Gaussian prior on the weights, the basis function representation can be interpreted as a GP. Further, the Wendland functions have been shown to be of minimal degree, as well as being continuously differentiable both at the origin and the cut-off radius c .

In the remainder of the paper, ϕ is chosen as the particular Wendland function

$$\phi_i^j(r_{k,i}^j) = (1 - r_{k,i}^j)_+^6 (35(r_{k,i}^j)^2 + 18r_{k,i}^j + 3)/3, \quad (18)$$

where $(\cdot)_+ = \max(0, \cdot)$. Clearly, $\phi_i^j(r_{k,i}^j) \equiv 0$ if $r_{k,i}^j \geq 1$. To get a larger or smaller support, $r_{k,i}^j$ can be scaled appropriately. Henceforth, $r_{k,i}^j$ is assumed to be scaled by $1/\alpha$, which, for e.g., $\alpha = 2$, increases the support by a factor two. Lastly, the CSRBF (18), is positive definite in $d = 2$ and 4 times continuously differentiable [28].

By considering any function which satisfies (17), the dynamics of $\mathbf{x}_k$ will now depend on just a few of the basis functions at every time instance k . Let

$$\mathbf{A}_k^j(\mathbf{x}_k) \in \mathbb{R}^{\tilde{n}_\theta^j \times n_\theta^j}$$

be a matrix-valued function that takes $\mathbf{x}_k$ as input and produces an indicator matrix such that only the *active*, or non-zero, basis functions and weights remain (*i.e.*, the products $\mathbf{A}_k^j \phi_k^j$, $\mathbf{A}_k^j \theta_k^j$ contain only non-zero basis functions and the corresponding weights). Moreover, note that the products produce a vector of size $\tilde{n}_\theta^j$. The expansion (2) can then be rewritten as

$$[\mathbf{g}_k]_j = (\mathbf{A}_k^j \phi_k^j)^\top \mathbf{A}_k^j \theta_k^j = (\phi_k^j)^\top (\mathbf{A}_k^j)^\top \mathbf{A}_k^j \theta_k^j. \quad (19)$$

Now, if $\tilde{n}_\theta^j \ll n_\theta^j$, the computational complexity is greatly reduced, since only these weights are needed in the recursions of the filter. Note that $\mathbf{A}_k^j$ is time-varying, as is $\tilde{n}_\theta^j$. Hence, the number of active basis functions varies over time, but is bounded from above. If the set of basis function centers Ξ is assumed to be a regular grid, the upper bound on the number of active basis functions is completely determined by the density of Ξ , as well as the support of the basis functions. For an RBF expansion to be viable in large scale online applications, there are a few more things one needs to consider. First of all, the size and density of the grid of basis functions needs to be chosen in such a way that the entire state-space region of interest is covered. In the CSRBF case, one also needs to consider the scaling factor α , *i.e.*, the support of the basis functions. This is application-specific and the parameters can be found via maximum likelihood estimation, see *e.g.* [29], or they can be manually selected based on expert knowledge. Further, in the CSRBF case, one needs to consider how long spatial correlations to capture, as the support determines how long correlations that can be represented.

In [29], multiple grids with varying resolution and scaling parameters were considered to capture both short-range and long-range correlations. However, in a real-time scenario, the approach taken there is prohibitive, as all of the resolutions are active simultaneously, and thus, the computational complexity is still restrictive. Another multi-resolution basis function expansion was considered in [30], where the multiple resolutions were instead used to vary the fineness of the grid. In this paper, we shall focus on manually selected grid resolutions and leave adaptivity or multi-resolution for future research. There are also a few more nuanced aspects to making the expansion real-time applicable, which we shall discuss next.

B. Finding Necessary Basis Functions

Assuming that the basis function expansion uses RBFs with compact support, how can $\mathbf{A}_k^j$ be determined? The naïve approach is to calculate the basis function values and remove any that are identically 0. If n_θ^j is large, this might prove to be computationally prohibitive. Instead, consider the special case when Ξ is a Cartesian grid (*i.e.*, decomposes as $\Xi = \Xi_1 \times \Xi_2 \times \dots \times \Xi_P$, where Ξ_p are the center coordinates

Fig. 1: Basis functions used in the description of $\mathbf{g}_k(\mathbf{x}_k)$ when considering a grid $\Xi_1, \Xi_2 = [0, 4]$ and $\mathbf{x}_k = [2, 2]^T$. Blue dots are for the global RBF description. The orange dots indicate the active (non-zero) basis functions in the CSRBF description of $\mathbf{g}_k(\mathbf{x}_k)$. The dots within the black square indicates the basis functions selected by the fast CSRBF method.

along the p :th dimension and $\times$ denotes the Cartesian product). Then, if the basis functions is a product basis function, *i.e.*,

$$\phi_i^j(\mathbf{x}) = \prod_{p=1}^P \phi_i^j([x]_p), \quad (20)$$

where $[x]_p$ is the p :th element of $\mathbf{x}$, the basis function values can be computed in each separate dimension. The regression matrix Φ^j can then be expressed as a Kronecker product $\Phi^j = \Phi_1^j \otimes \dots \otimes \Phi_P^j$. For instance, the Gaussian RBF is an example of a product basis function and is separable across dimensions. In the GP description, this would correspond to using a product kernel, *i.e.*,

$$\kappa(\mathbf{x}, \mathbf{x}') = \prod_{l=1}^P \kappa([x]_p, [x']_p). \quad (21)$$

If a P -dimensional square grid is assumed (*i.e.*, Ξ_p is of size m , $\forall p$), using (20) can reduce the number of basis function evaluations from m^P to Pm . Unfortunately, factorizing compact RBFs in this way is usually difficult, which is easy to see from (18). It can, however, be approximately factorized in this way for finding relevant basis functions.

The discussion is now, without loss of generality, temporarily restricted to two dimensions. Essentially, the RBFs considered here are circles in 2D, *i.e.*, if both dimensions are considered simultaneously, the basis functions will yield a circle of non-zero values. If, instead, each dimension is evaluated separately, the basis functions will yield a square of non-zero values. As a square of side-length a will always enclose a circle of diameter a , the necessary basis functions can be approximately computed using (20). See Fig. 1 for an

Fig. 2: Evaluation time of $\mathbf{g}_k$ for increasing number of basis functions. Using the standard compact formulation sees no gain in evaluation time compared to the full RBF formulation since it is still necessary to evaluate all basis functions to determine which ones to use. The approximate compact method vastly improves over both other methods.

illustration of the suggested basis function selection method. The suggested method for finding active basis functions and computing their values, henceforth referred to as fast CSRBF, clearly reduces the amount of necessary basis function evaluations. The drawback is a slight increase in memory requirement as the correlations between these weights need to be stored, as well as a restriction to square grids. It also increases $\tilde{n}_\theta^j$ by the set difference of the CSRBF and fast CSRBF basis function selection methods. However, as this set difference is small compared to the total number of basis functions, a performance increase is still expected.

