

The Kernel Kalman Rule

Efficient Nonparametric Inference by Recursive Least-Squares and Subspace Projections

Gregor H. W. Gebhardt · Andras Kupcsik ·
Gerhard Neumann

Received: date / Accepted: date

Abstract Enabling robots to act in unstructured and unknown environments requires versatile state estimation techniques. While traditional state estimation methods require known models and make strong assumptions about the dynamics, such versatile techniques should be able to deal with high dimensional observations and non-linear, unknown system dynamics. The recent framework for nonparametric inference (Song et al 2013) allows to perform inference on arbitrary probability distributions. High-dimensional embeddings of distributions into reproducing kernel Hilbert spaces (RKHS) are manipulated by kernelized inference rules, most prominently the kernel Bayes' rule (KBR). However, the computational demands of the KBR do not scale with the number of samples. In this paper, we present two techniques to increase the computational efficiency of non-parametric inference. First, the kernel Kalman rule (KKR) is presented as an alternative to the KBR. The derivation of the KKR follows the clear objective of recursive least squares. Based on the KKR we present the kernel Kalman filter (KKF) that updates an embedding of the belief state and learns the system and observation models from data. We further derive the kernel forward backward smoother (KFBS) based on a forward and backward KKF and a smoothing update in Hilbert space. Second, we present the subspace conditional embedding operator as a sparsification technique that still leverages from the full data set. We apply this sparsification to the KKR and derive the corresponding sparse KKF and KFBS algorithms. We show on nonlinear state estimation tasks that our approaches provide a significantly improved estimation accuracy while the computational demands are considerably decreased.

Gregor H.W. Gebhardt
Technische Universität Darmstadt, Computational Learning for Autonomous Systems
Hochschulstr. 10, 64285 Darmstadt, Germany
E-mail: gregor@robot-learning.de Tel.: +49-6151-16-25371
0000-0002-8575-069X

Andras Kupcsik
Idiap Research Institute, Centre du Parc
Rue Marconi 19, PO Box 592, 1920 Martigny, Switzerland
E-mail: andras.kupcsik@idiap.ch

Gerhard Neumann
University of Lincoln, Lincoln Centre for Autonomous Systems
Brayford Pool, LN6 7TS, Lincoln, UK
E-mail: geri@robot-learning.de

Keywords nonparametric inference · kernel methods · state estimation · model learning

1 Introduction

The ability to reason about past, current and future states in continuous, partially observable stochastic processes is a fundamental stepstone towards fully autonomous and intelligent systems. Such models are required in many applications as for example state estimation in case of incomplete sensory data, smoothing noisy data from mediocre sensors, or predicting future states from past and current observations.

Traditional state estimation techniques usually require analytical models of the underlying system, are often limited to a set of models with a special structure, and require knowledge about the moments of the stochastic processes. When assuming linear Gaussian models with known mean and covariance for instance, the Kalman filter (Kalman 1960) yields the optimal solution. However, the requirement of linear Gaussian models with known statistics imposes a strong limitation to the applicability of this method. For more complex processes, approximate solutions have to be used instead. Examples are the extended Kalman filter (McElhoo 1966; Smith et al 1962) or the unscented Kalman filter (Julier and Uhlmann 1997; Wan and Van Der Merwe 2000). These solutions inherit the Gaussian representation of the belief state to which they apply the non-linear system dynamics. However, the Gaussian distribution with its unimodal nature is a strong assumption about the belief state which leads to poor results for systems that yield a more complex distribution over possible states. Moreover, both the Kalman filter but also its extensions to non-linear systems, require that the dynamics of the systems are known as analytical models. Yet, these analytical models are often hard to obtain or make simplifying assumptions about the system.

