

Numerical Gaussian process Kalman filtering

Armin Küper and Steffen Waldherr*

* *Department of Chemical Engineering, KU Leuven, Leuven, Belgium*
(e-mail: {armin.kuper, steffen.waldherr}@kuleuven.be).

Abstract: In this manuscript we introduce numerical Gaussian process Kalman filtering (GPKF). Numerical Gaussian processes have recently been developed to simulate spatiotemporal models. The contribution of this paper is to embed numerical Gaussian processes into the recursive Kalman filter equations. This embedding enables us to do Kalman filtering on infinite-dimensional systems using Gaussian processes. This is possible because i) we are obtaining a linear model from numerical Gaussian processes, and ii) the states of this model are by definition Gaussian distributed random variables. Convenient properties of the numerical GPKF are that no spatial discretization of the model is necessary, and manual setting up of the Kalman filter, that is fine-tuning the process and measurement noise levels by hand is not required, as they are learned online from the data stream. We showcase the capability of the numerical GPKF in a simulation study of the advection equation.

Keywords: Kalman filters, Gaussian processes, infinite-dimensional systems, state estimation.

1. INTRODUCTION

Monitoring physical, chemical, or biological systems requires measurements obtained from sensors. These measurements are subjected to stochastic noise. In addition, technical or financial restrictions might prevent measuring all properties of interest. State estimation methods help to overcome both of these obstacles. They filter signal from noise, and reconstruct latent properties by leveraging mathematical models of the process at hand. The most known state estimator is the Kalman filter, see Kalman (1960). It gives the optimal, in an expected squared error sense, state estimate for a linear system that is subjected to additive Gaussian noise. State estimation methods for infinite-dimensional systems, that is systems described through partial differential equations (PDEs), are not that well established. Usually, the PDEs are spatially discretized into large ODE systems that can be incorporated into the recursive Kalman filter equations.

In this work, we take a different route founded on Gaussian process (GP) regression, see Williams and Rasmussen (2006), and the recently introduced numerical GPs by Raissi et al. (2018). These are machine learning methods that unlike in a first principles derivation of a state space model, work with non-parametric models that adapt to the data. A GP is a Gaussian distribution over functions and it is fully defined by a mean function and a covariance function. When doing regression, a GP prior is placed on the regressor function. The posterior distribution of the GP is then calculated by conditioning on measurements and estimating the so-called hyper-parameters of the covariance function. An elegant extension to traditional GP regression to solve (non-)linear time-dependent PDEs, are numerical

GPs. In this method the PDE is first temporally discretized, hence the name numerical. From this a multi-output GP is formulated whose structure, i.e. covariance functions, is informed by the time-discretized PDE. Propagation through time is then simply done by iteratively forming the conditional distribution of this multi-output GP without having to do numerical integration.

Our contribution lies in embedding numerical GP regression into the recursive Kalman filter equations. This is possible because the outputs of the numerical GP are, by definition, Gaussian distributed, and they relate linearly to each other through the time-discretized PDE. We therefore obtain a Kalman filter for infinite-dimensional systems whose underlying process model is a first principles structured Gaussian process. Tedious manual tuning of the Kalman filter is not necessary anymore as the process and measurement noise variance are learned dynamically on the data stream.

This paper is structured as follows: we start by introducing Kalman filtering and (numerical) GP regression in the methods section. Next, we show how numerical GP regression can be embedded into the recursive Kalman filter equations to obtain the numerical Gaussian process Kalman filter (GPKF). Before giving a conclusion, we showcase the numerical GPKF using a simulation case study of the one-dimensional advection equation.

Related works: Infinite-dimensional Kalman filtering via Gaussian process regression has been done previously by Särkkä and Hartikainen (2012) and Särkkä et al. (2013). Therein, spatiotemporal GP regression setups are converted into infinite-dimensional state space models. These can then be used for infinite-dimensional Kalman filtering. In numerical Gaussian process regression, the spatiotemporal model is first discretized in time and then formulated as a multi-output GP with the spatial coordinates as inputs.

* The authors would like to thank the reviewers for their helpful comments.

Finite-dimensional filtering with Gaussian processes has been done, among others, by Deisenroth et al. (2009), Deisenroth and Ohlsson (2011), and Ko and Fox (2009).

2. METHODS

We give a brief introduction to Kalman filtering and (numerical) Gaussian process regression. We use the respective common notations for Kalman filtering (state space models) and GP regression. Both sections should therefore be read as separate entities. The notation of numerical GP regression is oriented after the original work of Raissi et al. (2018) where the focus lies on spatiotemporal models. In Section 3 we will unify both frameworks.

2.1 Kalman filtering

Kalman filtering can be used for stochastic filtering of a signal from noisy online measurements. Furthermore, it can also be employed to reconstruct non-measurable states. It does so by using a probabilistic process and measurement model. In case of linear process dynamics and an identical, independent Gaussian distributed noise acting on the process and measurements, the Kalman filter gives the optimal estimate of the states with respect to the expected squared error.

What follows is the Kalman filter from a Bayesian perspective as presented in Särkkä (2013).