To illustrate the reduction in computational time when evaluating $\mathbf{g}_k$ using the fast CSRBF method, a 2D basis function grid was constructed, *i.e.*, $P = 2$. A constant velocity model was then used together with a two-dimensional $\mathbf{g}_k$, *i.e.*, $J = 2$. The grid density was then gradually increased to increase the number of basis functions. The three approaches (*i.e.*, full RBF, CSRBF and fast CSRBF), were then used to compute $\mathbf{g}_k$ in a single point $\mathbf{x}_k$. The different evaluation times are shown in Fig. 2. It is clear that the fast CSRBF is superior, especially for large n_θ . The evaluation times of the CSRBF and full RBF descriptions are approximately the same, which is not surprising given that both of the methods are evaluating all of the basis functions.

The one-step-ahead prediction times for the complete model, *i.e.*, predicting the state and covariance at the next time step were also compared for the three methods, see Fig. 3. After $\sum_j n_\theta^j \approx 10000$, the full RBF caused memory overflow and was left out for further increases. The CSRBF and fast CSRBF descriptions have similar performance. However, the fast CSRBF description still has a clear advantage, because of the computationally efficient selection of active basis functions highlighted in Fig. 2. Of course, the fast CSRBF description is limited to regular grids. However, using a regular grid alleviates the need for precise prior knowledge of the geometry of the environment and is instead limited to knowledge about

Fig. 3: Computational time as a function of the number of basis functions when performing one-step ahead prediction with a two-dimensional basis function grid and a constant velocity dynamical model augmented with a two-dimensional $\mathbf{g}_k$.

the extent of the scenario. Hence, regular grids are potentially more practically useful than non-regular grids. Furthermore, both the fast CSRBF method as well as the ordinary CSRBF methods are nearly invariant to the total number of basis functions, see Fig. 3. Hence, it is mainly the amount of available memory that limits the number of basis functions to be used.

C. Choice of Basis Function Sets

The set of basis function centers Ξ can be chosen rather arbitrarily. Let $\mathbf{x}_k \in \mathcal{X}$, $\forall k$ denote the input to ϕ_i^j . Then, the ξ_i^j are typically chosen from $\mathcal{X}$, *i.e.*, $\Xi \subseteq \mathcal{X}$. As the method described herein is concerned with basis functions with compact support, it is important that the basis function grid covers the entire domain of $\mathbf{g}_k$, as the function will be identically equal to 0 outside the support of the basis functions.

Further, if $\mathbf{g}_k$ is multi-dimensional, then several basis function expansions are necessary. If the effects being modeled are similar in nature, such as the two-dimensional globally aligned acceleration of a vehicle, it is reasonable to use the same basis functions for all of the expansions (*i.e.* $\phi_k^1 = \dots = \phi_k^J$). Since RBFs only depend on $\|\mathbf{x}_k - \xi_i^j\|$, the same basis function values can be used for each separate dimension j of $\mathbf{g}_k$. Hence, the number of basis function evaluations are reduced from n_θ to n_θ^j .

D. Basis Function Weight Ordering

Assuming $\phi_k^1 = \dots = \phi_k^J = \phi_k$ (*i.e.*, the same set of basis functions are used for all dimensions of $\mathbf{g}_k$), then two distinct options of weight ordering (*i.e.*, the ordering of θ_k) exist. The first option is to stack all of the weights as $\theta_k = ((\theta_k^1)^\top \dots (\theta_k^J)^\top)^\top$. This leads to

$$\mathbf{g}_k = (I \otimes \phi_k^\top) \theta_k. \quad (22)$$

The second option is to stagger the weights (*i.e.*, keep all weights $\theta_{k,i}^j$ associated with ϕ_i together in θ_k). Then,

$$\mathbf{g}_k = (\phi_k^\top \otimes \mathbf{I})\theta_k. \quad (23)$$

Mathematically, the representations in (22) and (23) are equivalent. However, depending on how the memory access is handled by the used programming language and hardware, these two representations result in significantly different execution times. For example, the representation in (23) results in a $3 \times$ faster function evaluation in Python.

E. Memory Usage

Thus far, we have mainly analyzed aspects relating to the computational complexity. However, the necessary storage also needs to be considered. Here, we will only focus on the memory S needed for the representation of $\mathbf{g}_k$ and leave out the memory requirements relating to $\mathbf{f}_k$, which are application-specific. Assuming that all numbers are stored with d bits and an equally spaced grid is used so that $n_\theta^1 = \dots = n_\theta^J$, the mean vector $\hat{\theta}$ will require $d \cdot n_\theta^j \cdot J$ bits, and if the symmetry is ignored, the covariance matrix $\mathbf{P}^{\theta\theta}$ will require

$$S = d \cdot (n_\theta^j \cdot J)^2 \quad (24)$$

bits. Unfortunately, if the dynamical systems affected by $\mathbf{g}_k$ are allowed to travel over the entire support of $\mathbf{g}_k$, the matrices $\mathbf{P}^{\theta\theta}$, $\mathbf{P}^{x\theta}$ and $\mathbf{P}^{\theta x}$ will be dense and hence, the storage required is fixed. Thus, the compact basis functions offer no gains in terms of storage complexity.

IV. COMPLETE MODEL AND ESTIMATION

To summarize the suggested complexity reduction, explicit design choices must be made. Assume that the same set of basis functions are used for each dimension of $\mathbf{g}_k$, *i.e.*, $\phi_k^1 = \dots = \phi_k^J$. This implies that the indicator matrices $\mathbf{A}_k^1 = \dots = \mathbf{A}_k^J = \mathbf{A}_k$. Also, assume that the weights are stacked according to (22). The function $\mathbf{g}_k$ is then described by

$$\mathbf{g}_k = \Phi_k \theta_k, \quad (25)$$

where $\Phi_k = (\mathbf{I} \otimes \phi_k^\top \mathbf{A}_k^\top \mathbf{A}_k)$. There are two remaining aspects necessary to consider. First, the dependency of Φ_k on $\mathbf{x}_k$ needs to be specified. Here, consider a general transformation of $\mathbf{x}_k$, *i.e.*,

$$\mathbf{g}_k = \mathbf{g}_k(\psi(\mathbf{x}_k), \mathbf{u}_k), \quad (26)$$

where ψ is some transformation of $\mathbf{x}_k$, for instance, $\psi(\mathbf{x}_k) = \mathbf{D}_k \mathbf{x}_k$ if the dependency is linear. The transformation ψ is assumed known and is specified a priori. The final model is then given by

