Numerical Gaussian process Kalman filtering for spatiotemporal systems

Armin Küper and Steffen Waldherr

arXiv:2105.02079v1 [eess.SY] 5 May 2021

Abstract— **We present a novel Kalman filter for spatiotemporal systems called the numerical Gaussian process Kalman filter (GPKF). Numerical Gaussian processes have recently been introduced as a physics informed machine learning method for simulating time-dependent partial differential equations without the need for spatial discretization. We bring numerical GPs into probabilistic state space form. This model is linear and its states are Gaussian distributed. These properties enable us to embed the numerical GP state space model into the recursive Kalman filter algorithm. We showcase the method using two case studies.**

Index Terms— **Kalman filtering, Gaussian processes, Spatiotemporal systems, Machine learning**

I. INTRODUCTION

State estimators are algorithms for reconstructing a system's state from a stream of noisy online measurements and model predictions. The model is usually in the form of a probabilistic state space model. Given linear finite-dimensional dynamics, the process and measurement equation read as

$$x_t = A x_{t-1} + q_{t-1}, \tag{1}$$
$$y_t = C x_t + r_t, \tag{2}$$

with the state at time t as $x_t \in \mathbb{R}^{d_x}$, and the measurements as $y_t \in \mathbb{R}^{d_y}$. Dynamics are described by matrix A while measurements are obtained through C. Process q_{t-1} and measurement noise r_t are modeled as Gaussian distributed zero-mean white noise with covariance matrix Q and R, respectively. Process and measurement equations $(1)-(2)$ can also be represented by probability density functions $p(\cdot)$ that take the form of Gaussian distributions $\mathrm{N}(\cdot)$ due to the nature of the noise processes. In particular we have

$$p(x_t|x_{t-1}) = \mathrm{N}\left(x_t|A x_{t-1}, Q\right), \tag{3}$$
$$p(y_t|x_t) = \mathrm{N}\left(y_t|C x_t, R\right). \tag{4}$$

The following question arises: How should the measurement y_t be used to correct the model prediction of x_t in real time?

One answer to this question is the Kalman filter. In fact, it gives the optimal estimate m_t^{opt} of the current state x_t under consideration of the measurement history up to the present $y_{1:t} = \{y_1, \dots, y_t\}$ [1]. It is optimal in the sense that the expected squared error between estimated and true state conditioned on the measurement history is minimized.

Armin Küper and Steffen Waldherr are with the CREaS group of the Chemical Engineering Department, KU Leuven, 3001 Leuven, Belgium (e-mail:{armin.kuper, steffen.waldherr}@kuleuven.be).

The Kalman filter estimates the state in a recursive manner. In Bayesian statistical terms this can be written as

$$p(x_t|y_{1:t}) = \frac{p(y_t|x_t)p(x_t|y_{1:t-1})}{p(y_t|y_{1:t-1})}. \tag{5}$$

Since the recursion starts with a Gaussian distribution and the dynamics are linear, all the above distributions remain Gaussian, and the posterior estimate can be obtained in closed form by calculating its mean and covariance.

Remember that the above model equations described finite-dimensional systems. Infinite-dimensional filtering theory certainly exists [2], but direct practical infinite-dimensional Kalman filtering does not to the best knowledge of the authors. Finite-dimensional approximations are therefore needed. Using spatial discretization methods, e.g. the finite volume method [3], one usually obtains a large system of ordinary differential equations which can be formulated as the above state space model $(1)-(2)$.

In this paper, we will take a different route. We will use so-called numerical Gaussian processes [4] to solve spatiotemporal models. Gaussian processes (GPs) are a probabilistic machine learning method that can be used for regression [6]. Numerical GPs are structured by the spatiotemporal model in the form of a (linear) time-dependent partial differential equation

$$\frac{\partial}{\partial t} f(t, x) = \mathcal{L}_x f(t, x), \quad x \in \Omega \subset \mathbb{R}^{d_x} \tag{6}$$
$$f(t, x_b) = \mathcal{B}_x f(t, x), \quad x_b \in \partial\Omega, \tag{7}$$

where $\mathcal{L}_x : L_2(\mathbb{R}^{d_x}, \mathbb{R}) \to L_2(\mathbb{R}^{d_x}, \mathbb{R})$, $f(t, x) \mapsto \mathcal{L}_x f(t, x)$ is a linear (integro-)differential operator determining the dynamics, and $\mathcal{B}_x$ is a linear functional imposing the boundary condition.

Due to the probabilistic nature of GPs, they provide us with uncertainty quantification. This will be of much benefit for their use in state estimation.

The contributions of this paper are as follows:

- we extend numerical GPs with an output channel for online measurements and formulate this as a probabilistic linear state space model which is by definition Gaussian distributed, see Section III;
- this allows us to put numerical GPs through the Kalman filter algorithm and therefore derive the numerical Gaussian process Kalman filter (GPKF), see Section IV;
- using numerical case studies of spatiotemporal systems, we showcase the numerical GPKF, see Section V.
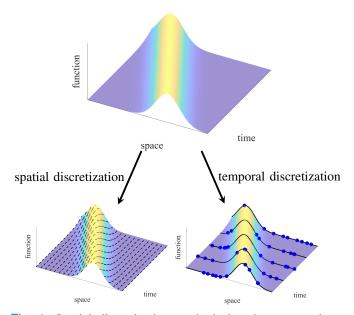
Fig. 1: Spatial discretization methods lay down a mesh on the spatial domain and describe how neighboring mesh cells should interact with each other dictated by the PDE. This results in a system of ODEs that is simulated with traditional solvers. Numerical Gaussian processes [4] on the other hand start with a temporal discretization. The resulting time discretized PDE is used to build up a structured Gaussian processes that is evaluated on arbitrarily chosen regression points. To simulate, one recursively formulates the so-called posterior Gaussian process. Figure adapted from [5].

Preliminary results were presented in the conference paper [7]. This article significantly extends this previous work. In contrast to [7], the preliminaries now contain a didactical example to streamline the section. Furthermore, numerical GPs are explicitly expressed as probabilistic state space models in this paper. Moreover, boundary conditions are treated as part of the model and not as measurements. This affects the KF algorithm as such that boundary conditions enter through the prediction step, rather than the update step of the Kalman filter. Finally, more sophisticated case studies are presented.

Related works: Infinite-dimensional Kalman filtering with Gaussian processes has previously been done by [8], [9]. Therein, certain covariance functions are shown to be convertible via Fourier-transforms to infinite-dimensional state space models. Using these models, infinite-dimensional Kalman filtering can then be done. This approach does not need a structured PDE to inform the GPs, contrary to numerical GPs. Of course, this also means that the GP has to learn the dynamics from the full state vector which might not be accessible in the first place.

Another fully non-parametric approach is given in [10]. Assuming space-time separability of the covariance function, a finite-dimensional discrete-time state space model is built up that can be used for Kalman filtering. Again, the full states need to be accessible for this method to work.

Finite-dimensional filtering with Gaussian processes has been done by [11]. Different Bayesian filters such as particle filters, extended and unscented Kalman filters are recovered. Closely related is the work of [12], wherein an analytic moment-based Gaussian process filter is presented. A general perspective on finite-dimensional Gaussian filtering is given in [13].

Mathematical notation: Scalars are lowercase non-bold symbols x, while vectors are bold $\boldsymbol{x}$, and matrices are uppercase and bold $\boldsymbol{X}$.

A random variable x stemming from probability density function $p(\cdot)$ is symbolically written as $x \sim p(x)$. Gaussian random variables and their distributions are denoted as $x \sim \mathrm{N}(x|m, P)$ with mean m and variance P.

A covariance matrix $\boldsymbol{K}(\boldsymbol{X}, \boldsymbol{X}) \in \mathbb{R}^{M \times M}$ in the Gaussian process regression context is built up by evaluating the underlying covariance function element-wise $\boldsymbol{K}_{ij} = k(\boldsymbol{x}_i, \boldsymbol{x}_j)$ using the data matrix $\boldsymbol{X} = \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_M\} \in \mathbb{R}^{d_x \times M}$. A cross-covariance matrix $\boldsymbol{K}(\boldsymbol{X}, \boldsymbol{X}_*) \in \mathbb{R}^{M \times S}$ can also be built up with $\boldsymbol{X}_* \in \mathbb{R}^{d_x \times S}$.

II. Preliminaries

In this section we will introduce Gaussian process regression, linear operators in combination with GP regression, and numerical GPs for time-dependent partial differential equations.

Gaussian process regression

A Gaussian process is a stochastic process $f(\boldsymbol{x})$ that is fully defined by its mean function $m(\boldsymbol{x})$ and covariance function $k(\boldsymbol{x}, \boldsymbol{x}')$

$$m(\boldsymbol{x}) = \mathrm{E}\left[f(\boldsymbol{x})\right] \tag{8}$$

$$k(\boldsymbol{x}, \boldsymbol{x}') = \mathrm{E}\left[\left(f(\boldsymbol{x}) - m(\boldsymbol{x})\right)\left(f(\boldsymbol{x}') - m(\boldsymbol{x}')\right)^T\right]. \tag{9}$$

What makes GPs usable in practice is their property that any finite dimensional collection of random variables $\boldsymbol{f}(\boldsymbol{X}) = \{f(\boldsymbol{x}_1), \ldots, f(\boldsymbol{x}_M)\}$ is jointly Gaussian distributed

$$\boldsymbol{f}(\boldsymbol{X}) \sim \mathrm{N}\left(\boldsymbol{m}(\boldsymbol{X}), \boldsymbol{K}(\boldsymbol{X}, \boldsymbol{X})\right), \tag{10}$$

with

$$\boldsymbol{m}(\boldsymbol{X}) = \begin{pmatrix} m(\boldsymbol{x}_1) & \cdots & m(\boldsymbol{x}_M) \end{pmatrix}^T \tag{11}$$

$$\boldsymbol{K}(\boldsymbol{X}, \boldsymbol{X}) = \begin{pmatrix} k(\boldsymbol{x}_1, \boldsymbol{x}_1) & \cdots & k(\boldsymbol{x}_1, \boldsymbol{x}_M) \\ \vdots & \ddots & \vdots \\ k(\boldsymbol{x}_M, \boldsymbol{x}_1) & \cdots & k(\boldsymbol{x}_M, \boldsymbol{x}_M) \end{pmatrix}. \tag{12}$$
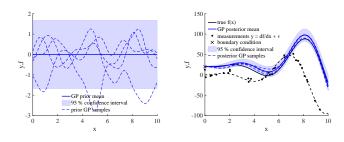
In regression, we wish to infer the process $f(\cdot)$ based on its inputs $\boldsymbol{x}$ and its (possibly noisy) output y. The output is not restricted to algebraic expressions such as $y = f(\boldsymbol{x}) + noise$, but can also be generalized to linear transforms of the process $y = \mathcal{L}_x f(\boldsymbol{x}) + noise$.

We will illustrate GP regression using the following differential example

$$y = \frac{\mathrm{d}f}{\mathrm{d}x} + \epsilon, \quad \text{with } \epsilon \sim \mathrm{N}\left(0, \sigma_\epsilon^2\right), \tag{13}$$

$$f^b = \mathcal{B}_x f, \tag{14}$$

(a) Gaussian process prior consisting of a zero mean and the squared exponential covariance function.

(b) Corresponding posterior Gaussian process.