The recently introduced framework for nonparametric inference (Song et al 2013; Fukumizu et al 2013) alleviates the problems of traditional state estimation methods for nonlinear systems. The basic idea of these methods is to embed the probability distributions into reproducing kernel Hilbert spaces (RKHS). These embeddings allow the representation of arbitrary probability distributions using empirical estimators. With the kernelized versions of the sum rule, the chain rule, and the Bayes' rule, inference on the embedded distribution can then be performed efficiently and entirely in the RKHS. Additionally, Song et al (2013) use the kernel sum rule and the kernel Bayes' rule to construct the kernel Bayes' filter (KBF). The KBF learns the transition and observation models from observed samples and can be applied to nonlinear systems with high-dimensional observations. However, the computational complexity of the KBF update scales poorly with the number of samples such that hyper-parameter optimization becomes prohibitively expensive. Moreover, the KBF requires mathematical tricks that may cause numerical instabilities and also render the objective which is optimized by the KBF unclear.

We propose two additions to the framework for nonparametric inference to overcome the limitations named above. First, we introduce the *subspace conditional embedding operator*. In contrast to the conditional embedding operator (Song et al 2009), this operator allows to estimate its empirical estimator with a much larger data set while maintaining computational efficiency. We further apply this subspace conditional embedding operator to the kernel sum rule, kernel chain rule and kernel Bayes rule to derive their subspace versions. We presented these results in a workshop paper at the large-scale kernel learning workshop at ICML 2015 (Gebhardt et al 2015).

Furthermore, we present the *kernel Kalman rule* (KKR) as alternative to the kernel Bayes' rule. Our derivations closely follow the derivations of the innovation update used

in the Kalman filter and are based on a recursive least squares minimization objective in a reproducing kernel Hilbert space (RKHS). We show that our update of the mean embedding is unbiased and has minimum variance. While the update equations are formulated in a potentially infinite dimensional RKHS, we derive, through the application of the kernel trick and by virtue of the representer theorem, an algorithm using only products and inversions of finite kernel matrices. We employ the kernel Kalman rule together with the kernel sum rule for filtering, which results in the *kernel Kalman filter* (KKF). In contrast to filtering techniques that rely on the KBR, the KKF allows to precompute expensive matrix inversions which significantly reduces the computational complexity and which also allows us to apply hyper-parameter optimization for the KKF. This work has been presented at AAAI 2017 (Gebhardt et al 2017).

In addition to the KKF, we introduce the *kernel forward backward smoother* (KFBS) which computes the embedding of the belief state given all available observations from the past and the future. The kernel forward backward smoother combines the belief state embeddings of a forward pass and a backward pass into smoothed embeddings using Hilbert space operations. Both, the forward and the backward pass are realized by a KKF, where the backward KKF operates backwards in time starting at the last observation. To scale gracefully with larger data sets, we rederive the KKR, the KKF and the KFBS with the subspace conditional operator (Gebhardt et al 2015).

We compare our approach to different versions of the KBR and demonstrate its improved estimation accuracy and computational efficiency. Furthermore, we evaluate the KKR on a simulated 4-link pendulum task, on a human motion capture data set (Wojtusich and von Stryk 2015) and on data from a table-tennis setup (Gomez-Gonzalez et al 2016).

1.1 Related Work

To the best of our knowledge the kernel Bayes' rule exists in three different versions. It was first introduced in its original version by Fukumizu et al (2013). Here, the KBR is derived, similar to the conditional operator, using prior modified covariance operators. These prior-modified covariance operators are approximated by weighting the feature mappings with the weights of the embedding of the prior distribution. Since these weights are potentially negative, the covariance operator might become indefinite, and thus, rendering its inversion impossible. To overcome this drawback, the authors have to apply a form of the Tikhonov regularization that significantly decreases accuracy and increases the computational costs. A second version of the KBR was introduced by Song et al (2013) in which they use a different approach to approximate the prior-modified covariance operator. In the experiments conducted for this paper, this second version often leads to more stable algorithms than the first version. Boots et al (2013) introduced a third version of the KBR where they apply only the simple form of the Tikhonov regularization. However, this rule requires the inversion of a matrix that is often indefinite, and therefore, high regularization constants are required, which again degrades the performance. In our experiments, we refer to these different versions with KBR(b) for the first, KBR(a) for the second (both adapted from the literature), and KBR(c) for the third version. Song et al (2013) propose in their framework for nonparametric inference to combine the KBR with the kernel sum rule to obtain the kernel Bayes filter (KBF). Our kernel Kalman filter is closely related to this, as we simply replace the KBR with the KKR. We compare to the KBF in our experiments. Nishiyama et al (2016) recently proposed the nonparametric kernel Bayes smoother. This approach builds on top of the kernel Bayes filter, which is used