$$\mathbf{x}_{k+1}^e \triangleq \mathbf{f}_k^e(\mathbf{x}_k^e, \mathbf{u}_k, \mathbf{w}_k^e, \mathbf{g}_k) \quad (27a)$$

$$\mathbf{f}_k^e \triangleq \begin{pmatrix} \mathbf{f}_k(\mathbf{x}_k, \mathbf{u}_k, \mathbf{w}_k, \mathbf{g}_k) \\ \theta_k + \hat{\mathbf{w}}_k \end{pmatrix} \quad (27b)$$

$$\mathbf{g}_k = (\mathbf{I} \otimes \phi_k^\top(\psi(\mathbf{x}_k)) \mathbf{A}_k^\top) \mathbf{A}_k \theta_k \triangleq \Phi_k \theta_k \quad (27c)$$

$$\mathbf{y}_k = \mathbf{h}_k(\mathbf{x}_k, \mathbf{u}_k, \mathbf{e}_k), \quad (27d)$$

where $\mathbf{e}_k$, $\mathbf{w}_k$ and $\hat{\mathbf{w}}_k$ are mutually independent white noise processes with covariance matrices $\mathbf{R}_k$, $\mathbf{Q}_k$, and Σ_k , respectively.

A. Filter Modifications

The EKF described in Section II-B can be used to estimate $\mathbf{x}_k$ and θ_k for every k . However, to fully exploit the new model structure and the computational complexity reductions it enables, the filter equations must be altered. Starting with the time update, note that from (27a), the mean update essentially remains the same apart from $\mathbf{g}_k$ now being defined by (27c). The covariance update (11) is altered according to

$$\begin{aligned} \mathbf{P}_{k+1|k}^{xx} &= \mathbf{F}_x \mathbf{P}_{k|k}^{xx} \mathbf{F}_x^\top + \mathbf{F}_x \mathbf{P}_{k|k}^{x\theta} \mathbf{A}_k^\top \mathbf{F}_\theta^\top \\ &\quad + \mathbf{F}_\theta \mathbf{A}_k \mathbf{P}_{k|k}^{\theta x} \mathbf{F}_x^\top + \mathbf{F}_\theta \mathbf{A}_k \mathbf{P}_{k|k}^{\theta\theta} \mathbf{A}_k^\top \mathbf{F}_\theta^\top + \mathbf{Q}_k \end{aligned} \quad (28a)$$

$$\mathbf{P}_{k+1|k}^{x\theta} = \mathbf{F}_x \mathbf{P}_{k|k}^{x\theta} + \mathbf{F}_\theta \mathbf{A}_k \mathbf{P}_{k|k}^{\theta\theta}. \quad (28b)$$

Hence, the bulk of the computational complexity relating to the product $\mathbf{F}_\theta \mathbf{P}^{\theta\theta} \mathbf{F}_\theta^\top$ has been reduced to $\mathcal{O}(\tilde{n}_\theta^2) \ll \mathcal{O}(n_\theta^2)$. The measurement update alterations are a bit more involved as some approximations are involved. First, we must begin by noting that $\mathbf{K}_k^\theta = \mathbf{P}_{k|k-1}^{x\theta} \mathbf{H}_k^{x\top} \mathbf{S}_k^{-1}$, and as such, the computational complexity can not be reduced while remaining with an exact algorithm as $\mathbf{P}_{k|k-1}^{x\theta}$ is dense after some time (given that the system has been influenced by all basis functions). Hence, approximations are needed to reduce the computational complexity. Here, we draw inspiration from Simultaneous Localization and Mapping, and more precisely [31], where a sparse weight Kalman filter is derived. Particularly, a gain matrix $\mathbf{K}_k$, which updates but a few of the states, is derived and shown to minimize the trace of the resulting covariance matrix. Here, the measurement update is restricted to recently active basis functions, *i.e.*,

$$\mathbf{K} = [\mathbf{K}_k^\top \quad (\mathbf{A}_k^\top \mathbf{A}_k \mathbf{K}_k^\theta)^\top]^\top. \quad (29)$$

Hence, $\mathbf{K}_k$ is sparse and all that remains is to alter the measurement equations. Essentially, all that is required is to replace $\mathbf{K}_k^\theta$ with $\mathbf{A}_k^\top \mathbf{A}_k \mathbf{K}_k^\theta$ which will leave only the elements of $\mathbf{K}_k^\theta$ that are non-zero in $\mathbf{A}_k$. To summarize, the new measurement update of the EKF becomes

$$\hat{\theta}_{k|k} = \hat{\theta}_{k|k-1} + \mathbf{A}_k^\top \mathbf{A}_k \mathbf{K}_k^\theta (\mathbf{y}_k - \mathbf{H}_k^x \hat{\mathbf{x}}_{k|k-1}) \quad (30a)$$

$$\begin{aligned} \mathbf{P}_{k|k}^{x\theta} &= (\mathbf{I} - \mathbf{K}_k^x \mathbf{H}_k^x) \mathbf{P}_{k|k-1}^{xx} - \mathbf{A}_k^\top \mathbf{A}_k \mathbf{K}_k^\theta \mathbf{H}_k^{x\top} \\ &\quad + (\mathbf{I} - \mathbf{K}_k^x \mathbf{H}_k^x) \mathbf{P}_{k|k-1}^{x\theta} + \mathbf{K}_k^x \mathbf{R}_k \mathbf{K}_k^{\theta\top} \mathbf{A}_k^\top \mathbf{A}_k \end{aligned} \quad (30b)$$

$$\begin{aligned} \mathbf{P}_{k|k}^{\theta\theta} &= \mathbf{P}_{k|k-1}^{\theta\theta} + \mathbf{A}_k^\top \mathbf{A}_k \mathbf{K}_k^\theta \mathbf{H}_k^{xx} \mathbf{P}_{k|k-1}^{xx} \mathbf{H}_k^{x\top} \mathbf{K}_k^{\theta\top} \mathbf{A}_k \mathbf{A}_k^\top \\ &\quad - \mathbf{A}_k^\top \mathbf{A}_k \mathbf{K}_k^\theta \mathbf{H}_k^{x\theta} \mathbf{P}_{k|k-1}^{x\theta} - \mathbf{P}_{k|k-1}^{\theta x} \mathbf{H}_k^{x\top} \mathbf{K}_k^{\theta\top} \mathbf{A}_k \mathbf{A}_k^\top \\ &\quad + \mathbf{A}_k^\top \mathbf{A}_k \mathbf{K}_k^\theta \mathbf{R}_k \mathbf{K}_k^{\theta\top} \mathbf{A}_k \mathbf{A}_k^\top. \end{aligned} \quad (30c)$$

Hence, similar to the time-update, the computational complexity relating to the product $\mathbf{K}_k^\theta \mathbf{R}_k \mathbf{K}_k^{\theta\top}$ has been reduced to $\mathcal{O}(\tilde{n}_\theta^2) \ll \mathcal{O}(n_\theta^2)$. Note that the reduction in computational complexity hinges on how $\mathbf{A}_k$ is used. In a practical implementation, the direct use of $\mathbf{A}_k$ *increases* the complexity; therefore, the user must instead devise a way of efficiently indexing into ϕ_k and θ_k and thus indirectly use $\mathbf{A}_k$. However, this has been left out as this depends on the choice of implementation language.