Suppose the process and measurement equations are given as a state space model

$$\begin{aligned} \mathbf{x}_t &= \mathbf{A}\mathbf{x}_{t-1} + \mathbf{q}_{t-1} \\ \mathbf{y}_t &= \mathbf{C}\mathbf{x}_t + \mathbf{r}_t. \end{aligned} \quad (1)$$

Here, $\mathbf{x}_t \in \mathbb{R}^{d_x}$ is the state at time t and $\mathbf{y}_t \in \mathbb{R}^{d_y}$ is the measurement at time t . Process dynamics of the model are given by $\mathbf{A} \in \mathbb{R}^{d_x \times d_x}$ and the measurement model matrix is $\mathbf{C} \in \mathbb{R}^{d_y \times d_x}$. Process noise $\mathbf{q}_{t-1} \in \mathbb{R}^{d_x}$ and measurement noise $\mathbf{r}_t \in \mathbb{R}^{d_y}$ are both modeled as white, additive Gaussian noise with diagonal covariance matrices $\mathbf{Q} \in \mathbb{R}^{d_x \times d_x}$ and $\mathbf{R} \in \mathbb{R}^{d_y \times d_y}$, respectively. We formulate (1) as probability density functions from which the Bayesian viewpoint of Kalman filtering naturally arises

$$\begin{aligned} p(\mathbf{x}_t|\mathbf{x}_{t-1}) &= \mathcal{N}(\mathbf{x}_t|\mathbf{A}\mathbf{x}_{t-1}, \mathbf{Q}), \\ p(\mathbf{y}_t|\mathbf{x}_t) &= \mathcal{N}(\mathbf{y}_t|\mathbf{C}\mathbf{x}_t, \mathbf{R}). \end{aligned} \quad (2)$$

Here the first equation describes the stochastic dynamics of the system while the second equation gives the distribution of the current measurement given the current state. These distributions are Gaussian for all times because the noise terms are Gaussian random variables, and the process and measurement equations are linear.

The goal in Bayesian filtering is to compute the marginal posterior distribution of the state \mathbf{x}_t at each time step given all measurements up to the current time step. The term marginal refers here to the marginalization over the previous state \mathbf{x}_{t-1} . Using Bayes' theorem we have

$$p(\mathbf{x}_t|\mathbf{y}_{1:t}) = \frac{p(\mathbf{y}_t|\mathbf{x}_t)p(\mathbf{x}_t|\mathbf{y}_{1:t-1})}{\int p(\mathbf{y}_t|\mathbf{x}_t)p(\mathbf{x}_t|\mathbf{y}_{1:t-1})d\mathbf{x}_t}. \quad (3)$$

With a constant data stream of measurements, this quickly becomes intractable as the measurement history grows ever larger. The Kalman filter circumvents this problem

by solving this equation recursively starting from a prior mean \mathbf{m}_0 and covariance \mathbf{P}_0 . The predictive, posterior, and normalizing distribution of (3) can be calculated in closed form. They are

$$p(\mathbf{x}_t|\mathbf{y}_{1:t-1}) = \mathcal{N}(\mathbf{x}_t|\mathbf{m}_t^-, \mathbf{P}_t^-), \quad (4)$$

$$p(\mathbf{x}_t|\mathbf{y}_{1:t}) = \mathcal{N}(\mathbf{x}_t|\mathbf{m}_t, \mathbf{P}_t), \text{ and} \quad (5)$$

$$p(\mathbf{y}_t|\mathbf{y}_{1:t-1}) = \mathcal{N}(\mathbf{y}_t|\mathbf{C}\mathbf{m}_t^-, \mathbf{S}_t), \quad (6)$$

respectively. Moments of the above distributions are calculated in a prediction step

$$\mathbf{m}_t^- = \mathbf{A}\mathbf{m}_{t-1}, \quad (7)$$

$$\mathbf{P}_t^- = \mathbf{A}\mathbf{P}_{t-1}\mathbf{A}^T + \mathbf{Q}, \quad (8)$$

and an update step

$$\begin{aligned} \mathbf{v}_t &= \mathbf{y}_t - \mathbf{C}\mathbf{m}_t^-, & \mathbf{m}_t &= \mathbf{m}_t^- + \mathbf{K}_t\mathbf{v}_t, \\ \mathbf{S}_t &= \mathbf{C}\mathbf{P}_t^-\mathbf{C}^T + \mathbf{R}, & \mathbf{P}_t &= \mathbf{P}_t^- - \mathbf{K}_t\mathbf{S}_t\mathbf{K}_t^T. \\ \mathbf{K}_t &= \mathbf{P}_t^-\mathbf{C}^T\mathbf{S}_t^{-1}, \end{aligned} \quad (9)$$

One can derive these equations by formulating the required conditional, joint, and marginal distributions of the states and measurements. This procedure can be found in Särkkä (2013) and we will use it later to derive the numerical Gaussian process Kalman filter equations.

2.2 Gaussian process regression

Here we give a brief introduction into Gaussian process regression where we take the function-space view presented in Williams and Rasmussen (2006). Afterwards we explain how numerical GP regression by Raissi et al. (2018) is used to solve time-dependent partial differential equations.

A GP is a Gaussian distribution over a random function $f(\mathbf{x})$. It is fully defined through its mean function $m(\mathbf{x})$ and covariance function $k(\mathbf{x}, \mathbf{x}')$

$$\begin{aligned} m(\mathbf{x}) &= \mathbb{E}[f(\mathbf{x})] \\ k(\mathbf{x}, \mathbf{x}') &= \mathbb{E}[(f(\mathbf{x}) - m(\mathbf{x}))(f(\mathbf{x}') - m(\mathbf{x}'))^T]. \end{aligned} \quad (10)$$

Any finite dimensional collection of random variables $f(\mathbf{x}_1), \dots, f(\mathbf{x}_l)$ is jointly Gaussian distributed

$$\begin{pmatrix} f(\mathbf{x}_1) \\ \vdots \\ f(\mathbf{x}_l) \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} m(\mathbf{x}_1) \\ \vdots \\ m(\mathbf{x}_l) \end{pmatrix}, \begin{pmatrix} k(\mathbf{x}_1, \mathbf{x}_1) & \cdots & k(\mathbf{x}_1, \mathbf{x}_l) \\ \vdots & \ddots & \vdots \\ k(\mathbf{x}_l, \mathbf{x}_1) & \cdots & k(\mathbf{x}_l, \mathbf{x}_l) \end{pmatrix} \right). \quad (11)$$