to compute the estimates of a normal forward pass. The smoothing update is then obtained by propagating the embeddings backwards in time without performing a second filtering pass.

For filtering tasks with known linear system equations and Gaussian noise, the Kalman filter (KF) yields the solution that minimizes the squared error of the estimate to the true state. Two widely known and applied approaches to extend the Kalman filter to non-linear systems are the *extended Kalman filter* (EKF) (McElhoe 1966; Smith et al 1962) and the *unscented Kalman filter* (UKF) (Wan and Van Der Merwe 2000; Julier and Uhlmann 1997). Both, the EKF and the UKF, assume that the non-linear system dynamics are known and use them to update the prediction mean. Yet, updating the prediction covariance is not straightforward. In the EKF the system dynamics are linearized at the current estimate of the state, and in the UKF the covariance is updated by applying the system dynamics to a set of sample-points (sigma points). While these approximations make the computations tractable, they can significantly reduce the quality of the state estimation, in particular for high-dimensional systems.

Hsu et al (2012) recently proposed an algorithm for learning Hidden Markov Models (HMMs) by exploiting the spectral properties of observable measures to derive an observable representation of the HMM (Jaeger 2000). An RKHS embedded version thereof was presented in (Song et al 2010). While this method is applicable for continuous state spaces, it still assumes a finite number of discrete hidden states.

Other closely related algorithms to our approach are the kernelized version of the Kalman filter and Kalman smoother by Ralaivola and d'Alche Buc (2005) and the kernel Kalman filter based on the conditional embedding operator (KKF-CEO) by Zhu et al (2014). The former approach formulates the Kalman filter in a sub-space of the infinite feature space that is defined by the kernels. Hence, this approach does not fully leverage the kernel idea of using an infinite feature space. In contrast, the KKF-CEO approach embeds the belief state also in an RKHS. However, they require that the observation is a noisy version of the full state of the system, and thus, they cannot handle partial observations. Moreover, they also deviate from the standard derivation of the Kalman filter, which, as our experiments show, decreases the estimation accuracy. The full observability assumption is needed in order to implement a simplified version of the innovation update of the Kalman filter in the RKHS. Our algorithm, the KKF does not suffer from this restriction. It also provides update equations that are much closer to the original Kalman filter and outperforms the KKF-CEO algorithm as shown in our experiments. Another approach to state estimation is presented in (Kawahara et al 2007). In this paper the authors propose to estimate low-dimensional state vectors based on kernel canonical correlation analysis and then regress a linear transition model of the estimated state vectors and the nonlinear features of the input.

Learning predictors in the space of predictive state representations to perform filtering has been proposed in (Sun et al 2016b) and later extended to smoothing in (Sun et al 2016a). They introduce predictive state inference machines (PSIM) which are (nonlinear) regressors on predictive states learned from data to perform filtering. With the smoothing machine (SMACH) they extend this concept for smoothing.

2 Preliminaries

Our work is based on the recent formulations of embedding distributions into reproducing kernel Hilbert spaces (Smola et al 2007; Song et al 2013). These embeddings allow to represent arbitrary probability distributions non-parametrically by a potentially infinite dimensional feature vector. Through the application of derived operators (Song et al 2009;

Fukumizu et al 2011) it is furthermore possible to apply inference rules entirely in the Hilbert space. In the first part of this section, we want to give the reader an introduction into this technology and define the notation we will use throughout this article. We will furthermore show a modification of this framework which allows to learn the embeddings and inference rules with much larger datasets while maintaining a low computational complexity.