Lastly, to quantify the number of basis functions used in each iteration, assume that the basis function grid is regular

and of dimension P with grid density δ_c in all dimensions. Further, assume that the support of the basis functions is α , and that the current state $\mathbf{x}_k$ is equal to one of the basis function centers. Then, the number of active basis functions with the fast CSRBF selection method is found by noting that the active region is an n -dimensional cube with side length α . The upper bound of active number of basis functions is thus

$$\tilde{n}_\theta \leq \left(\frac{2\alpha}{\delta_c} + 1 \right)^P, \quad (31)$$

which is useful for determining a reasonable tuning of α and δ_c , depending on the available computational resources. Note that if $\mathbf{x}_k$ is not equal to one of the basis function centers, or if $\mathbf{x}_k$ is close to the boundary of the basis function grid, then the inequality is strict.

B. Parameter Selection

The parameter selection in the model is a bit nuanced as it has two distinct effects: (i) it influences the state error, and (ii) it affects what types of functions can be represented as well as the convergence rate of the basis function expansion. The parameters to choose are the prior motion model, the domain of $\mathbf{g}$, the basis function grid density δ_c and support α , the noise covariances $\mathbf{Q}_k$, $\mathbf{R}_k$, and Σ_k , and the priors $\hat{\mathbf{x}}_0$, $\mathbf{P}_0^{xx}$, $\mathbf{P}_0^{x\theta}$, θ_0 , and $\mathbf{P}_0^{\theta\theta}$. The parameters are problem dependent and it is hard to give general advice for the selection. Most of the parameters, excluding the prior motion model and the domain of $\mathbf{g}$, can be found through, *e.g.*, maximum likelihood methods applied to historical data, or they can be manually selected based on expert knowledge. An intuition about the parameters is given here, and we later illustrate the selection process in connection to the numerical examples.

The prior motion model should generally be chosen to, as accurately as possible, represent the system at hand, so that the basis function expansion does not need to represent already known system properties. The basis function support α and grid density δ_c affect the computational complexity of the algorithm, but also what types of functions that can be represented. The support α completely determines the range of the spatial correlations that can be captured by the model. The density δ_c determines how quickly the learned function can vary with the input. The noise covariances $\mathbf{Q}_k$, $\mathbf{R}_k$, and Σ_k affect both the tracking error and the learning rate of the basis function weights. Further, Σ_k also determines how fast the model “forgets” past information. Lastly, the prior $\mathbf{P}_0^{\theta\theta}$ acts as a regularizer which keeps the basis function weights close to the prior and is key to keeping the learning rate slow.

Some of the parameters can also possibly be learned online, for instance through further random walk assumptions. Particularly interesting, as alluded to in Section III-A, is the BF support α and the grid density δ_c . These two parameters completely determine the computational complexity of the method. Hence, if these are to be learned online, it will be important to impose restrictions on the values they can take, to remain with a real-time applicable algorithm. Another possibility is to adapt the noise covariances $\mathbf{Q}_k$ and Σ_k online,

as in, *e.g.*, [32], [33], where an expectation maximization approach was taken to estimate the noise covariances over time.

V. NUMERICAL EXAMPLES

To illustrate the application of the proposed joint state inference and model learning approach, we will present three simulation examples. The first example highlights a key property of the presented method, namely the effect of including a prior motion model. The second example estimates the longitudinal tire-friction of a car and is designed to quantitatively compare the estimation performance of a standard EKF using the SSM+RBF model and the sparsely weighted EKF with the SSM+CSRBF model. The last example estimates the accelerations of vehicles throughout a three-way intersection and is designed to illustrate the reduction in computational complexity when using the fast CSRBF description compared to the standard RBF description of $\mathbf{g}_k$ in a larger setting. In the two latter examples, the method is compared to an EKF with the SSM+RBF model which in all essence is identical to the SSM+GP model used in [5].

A. Influence of Prior Motion Model

One particular property of the model (7) is that it inherits an initial estimation performance, in a *root mean squared error* (RMSE) sense, from the chosen prior motion model. Further, the model can improve upon that performance in cases where the prior motion model does not accurately describe the system. These properties are highlighted in the following example.

1) *Modeling*: Three different models, $\mathbf{f}_k^{(a)}$, $\mathbf{f}_k^{(b)}$, and $\mathbf{f}_k^{(c)}$ are considered, which are of the following form

$$\mathbf{x}_{k+1}^{(a)} = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \mathbf{x}_k^{(a)} + \begin{bmatrix} 1/2 \\ 1 \end{bmatrix} \mathbf{w}_k^{(a)} \quad (32a)$$

$$\mathbf{x}_{k+1}^{(b)} = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \mathbf{x}_k^{(b)} + \begin{bmatrix} 1/2 \\ 1 \end{bmatrix} (\mathbf{g}_k^{(b)}(\mathbf{x}_k^{(b)}) + \mathbf{w}_k^{(b)}) \quad (32b)$$

$$\mathbf{x}_{k+1}^{(c)} = \mathbf{g}_k^{(c)}(\mathbf{x}_k^{(c)}) + \mathbf{w}_k^{(c)}. \quad (32c)$$

Also, assume that the first state component (the position) is measurable, *i.e.*, the sensor model is

$$y = [x_k^{(i)}]_1 + e_k^{(i)}, \quad i \in [a, b, c]. \quad (33)$$

Thus, the models (32a)–(32c) are a CV model, a CV model with basis function expansion augmentation, and a basis function expansion description, respectively.

2) *Parameter Selection*: The three models (32a)–(32c) are tuned identically. The domain of $\mathbf{g}$ is chosen to cover the state-space of the simulations. The grid density $\delta_c = 1$ m and the basis function support is chosen as $\alpha = 10$ to have at least a few basis functions active in each time step. The measurement and process noise have variances $\mathbf{Q} = 0.01$ and $\mathbf{R} = 0.01$, respectively. The priors $\mathbf{P}_0^{\theta\theta} = 0.1\mathbf{I}$ and $\mathbf{P}_0^{xx} = \mathbf{I}$.