Here, we introduce the notation $\mathbf{K} = \mathbf{K}(\mathbf{X}, \mathbf{X})$, with $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_l)$ and $\mathbf{K}_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$. In regression we want to learn the function $f(\mathbf{x})$ from possibly noisy observations of its outputs at known inputs

$$y(\mathbf{x}_i) = f(\mathbf{x}_i) + \epsilon_i, \quad \epsilon_i \sim \mathcal{N}(0, \sigma_\epsilon^2). \quad (12)$$

For GP regression we place a GP prior on $f(\mathbf{x}) \sim \text{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}'))$ and formulate the conditional distribution of the function values to be predicted $\mathbf{f}(\mathbf{X}_*) = (f(\mathbf{x}_{*1}), \dots, f(\mathbf{x}_{*p}))$ at test points \mathbf{X}_* , given the observations $\mathbf{y}(\mathbf{X}) = (y(\mathbf{x}_1), \dots, y(\mathbf{x}_l))$. To obtain the conditional distribution, we first formulate the joint Gaussian distribution of the prediction and observation

$$\begin{pmatrix} \mathbf{f}(\mathbf{X}_*) \\ \mathbf{y}(\mathbf{X}) \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} m(\mathbf{X}_*) \\ m(\mathbf{X}) \end{pmatrix}, \begin{pmatrix} \mathbf{K}_{**} & \mathbf{K}_* \\ \mathbf{K}_*^T & \mathbf{K}_y \end{pmatrix} \right), \quad (13)$$

from which we get the conditional distribution using Lemma 2 (see Appendix)

$$\begin{aligned} \mathbf{f}(\mathbf{X}_*)|\mathbf{y}(\mathbf{X}) &\sim \mathcal{N}(\mathbf{m}(\mathbf{X}_*) + \mathbf{K}_* \mathbf{K}_y^{-1}(\mathbf{y}(\mathbf{X}) - \mathbf{m}(\mathbf{X})), \\ &\quad \mathbf{K}_{**} - \mathbf{K}_* \mathbf{K}_y^{-1} \mathbf{K}_*^T). \end{aligned} \quad (14)$$

Here, measurement noise is accounted for inside $\mathbf{K}_y = \mathbf{K} + \sigma_\epsilon^2 \mathbf{I}$. The covariance matrices in (13) and (14) are $\mathbf{K}_{**} = \mathbf{K}(\mathbf{X}_*, \mathbf{X}_*) \in \mathbb{R}^{p \times p}$ and $\mathbf{K}_* = \mathbf{K}(\mathbf{X}_*, \mathbf{X}) \in \mathbb{R}^{p \times l}$.

Although GP regression is non-parametric, meaning that we do not need to define a structure for $f(\mathbf{x})$, we incorporate prior information through our choice of the kernel $k(\mathbf{x}, \mathbf{x}'; \boldsymbol{\theta})$ and its hyper-parameters $\boldsymbol{\theta}$. The squared exponential kernel (SE) finds broad application. For scalar inputs it is

$$k(x, x') = \sigma^2 \exp\left(-\frac{(x - x')^2}{2l^2}\right). \quad (15)$$

The hyper-parameters are the variance σ^2 and the length-scale l .

The hyper-parameters $\boldsymbol{\theta}$ can be estimated from the measurement data by minimizing the negative log marginal likelihood function

$$-\log p(\mathbf{y}|\mathbf{X}) = +\frac{1}{2} \mathbf{y}^T (\mathbf{K}(\mathbf{X}, \mathbf{X}; \boldsymbol{\theta}) + \sigma_\epsilon^2 \mathbf{I})^{-1} \mathbf{y} + \frac{1}{2} \log |\mathbf{K}(\mathbf{X}, \mathbf{X}; \boldsymbol{\theta}) + \sigma_\epsilon^2 \mathbf{I}| + \frac{l}{2} \log 2\pi.$$

This provides a good compromise between data fit, first term, and model complexity, second term.

Numerical Gaussian processes: In Raissi et al. (2018) numerical GPs have been introduced and used to solve spatiotemporal models. Numerical GPs combine the data-driven machine learning nature of GPs with first principles knowledge from a spatiotemporal model. Numerical refers to the fact that the spatiotemporal model has to be time-discretized. For an easy to follow introduction of numerical GPs, we will be using the explicit Euler discretization here. For a general formulation, regardless of the discretization method, readers are referred to the original work of Raissi et al. (2018). To start, consider a linear partial differential equation

$$\frac{\partial n}{\partial t}(t, x) = \mathcal{L}_x n(t, x), \quad (16)$$

where \mathcal{L}_x is a linear operator acting on $n(t, x)$ with respect to $x \in \mathcal{R}$. Discretization of (16) in time with the explicit Euler scheme yields

$$\begin{aligned} n_t &= n_{t-1} + \Delta t \mathcal{L}_x n_{t-1} \\ &= \mathcal{Q}_x n_{t-1}. \end{aligned} \quad (17)$$