Fig. 2: Inferring $f(x)$ from noisy differential observations $y = \frac{df}{dx} + \epsilon$ and boundary condition $f^b = f(x_b)$ using Gaussian process regression. The true function is $f_{true} = x^2 \sin(x) + 0.2x^{3/2} + 20$.

with a Dirichlet boundary condition rendering the boundary operator to the identity operator $\mathcal{B}_x = \mathcal{I}_x$, and therefore $f^b = f(x_b = 0)$. To start, we place a GP prior on $f(x)$ and write

$$f(x) \sim \text{GP}\left(0, k^{ff}(x, x')\right), \tag{15}$$

where we chose the prior consisting of a zero-mean function and a squared exponential covariance function $k(x, x') = \sigma^2 \exp\left(-\frac{(x-x')^2}{2l^2}\right)$. Here, σ^2 and l are hyper-parameters that permit to adjust the shape of the covariance function, and therefore the behavior of $f(x)$, to the data.

Next, we make use of the fact that a linear transform of a GP is another GP that is structured by the linear transform, see [14] and the references therein. We therefore have

$$y \sim \text{GP}\left(0, \underbrace{\mathcal{L}_x \mathcal{L}_{x'} k^{ff}(x, x')}_{=k^{yy}}\right), \tag{16}$$

$$f^b \sim \text{GP}\left(0, \underbrace{\mathcal{B}_x \mathcal{B}_{x'} k^{ff}(x, x')}_{=k^{bb}}\right). \tag{17}$$

For our example, k^{ff} and k^{bb} are identical due to the Dirichlet boundary condition. The covariance function k^{yy} can easily be derived under use of the formal definition of a covariance function (9). Note the important distinction between $\mathcal{L}_x$ and $\mathcal{L}_{x'}$, meaning that the operator should be applied with respect to either the first or second argument, that is x or x'.

We continue by formulating the evaluated joint model $p(\boldsymbol{f}, f^b, \boldsymbol{y})$ at test locations $\boldsymbol{x}_* = \{x_{*1}, \ldots, x_{*S}\}$, the boundary x_b, and observation locations $\boldsymbol{x} = \{x_1, \ldots, x_M\}$

$$\begin{pmatrix} \boldsymbol{f}(\boldsymbol{x}_*) \\ f^b(x_b) \\ \boldsymbol{y}(\boldsymbol{x}) \end{pmatrix} \sim \text{N}(\boldsymbol{0}, \boldsymbol{K}), \tag{18}$$

with covariance matrix

$$\boldsymbol{K} = \begin{bmatrix} \boldsymbol{K}^{ff}(\boldsymbol{x}_*, \boldsymbol{x}_*) & \boldsymbol{k}^{fb}(\boldsymbol{x}_*, x_b) & \boldsymbol{K}^{fy}(\boldsymbol{x}_*, \boldsymbol{x}) \\ \boldsymbol{k}^{bf}(x_b, \boldsymbol{x}_*) & k^{bb}(x_b, x_b) & \boldsymbol{k}^{by}(x_b, \boldsymbol{x}) \\ \boldsymbol{K}^{yf}(\boldsymbol{x}, \boldsymbol{x}_*) & \boldsymbol{k}^{yb}(\boldsymbol{x}, x_b) & \boldsymbol{K}^{yy}(\boldsymbol{x}, \boldsymbol{x}) \end{bmatrix}. \tag{19}$$

Before doing inference, we learn the hyper-parameters of the covariance function by minimizing the negative log marginal likelihood (NLML) $-\log p(f^b, \boldsymbol{y})$. We will touch on this more thoroughly later on.

Inference is now done by calculating the conditional distribution $p(\boldsymbol{f}|f^b, \boldsymbol{y})$. For Gaussian distributions, this results in another Gaussian with conditional mean $\text{E}[\boldsymbol{f}|f^b, \boldsymbol{y}]$ and covariance $\text{V}[\boldsymbol{f}|f^b, \boldsymbol{y}]$ as

$$\text{E}[\boldsymbol{f}|f^b, \boldsymbol{y}] = \begin{bmatrix} \boldsymbol{k}^{fb} & \boldsymbol{K}^{fy} \end{bmatrix} \begin{bmatrix} k^{bb} & \boldsymbol{k}^{by} \\ \boldsymbol{k}^{yb} & \boldsymbol{K}^{yy} + \sigma_n^2 \boldsymbol{I} \end{bmatrix}^{-1} \begin{pmatrix} f^b \\ \boldsymbol{y} \end{pmatrix} \tag{20}$$

$$\text{V}[\boldsymbol{f}|f^b, \boldsymbol{y}] = \boldsymbol{K}^{ff}$$
$$- \begin{bmatrix} \boldsymbol{k}^{fb} & \boldsymbol{K}^{fy} \end{bmatrix} \begin{bmatrix} k^{bb} & \boldsymbol{k}^{by} \\ \boldsymbol{k}^{yb} & \boldsymbol{K}^{yy} + \sigma_n^2 \boldsymbol{I} \end{bmatrix}^{-1} \begin{bmatrix} \boldsymbol{k}^{bf} \\ \boldsymbol{K}^{yf} \end{bmatrix}. \tag{21}$$

Here, another hyper-parameter σ_n^2 has been introduced to model measurement noise ϵ as additive zero mean Gaussian white noise. This hyper-parameter is also learned via the NLML. As can be seen in Fig. 2, the structured GP (18) is able to infer f from derivative observations.

To summarize, we give a cooking recipe for probabilistic inference with GPs:

1) Carefully place a GP prior to prevent inversion of linear operators (prior on f in our example to prevent inversion of $\frac{d}{dx}$).
2) Build up output GPs that are linear transforms of the GP prior (y, f^b in our example).
3) Construct the joint model $p(f, f^b, y)$ and learn the hyper-parameters via the marginal distribution $p(y, f^b)$.
4) Do inference by conditioning $p(f|f^b, y)$.

Numerical Gaussian processes for time-dependent partial differential equations

So far we looked at static examples. We are however interested in spatiotemporal systems described through time-dependent partial differential equations (PDEs). We assume that the PDE model is known and linear. Given an initial condition and boundary conditions, we wish to solve the PDE via GPs. We can do this with the recently introduced numerical GPs [4]. We will present the essential idea here, referring the reader to the original paper for further details.

Numerical GPs are built on the fact that linear transforms of a GP result in another GP. In fact, numerical GPs use a very similar procedure as shown earlier.

The first step is to discretize the time-dependent PDE (7) in time using one of the many existing methods (explicit Euler for simplicity here)

$$f_t(\boldsymbol{x}) = f_{t-1}(\boldsymbol{x}) + \Delta t \mathcal{L}_x f_{t-1}(\boldsymbol{x}) \tag{22}$$
$$= \mathcal{A}_x f_{t-1}(\boldsymbol{x}). \tag{23}$$

Next, a GP prior is placed on $f_{t-1} \sim \text{GP}\left(0, k_{t-1,t-1}^{ff}(\boldsymbol{x}, \boldsymbol{x}')\right)$

and therefore f_t is also a GP defined as

$$f_t(\boldsymbol{x}) \sim \text{GP}\left(0, \underbrace{\mathcal{A}_x \mathcal{A}_{x'} k_{t-1,t-1}^{ff}}_{=k_{t,t}^{ff}}\right), \qquad (24)$$

and so is the boundary condition $f_t^b = \mathcal{B}_x f_t$

$$f_t^b(\boldsymbol{x}) \sim \text{GP}\left(0, \underbrace{\mathcal{B}_x \mathcal{B}_{x'} \mathcal{A}_x \mathcal{A}_{x'} k_{t-1,t-1}^{ff}}_{=k_{t,t}^{f^b f^b}}\right). \qquad (25)$$

Then, the evaluated joint model $p(\boldsymbol{f}_t, \boldsymbol{f}_t^b, \boldsymbol{f}_{t-1})$ on arbitrary test points $\boldsymbol{X}_t$, the boundaries $\boldsymbol{X}_b$, and previous points $\boldsymbol{X}_{t-1}$ is formulated as

$$\begin{pmatrix} \boldsymbol{f}_t \\ \boldsymbol{f}_t^b \\ \boldsymbol{f}_{t-1} \end{pmatrix} \sim \text{N}\left(\boldsymbol{0}, \begin{bmatrix} \boldsymbol{K}_{t,t}^{ff} & \boldsymbol{K}_{t,t}^{ff^b} & \boldsymbol{K}_{t,t-1}^{ff} \\ \boldsymbol{K}_{t,t}^{f^b f} & \boldsymbol{K}_{t,t}^{f^b f^b} & \boldsymbol{K}_{t,t-1}^{f^b f} \\ \boldsymbol{K}_{t-1,t}^{ff} & \boldsymbol{K}_{t-1,t}^{ff^b} & \boldsymbol{K}_{t-1,t-1}^{ff} \end{bmatrix}\right). \qquad (26)$$

Here, the inputs to the individual covariance matrices are omitted. They can be inferred from the sub- and superscripts, e.g. $\boldsymbol{K}_{t-1,t}^{ff^b}(\boldsymbol{X}_{t-1}, \boldsymbol{X}_b)$.

To simulate the spatiotemporal model, we now recursively compute the conditional distribution $p(\boldsymbol{f}_t|\boldsymbol{f}_t^b, \boldsymbol{f}_{t-1})$. Here the solution from the previous simulation step $\boldsymbol{f}_{t-1}$ acts as an artificial measurement. Hyper-parameters are optimized in each simulation step by minimizing $-\log p(\boldsymbol{f}_t^b, \boldsymbol{f}_{t-1})$.

III. NUMERICAL GAUSSIAN PROCESSES AS PROBABILISTIC STATE SPACE MODELS

In this section we will derive a probabilistic state space model from numerical Gaussian processes. This probabilistic state space model is linear and its random states are, by definition of GPs, Gaussian distributed. These model properties allow us in Section IV to go through the Kalman filter algorithm, thereby deriving the numerical Gaussian process Kalman filter.

We start by introducing an output channel for noisy online measurements that are linear transforms of $f(t, \boldsymbol{x})$. Mathematically this reads as

$$f_t^y(\boldsymbol{y}) = (\mathcal{H}_x f_t)(\boldsymbol{x}) + r_t(\boldsymbol{y}). \qquad (27)$$

The measurement operator is $\mathcal{H}_x : L_2(\mathbb{R}^{d_x}, \mathbb{R}) \to L_2(\mathbb{R}^{d_y}, \mathbb{R}), f_t \mapsto f_t^y$. Obtained measurements are point evaluations of f_t^y at locations $\boldsymbol{y}_i$.

In this work, we consider measurement noise r_t as zero-mean spatiotemporal white noise with constant variance over $\boldsymbol{y}$. To model this, we place a GP prior on r_t that is independent of all the other GP priors. We write

$$r_t \sim \text{GP}\left(0, k_{t,t}^{rr}(\boldsymbol{y}, \boldsymbol{y}')\right), \qquad (28)$$

where the covariance function is $k_{t,t}^{rr} = \sigma_r^2 \delta(\boldsymbol{y} - \boldsymbol{y}')$ and $E[r_t r_{t+\tau}] = 0$. Other additive noise models are conceivable. Particularly the spatial white noise assumption could be relaxed to something more suitable under use of an appropriate covariance function [6].

Since $f_t(x)$ is a GP, so is the measurement output

$$f_t^y \sim \text{GP}\left(0, \underbrace{\left(\mathcal{H}_x \mathcal{H}_{x'} k_{t,t}^{ff}\right)(\boldsymbol{x}, \boldsymbol{x}') + k_{t,t}^{rr}(\boldsymbol{y}, \boldsymbol{y}')}_{=k_{t,t}^{f^y f^y}}\right). \qquad (29)$$

Regarding the PDE itself, we add spatiotemporal process noise to the PDE to obtain

$$\frac{\partial}{\partial t} f(t, \boldsymbol{x}) = \mathcal{L}_x f(t, \boldsymbol{x}) + q(t, \boldsymbol{x}). \qquad (30)$$

Discretization in time yields

$$f_t(\boldsymbol{x}) = \mathcal{A}_x f_{t-1}(\boldsymbol{x}) + \Delta t q_{t-1}(\boldsymbol{x}), \qquad (31)$$

with process noise modeled as an independent zero-mean GP prior with a white noise kernel $k_{t-1,t-1}^{qq}(x, x') = \sigma_q^2 \delta(x - x')$. Of course, the above discretization is mathematically not rigorous because (30) is not continuous due to the spatiotemporal white noise. A more rigorous treatment would require Itô–Calculus [15]. We could formally circumvent this by first discretizing and then adding process noise $q_{t-1}(\boldsymbol{x})$ (without Δt), as done in the finite-dimensional state space models. However, we found that exclusion of Δt in the process noise term drastically reduced estimation performance. Moreover, above we recovered a time-discrete form that is similar to the Euler-Maruyama method. This brief discussion lays bare the gap between proper theoretical treatment and actual implementation, as lamented as early as in [16].