TABLE I
MEAN RMSE FOR SECTION V-A

Model	Mean RMSE	
	Scenario 1)	Scenario 2)
(32a)	0.09	0.18
(32b)	0.09	0.09
(32c)	4.14	12.69

3) *Results:* To illustrate the benefits of including the prior motion model in (32b), two different scenarios are considered. The objective is to estimate the position of the system with minimal error and the performance is compared through the positional RMSE. A Kalman filter is applied to (32a) and an EKF to (32b) and (32c), respectively. The two considered scenarios are:

- 1) The “true” system is assumed to be given by (32a). It is simulated 100 time steps with a process noise with variance $\mathbf{Q} = 0.01$, and the measurement error has variance $\mathbf{R} = 0.01$.
- 2) The “true” system is assumed to be given by a model similar to (32b), where the “true” $\mathbf{g}_k^{(2)}$ is assumed to be given by $\frac{1}{2} \sin(\frac{1}{25} \pi [\mathbf{x}_k]_0) + 0.01$. This choice is to keep the states somewhat close to the origin so that the number of basis functions necessary to describe the system using (32c) is kept low.

The simulations were run 50 times and the mean positional RMSE is reported in Table I. In the first scenario, the models (32a) and (32b) have essentially identical estimation performance whereas (32c) has a higher RMSE. This indicates that the prior motion model included in (32b) provides an initial estimation performance, as opposed to (32c). For the second scenario, model (32c) performs worst. The model (32b) performs better than both (32a) and (32c), which again indicates an initial estimation performance, but also that the model improves over time, as opposed to (32a). Clearly, the prior motion model provides a minimal estimation performance, in an RMSE sense, while the basis function expansion augmentation provides a means of improving over time. Note that this does not mean that including a prior motion model will always perform better than just using a basis function expansion. Depending on how well the assumptions of the model on the function $\mathbf{g}_k$ fit the actual scenario, one model or the other may perform better.

B. Longitudinal Tire-Friction Estimation

The second example deals with longitudinal tire-friction estimation in a two-wheel drive vehicle. The scenario is chosen to illustrate that the approximate algorithm presented here is competitive, in a function estimation RMSE sense, compared to using an exact algorithm, see [5].

1) *Modeling:* The motion model is chosen as the commonly used single-track model, see e.g. [34] for details. A detailed derivation of the motion model can be found in the supplementary material and it is summarized here. The following assumptions are made

TABLE II
LONGITUDINAL FRICTION MODEL/SIMULATION PARAMETERS

Parameter	Value	Description
δ_c	0.025	Grid spacing
α	0.15	CSRBF support
l	0.01	RBF length scale
$\mathbf{R}$	diag[0.1, 0.01]	Measurement noise covariance
$\hat{x}_0$	0	Prior state estimate
$\mathbf{P}_0^{xx}$	10^{-6}	Prior state error covariance
$\mathbf{P}_0^{\theta\theta}$	$10^{-5} \mathbf{I}$	Prior weight error covariance
q	1	Process noise variance
Σ	$10^{-8} \mathbf{I}$	Weight noise covariance
T_s	40 ms	Sampling interval
l_r	1.6 m	Length from CG to rear axle
l_f	1.4 m	Length from CG to front axle
m	1000 kg	Vehicle mass
g_0	9.81 m/s ²	Nominal gravity
B	11.7	Pacejka stiffness parameter
C	1.69	Pacejka shape parameter
D	1.2	Pacejka peak parameter
E	0.377	Pacejka curvature parameter

- The steering angle is small, such that the lateral and angular velocities are negligible.
- The pitch dynamics is negligible, so that the vertical load on the wheels is constant.
- The angular velocity of the wheels is assumed measurable.
- The vehicle is assumed to be front-wheel drive, so the longitudinal velocity is measurable through the free-rolling rear wheels.
- The vehicle is equipped with an accelerometer capable of measuring the longitudinal acceleration, a commonly occurring automotive sensor [12].

The tire friction is modeled as a basis function expansion and the complete, discretized, model is thus given by

$$x_{k+1} = x_k + T_s \underbrace{\frac{g_0 l_f}{l_r + l_f}}_G g_{f,k} + w_k \quad (34a)$$

$$g_{f,k} = \phi^\top(s_k(x_k)) \mathbf{A}_k^\top \mathbf{A}_k \theta_k \quad (34b)$$

$$s_k(x_k) = \frac{r_w \omega_{f,k} - x_k}{x_k} \quad (34c)$$

$$\theta_{k+1} = \theta_k + \dot{\mathbf{w}}_k \quad (34d)$$

$$\mathbf{y}_k = \begin{bmatrix} G g_{f,k} \\ x_k \end{bmatrix} + \mathbf{e}_k. \quad (34e)$$

Here, x_k is the longitudinal velocity of the vehicle, g_0 is the nominal gravity, and l_r and l_f are the distances from the *center of gravity* (CG) to the rear and front axle, respectively. Further, r_w and $\omega_{f,k}$ are the radius and angular velocity of the front wheels, respectively. The wheel slip is denoted s_k and T_s is the sampling interval, which is here $T_s = 40$ ms. Furthermore, w_k , $\dot{\mathbf{w}}_k$ and $\mathbf{e}_k$ are mutually independent white noise processes with covariances q , Σ and $\mathbf{R}$, respectively. Lastly, the static vehicle parameters are given in Table II.

Fig. 4: Longitudinal friction force estimate with the proposed method and the SSM+RBF method described in [5]. The frequency of the true slip is plotted as a black histogram, displaying in what regions the methods “should” have data available. The true slip is not known but estimated online. The confidence intervals is a 3σ interval constructed from the 50 MC realizations.

2) *Parameter Selection:* As the motion model has been fixed, the basis function grid can be chosen. The interest here lies in estimating the main characteristics of the tire-friction curve and the domain of $\mathbf{g}_k$ is thus chosen as $s_k \in [-0.5, 0.5]$. The noise covariance $\mathbf{R}_k = \mathbf{R}$ is chosen according to the simulated noise levels in the sensors. The prior $\hat{x}_0 = 0$ and $\mathbf{P}_0^{xx} = 10^{-6}$ as it is known that the vehicle starts from a standstill. Since the unknown function $g_{f,k}$ influences both the motion model and the sensor model, the prior $\mathbf{P}_0^{\theta\theta}$ is chosen to be small, to keep the model from adapting too quickly, which can otherwise lead to divergence. To not place too strict assumptions on the smoothness of the tire friction curve, the grid density $\delta_c = 0.025$ and basis function support $\alpha = 0.15$, which keeps the measurements from influencing function estimates far from the current slip value. The RBF length scale is chosen as $l = 0.01$, as this proved to yield the best performance in that case. The covariances q and Σ are finally tuned to achieve a low RMSE w.r.t. to the function estimate. The two methods are tuned independently for the lowest possible RMSE and hence, the parameters differ slightly. The complete set of parameters for the SSM+CSRBF model is presented in Table II.