We now place a GP prior of our choice on n_{t-1}

$$n_{t-1} \sim \text{GP}(0, k_{t-1, t-1}^{nn}(x, x')). \quad (18)$$

It follows that

$$n_t = \mathcal{Q}_x n_{t-1} \sim \text{GP}(0, k_{t,t}^{nn}(x, x')), \quad (19)$$

since a linear transformation of a GP is still a GP but with a different kernel that is informed by the linear transformation, see Särkkä (2011) for example. Using the definition of a covariance function (10), we can derive

$$\begin{aligned} k_{t,t} &= \mathbb{E} \left[n_t(x) (n_t(x'))^T \right] \\ &= \mathbb{E} \left[\mathcal{Q}_x n_{t-1}(x) (\mathcal{Q}_{x'} n_{t-1}(x'))^T \right] \\ &= \mathcal{Q}_x k_{t-1, t-1}^{nn} \mathcal{Q}_{x'}^T, \end{aligned} \quad (20)$$

and

$$\begin{aligned} k_{t,t-1}^{nn} &= \mathbb{E} \left[n_t(x) (n_{t-1}(x'))^T \right] \\ &= \mathbb{E} \left[\mathcal{Q}_x n_{t-1}(x) (n_{t-1}(x'))^T \right] \\ &= \mathcal{Q}_x k_{t-1, t-1}^{nn}. \end{aligned} \quad (21)$$

Note that if we place the prior on n_t , we would have to invert the linear operator. Choosing where to place the GP prior is therefore crucial.

To perform temporal propagation, we first formulate the following multi-output GP

$$\begin{bmatrix} n_t \\ n_{t-1} \end{bmatrix} \sim \text{GP} \left(0, \begin{bmatrix} k_{t,t}^{nn} & k_{t,t-1}^{nn} \\ k_{t-1,t}^{nn} & k_{t-1,t-1}^{nn} \end{bmatrix} \right). \quad (22)$$

Starting from the initial condition, we now recursively calculate the conditional posterior distribution

$$p(n_t | n_{t-1}) = \mathcal{N}(\mu_t, \Sigma_{t,t}). \quad (23)$$

The posterior mean and covariance are calculated as in (14).

Before formulating the conditional posterior distribution, the hyper-parameters have to be learned. In the first time step, this is done on the initial and boundary data, and in the succeeding steps on artificially generated data, i.e. test points $\mathbf{X}_{t-1,*}$, $\mathbf{X}_{t,*}$, and the current boundary data. This artificial data has to be marginalized out from the posterior distribution to assure correct propagation of uncertainty. The reader is referred to Raissi et al. (2018) for exact details of this procedure.

Numerical GPs can handle explicit and implicit numerical methods. Merely (22) changes depending on the outputs and kernels introduced through the discretization method. Furthermore, numerous types of boundary conditions and their combinations can be handled within this framework.

3. NUMERICAL GAUSSIAN PROCESS KALMAN FILTERING

We start by extending numerical GPs to include online measurements and stochastic noise. Then, we embed this structure into the traditional Kalman filter equations.

3.1 Inclusion of a measurement equation and noise

Consider a linear stochastic partial differential equation (PDE)

$$\frac{\partial n}{\partial t}(t, x) = \mathcal{L}_x n(t, x) + q(t, x), \quad (24)$$

with \mathcal{L}_x being a linear operator acting on the solution of the PDE $n(t, \cdot) \in L_2(\mathbb{R}^{d_x}, \mathbb{R})$, $x \mapsto n(t, x)$, and $x \in \mathbb{R}^{d_x}$. Process noise $q(t, x)$ is white, additive Gaussian noise. To make the derivation easily accessible and without loss of generality, we use the explicit Euler time stepping scheme to derive the numerical GPKF. Forward Euler discretization of (24) results in

$$\begin{aligned} n_t(x) &= n_{t-1}(x) + \Delta t \mathcal{L}_x n_{t-1}(x) + \Delta t q_{t-1}(x), \\ &= \mathcal{F}_x n_{t-1}(x) + \Delta t q_{t-1}(x), \end{aligned} \quad (25)$$

where \mathcal{F}_x is again a linear operator summarizing the structure of the discretized PDE. The right-hand side of (24) is not continuous, and one should therefore use the Euler-Maruyama method for discretization instead. However, the covariance of the noise term q_{t-1} will enter the algorithm by being multiplied with the step size Δt , much

Table 1. Kernels of the multi-output explicit Euler GP (30). For readability we write k instead of $k_{t-1,t-1}^{nn}$.

Kernel	Structure
$k_{t-1,t-1}^{nn}$	k
$k_{t,t-1}^{nn}$	$\mathcal{F}_x k$
$k_{t,t}^{nn}$	$\mathcal{F}_x k \mathcal{F}_x^T + \Delta t^2 k_{t-1,t-1}^{qq}$
$k_{t,t-1}^{ny}$	$\mathcal{H}_x \mathcal{F}_x k$
$k_{t,t}^{ny}$	$\mathcal{F}_x k (\mathcal{F}_x' \mathcal{H}_x')^T + \Delta t^2 k_{t-1,t-1}^{qq}$
$k_{t,t}^{ny}$	$\mathcal{H}_x \mathcal{F}_x k (\mathcal{F}_x' \mathcal{H}_x')^T + \Delta t^2 \mathcal{H}_x k_{t-1,t-1}^{qq} \mathcal{H}_x'^T + k_{t,t}^{rr}$

like in the Euler-Maruyama method, see $k_{t,t}^{nn}$ in Table 1. We place mutually independent GP priors on n_{t-1} and q_{t-1}

$$n_{t-1} \sim \text{GP}(0, k_{t-1,t-1}^{nn}(x, x')), \quad (26)$$

$$q_{t-1} \sim \text{GP}(0, k_{t-1,t-1}^{qq}(x, x')), \quad (27)$$

with a kernel of our choice for n_{t-1} and a white noise kernel $k_{t-1,t-1}^{qq}(x, x') = \delta(x - x')\sigma_q^2$ for q_{t-1} . Other noise kernel choices are possible, but not sensible given the linear Gaussian setting.