The updated GP for $f_t(x)$ is

$$f_t(x) \sim \text{GP}\left(0, \underbrace{\mathcal{A}_x \mathcal{A}_{x'} k_{t-1,t-1}^{ff} + \Delta t^2 k_{t-1,t-1}^{qq}}_{=k_{t,t}^{ff}}\right). \qquad (32)$$

The complete multi output GP now reads as

$$\begin{pmatrix} f_t \\ f_t^y \\ f_t^b \\ f_{t-1} \end{pmatrix} \sim \text{GP}\left(\boldsymbol{0}, \underbrace{\begin{bmatrix} k_{t,t}^{ff} & k_{t,t}^{f^y f} & k_{t,t}^{ff^b} & k_{t,t-1}^{ff} \\ k_{t,t}^{f^y f} & k_{t,t}^{f^y f^y} & k_{t,t}^{f^y f^b} & k_{t,t-1}^{f^y f} \\ k_{t,t}^{f^b f} & k_{t,t}^{f^b f^y} & k_{t,t}^{f^b f^b} & k_{t,t-1}^{f^b f} \\ k_{t-1,t}^{ff} & k_{t-1,t}^{ff^y} & k_{t-1,t}^{ff^b} & k_{t-1,t-1}^{ff} \end{bmatrix}}_{=\boldsymbol{K}_{GP}}\right). \tag{33}$$

All covariance functions inside $\boldsymbol{K}_{GP}$ are known as they can be built up from the prior covariance function $k_{t-1,t-1}^{ff}$ (for the explicit Euler) and the respective linear operators $\mathcal{A}_x, \mathcal{H}_x$ or functionals $\mathcal{B}_x$, see Appendix C.

We can recover a probabilistic state space model from (33). For this purpose we assume that (33) has been evaluated at measurement locations $\boldsymbol{Y}$, boundaries $\boldsymbol{X}_b$ and arbitrarily chosen test locations $\boldsymbol{X}$. The covariance functions therefore become covariance matrices, e.g. $k_{t,t}^{f^y f^b}(\boldsymbol{y}, \boldsymbol{x}_b)$ becomes $\boldsymbol{K}_{t,t}^{f^y f^b}(\boldsymbol{Y}, \boldsymbol{X}_b)$.

The process equation is

$$p\left(\boldsymbol{f}_t|\boldsymbol{f}_t^b, \boldsymbol{f}_{t-1}\right) = \text{N}\left(\boldsymbol{f}_t|\boldsymbol{A}_t\begin{pmatrix} \boldsymbol{f}_t^b \\ \boldsymbol{f}_{t-1} \end{pmatrix}, \boldsymbol{P}_t^{GPf}\right), \qquad (34)$$

with the covariance matrix

$$\boldsymbol{A}_t = \begin{bmatrix} \boldsymbol{K}_{t,t}^{ff^b} & \boldsymbol{K}_{t,t-1}^{ff} \end{bmatrix} \begin{bmatrix} \boldsymbol{K}_{t,t}^{f^b f^b} & \boldsymbol{K}_{t,t-1}^{f^b f} \\ \boldsymbol{K}_{t-1,t}^{ff^b} & \boldsymbol{K}_{t-1,t-1}^{ff} \end{bmatrix}^{-1} \quad (35)$$

taking the role of the dynamic matrix in a state space model. Here, $\boldsymbol{A}_t$ is time-varying, indicated by the subscript, due to changing hyper-parameter values and possibly changing regression points. The process noise covariance matrix of the state space model, usually denoted $\boldsymbol{Q}$, is given by

$$\boldsymbol{P}_t^{GPf} = \boldsymbol{K}_{t,t}^{ff} - \boldsymbol{A}_t \begin{bmatrix} \boldsymbol{K}_{t,t}^{ff^b} & \boldsymbol{K}_{t,t-1}^{ff} \end{bmatrix}^T. \quad (36)$$

The measurement equation reads as

$$p\left(\boldsymbol{f}_t^y | \boldsymbol{f}_t\right) = \mathrm{N}\left(\boldsymbol{f}_t^y \middle| \boldsymbol{C}_t \boldsymbol{f}_t, \, \boldsymbol{P}_t^{GP,f^y}\right), \quad (37)$$

with

$$\boldsymbol{C}_t = \boldsymbol{K}_{t,t}^{f^y f} \left(\boldsymbol{K}_{t,t}^{ff}\right)^{-1} \quad (38)$$

taking the role of the measurement matrix in a state space model. The measurement noise covariance matrix of the state space model, usually denoted $\boldsymbol{R}$, is given by

$$\boldsymbol{P}_t^{GP,f^y} = \boldsymbol{K}_{t,t}^{f^y f^y} - \boldsymbol{C}_t \boldsymbol{K}_{t,t}^{ff^y}. \quad (39)$$

The noise covariance matrices $\boldsymbol{Q}$ and $\boldsymbol{R}$ in a classical state space model are usually static matrices. The GP model however naturally allows for an adaptive noise description, due to online adaptation of the hyper-parameters, thereby drastically reducing the usually required tedious fine-tuning of the Kalman filter. This will be showcased with case studies in Section V and discussed in Section VI.

For the recursive Kalman filter algorithm we will also need the following

$$p(\boldsymbol{f}_t^b | \boldsymbol{f}_{t-1}) = \mathrm{N}\left(\boldsymbol{f}_t^b | \boldsymbol{A}_t^{f^b} \boldsymbol{f}_{t-1}, \boldsymbol{P}_t^{GP,f^b}\right), \quad (40)$$

with

$$\boldsymbol{A}_t^{f^b} = \boldsymbol{K}_{t,t-1}^{f^b f} \left(\boldsymbol{K}_{t-1,t-1}^{nn}\right)^{-1}, \quad (41)$$

$$\boldsymbol{P}_t^{GP,f^b} = \boldsymbol{K}_{t,t}^{f^b f^b} - \boldsymbol{A}_t^{f^b} \boldsymbol{K}_{t-1,t}^{ff^b}. \quad (42)$$

Equipped with a probabilistic state space model, we are now ready to write down the recursive Kalman filter equations.

IV. NUMERICAL GAUSSIAN PROCESS KALMAN FILTER DERIVATION

In this section we will embed the numerical GP state space model into the Kalman filter algorithm. Afterwards, we will briefly discuss the Kalman filter marginal likelihood for the use of hyper-parameter estimation.

There are two ways to derive the Kalman filter. The first derivation is rooted in recursive least squares regression, see e.g. [17]. The second derivation takes a Bayesian filtering perspective, see e.g. [18]. In this article, we take the latter perspective to derive the numerical Gaussian process Kalman filter since the Bayesian perspective neatly fits into the Gaussian process framework.
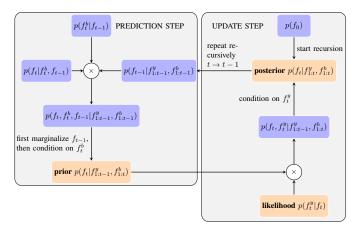

Fig. 3: The recursive Kalman filter algorithm for calculating the posterior state distribution of a spatiotemporal system.

Recall that in Kalman filtering, we want to calculate the probability distribution of a dynamic state given all the measurement history up to the present. For spatiotemporal systems, we additionally condition on the boundary data history. Boundary conditions are treated as part of the model and therefore enter through the prediction step of the Kalman filter. According to Bayes' rule, the posterior distribution of the state estimate is given by

$$p(\boldsymbol{f}_t | \boldsymbol{f}_{1:t}^y, \boldsymbol{f}_{1:t}^b) = \frac{p(\boldsymbol{f}_t^y | \boldsymbol{f}_t) p(\boldsymbol{f}_t | \boldsymbol{f}_{1:t-1}^y, \boldsymbol{f}_{1:t}^b)}{p(\boldsymbol{f}_t^y | \boldsymbol{f}_{1:t-1}^y, \boldsymbol{f}_{1:t}^b)}. \quad (43)$$

We make the usual assumptions for the Kalman filter [18] and an additional one for the boundary data:

Assumption 1) States are assumed to be Markovian, i.e. the current state $\boldsymbol{f}_t$ is conditionally independent of anything that happened before $t-1$ given the previous state $\boldsymbol{f}_{t-1}$.

Assumption 2) Given the current state $\boldsymbol{f}_t$, the current measurement $\boldsymbol{f}_t^y$ is conditionally independent of the measurement $\boldsymbol{f}_{1:t-1}^y$ and state histories $\boldsymbol{f}_{1:t-1}$, as well as the boundary history up to the present $\boldsymbol{f}_{1:t}^b$.

Assumption 3) Given the previous state $\boldsymbol{f}_{t-1}$, the current boundary condition $\boldsymbol{f}_t^b$ is conditionally independent of the measurement $\boldsymbol{f}_{1:t-1}^y$ and boundary history data $\boldsymbol{f}_{1:t-1}^b$.

In this article, only one-step discretization schemes such as the Euler method are considered. Using multi-step methods would require to extend the Markovian property to an accordingly higher order.

The individual terms of (43) can be calculated in closed form for Gaussian distributions as

- prior $p(\boldsymbol{f}_t | \boldsymbol{f}_{1:t-1}^y, \boldsymbol{f}_{1:t}^b) = \mathrm{N}(\boldsymbol{f}_t | \boldsymbol{m}_t^-, \boldsymbol{P}_t^-)$
- posterior $p(\boldsymbol{f}_t | \boldsymbol{f}_{1:t}^y, \boldsymbol{f}_{1:t}^b) = \mathrm{N}(\boldsymbol{f}_t | \boldsymbol{m}_t, \boldsymbol{P}_t)$
- marginal likelihood
  $p(\boldsymbol{f}_t^y | \boldsymbol{f}_{1:t-1}^y, \boldsymbol{f}_{1:t}^b) = \mathrm{N}(\boldsymbol{f}_t^y | \boldsymbol{C}_t \boldsymbol{m}_t^-, \boldsymbol{S}_t)$

The road map ahead to derive these expressions is shown in Fig. 3. To calculate the prior distribution $p(\boldsymbol{f}_t | \boldsymbol{f}_{1:t-1}^y, \boldsymbol{f}_{1:t}^b)$, we will first formulate the joint distribution between states $\boldsymbol{f}_t$,

$\boldsymbol{f}_{t-1}$, and current boundary condition $\boldsymbol{f}_t^b$, conditioned on measurement and boundary histories. Under use of Assumptions 1 and 3 [1], the prior joint is calculated as

$$p(\boldsymbol{f}_t, \boldsymbol{f}_t^b, \boldsymbol{f}_{t-1} | \boldsymbol{f}_{1:t-1}^y, \boldsymbol{f}_{1:t-1}^b) \tag{44}$$

$$= \underbrace{p(\boldsymbol{f}_t | \boldsymbol{f}_t^b, \boldsymbol{f}_{t-1})}_{\text{model prediction}} p(\boldsymbol{f}_t^b | \boldsymbol{f}_{t-1}) \underbrace{p(\boldsymbol{f}_{t-1} | \boldsymbol{f}_{1:t-1}^y, \boldsymbol{f}_{1:t-1}^b)}_{\text{previous posterior}} \tag{45}$$

$$= \mathrm{N} \left(\boldsymbol{f}_t | \boldsymbol{A}_t \begin{pmatrix} \boldsymbol{f}_t^b \\ \boldsymbol{f}_{t-1} \end{pmatrix}, \boldsymbol{P}_t^{GP,f} \right) \tag{46}$$

$$\times \mathrm{N} \left(\boldsymbol{f}_t^b | \boldsymbol{A}_t^{f^b} \boldsymbol{f}_{t-1}, \boldsymbol{P}_t^{GP,f^b} \right)$$

$$\times \mathrm{N} \left(\boldsymbol{f}_{t-1} | \boldsymbol{m}_{t-1}, \boldsymbol{P}_{t-1} \right)$$

$$= \mathrm{N} \left(\begin{pmatrix} \boldsymbol{f}_t \\ \boldsymbol{f}_t^b \\ \boldsymbol{f}_{t-1} \end{pmatrix} | \boldsymbol{m}', \boldsymbol{P}' \right). \tag{47}$$

The covariance matrices $\boldsymbol{A}_t$ and $\boldsymbol{P}_t^{GP,f}$ have been introduced in (34), while $\boldsymbol{A}_t^{f^b}$ and $\boldsymbol{P}_t^{GP,f^b}$ have been introduced in (42).