3) *Results:* The scenario was simulated in MATLAB using the Simscape Driveline testbed with the true tire friction given by the Pacejka tire model [35]

$$\mu = D \sin(C \arctan(Bs - E(Bs - \arctan(Bs)))), \quad (35)$$

where the parameters B, C, D , and E are chosen to correspond to an asphalt surface, see Table II. The scenario consists of five accelerations from standstill to a longitudinal velocity of $v \approx 20$ m/s to excite the system into the nonlinear region of the slip curve, since it is approximately linear for small slip values. 100 accelerations from standstill were simulated and Gaussian noise was added to the longitudinal acceleration and wheel angular velocity measurements, with standard deviation 0.1 and 0.01, respectively. Five of these accelerations are then sampled during evaluation to learn the function $g_{f,k}$ while jointly estimating the vehicle longitudinal velocity. This was done 50 times, and the resulting function estimates are presented in Fig. 4, where the confidence interval is a 3σ

TABLE III
RMSE OF THE ESTIMATED FUNCTION

Model	RMSE (std)
SSM+CSRBF	0.305 (± 0.0368)
SSM+RBF	0.278 (± 0.0046)

Fig. 5: Three-way intersection simulation scenario. Cars approach on the middle road and turn left or right with equal probability. The two paths are visualized as well as the centers of the basis functions.

interval constructed from the 50 realizations. The frequency of true slip is also visualized as histogram in the bottom of the figure. Note that the true slip is not known, but estimated online, and as such, the histogram is not a one-to-one mapping to the actual input to the basis function expansions. Neither of the methods recover the true friction force curve exactly, but both of them capture the general characteristics of the friction curve. The error was computed at the parts of the curve where there should be measurements, i.e., at the locations of the true slip. The mean RMSE for the two respective methods is presented in Table III.

C. Three-Way Intersection Simulation

The third example is similar to the one explored in [5]. The scenario is a simulation of the three-way intersection shown in Fig. 5. Cars approach from the bottom road, starting at $p_y = 0$ (outside the image), and turn left or right with equal probability. In [5], we used a SSM+GP model parameterized by inducing points to model the dynamics of the cars. The inducing points were placed only over the actual road area to reduce computational burden and make the model viable in that specific scenario. Hence, there was a requirement on exact prior knowledge of the environment. As pointed out in Section III-B, this requirement has been alleviated because the method described herein is *nearly* invariant to the total number of basis functions. Hence, we use a simpler regular grid, and the only prior information needed is the total extent of the intersection.

1) *Modeling*: The prior motion model is chosen as a CV motion model, *i.e.*, the complete motion model is

$$\mathbf{x}_{k+1} = \mathbf{F}_k \mathbf{x}_k + \mathbf{G}_k \mathbf{g}_k + \mathbf{G}_k \mathbf{w}_k \quad (36a)$$

$$\mathbf{F}_k = \begin{bmatrix} 1 & T_s \\ 0 & 1 \end{bmatrix} \otimes \mathbf{I} \quad \mathbf{G}_k = \begin{bmatrix} T_s^2/2 \\ T_s \end{bmatrix} \otimes \mathbf{I}, \quad (36b)$$

with $\mathbf{g}_k$ given by (27c). Moreover, T_s is the sampling interval for the system(s). Information about the input $\mathbf{u}_k$ is assumed unknown; hence, $\mathbf{u}_k = 0$, $\forall k$. Also, note that $\mathbf{g}_k$ is pre-multiplied by $\mathbf{G}_k$, which will render $\mathbf{g}_k$ to be interpretable as the vehicle acceleration throughout the intersection. Further, the necessary derivatives are given by

$$\mathbf{F}_x = \mathbf{F}_k + \frac{\mathbf{G}_k \Phi_k \hat{\boldsymbol{\theta}}_{k|k}}{\partial \mathbf{x}_{k|k}} \bigg|_{\mathbf{x}_{k|k} = \hat{\mathbf{x}}_{k|k}} \quad (37a)$$

$$\mathbf{F}_\theta = \mathbf{G}_k \Phi_k(\hat{\mathbf{x}}_{k|k}), \quad (37b)$$

where $\Phi_k(\hat{\mathbf{x}}_{k|k}) = (\mathbf{I} \otimes \phi_k^\top(\psi(\hat{\mathbf{x}}_{k|k}))) \mathbf{A}_k^\top \mathbf{A}_k$.

For the observation model, it is assumed that the position of the vehicle is measurable, albeit with some noise, *i.e.*,

$$\mathbf{h}_k = \mathbf{H}_k \mathbf{x}_k + \mathbf{e}_k \quad (38a)$$

$$\mathbf{H}_k = [\mathbf{I} \quad \mathbf{0}] \quad \mathbf{R}_k = \sigma_e^2 \mathbf{I}, \quad (38b)$$

where σ_e^2 is the measurement noise variance. Note that a linear observation model implies $\nabla_{\mathbf{x}_{k|k}} \mathbf{h}_k = \mathbf{H}_k$.

2) *Parameter selection*: The basis function grid spacing is selected $\delta_c = 1$ m primarily to allow the evaluations to run in reasonable time since the Gaussian RBFs otherwise cause a huge computational burden. Further, it could also be argued that the acceleration of a vehicle does not change rapidly enough to motivate a grid spacing less than $\delta_c = 1$ m, thus justifying the choice. The support of the CSRBFs and length scale of the Gaussian RBFs is chosen as $\alpha = 5$ and $l = 1$, respectively. The length scale of the Gaussian RBFs is justified by the same argument as the grid spacing. The support of the CSRBFs is chosen as $\alpha = 5$ to include at least a few basis functions in each time step. The measurement noise variance is set to $R = 0.2$, which can be achieved through, *e.g.*, differential GPS [36]. The prior $\hat{\mathbf{x}}_0$ is set to the true initial position and velocity and $\mathbf{P}_0^{xx} = 0.1 \cdot \mathbf{I}$. In a real scenario, these can be found through the first two measurements and a

simple numerical difference. The prior $\mathbf{P}_0^{\theta\theta} = 0.01 \mathbf{I}$, which is significantly higher than the tire friction example, but is made possible by the inclusion of a prior motion model. The process noise covariance $\boldsymbol{\Sigma}_k = \mathbf{0}$, as the vehicles travel at approximately the same velocity, and the acceleration is thus not expected to change. The process noise covariance $\mathbf{Q}_k = \mathbf{Q}$ is then tuned to achieve a low RMSE w.r.t. the position and velocity estimates and was finally chosen as $\mathbf{Q} = 0.1 \cdot \mathbf{I}$. Furthermore, let the modeled acceleration depend only on the position of the vehicle, *i.e.*,

$$\psi(\mathbf{x}_k) = [\mathbf{I} \quad \mathbf{0}] \mathbf{x}_k.$$

One might argue that normal acceleration also depends on the absolute velocity, but in this study, it is assumed that the vehicles travel at approximately the same speed throughout the intersection. Lastly, the sampling interval is set to $T_s = 0.2$. All of the simulation parameters are summarized in Table IV(a).