We introduce a measurement equation

$$n_t^y(y) = \mathcal{H}_x n_t(y) + r_t(y), \quad (28)$$

with the measurement operator $\mathcal{H}_x : L_2(\mathbb{R}^{d_x}, \mathbb{R}) \rightarrow L_2(\mathbb{R}^{d_y}, \mathbb{R})$, $n^t \mapsto n_t^y$, the measurement $y \in \mathbb{R}^{d_y}$, and additive, white noise $r_t(y)$ modeled as a GP with kernel $k_{t,t}^{rr}(y, y') = \delta(y - y')\sigma_r^2$. Since n_{t-1} is a GP, so is n_t , as already established, but so is also n_t^y . Boundary conditions are usually a linear transformation of the state and are hence treated in the same fashion

$$n_t^b(x_b) = \mathcal{B}_x n_t(x_b). \quad (29)$$

We write the complete multi-output GP as

$$\begin{pmatrix} n_t \\ n_t^y \\ n_t^b \\ n_{t-1} \end{pmatrix} \sim \text{GP} \left(\begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \begin{bmatrix} k_{t,t}^{nn} & k_{t,t}^{ny} & k_{t,t}^{nb} & k_{t,t-1}^{nn} \\ k_{t,t}^{ny} & k_{t,t}^{ny} & k_{t,t}^{ny} & k_{t,t-1}^{ny} \\ k_{t,t}^{nb} & k_{t,t}^{nb} & k_{t,t}^{nb} & k_{t,t-1}^{nb} \\ k_{t-1,t}^{nn} & k_{t-1,t}^{ny} & k_{t-1,t}^{nb} & k_{t-1,t-1}^{nn} \end{bmatrix} \right). \quad (30)$$

These kernels, except for the boundary kernels, are listed in Table 1. With this structure we are equipped to write down the recursive Kalman filter equations. This is possible because we have a linear model whose states are Gaussian distributed random variables, i.e. Gaussian process distributed.

3.2 The numerical Gaussian process Kalman filter

In this subsection we only work with finite-dimensional collections of the Gaussian process random variables. We therefore write $\text{N}(\cdot, \cdot)$ instead of $\text{GP}(\cdot, \cdot)$. Furthermore, we use bold lowercase letters for the random variables and kernels become covariance matrices, indicated by bold capital letters, as already introduced in subsection 2.2.

Recall that in Kalman filtering we want to calculate the posterior distribution of a dynamic state given model predictions and measurements up to the current time. According to Bayes' rule the posterior is given by

$$p(\mathbf{n}_t | \mathbf{n}_{1:t}^y) = \frac{p(\mathbf{n}_t^y | \mathbf{n}_t) p(\mathbf{n}_t | \mathbf{n}_{1:t-1}^y)}{p(\mathbf{n}_t^y | \mathbf{n}_{1:t-1}^y)}. \quad (31)$$

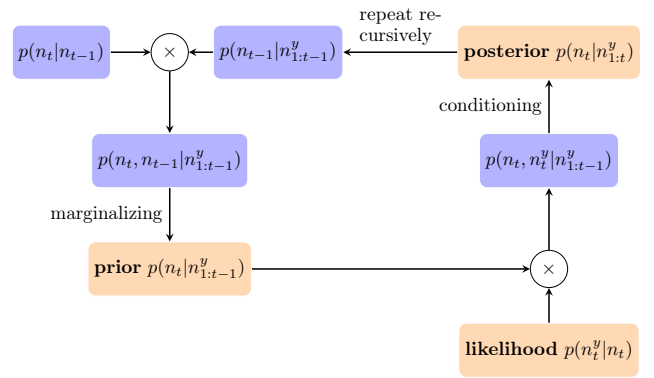


Fig. 1. Recursive posterior distribution calculation. Boundary data has been omitted for readability as it is treated in the same fashion as measurements.

Here we omitted the boundary data $\mathbf{n}_t^b(x_b)$ for readability. We assume the states to be Markovian, i.e. the current state \mathbf{n}_t is conditionally independent of anything that happened before $t - 1$. Furthermore, the current measurement given the current state is conditionally independent of the measurement and state histories. The individual terms can be calculated in closed form

- prior $p(\mathbf{n}_t | \mathbf{n}_{t-1}^y) = \text{N}(\mathbf{n}_t | \mathbf{m}_t^-, \mathbf{P}_t^-)$,
- posterior $p(\mathbf{n}_t | \mathbf{n}_t^y) = \text{N}(\mathbf{n}_t | \mathbf{m}_t, \mathbf{P}_t)$,
- normalizing constant $p(\mathbf{n}_t^y | \mathbf{n}_{t-1}^y) = \text{N}(\mathbf{n}_t^y | \mathcal{H}_x \mathbf{m}_t^-, \mathbf{P}_t)$.

An illustration of the road map ahead is shown in Fig. 1.

To calculate the prior distribution, we first calculate the joint distribution of the states \mathbf{n}_{t-1} and \mathbf{n}_t conditioned on the measurement and boundary data history

$$\begin{aligned} p(\mathbf{n}_{t-1}, \mathbf{n}_t | \mathbf{n}_{1:t-1}^y, \mathbf{n}_{1:t-1}^b) &= \underbrace{p(\mathbf{n}_t | \mathbf{n}_{t-1})}_{\text{model prediction}} \\ &\quad \times \underbrace{p(\mathbf{n}_{t-1} | \mathbf{n}_{1:t-1}^y, \mathbf{n}_{1:t-1}^b)}_{\text{previous posterior}} \\ &= \text{N}(\mathbf{n}_t | \mathbf{A} \mathbf{n}_{t-1}, \mathbf{P}_t^{GP,nn}) \\ &\quad \times \text{N}(\mathbf{n}_{t-1} | \mathbf{m}_{t-1}, \mathbf{P}_{t-1}) \\ &= \text{N} \left(\begin{pmatrix} \mathbf{n}_{t-1} \\ \mathbf{n}_t \end{pmatrix} | \mathbf{m}', \mathbf{P}' \right). \quad (32) \end{aligned}$$