Using Lemma A two times, the joint mean $\boldsymbol{m}'$ of (47) is

$$\boldsymbol{m}' = \begin{pmatrix} \boldsymbol{A}_t \begin{pmatrix} \boldsymbol{A}_t^{f^b} \boldsymbol{m}_{t-1} \\ \boldsymbol{m}_{t-1} \end{pmatrix} \\ \boldsymbol{A}_t^{f^b} \boldsymbol{m}_{t-1} \\ \boldsymbol{m}_{t-1} \end{pmatrix} \tag{48}$$

and the covariance is

$$\boldsymbol{P}' = \begin{bmatrix} \boldsymbol{A}_t \tilde{\boldsymbol{P}}_{t-1} \boldsymbol{A}_t^T + \boldsymbol{P}_t^{GP,f} & \boldsymbol{A}_t \begin{bmatrix} \boldsymbol{S}_t^{f^b} \\ \boldsymbol{P}_{t-1}(\boldsymbol{A}_t^{f^b})^T \end{bmatrix} & \boldsymbol{A}_t \begin{bmatrix} \boldsymbol{A}_t^{f^b} \boldsymbol{P}_{t-1} \\ \boldsymbol{P}_{t-1} \end{bmatrix} \\ \left(\boldsymbol{A}_t \begin{bmatrix} \boldsymbol{S}_t^{f^b} \\ \boldsymbol{P}_{t-1}(\boldsymbol{A}_t^{f^b})^T \end{bmatrix} \right)^T & \boldsymbol{S}_t^{f^b} & \boldsymbol{A}_t^{f^b} \boldsymbol{P}_{t-1} \\ \left(\boldsymbol{A}_t \begin{bmatrix} \boldsymbol{A}_t^{f^b} \boldsymbol{P}_{t-1} \\ \boldsymbol{P}_{t-1} \end{bmatrix} \right)^T & \left(\boldsymbol{A}_t^{f^b} \boldsymbol{P}_{t-1} \right)^T & \boldsymbol{P}_{t-1} \end{bmatrix}. \tag{49}$$

Here we introduced

$$\boldsymbol{S}_t^{f^b} = \boldsymbol{A}_t^{f^b} \boldsymbol{P}_{t-1} \left(\boldsymbol{A}_t^{f^b} \right)^T + \boldsymbol{P}_t^{GP,f^b}, \tag{50}$$

$$\tilde{\boldsymbol{P}}_{t-1} = \begin{bmatrix} \boldsymbol{S}_t^{f^b} & \boldsymbol{A}_t^{f^b} \boldsymbol{P}_{t-1} \\ \left(\boldsymbol{A}_t^{f^b} \boldsymbol{P}_{t-1} \right)^T & \boldsymbol{P}_{t-1} \end{bmatrix}. \tag{51}$$

The prior distribution

$$p(\boldsymbol{f}_t | \boldsymbol{f}_{1:t-1}^y, \boldsymbol{f}_{1:t}^b) = \mathrm{N} \left(\boldsymbol{f}_t | \boldsymbol{m}_t^-, \boldsymbol{P}_t^- \right), \tag{52}$$

is now obtained from (47) by first marginalizing over $\boldsymbol{f}_{t-1}$ and then conditioning on $\boldsymbol{f}_t^b$. Prior mean $\boldsymbol{m}_t^-$ and covariance $\boldsymbol{P}_t^-$ are

$$\boldsymbol{m}_t^- = \boldsymbol{A}_t \left(\boldsymbol{m}_{t-1} + \boldsymbol{P}_{t-1} \left(\boldsymbol{A}_t^{f^b} \right)^T \left(\boldsymbol{S}_t^{f^b} \right)^{-1} \begin{pmatrix} \boldsymbol{f}_t^b \\ \left(\boldsymbol{f}_t^b - \boldsymbol{A}_t^{f^b} \boldsymbol{m}_{t-1} \right) \end{pmatrix} \right), \tag{53}$$

$$\boldsymbol{P}_t^- = \boldsymbol{A}_t \tilde{\boldsymbol{P}}_{t-1} \boldsymbol{A}_t^T + \boldsymbol{P}_t^{GP,f} - \boldsymbol{A}_t \begin{bmatrix} \boldsymbol{S}_t^{f^b} \\ \boldsymbol{P}_{t-1} \left(\boldsymbol{A}_t^{f^b} \right)^T \end{bmatrix}$$

$$\times \left(\boldsymbol{S}_t^{f^b} \right)^{-1} \begin{bmatrix} \boldsymbol{S}_t^{f^b} \\ \boldsymbol{P}_{t-1} \left(\boldsymbol{A}_t^{f^b} \right)^T \end{bmatrix}^T \boldsymbol{A}_t^T. \tag{54}$$

[1] regarding the boundary, one could also think of assuming that the current boundary $\boldsymbol{f}_t^b$ and the previous state $\boldsymbol{f}_{t-1}$ are independent but this would be in conflict with the numerical GP state space model.

Equipped with this, we can calculate the joint distribution of the current state (prior distribution) and measurements (likelihood). This will allow us to write down the posterior distribution later on. Using Assumption 2, we have

$$p(\boldsymbol{f}_t, \boldsymbol{f}_t^y | \boldsymbol{f}_{1:t-1}^y, \boldsymbol{f}_{1:t}^b)$$

$$= \underbrace{p(\boldsymbol{f}_t^y | \boldsymbol{f}_t)}_{\text{likelihood}} \underbrace{p(\boldsymbol{f}_t | \boldsymbol{f}_{1:t-1}^y, \boldsymbol{f}_{1:t}^b)}_{\text{prior}}$$

$$= \mathrm{N}(\boldsymbol{f}_t^y | \boldsymbol{C}_t \boldsymbol{f}_t, \boldsymbol{P}_t^{GP,f^y}) \mathrm{N} \left(\boldsymbol{f}_t | \boldsymbol{m}_t^-, \boldsymbol{P}_t^- \right)$$

$$= \mathrm{N} \left(\begin{pmatrix} \boldsymbol{f}_t \\ \boldsymbol{f}_t^y \end{pmatrix} | \boldsymbol{m}'', \boldsymbol{P}'' \right). \tag{55}$$

The covariance matrices $\boldsymbol{C}_t$ and $\boldsymbol{P}_t^{GP,f^y}$ were introduced in (37). We use Lemma A one more time to get the joint mean as

$$\boldsymbol{m}'' = \begin{pmatrix} \boldsymbol{m}_t^- \\ \boldsymbol{C}_t \boldsymbol{m}_t^- \end{pmatrix}, \tag{56}$$

and the covariance as

$$\boldsymbol{P}'' = \begin{bmatrix} \boldsymbol{P}_t^- & \boldsymbol{P}_t^- \boldsymbol{C}_t^T \\ \boldsymbol{C}_t \boldsymbol{P}_t^- & \boldsymbol{C}_t \boldsymbol{P}_t^- \boldsymbol{C}_t^T + \boldsymbol{P}_t^{GP,f^y} \end{bmatrix}. \tag{57}$$

The posterior distribution

$$p(\boldsymbol{f}_t | \boldsymbol{f}_t^y, \boldsymbol{f}_{1:t-1}^y, \boldsymbol{f}_{1:t}^b) = p(\boldsymbol{f}_t | \boldsymbol{f}_{1:t}^y, \boldsymbol{f}_{1:t}^b)$$

$$= \mathrm{N} \left(\boldsymbol{m}_t, \boldsymbol{P}_t \right) \tag{58}$$

is obtained by conditioning the joint distribution (55) on the current measurement using Lemma B. The posterior mean is

$$\boldsymbol{m}_t = \boldsymbol{m}_t^- + \boldsymbol{P}_t^- \boldsymbol{C}_t^T$$

$$\times \left(\boldsymbol{C}_t \boldsymbol{P}_t^- \boldsymbol{C}_t^T + \boldsymbol{P}_t^{GP,f^y} \right)^{-1} \left(\boldsymbol{f}_t^y - \boldsymbol{C}_t \boldsymbol{m}_t^- \right), \tag{59}$$

and the corresponding posterior variance is

$$\boldsymbol{P}_t = \boldsymbol{P}_t^- - \boldsymbol{P}_t^- \boldsymbol{C}_t^T \left(\boldsymbol{C}_t \boldsymbol{P}_t^- \boldsymbol{C}_t^T + \boldsymbol{P}_t^{GP,f^y} \right)^{-1} \boldsymbol{C}_t \boldsymbol{P}_t^-. \tag{60}$$

To summarize the numerical GPKF, we have the prediction step as

$$\boldsymbol{m}_t^- = \boldsymbol{A}_t \left(\boldsymbol{m}_{t-1} + \boldsymbol{P}_{t-1} \left(\boldsymbol{A}_t^{f^b} \right)^T \left(\boldsymbol{S}_t^{f^b} \right)^{-1} \begin{pmatrix} \boldsymbol{f}_t^b \\ \left(\boldsymbol{f}_t^b - \boldsymbol{A}_t^{f^b} \boldsymbol{m}_{t-1} \right) \end{pmatrix} \right), \tag{61}$$

$$\boldsymbol{P}_t^- = \boldsymbol{A}_t \tilde{\boldsymbol{P}}_{t-1} \boldsymbol{A}_t^T + \boldsymbol{P}_t^{GP,f} - \boldsymbol{A}_t \begin{bmatrix} \boldsymbol{S}_t^{f^b} \\ \boldsymbol{P}_{t-1} \left(\boldsymbol{A}_t^{f^b} \right)^T \end{bmatrix}$$

$$\times \left(\boldsymbol{S}_t^{f^b} \right)^{-1} \begin{bmatrix} \boldsymbol{S}_t^{f^b} \\ \boldsymbol{P}_{t-1} \left(\boldsymbol{A}_t^{f^b} \right)^T \end{bmatrix}^T \boldsymbol{A}_t^T. \tag{62}$$

and the update step as

$$\boldsymbol{v}_t = \boldsymbol{f}_t^y - \boldsymbol{C}_t \boldsymbol{m}_t^-, \tag{63}$$

$$\boldsymbol{S}_t = \boldsymbol{C}_t \boldsymbol{P}_t^- \boldsymbol{C}_t^T + \boldsymbol{P}_t^{GP,f^y}, \tag{64}$$

$$\boldsymbol{K}_t = \boldsymbol{P}_t^- \boldsymbol{C}_t^T (\boldsymbol{S}_t)^{-1}, \tag{65}$$

$$\boldsymbol{m}_t = \boldsymbol{m}_t^- + \boldsymbol{K}_t \boldsymbol{v}_t, \tag{66}$$

$$\boldsymbol{P}_t = \boldsymbol{P}_t^- - \boldsymbol{K}_t \boldsymbol{S}_t (\boldsymbol{K}_t)^T. \tag{67}$$

Due to having to include boundary conditions, the prediction step doesn't perfectly align with the traditional KF. It is not possible to simply write the prior mean prediction as the previous posterior mean mapped through the dynamic matrix $\boldsymbol{A}_t$. Treating boundary conditions as measurements results in much simpler expressions that smoothly align with the traditional KF, see [7]. Of course this means that prediction steps do not abide boundary conditions which can be problematic when multiple prediction steps have to be computed between update steps.