One of the main contributions of our paper is a novel method for performing Bayesian updates on a distribution embedded into an RKHS. The derivations of this update rule are based on a least-squares objective and inspired by the derivations of the Kalman filter update, thus we name this method *kernel Kalman rule*. In the second part of this section, we will recapitulate the classical Kalman filter equations and review the derivations of the innovation update based on the least-squares objective.

2.1 Nonparametric Inference with Hilbert Space Embeddings of Distributions

Intuitively, a Hilbert space is an extension of the well known two- or three-dimensional Euclidean vector space to arbitrary many dimensions, specifically including infinite dimensional vector spaces. Such infinite dimensional Hilbert spaces include spaces where the single elements are functions, i.e., infinite dimensional vectors that contain for each element of the domain the corresponding function value in the image. In addition, a Hilbert space has an inner product that allows to measure distances and angles between its elements. For a reproducing kernel Hilbert space \mathcal{H}_k , this inner product $\langle \cdot, \cdot \rangle$ is implicitly defined by a reproducing kernel, with $k : \Omega \times \Omega \rightarrow \mathbb{R}$ and $k(\mathbf{x}, \mathbf{y}) =: \langle \boldsymbol{\varphi}(\mathbf{x}), \boldsymbol{\varphi}(\mathbf{y}) \rangle$, where $\boldsymbol{\varphi}(\mathbf{x})$ is a feature mapping into a possibly infinite dimensional space, intrinsic to the kernel function. For example the Gaussian kernel computes the inner product of the feature mappings of its inputs where the feature mappings itself cannot be written down explicitly as they are into an infinite dimensional space. Due to the reproducing property of the kernel, all elements f of the RKHS can be reproduced by k in the sense that the outcome $f(x)$ of the function for a specific value x can be obtained by an evaluation of the kernel function (Aronszajn 1950), i.e., $f(\mathbf{y}) = \langle f, \boldsymbol{\varphi}(\mathbf{y}) \rangle$ for any $f \in \mathcal{H}_k$.

In a practical setting, we want to embed probability distributions in an RKHS spanned by samples $\mathcal{D} = [(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_n, \mathbf{y}_n)]$. Based on the representer theorem (Schölkopf et al 2001) and the reproducing property, the elements f of an RKHS \mathcal{H}_k can then be written as

$$f(\cdot) = \sum_{i=1}^n \alpha_i k(\mathbf{x}_i, \cdot) = \sum_{i=1}^n \alpha_i \langle \boldsymbol{\varphi}(\mathbf{x}_i), \boldsymbol{\varphi}(\cdot) \rangle = \boldsymbol{\alpha}^\top \boldsymbol{\Upsilon}_x^\top \boldsymbol{\varphi}(\cdot), \quad (1)$$

with the weights $\alpha_i \in \mathbb{R}$ and where we denote the feature matrix of samples \mathbf{x}_i by $\boldsymbol{\Upsilon}_x = [\boldsymbol{\varphi}(\mathbf{x}_1), \dots, \boldsymbol{\varphi}(\mathbf{x}_n)]$. In the following paragraphs we will show how probability distributions can be represented as an embedding in such a reproducing kernel Hilbert space and how the operators for performing inference in the RKHS can be derived.