Using Lemma 1, the joint mean of (32) is

$$\mathbf{m}' = \begin{pmatrix} \mathbf{m}_{t-1} \\ \mathbf{A} \mathbf{m}_{t-1} \end{pmatrix} \quad (33)$$

and the covariance is

$$\mathbf{P}' = \begin{bmatrix} \mathbf{P}_{t-1} & \mathbf{P}_{t-1} \mathbf{A}^T \\ \mathbf{A} \mathbf{P}_{t-1} & \mathbf{A} \mathbf{P}_{t-1} \mathbf{A}^T + \mathbf{P}_t^{GP,nn} \end{bmatrix}. \quad (34)$$

We introduced

$$\mathbf{A} = \mathbf{K}_{t,t-1}^{nn} (\mathbf{K}_{t-1,t-1}^{nn})^{-1}, \quad (35)$$

and

$$\mathbf{P}_t^{GP,nn} = \mathbf{K}_{t,t}^{nn} - \mathbf{K}_{t,t-1}^{nn} (\mathbf{K}_{t-1,t-1}^{nn})^{-1} \mathbf{K}_{t-1,t}^{nn}. \quad (36)$$

These two equations are obtained by forming the conditional $\mathbf{n}_t | \mathbf{n}_{t-1}$ of the multi-output GP (22). We can think of (35) as the equivalent to the dynamic matrix in a state space model with the difference that its entries depend on the test points $\mathbf{A}(\mathbf{X}_{t,*}, \mathbf{X}_{t-1,*})$. Now, the prior distribution is obtained by marginalizing over \mathbf{n}_{t-1} in (32)

$$p(\mathbf{n}_t | \mathbf{n}_{1:t-1}^y, \mathbf{n}_{1:t-1}^b) = \text{N}(\mathbf{n}_t | \mathbf{m}_t^-, \mathbf{P}_t^-), \quad (37)$$

with prior mean and covariance

$$\mathbf{m}_t^- = \mathbf{A}\mathbf{m}_{t-1}, \quad (38)$$

$$\mathbf{P}_t^- = \mathbf{A}\mathbf{P}_{t-1}\mathbf{A}^T + \mathbf{P}_t^{GP,nn}. \quad (39)$$

Equipped with this, we can calculate the joint distribution $p(\mathbf{n}_t, \mathbf{n}_t^y, \mathbf{n}_t^b | \mathbf{n}_{1:t-1}^y, \mathbf{n}_{1:t-1}^b)$. This will allow us to write down the posterior distribution later. We start by noting that

$$\begin{aligned} p(\mathbf{n}_t, \mathbf{n}_t^y, \mathbf{n}_t^b | \mathbf{n}_{1:t-1}^y, \mathbf{n}_{1:t-1}^b) &= p(\mathbf{n}_t^y, \mathbf{n}_t^b | \mathbf{n}_t) \\ &\quad \times p(\mathbf{n}_t | \mathbf{n}_{1:t-1}^y, \mathbf{n}_{1:t-1}^b) \\ &= \mathbf{N}(\mathbf{n}_t^y, \mathbf{n}_t^b | \mathbf{C}\mathbf{n}_t, \mathbf{P}_t^{GP,n^y n^y}) \\ &\quad \times \mathbf{N}(\mathbf{n}_t | \mathbf{m}_t^-, \mathbf{P}_t^-) \\ &= \mathbf{N}\left(\begin{pmatrix} \mathbf{n}_t \\ \mathbf{n}_t^y \\ \mathbf{n}_t^b \end{pmatrix} | \mathbf{m}'', \mathbf{P}''\right). \end{aligned} \quad (40)$$

Here we introduced

$$\mathbf{C} = \begin{bmatrix} \mathbf{K}_{t,t}^{n^y n^y} \\ \mathbf{K}_{t,t}^{n^b n^y} \end{bmatrix} (\mathbf{K}_{t,t}^{nn})^{-1}, \quad (41)$$

and

$$\begin{aligned} \mathbf{P}_t^{GP,n^y n^y} &= \begin{bmatrix} \mathbf{K}_{t,t}^{n^y n^y} & \mathbf{K}_{t,t}^{n^y n^b} \\ \mathbf{K}_{t,t}^{n^b n^y} & \mathbf{K}_{t,t}^{n^b n^b} \end{bmatrix} \\ &\quad - \begin{bmatrix} \mathbf{K}_{t,t}^{n^y n} \\ \mathbf{K}_{t,t}^{n^b n} \end{bmatrix} (\mathbf{K}_{t,t}^{nn})^{-1} \begin{bmatrix} \mathbf{K}_{t,t}^{n^y n^y} & \mathbf{K}_{t,t}^{n^y n^b} \end{bmatrix}. \end{aligned} \quad (42)$$

These two expressions are obtained in the similar fashion as \mathbf{A} and $\mathbf{P}_t^{GP,nn}$ by forming $\mathbf{n}_t^y, \mathbf{n}_t^b | \mathbf{n}_t$ from (30). Again, we can think of \mathbf{C} as the measurement matrix equivalent in a state space model, albeit its entries depend on the measurement, boundary, and test point locations $\mathbf{C}(\mathbf{Y}_t, \mathbf{X}_{t,b}, \mathbf{X}_{t,*})$. Using Lemma 1 one more time, the joint mean is

$$\mathbf{m}'' = \begin{pmatrix} \mathbf{m}_t^- \\ \mathbf{C}\mathbf{m}_t^- \end{pmatrix}. \quad (43)$$