Hyper-parameter estimation via the marginal likelihood

Hyper-parameters are estimated in each update step. In GP regression the hyper-parameters are usually estimated by minimizing the negative log marginal likelihood (NLML). The marginal likelihood of the Kalman filter, marginal with respect to the current state f_t, is found in the recursive Bayes denominator (43) as

$$p(\boldsymbol{f}_t^y | \boldsymbol{f}_{1:t-1}^y, \boldsymbol{f}_{1:t}^b) = \mathrm{N}\left(\boldsymbol{C}_t \boldsymbol{m}_t^-, \boldsymbol{S}_t\right). \tag{68}$$

The corresponding NLML is

$$-\log p(\boldsymbol{f}_t^y | \boldsymbol{f}_{1:t-1}^y, \boldsymbol{f}_{1:t}^b) \tag{69}$$
$$= \frac{1}{2}\left(\boldsymbol{f}_t^y - \boldsymbol{C}_t \boldsymbol{m}_t^-\right)^T \boldsymbol{S}_t^{-1}\left(\boldsymbol{f}_t^y - \boldsymbol{C}_t \boldsymbol{m}_t^-\right)$$
$$+ \frac{1}{2}\log(\det(\boldsymbol{S}_t)) + \frac{N_y}{2}\log(2\pi), \tag{70}$$

with N_y being the number of measurements. The partial derivatives of this NLML required for minimization can be calculated analytically and are shown in Appendix D.

V. Simulation case studies

In this section we will showcase the numerical Gaussian process Kalman filter (GPKF) using different case studies. For all case studies the relative error between posterior mean estimate $\boldsymbol{m}_t$ and analytical solution $\boldsymbol{f}_{t,ref}$ evaluated on the same points $\boldsymbol{X}$ is defined as

$$\text{relative error at time } t = \frac{\|\boldsymbol{f}_{t,ref} - \boldsymbol{m}_t\|}{\|\boldsymbol{f}_{t,ref}\|}, \tag{71}$$

with euclidean distance $\|\cdot\|$. All case studies have been implemented in MATLAB.

A. Step shaped one-dimensional advection equation

We study the one-dimensional advection equation

$$\frac{\partial}{\partial t} f(t, x) = -g \frac{\partial f}{\partial x}, \tag{72}$$

with initial and boundary condition as

$$f(0, x) = 2H(x) - 1, \tag{73}$$
$$f(t, 0) = 0. \tag{74}$$

Here, $H(x)$ is the Heaviside step function. We solve (72) with the implicit Euler and a step size of $\Delta t = 0.005$. The GP prior

is $f_t \sim \mathrm{GP}(0, k_{t,t}^{ff}(x, x'))$ with a neural network covariance function

$$k_{t,t}^{ff}(x, x') = \frac{2}{\pi} \sin^{-1}\left(\frac{2(\sigma_0^2 + \sigma^2 x x')}{\sqrt{(1 + 2(\sigma_0^2 + \sigma^2 x^2))(1 + 2(\sigma_0^2 + \sigma^2 x'^2))}}\right). \tag{75}$$

This covariance function is capable of capturing discontinuities. The hyper-parameters are σ_0^2 and σ^2. Measurements are noisy point evaluations so that the measurement equation reads as

$$f_t^y(y) = f_t(x) + r_t, \tag{76}$$

with $r_t \sim \mathrm{N}(0, \sigma_{r,true}^2)$ being white noise. Measurements are received every third model step size so that between each KF update step there are three KF prediction steps. Hyper-parameters are learned by minimizing the NLML (70) in each update step.

Estimation results are illustrated in Fig. 4. The initial estimate (blue curve with uncertainty band in Fig. 4a) is far off from the true solution (black line). However, rather quickly the estimate converges closer, as indicated by a dropping relative error (black curve in Fig. 4b). The relative error drops to a value of around 0.1 before it slowly starts to increase. This divergence is not uncommon in Kalman filtering [19]. One possible cause, that is also present in Fig. 4a, is a low posterior covariance matrix $\boldsymbol{P}_t$. As a result, the KF puts less importance on measurements and more on model predictions, which in this case aren't perfect due to the temporal discretization required for numerical GPs. Remedies such as the Kalman filter with fading memory exist [17].

The measurement noise level σ_r is correctly learned online to be $\sigma_{r,true}$. Estimated process noise σ_q is non-zero, likely reflecting the aforementioned model imperfection. For a few time steps at the beginning its estimation gets exceedingly large. This is cut off from the plot.

B. Liouville's equation

We study Liouville's equation

$$\frac{\partial}{\partial t} f(t, \boldsymbol{x}) = -\mathrm{div}\left(f(t, \boldsymbol{x})\dot{\boldsymbol{x}}\right) \tag{77}$$
$$= -\dot{x}_1 \frac{\partial f}{\partial x_1} - f \frac{\partial \dot{x}_1}{\partial x_1} - \dot{x}_2 \frac{\partial f}{\partial x_2} - f \frac{\partial \dot{x}_2}{\partial x_2}, \tag{78}$$

wherein the vector field $\dot{\boldsymbol{x}}$ is governed by

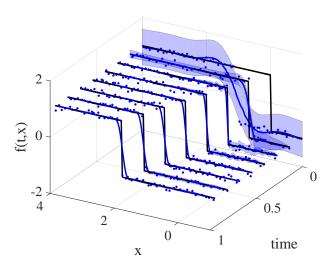$$\dot{\boldsymbol{x}} = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix} \boldsymbol{x}. \tag{79}$$
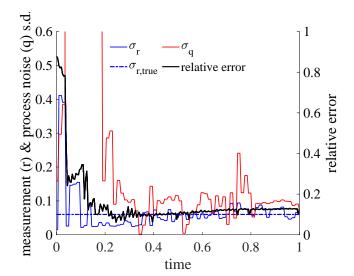
Initial and boundary conditions are

$$f(0, x_1, x_2) = \mathrm{N}\left(\boldsymbol{\mu}_1, \Sigma\right) + \mathrm{N}\left(\boldsymbol{\mu}_2, \Sigma\right), \tag{80}$$
$$f(t, 0, x_2) = 0, \tag{81}$$
$$f(t, x_1, 0) = 0. \tag{82}$$

Here, the initial condition consists of two Gaussian bumps sitting opposite of each other at $\boldsymbol{\mu}_1 = (0, -2)^T$ and $\boldsymbol{\mu}_2 = (0, 2)^T$ with covariance matrix $\Sigma = \mathrm{diag}(0.4, 0.65)$. The initial estimate communicated to the Kalman filter consists of only one Gaussian bump located at $\boldsymbol{\mu}_1$ with increased variance $\Sigma + \mathrm{diag}(0.4, 0.4)$.

(a) Snapshots for the step shaped one-dimensional advection equation. Noisy measurements are blue dots, while the true solution is represented by black lines. Posterior mean estimates along with their 95 % confidence intervals are shown in blue.

(b) Hyper-parameter estimation and relative error (71) over time. Estimated process noise σ_q gets exceedingly large at the beginning. This is cut off from the plot.

Fig. 4: Estimation results for the step shaped one-dimensional advection equation (72).

We discretize (78) in time using the implicit Euler method with a step size of $\Delta t = 0.005$, and place a GP prior on $f_t \sim \mathrm{GP}(0, k_{t,t}^{ff}(x, x'))$ with the following squared exponential covariance function

$$k_{t,t}^{ff}(x, x') = \sigma_n^2 \exp\left(-\frac{(x_1 - x_1')^2}{2l_1^2} - \frac{(x_2 - x_2')^2}{2l_2^2}\right). \quad (83)$$

Here, the hyper-parameters are σ_n^2, l_1, and l_2. Through noisy measurements we receive the marginal distribution of x_1 as

$$f_t^y(y = x_1) = \int_{x_{2,min}}^{x_{2,max}} f_t(x_1, x_2)\mathrm{d}x_2 + r_t, \quad (84)$$

with $r_t \sim \mathrm{N}(0, \sigma_{r,true}^2)$ being white noise.

The vector field $\dot{x}$ describes harmonic oscillations, plotting the solution $f(t, \boldsymbol{x})$ results therefore in circular motions. The simulation runs for one full rotation $T_{sim} = 2\pi$. Although measurements are formally the x_1-marginal, over time measurements will be marginals of the initial distribution from all angles due to the circular motion of the joint distribution $f(t, \boldsymbol{x})$. Interested readers are referred to [20] for further details. As discovered and explained by [21], this case study illustrates the relatedness to tomography. Measurements are received every $\frac{T_{sim}}{64}$, which translates to 19 KF predictions steps for every KF update step.

State estimation results are displayed in Fig. 5 while noise level estimations are displayed in Fig. 6. Through measurements of the x_1-marginal $f(t, x_1)$ the numerical GPKF is able to reconstruct the latent x_2-marginal distribution $f(t, x_2)$ and the joint distribution $f(t, x_1, x_2)$. Noise level estimation σ_r oscillates around its true level $\sigma_{r,true}$, while process noise σ_q is set to zero by the numerical GPKF. The relative estimation error drops significantly before settling in at around 0.16.

Hyper-parameter optimization requires an ad-hoc approach for this case study. Simply using the marginal likelihood $p(\boldsymbol{f}_t^y | \boldsymbol{f}_{1:t-1}^y, \boldsymbol{f}_{1:t}^b)$ of the Kalman filter as before will result in

ever larger hyper-parameter values of the marginalized dimension, i.e. l_2. We conjecture that due to the marginalization of x_2, its impact on the marginal likelihood has been largely nullified, rendering l_2 unidentifiable. A remedy to this would be to use the Radon-transform and thereby treat this case study as a problem of tomographic type. Through a change of variables, known as input warping in the GP community, the Radon-transform would convolute the dimensions x_1 and x_2, and therefore also both length-scales l_1 and l_2, thereby preventing cancellation of l_2 through the measurement marginalization. See [22] for the use of GPs in tomographic inverse problems.

Instead, we successfully employed the function

$$
\begin{aligned}
&- \log p(\boldsymbol{f}_t^y, \boldsymbol{f}_t^b, \boldsymbol{f}_{t-1}) \\
&= \frac{1}{2} \begin{pmatrix} \boldsymbol{f}_t^y \\ \boldsymbol{f}_t^b \\ \boldsymbol{m}_{t-1} \end{pmatrix}^T \underbrace{\begin{bmatrix} \boldsymbol{K}_{t,t}^{f^y f^y} & \boldsymbol{K}_{t,t}^{f^y f^b} & \boldsymbol{K}_{t,t-1}^{ff} \\ \boldsymbol{K}_{t,t}^{f^b f^y} & \boldsymbol{K}_{t,t}^{f^b f^b} & \boldsymbol{K}_{t,t-1}^{f^b f} \\ \boldsymbol{K}_{t-1,t}^{ff^y} & \boldsymbol{K}_{t-1,t}^{ff^b} & \boldsymbol{K}_{t-1,t-1}^{ff} \end{bmatrix}^{-1}}_{=\boldsymbol{K}} \begin{pmatrix} \boldsymbol{f}_t^y \\ \boldsymbol{f}_t^b \\ \boldsymbol{m}_{t-1} \end{pmatrix} \\
&\quad + \frac{1}{2} \log(\det(\boldsymbol{K})) + \frac{N}{2} \log(2\pi) \quad (85)
\end{aligned}
$$

as the objective function to minimize. Here, we used the previous posterior mean $\boldsymbol{m}_{t-1}$ as fixed targets for $\boldsymbol{f}_{t-1}$. The number of data points is N. This objective function now contains both length-scales l_1 and l_2 explicitly and therefore works.