2.1.1 Embeddings of Marginal and Joint Distributions

The embedding of a marginal density $P(X)$ over the random variable X is defined as the expected feature mapping $\boldsymbol{\mu}_X := \mathbb{E}_X [\boldsymbol{\varphi}(X)]$, also called the *mean map* (Smola et al 2007). Using a finite set of samples from $P(X)$, the mean map can be estimated as

$$\hat{\boldsymbol{\mu}}_X = \frac{1}{n} \sum_{i=1}^n \boldsymbol{\varphi}(\mathbf{x}_i) = \frac{1}{n} \boldsymbol{\Upsilon}_x^\top \mathbf{1}_n, \quad (2)$$

where \mathbf{Y}_x is the feature matrix of the samples from the distribution and $\mathbf{1}_n \in \mathbb{R}^n$ is an n dimensional vector of ones. Because of the reproducing property of the kernel function, computing the expectation of a function which is an element of the same RKHS resolves to simple matrix operations. On the other hand, the probability of a single outcome or higher order statistics of the distributions are not straight forward to obtain.

Alternatively, a distribution can be embedded in a tensor product RKHS $\mathcal{H}_k \times \mathcal{H}_k$ as the expected tensor product of the feature mappings (Smola et al 2007)

$$C_{XX} := \mathbb{E}_{XX} [\varphi(X) \otimes \varphi(X)] - \mu_X \otimes \mu_X, \quad (3)$$

where we use \otimes to denote the tensor product (or outer product) of two vectors. This embedding is also called the *centered covariance operator*. The finite sample estimator is given by

$$\hat{C}_{XX} = \frac{1}{m} \sum_{i=1}^m \varphi(x_i) \otimes \varphi(x_i) - \hat{\mu}_X \otimes \hat{\mu}_X. \quad (4)$$

Similarly, we can define the *uncentered cross-covariance operator* for a joint distribution $p(X, Y)$ of two variables X and Y as $\hat{C}_{XY} = \frac{1}{m} \sum_{i=1}^m \varphi(x_i) \otimes \phi(y_i)$. Based on these two definitions for Hilbert space embeddings of probability distributions and with the conditional embedding operator discussed in the next section, all the rules of the framework for non-parametric inference (Song et al 2013) can be derived.

2.1.2 The Conditional Embedding Operator

The embedding of a conditional distribution $P(Y|X)$ is not like the mean map a single element of the RKHS, but rather a family of embeddings where there is one embedding for each realization of the conditioning variable X . To obtain the conditional distribution for a specific value $X = x_*$, Song et al (2009) defined the *conditional embedding operator* $C_{Y|X}$ which, if applied to the feature mapping of \mathbf{x} returns the embedding of $P(Y|X = x_*)$

$$\mu_{Y|x} := \mathbb{E}_{Y|x} [\phi(Y)] = C_{Y|X} \varphi(\mathbf{x}). \quad (5)$$

Given a finite set of samples $\mathcal{D} = [(x_1, y_1), \dots, (x_n, y_n)]$ from the conditional distribution $P(Y|X)$, the conditional embedding operator can be derived from a least-squares objective (Grünwälder et al 2012) as

$$\hat{C}_{Y|X} = \Phi_y (\mathbf{K} + \lambda \mathbf{I}_n)^{-1} \mathbf{Y}_x^\top, \quad (6)$$

with the feature matrices $\Phi_y := [\phi(y_1), \dots, \phi(y_n)]$ and $\mathbf{Y}_x := [\varphi(x_1), \dots, \varphi(x_n)]$, the Gram matrix $\mathbf{K} = \mathbf{Y}_x^\top \mathbf{Y}_x \in \mathbb{R}^{n \times n}$, the regularization parameter λ , and the identity matrix $\mathbf{I}_n \in \mathbb{R}^{n \times n}$. With the feature mapping of the realization x_* this results in

$$\mu_{Y|x_*} = \hat{C}_{Y|X} \varphi(x_*) = \Phi_y (\mathbf{K} + \lambda \mathbf{I}_n)^{-1} \mathbf{Y}_x^\top \varphi(x_*) = \Phi_y (\mathbf{K} + \lambda \mathbf{I}_n)^{-1} \mathbf{k}_{x_*}, \quad (7)$$

where \mathbf{k}_{x_*} is the kernel vector of the samples $[x_1, \dots, x_n]$ and the realization x_* . As the kernel matrices in the inverse and the kernel vector of the realization are finite, the embedding of the conditional distribution can be represented as a weighted sum of feature mappings

$$\mu_{Y|x_*} = \Phi_y \alpha = \sum_{i=1}^n \alpha_i \phi(y_i), \quad (8)$$

with the feature mappings of the samples $[y_1, \dots, y_n]$ and the finite weight vector $\alpha \in \mathbb{R}^n$.