The covariance is

$$\mathbf{P}'' = \begin{bmatrix} \mathbf{P}_t^- & \mathbf{P}_t^- \mathbf{C}^T \\ \mathbf{C}\mathbf{P}_t^- & \mathbf{C}\mathbf{P}_t^- \mathbf{C}^T + \mathbf{P}_t^{GP,n^y n^y} \end{bmatrix}. \quad (44)$$

To get the posterior $p(\mathbf{n}_t | \mathbf{n}_{1:t}^y, \mathbf{n}_{1:t}^b)$, we need to condition the above joint distribution $p(\mathbf{n}_t, \mathbf{n}_t^y, \mathbf{n}_t^b | \mathbf{n}_{1:t-1}^y, \mathbf{n}_{1:t-1}^b)$ on the current measurement and boundary data using Lemma 2

$$\begin{aligned} p(\mathbf{n}_t | \mathbf{n}_t^y, \mathbf{n}_t^b, \mathbf{n}_{1:t-1}^y, \mathbf{n}_{1:t-1}^b) &= p(\mathbf{n}_t | \mathbf{n}_{1:t}^y, \mathbf{n}_{1:t}^b) \\ &= \mathbf{N}(\mathbf{m}_t, \mathbf{P}_t). \end{aligned} \quad (45)$$

The posterior mean is ¹

$$\begin{aligned} \mathbf{m}_t &= \mathbf{m}_t^- + \mathbf{P}_t^- \mathbf{C}^T \\ &\quad \times \left(\mathbf{C}\mathbf{P}_t^- \mathbf{C}^T + \mathbf{P}_t^{GP,n^y n^y} \right)^{-1} \left([\mathbf{n}_t^y \ \mathbf{n}_t^b]^T - \mathbf{C}\mathbf{m}_t^- \right), \end{aligned} \quad (46)$$

and the corresponding variance is

$$\mathbf{P}_t = \mathbf{P}_t^- - \mathbf{P}_t^- \mathbf{C}^T \left(\mathbf{C}\mathbf{P}_t^- \mathbf{C}^T + \mathbf{P}_t^{GP,n^y n^y} \right)^{-1} \mathbf{C}\mathbf{P}_t^-. \quad (47)$$

¹ Careful readers will recognize the structure of this posterior KF estimate from the posterior GP mean and covariance, compare with (14). This shouldn't be a surprise, as in both cases a conditional Gaussian is formed.




Fig. 2. Temporal snapshots of the analytical solution (black) and the posterior estimate of the numerical GPKF (blue) along with two times its standard deviation. Blue circles indicate measurements. 41 training points are sampled from the initial estimate, from then on 5 noisy measurements are available in each time step $\Delta t = 5 \cdot 10^{-3}$.

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

$$\mathbf{m}_t^- = \mathbf{A}\mathbf{m}_{t-1}^-, \quad (48)$$

$$\mathbf{P}_t^- = \mathbf{A}\mathbf{P}_{t-1}\mathbf{A}^T + \mathbf{P}_t^{GP,nn}, \quad (49)$$

and the update step as

$$\mathbf{v}_t = \mathbf{n}_t^y - \mathbf{C}\mathbf{m}_t^-, \quad (50)$$

$$\mathbf{S}_t = \mathbf{C}\mathbf{P}_t^- \mathbf{C}^T + \mathbf{P}_t^{GP,n^y n^y}, \quad (51)$$

$$\mathbf{K}_t = \mathbf{P}_t^- \mathbf{C}^T (\mathbf{S}_t)^{-1}, \quad (52)$$

$$\mathbf{m}_t = \mathbf{m}_t^- + \mathbf{K}_t \mathbf{v}_t, \quad (53)$$

$$\mathbf{P}_t = \mathbf{P}_t^- - \mathbf{K}_t \mathbf{S}_t (\mathbf{K}_t)^T. \quad (54)$$

The covariances of process and measurement noise are added to the kernels $\mathbf{K}_{t,t}^{nn}$ and $\mathbf{K}_{t,t}^{n^y n^y}$, as shown in Table 1. They therefore enter the algorithm as in the traditional Kalman filter.

By minimizing the negative log-marginal likelihood, hyperparameters are estimated in every time step from the current measurement, boundary, and previous test data.

4. SIMULATION STUDY

We will look at the one-dimensional advection equation

$$\frac{\partial n}{\partial t}(t, x) + gn(t, x) = 0. \quad (55)$$

Here, the advection speed is $g = 3$. We impose a no-flux boundary condition $n(t, x_b = 0)g = 0$. The initial state is a bimodal Gaussian distribution $n(0, x) = \mathbf{N}(\mu_{0,1} = 2, \sigma_{0,1}^2 = 0.45^2) + \mathbf{N}(\mu_{0,2} = 3.75, \sigma_{0,2}^2 = 0.6^2)$.

For this numerical case study, we assume to obtain noise corrupted point evaluations of the analytical solution

$$n_t^y(y_i) = n_t(y_i) + \epsilon_i. \quad (56)$$

Here ϵ is generated from a zero mean Gaussian with variance $\sigma_\epsilon^2 = 0.06^2$. Process noise is zero.




Fig. 3. Process and measurement noise standard deviation estimation, as well as the mean integrated squared error between posterior state estimation and ground truth.

We discretize (55) using the implicit Euler method with a step size of $\Delta t = 5 \cdot 10^{-3}$. This is also used as the sampling time for measurements. A GP prior with a zero-mean function and squared exponential kernel is placed on $n_t(x)$. From this, the entire structure of the numerical GPKF is built up.