VI. DISCUSSION

In this section we elaborate on computational load, optimality, observability, and distinguishing features of the numerical Gaussian process Kalman filter.
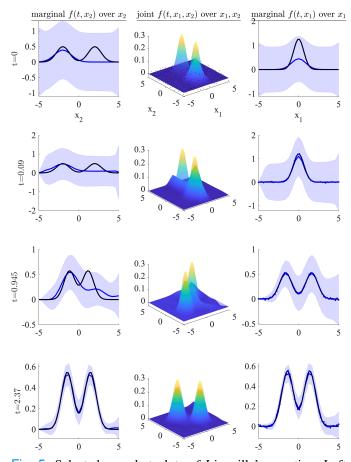
Fig. 5: Selected snapshot plots of Liouville's equation. Left and right plots show the marginal distributions $f(t, x_2)$ and $f(t, x_1)$, respectively. Posterior mean estimates of the marginals are shown in blue along with their 95 % confidence intervals (CI), while black curves show the analytical solution. Right plots also show noisy measurements of $f(t, x_1)$ as blue circles. Middle plots show the posterior mean estimate for the joint distribution $f(t, x_1, x_2)$, as well as the analytical solution (transparent). Uncertainty bands have been omitted. The initial estimate is based on the black dots (first row, middle plot). Online measurements of $f(t, x_1)$ shown as blue dots (right plots).

On computational cost

The main computational cost lies in the hyper-parameter optimization which requires the inversion of the measurement covariance matrix. This is inherent to GPs and scales cubically with the number of training points. There are however different approximations to reduce this to less than cubic scaling, see e.g. [23], [24].

Actual simulation of the model with numerical GPs, i.e. propagation through time, does not require numerical integration and instead only involves algebraic manipulations as given by (21). The two dimensional simulation ran for 111 min, while the one dimensional simulation ran for 9.5 min on a 2020 13" MacBook Pro $2, 3$ GHz quad-core i7.

In addition to the number of dimensions $\boldsymbol{x} \in \mathbb{R}^{d_x}$, computational load is also influenced by the prior covariance function and the PDE structure. Both can lead to less or more
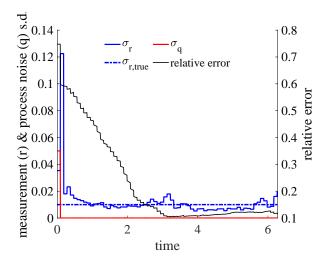

Fig. 6: Noise hyper-parameter estimation and relative estimation error for Liouville's equation.

difficult to handle contours of the NLML, thereby influencing the computational load via the hyper-parameter optimization. For example, we found simulations of the same PDE using a squared exponential covariance function to run faster and more robustly than using a neural network covariance function.

On optimality

For a linear and Gaussian system, the Kalman filter gives the optimal estimate of the state $\boldsymbol{m}_t^{opt}$ given the measurement history $\boldsymbol{n}_{1:t}^y$. For spatiotemporal systems, we also condition on the boundary history $\boldsymbol{n}_{1:t}^b$. Optimal is defined in this context as

$$\boldsymbol{m}_t^{opt}(\boldsymbol{n}_{1:t}^y, \boldsymbol{n}_{1:t}^b, \boldsymbol{\theta})$$
$$= \arg \min_{\boldsymbol{m}_t} \mathrm{E}\left[(\boldsymbol{n}_t - \boldsymbol{m}_t)^T (\boldsymbol{n}_t - \boldsymbol{m}_t) | \boldsymbol{n}_{1:t}^y, \boldsymbol{n}_{1:t}^b, \boldsymbol{\theta}\right]. \quad (86)$$

Here, the estimated state is $\boldsymbol{m}_t$ and the true (stochastic) state is $\boldsymbol{n}_t$. The hyper-parameters of the Gaussian process are denoted as $\boldsymbol{\theta}$. They influence the state estimate through i) the covariance function itself and ii) their estimated values.

Although working with spatiotemporal systems, the input space of numerical GPs is solely the spatial domain. Choosing an appropriate covariance function should therefore be based upon the solution shape with respect to spatial dimensions. For the step shaped one-dimensional advection equation for example, we chose the neural network covariance function due to its ability to handle discontinuities. In the original numerical GP work [4] the Burgers equation is solved with the neural network covariance function as well because its solution develops a shock front.

As showcased and argued in [6], multiple optima of the marginal likelihood can certainly exist, but are usually not a problem. Indeed, the different optima represent different interpretations of the data.

On observability

We showed that numerical Gaussian processes can be brought into probabilistic state space model form. By design

of numerical GPs, this state space model is structured by the underlying time-dependent partial differential equation. Showing observability of the PDE might therefore be transferable to the numerical GP state space model (34), (37).

Closely connected to observability is the concept of identifiability. This relates to the estimation of (hyper-)parameters. As seen and explained in the second case study involving Liouville's equation, an ad-hoc replacement for the marginal likelihood might be necessary to make hyper-parameter estimation work.

On distinguishing features

Kalman filters require fine-tuning before they deliver acceptable performance. This fine-tuning involves adjusting the process Q and measurement noise covariances R, as well as the initial error covariance matrix P_0.

Fine-tuning of numerical GPKFs is either much less tedious or not required at all. Once initialized with reasonable values, process and measurement noise level hyper-parameters σ_q and σ_r are adapted online to the measurement stream.

As for the initial error covariance matrix P_0, no manual fine-tuning was required for the case studies due to the probabilistic nature of GPs. P_0 was simply obtained by calculating the posterior GP variance (21) for given initial data. This was done without optimizing the hyper-parameters, as can be seen by the bad initial fit to the step function in Fig. 4a.

VII. Conclusions

In this article, we presented a new type of Kalman filter for spatiotemporal systems. Building on numerical Gaussian processes, we developed a probabilistic state space model that is linear and has Gaussian distributed states. These properties made it possible to derive the Kalman filter algorithm based on numerical Gaussian process state space models.

The resulting numerical Gaussian process Kalman filter was showcased with a step shaped one-dimensional advection equation where measurements were noisy point evaluations of the solution. In another case study Liouville's equation with integral measurements was studied. In both case studies the numerical GPKF was able to reconstruct the latent state and also estimate the true measurement noise levels through its hyper-parameters.

Appendix

Lemmata A and B are taken from [18].

A. Joint distribution of Gaussian variables

If random variables $x \in \mathbb{R}^{d_x}$ and $y \in \mathbb{R}^{d_y}$ have the Gaussian probability distributions

$$x \sim \mathrm{N}\left(m, P\right) \tag{87}$$

$$y|x \sim \mathrm{N}\left(Hx + u, R\right), \tag{88}$$

than the joint distribution of x, y and the marginal distribution of y are given as

$$\begin{pmatrix} x \\ y \end{pmatrix} \sim \mathrm{N}\left(\begin{pmatrix} m \\ Hm + u \end{pmatrix}, \begin{pmatrix} P & PH^T \\ HP & HPH^T + R \end{pmatrix}\right), \tag{89}$$

$$y \sim \mathrm{N}\left(Hm + u, HPH^T + R\right). \tag{90}$$

B. Conditional distribution of Gaussian variables

If the random variables x and y have the joint Gaussian probability distribution

$$\begin{pmatrix} x \\ y \end{pmatrix} \sim \mathrm{N}\left(\begin{pmatrix} a \\ b \end{pmatrix}, \begin{pmatrix} A & C \\ C^T & B \end{pmatrix}\right), \tag{91}$$

than the conditional distribution is

$$x|y \sim \mathrm{N}\left(a + CB^{-1}(y - b), A - CB^{-1}C^T\right). \tag{92}$$

C. Numerical Gaussian process state space model precursor

We can derive the covariance functions of the joint model (33), here shown using the implicit Euler, as

$$K_{GP} = \mathrm{E}\left[\begin{pmatrix} f_t(x) \\ f_t^y(y) \\ f^b t(x_b) \\ f_{t-1}(x) \end{pmatrix} \begin{pmatrix} f_t(x') \\ f_t^y(y') \\ f_t^b(x_b') \\ f_{t-1}(x') \end{pmatrix}^T\right] \tag{93}$$

$$= \mathrm{E}\left[\begin{pmatrix} f_t \\ \mathcal{H}_x f_t + r_t \\ \mathcal{B}_x f_t \\ \mathcal{A}_x f_t - \Delta t q_{t-1} \end{pmatrix} \begin{pmatrix} f_t \\ \mathcal{H}_{x'} f_t + r_t \\ \mathcal{B}_{x'} f_t \\ \mathcal{A}_{x'} f_t - \Delta t q_{t-1} \end{pmatrix}^T\right]$$

$$= \begin{bmatrix} k & \mathcal{H}_{x'} k & \mathcal{B}_{x'} k & \mathcal{A}_{x'} k \\ \mathcal{H}_x k & \mathcal{H}_x \mathcal{H}_{x'} k + k_{t,t}^{rr} & \mathcal{H}_x \mathcal{B}_{x'} k & \mathcal{H}_x \mathcal{A}_{x'} k \\ \mathcal{B}_x k & \mathcal{B}_x \mathcal{H}_{x'} k & \mathcal{B}_x \mathcal{B}_{x'} k & \mathcal{B}_x \mathcal{A}_{x'} k \\ \mathcal{A}_x k & \mathcal{A}_x \mathcal{H}_{x'} k & \mathcal{A}_x \mathcal{B}_{x'} k & \mathcal{A}_x \mathcal{A}_{x'} k + \Delta t^2 k_{t,t}^{qq} \end{bmatrix}. \tag{94}$$

Here, the prior covariance function $k_{t,t}^{ff}$ has been abbreviated as k.