3) *Results*: The intersection simulation was constructed and run in the SUMO microscopic traffic simulator [37]. A wide variety of vehicles were simulated and then sampled during estimation/learning, see Table IV(b) for a summary of the vehicle parameters. 150 vehicles were used for each of the 10 simulations. Further, as a tracking comparison, a common CV model is used with identical parameter tuning to both the SSM+CSRBF model as well as the SSM+RBF model.

The execution time of single time and measurement updates are shown in Table V. It is clear that the proposed SSM+CSRBF improves over the SSM+RBF. On average, the time update takes approximately 0.006 s, and the measurement update roughly 0.034 s. Hence, the SSM+CSRBF filter should be able to run at a rate of approximately 25 Hz, with this particular tuning of the parameters. Further, the RMSE over number of vehicles that have passed through the intersection is visualized in Fig. 6. The RMSE with respect to position and velocity is visualized separately for each path. As the number of vehicles increases, the RMSE both with respect to position and velocity approaches that of the SSM+RBF model, for both of the paths. This indicates that the proposed method does not lose any significant estimation performance (asymptotically), in the RMSE sense, at least in this particular example. Interestingly, the SSM+RBF method has a clear RMSE advantage already from the start. This is due to the use of an approximate Kalman gain in the SSM+CSRBF method, which is one of the keys to the computational complexity reduction. Thus, there is a trade-off between computational benefits and convergence rate in an RMSE sense.

VI. CONCLUSION AND FUTURE WORK

An extended Kalman filter method for online joint state inference and model learning has been presented and its application illustrated via two numerical examples. The underlying system dynamics are modeled using an SSM composed of a prior, known, model and an RBF expansion capturing unknown model properties. By using RBFs with compact support and an approximate Kalman gain in the filter it is shown that the method has a complexity of $\mathcal{O}(\tilde{n}_\theta^2)$, where $\tilde{n}_\theta$ depends

Fig. 6: RMSE with an increasing number of vehicles, *i.e.*, observations. The left column is the RMSE w.r.t. the position and the right column w.r.t. the velocity. The first row is w.r.t. the right path and the second row w.r.t. the left path. The SSM+CSRBF method clearly approaches the SSM+RBF method over time.

on the support of the compact basis functions as well as the basis function grid density. This should be compared to the SSM+RBF method which has a complexity of $\mathcal{O}(n_\theta^2)$, where n_θ is the number of basis functions. Note that $\tilde{n}_\theta \leq n_\theta$ and hence, the method approaches the SSM+RBF formulation if the support of the basis functions is large. The drawback with the proposed method is that correlation between distant points in the state-space domain can not be represented. Still, in many applications the expected correlations are short-range and hence, the proposed method provides a computationally efficient method for joint state inference and model learning, and enables large-scale real-time applications.

Future directions concern a mathematical analysis of the model's statistical properties, *i.e.*, under what conditions convergence can be guaranteed. Further, adaptive grid densities are of interest as this makes it possible to concentrate basis functions where they are needed. Lastly, methods for estimating multi-modal functions need to be addressed for this to be useful in a wider range of applications.

REFERENCES

- [1] S. Särkkä, *Bayesian Filtering and Smoothing*. Cambridge University Press, 2010.
- [2] X. R. Li and V. P. Jilkov, "Survey of Maneuvering Target Tracking. Part I: Dynamic Models," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 39, no. 4, pp. 1333–1364, 2003.
- [3] L. Ljung, *System Identification: Theory for the User*, 2nd ed. Pearson Education, 1998.
- [4] C. Veibäck, J. Olofsson, T. R. Lauknes, and G. Hendeby, "Learning Target Dynamics While Tracking Using Gaussian Processes," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 9251, pp. 1–10, 2019.
- [5] A. Kullberg, I. Skog, and G. Hendeby, "Learning Driver Behaviors Using a Gaussian Process Augmented State-Space Model," in *FUSION 2020*, Virtual, Jul. 2020.
- [6] A. Svensson, T. B. Schön, A. Solin, and S. Särkkä, "Nonlinear State Space Model Identification Using a Regularized Basis Function Expansion," in *CAMSAP 2015*, Cancun, Mexico, Dec. 2015, pp. 481–484.

TABLE IV
THREE-WAY INTERSECTION SIMULATION PARAMETERS

(a) MODEL PARAMETERS

Parameter	Value	Description
δ_c	1 m	Grid density
l	1	RBF length scale
α	5	CSRBF support
R	0.2 m/s ²	Measurement noise variance
$\hat{x}_0$	Ground truth	Prior state estimate
$\mathbf{P}_0^{xx}$	0.1 · $\mathbf{I}$	Prior state error covariance
$\mathbf{P}_0^{\theta\theta}$	0.01 · $\mathbf{I}$	Prior weight error covariance
$\mathbf{Q}$	0.1 $\mathbf{I}$	Process noise covariance
Σ	$\mathbf{0}$	Weight noise covariance
T_s	0.2 s	Sampling interval

(b) SUMO PARAMETERS

Parameter	Value
Vehicle length	$U(4.5, 5)$
Vehicle max speed	10 m/s
Vehicle speed deviation	5%
Vehicle follow model	ACC

TABLE V
EXECUTION TIMES FOR EACH SEPARATE FILTERING STEP

Model	Average Execution Time [s] (std)	
	Time update	Measurement update
SSM+CSRBF	0.006 (±0.0027)	0.034 (±0.009)
SSM+RBF	0.045 (±0.0273)	0.784 (±0.065)

- [7] R. Turner, M. P. Deisenroth, and C. E. Rasmussen, "State-space Inference and Learning with Gaussian Processes," *J. Mach. Learn. Res.*,