The initial state estimate is a shifted and diffused version of the true initial state $n(0, x) = N(\mu_0 + 0.5, \sigma_0^2 + 0.2)$. From the initial estimate, $N_{tr} = 41$ points are sampled and regular GP regression is performed. This gives the initial variance P_0 of the Kalman filter. Initial hyper-parameter values of the squared exponential are $l_{SE} = 0.5$ and $\sigma_{SE}^2 = 0.3^2$, and $\sigma_q^2 = 0.1^2$ for the process white noise kernel, and $\sigma_r^2 = 0.2^2$ for the measurement white noise kernel.

Fig. 2 shows temporal snapshots of the analytical solution and the numerical GPKF posterior estimate along with two times its standard deviation. State estimation convergence towards ground truth in terms of the mean integrated squared error (MISE) can be seen in Fig. 3. Therein, the adaptation of the noise hyper-parameters is shown as well.

5. CONCLUSION

In this paper, we introduced numerical Gaussian process Kalman filtering (GPKF) with which we can do state and noise estimation for infinite-dimensional systems. Numerical Gaussian processes by Raissi et al. (2018) solve spatiotemporal models with Gaussian process regression by discretization in time and leveraging the resulting structure for kernel design. We embedded this method into the well known recursive Kalman filter equations.

Learning the noise hyper-parameters online on the data stream could be a two-sided sword. In cases where the negative log-marginal likelihood function has a rugged contour, performance of the numerical GPKF could be drastically reduced by getting stuck in local minima. Further simulation studies with more complex dynamics should investigate this, as well as non-zero process noise.

REFERENCES

- Deisenroth, M.P., Huber, M.F., and Hanebeck, U.D. (2009). Analytic moment-based gaussian process filtering. In *Proceedings of the 26th annual international conference on machine learning*, 225–232. ACM.
- Deisenroth, M.P. and Ohlsson, H. (2011). A general perspective on gaussian filtering and smoothing: Explaining current and deriving new algorithms. In *Proceedings of the 2011 American Control Conference*, 1807–1812. IEEE.
- Kalman, R.E. (1960). A new approach to linear filtering and prediction problems.
- Ko, J. and Fox, D. (2009). Gp-bayesfilters: Bayesian filtering using gaussian process prediction and observation models. *Autonomous Robots*, 27(1), 75–90.
- Raissi, M., Perdikaris, P., and Karniadakis, G.E. (2018). Numerical gaussian processes for time-dependent and nonlinear partial differential equations. *SIAM Journal on Scientific Computing*, 40(1), A172–A198.
- Särkkä, S. (2011). Linear operators and stochastic partial differential equations in gaussian process regression. In *International Conference on Artificial Neural Networks*, 151–158. Springer.
- Särkkä, S. (2013). *Bayesian filtering and smoothing*, volume 3. Cambridge University Press.
- Särkkä, S. and Hartikainen, J. (2012). Infinite-dimensional kalman filtering approach to spatio-temporal gaussian process regression. In *International Conference on Artificial Intelligence and Statistics*, 993–1001.
- Särkkä, S., Solin, A., and Hartikainen, J. (2013). Spatiotemporal learning via infinite-dimensional bayesian filtering and smoothing: A look at gaussian process regression through kalman filtering. *IEEE Signal Process. Mag.*, 30(4), 51–61.
- Williams, C.K. and Rasmussen, C.E. (2006). *Gaussian processes for machine learning*, volume 2. MIT Press Cambridge, MA.

Appendix A. GAUSSIAN IDENTITIES

The following is taken from Särkkä (2013).

Lemma 1. (Joint distribution of Gaussian variables) If random variables $\mathbf{x} \in \mathbb{R}^{d_x}$ and $\mathbf{y} \in \mathbb{R}^{d_y}$ have the Gaussian probability distributions

$$\mathbf{x} \sim N(\mathbf{m}, \mathbf{P}) \quad (\text{A.1})$$

$$\mathbf{y}|\mathbf{x} \sim N(\mathbf{H}\mathbf{x} + \mathbf{u}, \mathbf{R}), \quad (\text{A.2})$$

then the joint distribution of \mathbf{x} , \mathbf{y} and the marginal distribution of \mathbf{y} are given as

$$\begin{pmatrix} \mathbf{x} \\ \mathbf{y} \end{pmatrix} \sim N\left(\begin{pmatrix} \mathbf{m} \\ \mathbf{H}\mathbf{m} + \mathbf{u} \end{pmatrix}, \begin{pmatrix} \mathbf{P} & \mathbf{P}\mathbf{H}^T \\ \mathbf{H}\mathbf{P} & \mathbf{H}\mathbf{P}\mathbf{H}^T + \mathbf{R} \end{pmatrix}\right), \quad (\text{A.3})$$

$$\mathbf{y} \sim N(\mathbf{H}\mathbf{m} + \mathbf{u}, \mathbf{H}\mathbf{P}\mathbf{H}^T + \mathbf{R}). \quad (\text{A.4})$$

Lemma 2. (Conditional distribution of Gaussian variables) If the random variables \mathbf{x} and \mathbf{y} have the joint Gaussian probability distribution

$$\begin{pmatrix} \mathbf{x} \\ \mathbf{y} \end{pmatrix} \sim N\left(\begin{pmatrix} \mathbf{a} \\ \mathbf{b} \end{pmatrix}, \begin{pmatrix} \mathbf{A} & \mathbf{C} \\ \mathbf{C}^T & \mathbf{B} \end{pmatrix}\right), \quad (\text{A.5})$$

then the conditional distribution is

$$\mathbf{x}|\mathbf{y} \sim N(\mathbf{a} + \mathbf{C}\mathbf{B}^{-1}(\mathbf{y} - \mathbf{b}), \mathbf{A} - \mathbf{C}\mathbf{B}^{-1}\mathbf{C}^T). \quad (\text{A.6})$$