D. Partial derivatives of the KF NLML

We had

$$-\frac{\partial}{\partial \theta_j} \log p(\boldsymbol{f}_t^y | \boldsymbol{f}_{1:t-1}^y, \boldsymbol{f}_{1:t}^b)$$

$$= \frac{1}{2} \frac{\partial}{\partial \theta_j} \left(\boldsymbol{f}_t^y - \boldsymbol{C}_t \boldsymbol{m}_t^-\right)^T \boldsymbol{S}_t^{-1} \left(\boldsymbol{f}_t^y - \boldsymbol{C}_t \boldsymbol{m}_t^-\right)$$

$$+ \frac{1}{2} \frac{\partial}{\partial \theta_j} \log(\det(\boldsymbol{S}_t)). \tag{95}$$

Expanding the first term and then using the product rule on it to evaluate all necessary partial derivatives gives

$$\frac{\partial}{\partial \theta_j} \left(\boldsymbol{f}_t^y - \boldsymbol{C}_t \boldsymbol{m}_t^- \right)^T \boldsymbol{S}_t^{-1} \left(\boldsymbol{f}_t^y - \boldsymbol{C}_t \boldsymbol{m}_t^- \right)$$

$$= \frac{\partial}{\partial \theta_j} \left((\boldsymbol{n}_t^y)^T \boldsymbol{S}_t^{-1} \boldsymbol{f}_t^y - (\boldsymbol{f}_t^y)^T \boldsymbol{S}_t^{-1} \boldsymbol{C}_t \boldsymbol{m}_t^- \right.$$

$$\left. - (\boldsymbol{C}_t \boldsymbol{m}_t^-)^T \boldsymbol{S}_t^{-1} \boldsymbol{f}_t^y + (\boldsymbol{C}_t \boldsymbol{m}_t^-)^T \boldsymbol{S}_t^{-1} \boldsymbol{C}_t \boldsymbol{m}_t^- \right)$$

$$= (\boldsymbol{f}_t^y)^T \frac{\partial \boldsymbol{S}_t^{-1}}{\partial \theta_j} \boldsymbol{f}_t^y$$

$$- (\boldsymbol{f}_t^y)^T \left(\frac{\partial \boldsymbol{S}_t^{-1}}{\partial \theta_j} \boldsymbol{C}_t \boldsymbol{m}_t^- + \boldsymbol{S}_t^{-1} \frac{\partial \boldsymbol{C}_t}{\partial \theta_j} \boldsymbol{m}_t^- + \boldsymbol{S}_t^{-1} \boldsymbol{C}_t \frac{\partial \boldsymbol{m}_t^-}{\partial \theta_j} \right)$$

$$- \left(\frac{\partial \left(\boldsymbol{m}_t^- \right)^T}{\partial \theta_j} \boldsymbol{C}_t^T \boldsymbol{S}_t^{-1} + (\boldsymbol{m}_t^-)^T \frac{\partial \boldsymbol{C}_t^T}{\partial \theta_j} \boldsymbol{S}_t^{-1} \right) \boldsymbol{f}_t^y$$

$$- (\boldsymbol{m}_t^-)^T \boldsymbol{C}_t^T \frac{\partial \boldsymbol{S}_t^{-1}}{\partial \theta_j} \boldsymbol{f}_t^y$$

$$+ \frac{\partial (\boldsymbol{m}_t^-)^T}{\partial \theta_j} \boldsymbol{C}_t^T \boldsymbol{S}_t^{-1} \boldsymbol{C}_t \boldsymbol{m}_t^- + (\boldsymbol{m}_t^-)^T \frac{\partial \boldsymbol{C}_t^T}{\partial \theta_j} \boldsymbol{S}_t^{-1} \boldsymbol{C}_t \boldsymbol{m}_t^-$$

$$+ (\boldsymbol{m}_t^-)^T \boldsymbol{C}_t^T \frac{\partial \boldsymbol{S}_t^{-1}}{\partial \theta_j} \boldsymbol{C}_t \boldsymbol{m}_t^- + (\boldsymbol{m}_t^-)^T \boldsymbol{C}_t^T \boldsymbol{S}_t^{-1} \frac{\partial \boldsymbol{C}_t}{\partial \theta_j} \boldsymbol{m}_t^-$$

$$+ (\boldsymbol{m}_t^-)^T \boldsymbol{C}_t^T \boldsymbol{S}_t^{-1} \boldsymbol{C}_t \frac{\partial \boldsymbol{m}_t^-}{\partial \theta_j}. \tag{96}$$

The second term in (95) can be calculated as

$$\frac{\partial}{\partial \theta_j} \log(\det(\boldsymbol{S}_t)) = \mathrm{tr} \left(\boldsymbol{S}_t^{-1} \frac{\partial \boldsymbol{S}_t}{\partial \theta_j} \right), \tag{97}$$

while the partial derivative of $\boldsymbol{S}_t^{-1}$ can be calculated as

$$\frac{\partial \boldsymbol{S}_t^{-1}}{\partial \theta_j} = -\boldsymbol{S}_t^{-1} \frac{\partial \boldsymbol{S}_t}{\partial \theta_j} \boldsymbol{S}_t^{-1}, \tag{98}$$

see [6] for both identities.

We still need the partial derivatives of $\boldsymbol{S}_t$, $\boldsymbol{C}_t$, and $\boldsymbol{m}_t^-$. Starting with the prior mean we have

$$\frac{\partial \boldsymbol{m}_t^-}{\partial \theta_j} = \frac{\partial}{\partial \theta_j} \boldsymbol{A}_t \underbrace{\left(\boldsymbol{m}_{t-1} + \boldsymbol{P}_{t-1} \left(\boldsymbol{A}_t^{f^b} \right)^T \left(\boldsymbol{S}_t^{f^b} \right)^{-1} \left(\boldsymbol{f}_t^b - \boldsymbol{A}_t^{f^b} \boldsymbol{m}_{t-1} \right) \right)}_{=\tilde{\boldsymbol{m}}_{t-1}}$$

$$= \frac{\partial \boldsymbol{A}_t}{\partial \theta_j} \tilde{\boldsymbol{m}}_{t-1} + \boldsymbol{A}_t \frac{\partial \tilde{\boldsymbol{m}}_{t-1}}{\partial \theta_j}, \tag{99}$$

with

$$\frac{\partial \boldsymbol{A}_t}{\partial \theta_j}$$

$$= \frac{\partial}{\partial \theta_j} \begin{bmatrix} \boldsymbol{K}_{t,t}^{ff^b} & \boldsymbol{K}_{t,t}^{ff} \end{bmatrix} \begin{bmatrix} \boldsymbol{K}_{t,t}^{f^b f^b} & \boldsymbol{K}_{t,t-1}^{f^b f} \\ \boldsymbol{K}_{t-1,t}^{ff^b} & \boldsymbol{K}_{t-1,t-1}^{ff} \end{bmatrix}^{-1}$$

$$= \begin{bmatrix} \dfrac{\partial \boldsymbol{K}_{t,t}^{ff^b}}{\partial \theta_j} & \dfrac{\partial \boldsymbol{K}_{t,t-1}^{ff}}{\partial \theta_j} \end{bmatrix} \begin{bmatrix} \boldsymbol{K}_{t,t}^{f^b f^b} & \boldsymbol{K}_{t,t-1}^{f^b f} \\ \boldsymbol{K}_{t-1,t}^{ff^b} & \boldsymbol{K}_{t-1,t-1}^{ff} \end{bmatrix}^{-1}$$

$$+ \begin{bmatrix} \boldsymbol{K}_{t,t}^{f^b} & \boldsymbol{K}_{t,t-1}^{ff} \end{bmatrix} \begin{bmatrix} \boldsymbol{K}_{t,t}^{f^b f^b} & \boldsymbol{K}_{t,t-1}^{f^b f} \\ \boldsymbol{K}_{t-1,t}^{ff^b} & \boldsymbol{K}_{t-1,t-1}^{ff} \end{bmatrix}^{-1}$$

$$\times \begin{bmatrix} \dfrac{\partial \boldsymbol{K}_{t,t}^{f^b f^b}}{\partial \theta_j} & \dfrac{\partial \boldsymbol{K}_{t,t-1}^{f^b f}}{\partial \theta_j} \\ \dfrac{\partial \boldsymbol{K}_{t-1,t}^{ff^b}}{\theta_j} & \dfrac{\partial \boldsymbol{K}_{t-1,t-1}^{ff}}{\theta_j} \end{bmatrix} \begin{bmatrix} \boldsymbol{K}_{t,t}^{f^b f^b} & \boldsymbol{K}_{t,t-1}^{f^b f} \\ \boldsymbol{K}_{t-1,t}^{ff^b} & \boldsymbol{K}_{t-1,t-1}^{ff} \end{bmatrix}^{-1}, \tag{100}$$

$$\frac{\partial \tilde{\boldsymbol{m}}_{t-1}}{\partial \theta_j} = \left[\boldsymbol{0}, \left(\boldsymbol{P}_{t-1} \frac{\partial \left(\boldsymbol{A}_t^{f^b} \right)^T}{\partial \theta_j} \left(\boldsymbol{S}_t^{f^b} \right)^{-1} \right. \right.$$

$$\left. + \boldsymbol{P}_{t-1} \left(\boldsymbol{A}_t^{f^b} \right)^T \frac{\partial \left(\boldsymbol{S}_t^{f^b} \right)^{-1}}{\partial \theta_j} \right)$$

$$\times \left(\boldsymbol{f}_t^b - \boldsymbol{A}_t^{f^b} \boldsymbol{m}_{t-1} \right)$$

$$\left. - \boldsymbol{P}_{t-1} \left(\boldsymbol{A}_t^{f^b} \right)^T \left(\boldsymbol{S}_t^{f^b} \right)^{-1} \frac{\partial \boldsymbol{A}_t^{f^b}}{\partial \theta_j} \boldsymbol{m}_{t-1} \right]^T, \tag{101}$$

$$\frac{\partial \boldsymbol{A}_t^{f^b}}{\partial \theta_j} = \frac{\partial}{\partial \theta_j} \left(\boldsymbol{K}_{t,t-1}^{f^b f} \left(\boldsymbol{K}_{t-1,t-1}^{ff} \right)^{-1} \right)$$

$$= \frac{\partial \boldsymbol{K}_{t,t-1}^{f^b f}}{\partial \theta_j} \left(\boldsymbol{K}_{t-1,t-1}^{ff} \right)^{-1} + \boldsymbol{K}_{t,t-1}^{f^b f} \frac{\partial \left(\boldsymbol{K}_{t-1,t-1}^{ff} \right)^{-1}}{\partial \theta_j}, \tag{102}$$

$$\frac{\partial \boldsymbol{S}_t^{f^b}}{\partial \theta_j} = \frac{\partial \boldsymbol{A}_t^{f^b}}{\partial \theta_j} \boldsymbol{P}_{t-1} \left(\boldsymbol{A}_t^{f^b} \right)^T + \frac{\partial \boldsymbol{P}_t^{GP,f^b}}{\partial \theta_j}$$

$$+ \boldsymbol{A}_t^{f^b} \boldsymbol{P}_{t-1} \frac{\partial \left(\boldsymbol{A}_t^{f^b} \right)^T}{\partial \theta_j}, \tag{103}$$

$$\frac{\partial \boldsymbol{P}_t^{GP,f^b}}{\partial \theta_j} = \frac{\partial \boldsymbol{K}_{t,t}^{f^b f^b}}{\partial \theta_j} - \frac{\partial \boldsymbol{A}_t^{f^b}}{\partial \theta_j} \boldsymbol{K}_{t-1,t}^{ff^b} - \boldsymbol{A}_t^{f^b} \frac{\partial \boldsymbol{K}_{t-1,t}^{ff^b}}{\partial \theta_j}. \tag{104}$$

For the measurement matrix we have

$$\frac{\partial \boldsymbol{C}_t}{\partial \theta_j} = \frac{\partial \boldsymbol{K}_{tt}^{f^y f} \left(\boldsymbol{K}_{tt}^{ff} \right)^{-1}}{\partial \theta_j}$$

$$= \frac{\partial \boldsymbol{K}_{tt}^{f^y f}}{\partial \theta_j} \left(\boldsymbol{K}_{tt}^{ff} \right)^{-1} + \boldsymbol{K}_{tt}^{f^y f} \frac{\partial \left(\boldsymbol{K}_{tt}^{ff} \right)^{-1}}{\partial \theta_j}$$

$$= \frac{\partial \boldsymbol{K}_{tt}^{f^y f}}{\partial \theta_j} \left(\boldsymbol{K}_{tt}^{ff} \right)^{-1} - \boldsymbol{K}_{tt}^{f^y f} \left(\boldsymbol{K}_{tt}^{ff} \right)^{-1} \frac{\partial \boldsymbol{K}_{tt}^{ff}}{\partial \theta_j} \left(\boldsymbol{K}_{tt}^{ff} \right)^{-1}. \tag{105}$$