- vol. 9, pp. 868–875, 2010.
- [8] R. Frigola, F. Lindsten, T. B. Schön, and C. E. Rasmussen, “Bayesian Inference and Learning in Gaussian Process State-Space Models with Particle MCMC,” in *Adv. Neural Inf. Process. Syst.*, Lake Tahoe, Dec. 2013, pp. 3156–3164.
 - [9] J. Ko and D. Fox, “GP-BayesFilters: Bayesian Filtering Using Gaussian Process Prediction and Observation Models,” *Auton. Robots*, vol. 27, no. 1, pp. 75–90, 2009.
 - [10] Y. Liu and P. M. Djuric, “Gaussian Process State-Space Models with Time-Varying Parameters and Inducing Points,” in *Eur. Signal Process. Conf.*, vol. 2021-Janua, Amsterdam, Jan. 2021, pp. 1462–1466.
 - [11] A. Svensson and T. B. Schön, “A Flexible State-Space Model for Learning Nonlinear Dynamical Systems,” *Automatica*, vol. 80, pp. 189–199, 2017.
 - [12] K. Berntorp, “Bayesian Tire-Friction Learning by Gaussian-Process State-Space Models,” in *18th Eur. Control Conf.* Napoli: EUCA, Jun. 2019, pp. 231–236.
 - [13] —, “Online Bayesian Inference and Learning of Gaussian-Process State-Space Models,” *Automatica*, vol. 129, 2021.
 - [14] A. Lederer, A. J. O. Conejo, K. Maier, W. Xiao, and S. Hirche, “Real-Time Regression with Dividing Local Gaussian Processes,” Jun. 2020. [Online]. Available: <http://arxiv.org/abs/2006.09446>
 - [15] M. Kok and A. Solin, “Scalable Magnetic Field SLAM in 3D Using Gaussian Process Maps,” in *FUSION 2018*, Cambridge, UK, Jul. 2018, pp. 1353–1360.
 - [16] J. Quiñero-Candela and C. E. Rasmussen, “A Unifying View of Sparse Approximate Gaussian Process Regression,” *J. Mach. Learn. Res.*, vol. 6, pp. 1939–1959, 2005.
 - [17] E. Snelson and Z. Ghahramani, “Sparse Gaussian Processes using Pseudo-inputs,” *Adv. Neural Inf. Process. Syst.* 18, pp. 1257–1264, 2006.
 - [18] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. MIT Press, 2006.
 - [19] F. Daum, “Nonlinear Filters: Beyond the Kalman Filter,” *IEEE Aerosp. Electron. Syst. Mag.*, vol. 20, no. 8 II, pp. 57–68, 2005.
 - [20] S. J. Julier and J. K. Uhlmann, “A New Extension of the Kalman Filter to Nonlinear Systems,” in *Signal Process. Sens. Fusion, Target Recognit.* VI, vol. 3068, Orlando, Apr. 1997, pp. 182–193.
 - [21] J. Steinbring and U. D. Hanebeck, “LRKF revisited: The Smart Sampling Kalman Filter (S2KF),” *J. Adv. Inf. Fusion*, vol. 9, no. 2, pp. 106–123, 2014.
 - [22] M. D. Buhmann, *Radial Basis Functions: Theory and Implementations*. Cambridge University Press, 2003.
 - [23] J. Park and I. W. Sandberg, “Universal Approximation Using Radial-Basis-Function Networks,” *Neural Comput.*, vol. 3, no. 2, pp. 246–257, 1991.
 - [24] E. Porcu, D. J. Daley, M. Buhmann, and M. Bevilacqua, “Radial Basis Functions with Compact Support for Multivariate Geostatistics,” *Stoch. Environ. Res. Risk Assess.*, vol. 27, no. 4, pp. 909–922, 2013.
 - [25] R. Furrer, M. G. Genton, and D. Nychka, “Covariance Tapering for Interpolation of Large Spatial Datasets,” *J. Comput. Graph. Stat.*, vol. 15, no. 3, pp. 502–523, 2006.
 - [26] P. L. Houtekamer and H. L. Mitchell, “A Sequential Ensemble Kalman Filter for Atmospheric Data Assimilation,” *Mon. Weather Rev.*, vol. 129, no. 1, pp. 123–137, 2001.
 - [27] T. M. Hamill, J. S. Whitaker, and C. Snyder, “Distance-dependent Filtering of Background Error Covariance Estimates in an Ensemble Kalman Filter,” *Mon. Weather Rev.*, vol. 129, no. 11, pp. 2776–2790, 2001.
 - [28] H. Wendland, “On the Smoothness of Positive Definite and Radial Functions,” *J. Comput. Appl. Math.*, vol. 101, no. 1-2, pp. 177–188, 1999.
 - [29] D. Nychka, S. Bandyopadhyay, D. Hammerling, F. Lindgren, and S. Sain, “A Multiresolution Gaussian Process Model for the Analysis of Large Spatial Datasets,” *J. Comput. Graph. Stat.*, vol. 24, no. 2, pp. 579–599, 2015.
 - [30] M. Katzfuss, “A Multi-Resolution Approximation for Massive Spatial Datasets,” *J. Am. Stat. Assoc.*, vol. 112, no. 517, pp. 201–214, 2017.
 - [31] S. J. Julier, “A Sparse Weight Kalman Filter Approach to Simultaneous Localisation and Map Building,” in *Int. Conf. Intell. Robot. Syst.*, Maui, Hawaii, Oct. 2001, pp. 1251–1256.
 - [32] V. A. Bavdekar, A. P. Deshpande, and S. C. Patwardhan, “Identification of Process and Measurement Noise Covariance for State and Parameter Estimation Using Extended Kalman Filter,” *J. Process Control*, vol. 21, no. 4, pp. 585–601, 2011.
 - [33] G. C. Goodwin and J. C. Agüero, “Approximate EM Algorithms for Parameter and State Estimation in Nonlinear Stochastic Models,” in *Proc. 44th IEEE Conf. Decis. Control. Eur. Control Conf. CDC-ECC ’05*. Seville: IEEE, Dec. 2005, pp. 368–373.
 - [34] R. Rajamani, *Vehicle Dynamics and Control*. Springer US, 2011.
 - [35] H. B. Pacejka, *Tire and Vehicle Dynamics*, 2nd ed. Elsevier Science, 2005.
 - [36] I. Skog and P. Händel, “In-car Positioning and Navigation Technologies Survey,” *IEEE Trans. Intell. Transp. Syst.*, vol. 10, no. 1, pp. 4–21, 2009.
 - [37] P. A. Lopez, M. Behrisch, L. Bieker-Walz, J. Erdmann, Y.-P. Flötteröd, R. Hilbrich, L. Lücken, J. Rummel, P. Wagner, and E. Wießner, “Microscopic Traffic Simulation using SUMO,” in *IEEE Intell. Transp. Syst. Conf.*, Maui, Nov. 2018.