The covariance matrix of the innovation is a bit more elaborate, we have

$$\frac{\partial \boldsymbol{S}_t}{\partial \theta_j} = \frac{\partial \boldsymbol{C}_t \boldsymbol{P}_t^- \boldsymbol{C}_t^T + \boldsymbol{P}_t^{GP,f^y f^y}}{\partial \theta_j}$$

$$= \frac{\partial \boldsymbol{C}_t}{\partial \theta_j} \boldsymbol{P}_t^- \boldsymbol{C}_t^T + \boldsymbol{C}_t \frac{\partial \boldsymbol{P}_t^-}{\partial \theta_j} \boldsymbol{C}_t^T$$

$$+ \boldsymbol{C}_t \boldsymbol{P}_t^- \frac{\partial \boldsymbol{C}_t^T}{\partial \theta_j} + \frac{\partial \boldsymbol{P}_t^{GP,f^y}}{\partial \theta_j}. \tag{106}$$

The partial derivatives in the second and last term need to be derived. For the prior covariance matrix of the state error we get

$$
\frac{\partial \boldsymbol{P}_t^-}{\partial \theta_j}
$$

$$
= \frac{\partial \boldsymbol{A}_t}{\partial \theta_j} \tilde{\boldsymbol{P}}_{t-1} \boldsymbol{A}_t^T + \boldsymbol{A}_t \tilde{\boldsymbol{P}}_{t-1} \frac{\partial \boldsymbol{A}_t^T}{\partial \theta_j} + \frac{\partial \boldsymbol{P}_t^{GP,f}}{\partial \theta_j}
$$

$$
- \boldsymbol{A}_t \begin{bmatrix} \dfrac{\partial \boldsymbol{S}_t^{f^b}}{\partial \theta_j} \\[2ex] \boldsymbol{P}_{t-1} \dfrac{\partial \left(\boldsymbol{A}_t^{f^b}\right)^T}{\partial \theta_j} \end{bmatrix} \left(\boldsymbol{S}_t^{f^b}\right)^{-1} \begin{bmatrix} \boldsymbol{S}_t^{f^b} \\[1ex] \boldsymbol{P}_{t-1}\left(\boldsymbol{A}_t^{f^b}\right)^T \end{bmatrix}^T \boldsymbol{A}_t^T
$$

$$
- \boldsymbol{A}_t \begin{bmatrix} \boldsymbol{S}_t^{f^b} \\[1ex] \boldsymbol{P}_{t-1}\left(\boldsymbol{A}_t^{f^b}\right)^T \end{bmatrix} \frac{\partial \left(\boldsymbol{S}_t^{f^b}\right)^{-1}}{\partial \theta_j} \begin{bmatrix} \boldsymbol{S}_t^{f^b} \\[1ex] \boldsymbol{P}_{t-1}\left(\boldsymbol{A}_t^{f^b}\right)^T \end{bmatrix}^T \boldsymbol{A}_t^T
$$

$$
- \boldsymbol{A}_t \begin{bmatrix} \boldsymbol{S}_t^{f^b} \\[1ex] \boldsymbol{P}_{t-1}\left(\boldsymbol{A}_t^{f^b}\right)^T \end{bmatrix} \left(\boldsymbol{S}_t^{f^b}\right)^{-1} \begin{bmatrix} \dfrac{\partial \boldsymbol{S}_t^{f^b}}{\partial \theta_j} \\[2ex] \boldsymbol{P}_{t-1} \dfrac{\partial \left(\boldsymbol{A}_t^{f^b}\right)^T}{\partial \theta_j} \end{bmatrix}^T \boldsymbol{A}_t^T
$$

$$
- \boldsymbol{A}_t \begin{bmatrix} \boldsymbol{S}_t^{f^b} \\[1ex] \boldsymbol{P}_{t-1}\left(\boldsymbol{A}_t^{f^b}\right)^T \end{bmatrix} \left(\boldsymbol{S}_t^{f^b}\right)^{-1} \begin{bmatrix} \boldsymbol{S}_t^{f^b} \\[1ex] \boldsymbol{P}_{t-1}\left(\boldsymbol{A}_t^{f^b}\right)^T \end{bmatrix}^T \frac{\partial \boldsymbol{A}_t^T}{\partial \theta_j}.
\tag{107}
$$

The partial derivatives of the dynamic matrix are already covered in (99). The partial derivative of the inherent GP uncertainty is

$$
\frac{\partial \boldsymbol{P}_t^{GP,f}}{\partial \theta_j}
$$

$$
= \frac{\partial \boldsymbol{K}_{t,t}^{ff} - \boldsymbol{K}_{t,t-1}^{ff} \left(\boldsymbol{K}_{t-1,t-1}^{ff}\right)^{-1} \boldsymbol{K}_{t-1,t}^{ff}}{\partial \theta_j}
$$

$$
= \frac{\partial \boldsymbol{K}_{t,t}^{ff}}{\partial \theta_j} - \frac{\partial \boldsymbol{K}_{t,t-1}^{ff}}{\partial \theta_j} \left(\boldsymbol{K}_{t-1,t-1}^{ff}\right)^{-1} \boldsymbol{K}_{t-1,t}^{ff}
$$

$$
+ \boldsymbol{K}_{t,t-1}^{ff} \left(\boldsymbol{K}_{t-1,t-1}^{ff}\right)^{-1} \frac{\partial \boldsymbol{K}_{t-1,t-1}^{ff}}{\partial \theta_j} \left(\boldsymbol{K}_{t-1,t-1}^{ff}\right)^{-1} \boldsymbol{K}_{t-1,t}^{ff}
$$

$$
- \boldsymbol{K}_{t,t-1}^{ff} \left(\boldsymbol{K}_{t-1,t-1}^{ff}\right)^{-1} \frac{\partial \boldsymbol{K}_{t-1,t}^{ff}}{\partial \theta_j}.
\tag{108}
$$

The partial derivative of the inherent GP uncertainty for the output values (last term of 106) is

$$
\frac{\partial \boldsymbol{P}_t^{GP,f^y}}{\partial \theta_j} = \frac{\partial \boldsymbol{K}_{t,t}^{f^y f^y} - \boldsymbol{K}_{t,t}^{f^y f} \left(\boldsymbol{K}_{t,t}^f\right)^{-1} \boldsymbol{K}_{t,t}^{ff^y}}{\partial \theta_j}
$$

$$
= \frac{\partial \boldsymbol{K}_{t,t}^{f^y f^y}}{\partial \theta_j} - \frac{\partial \boldsymbol{K}_{t,t}^{f^y f}}{\partial \theta_j} \left(\boldsymbol{K}_{t,t}^{ff}\right)^{-1} \boldsymbol{K}_{t,t}^{ff^y}
$$

$$
+ \boldsymbol{K}_{t,t}^{f^y f} \left(\boldsymbol{K}_{t,t}^{ff}\right)^{-1} \frac{\partial \boldsymbol{K}_{t,t}^{ff}}{\partial \theta_j} \left(\boldsymbol{K}_{t,t}^{ff}\right)^{-1} \boldsymbol{K}_{t,t}^{ff^y}
$$

$$
- \boldsymbol{K}_{t,t}^{f^y f} \left(\boldsymbol{K}_{t,t}^{ff}\right)^{-1} \frac{\partial \boldsymbol{K}_{t,t}^{ff^y}}{\partial \theta_j}.
\tag{109}
$$

REFERENCES

[1] R. E. Kalman, "A New Approach to Linear Filtering and Prediction Problems," *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, 03 1960. [Online]. Available: https://doi.org/10.1115/1.3662552

[2] R. F. Curtain and H. Zwart, *An introduction to infinite-dimensional linear systems theory*. Springer Science & Business Media, 2012, vol. 21.

[3] J. H. Ferziger, M. Perić, and R. L. Street, *Computational methods for fluid dynamics*. Springer, 2002, vol. 3.

[4] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Numerical Gaussian processes for time-dependent and nonlinear partial differential equations," *SIAM Journal on Scientific Computing*, vol. 40, no. 1, pp. A172–A198, 2018.

[5] A. Küper, N. Totis, and S. Waldherr, "Comparing cell population balance model simulation through Gaussian processes and discretisation," in *Proceedings of the 24th International Symposium on Mathematical Theory of Networks and Systems [not yet published]*. IFAC, 2020.

[6] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*. MIT Press Cambridge, MA, 2006, vol. 2, no. 3.

[7] A. Küper and S. Waldherr, "Numerical Gaussian process Kalman filtering," *IFAC-PapersOnLine*, vol. 53, no. 2, pp. 11 416–11 421, 2020, 21th IFAC World Congress. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2405896320308788

[8] S. Särkkä and J. Hartikainen, "Infinite-dimensional Kalman filtering approach to spatio-temporal Gaussian process regression," in *International Conference on Artificial Intelligence and Statistics*, 2012, pp. 993–1001.

[9] S. Särkkä, A. Solin, and J. Hartikainen, "Spatiotemporal learning via infinite-dimensional Bayesian filtering and smoothing: A look at Gaussian process regression through Kalman filtering," *IEEE Signal Processing Magazine*, vol. 30, no. 4, pp. 51–61, 2013.

[10] M. Todescato, A. Carron, R. Carli, G. Pillonetto, and L. Schenato, "Efficient spatio-temporal Gaussian regression via Kalman filtering," *Automatica*, vol. 118, p. 109032, 2020.

[11] J. Ko and D. Fox, "GP-BayesFilters: Bayesian filtering using Gaussian process prediction and observation models," *Autonomous Robots*, vol. 27, no. 1, pp. 75–90, 2009.

[12] M. P. Deisenroth, M. F. Huber, and U. D. Hanebeck, "Analytic moment-based Gaussian process filtering," in *Proceedings of the 26th annual international conference on machine learning*. ACM, 2009, pp. 225–232.

[13] M. P. Deisenroth and H. Ohlsson, "A general perspective on Gaussian filtering and smoothing: Explaining current and deriving new algorithms," in *Proceedings of the 2011 American Control Conference*. IEEE, 2011, pp. 1807–1812.

[14] S. Särkkä, "Linear operators and stochastic partial differential equations in Gaussian process regression," in *International Conference on Artificial Neural Networks*. Springer, 2011, pp. 151–158.

[15] S. Särkkä and A. Solin, *Applied stochastic differential equations*. Cambridge University Press, 2019, vol. 10.

[16] R. Curtain, "A survey of infinite-dimensional filtering," *Siam Review*, vol. 17, no. 3, pp. 395–411, 1975.

[17] D. Simon, *Optimal state estimation: Kalman, H infinity, and nonlinear approaches*. John Wiley & Sons, 2006.

[18] S. Särkkä, *Bayesian filtering and smoothing*. Cambridge University Press, 2013, vol. 3.

[19] R. Fitzgerald, "Divergence of the Kalman filter," *IEEE Transactions on Automatic Control*, vol. 16, no. 6, pp. 736–747, 1971.

[20] S. Zeng, "Sample-based population observers," *Automatica*, vol. 101, pp. 166–174, 2019.

[21] S. Zeng, S. Waldherr, C. Ebenbauer, and F. Allgöwer, "Ensemble observability of linear systems," *IEEE Transactions on Automatic Control*, vol. 61, no. 6, pp. 1452–1465, 2015.

[22] Z. Purisha, C. Jidling, N. Wahlström, T. B. Schön, and S. Särkkä, "Probabilistic approach to limited-data computed tomography reconstruction," *Inverse Problems*, vol. 35, no. 10, p. 105004, 2019.

[23] E. Snelson and Z. Ghahramani, "Sparse Gaussian processes using pseudo-inputs," *Advances in Neural Information Processing Systems*, vol. 18, pp. 1259–1266, 2006.

[24] A. Solin and S. Särkkä, "Hilbert space methods for reduced-rank Gaussian process regression," *Statistics and Computing*, vol. 30, no. 2, pp. 419–446, 2020.