2.1.3 The Kernel Sum Rule

The embedding of $Q(Y) = \sum_X P(Y|X)\pi(X)$ can be obtained from the *kernel sum rule* (Song et al 2013). To that end, the conditional operator is applied to the embedding $\mu_X^\pi = \Upsilon_x \alpha_\pi$ of the prior distribution $\pi(X)$,

$$\hat{\mu}_Y^\pi = \hat{C}_{Y|X} \mu_X^\pi = \Phi_y (\mathbf{K} + \lambda \mathbf{I}_m)^{-1} \mathbf{K} \alpha_\pi. \quad (9)$$

Here, we can see again that the result can be represented as a weighted sum over feature mappings. In order to obtain the distribution $Q(Y)$ as a covariance operator instead of a mean map, Song et al (2013) also proposed the kernel sum rule for tensor product features which yields the prior modified covariance operator C_{YY}^π as

$$C_{YY}^\pi = C_{(Y|X)} \mu_X^\pi \quad (10)$$

$$C_{YY}^\pi = \Phi_y \text{diag}((\mathbf{K} + \lambda \mathbf{I}_m)^{-1} \mathbf{K} \alpha) \Phi_y^\top, \quad (11)$$

where $C_{(Y|X)}$ is the conditional operator for tensor product features.

2.1.4 The Kernel Chain Rule

The *kernel chain rule* (Song et al 2013) yields an embedding of the joint distribution $Q(X, Y) = P(Y|X)\pi(X)$ as a prior modified covariance operator. There are two versions of the kernel chain rule. Both apply the conditional embedding operator of $P(Y|X)$ to an embedding of the prior distribution $\pi(X)$. In the first version the conditional operator is applied to a covariance embedding of the prior distribution. But, this covariance operator is not estimated directly from samples but approximated from the weight vector α_π of the mean map $\mu_X^\pi = \Upsilon_x \alpha_\pi$ as $\hat{C}_{XX}^\pi = \Upsilon_x \text{diag}(\alpha) \Upsilon_x^\top$. This yields Version (a) of the kernel chain rule as

$$\hat{C}_{YX}^\pi = C_{X|Y} \hat{C}_{XX}^\pi = \Phi_y (\mathbf{K} + \lambda \mathbf{I}_m)^{-1} \mathbf{K} \text{diag}(\alpha) \Upsilon_x^\top. \quad (12)$$

Version (b) of the kernel chain rule first computes the mean map μ_Y^π conditioned on the prior distribution $\pi(X)$ by applying the conditional embedding operator to the mean map μ_X^π . Afterwards, the prior-modified covariance operator of the joint distribution is constructed from the resulting weight vector which results in

$$\hat{C}_{YX}^\pi = \Phi_y \text{diag}((\mathbf{K} + \lambda \mathbf{I}_m)^{-1} \mathbf{K} \alpha) \Upsilon_x^\top. \quad (13)$$

Both versions of the kernel chain rule have been used to derive different versions of the kernel Bayes' rule as we will depict below.

2.1.5 The Kernel Bayes' Rule

Given the embedding of a prior distribution $\pi(X)$ and the feature mapping of an observation $\phi(\mathbf{y})$, the *kernel Bayes' rule* (KBR) infers the mean embedding of the posterior distribution $Q_\pi(X|Y = \mathbf{y})$. The idea is to construct a prior-modified conditional embedding operator that yields the mean map of the posterior if applied to the feature mapping of the observation (Fukumizu et al 2013)

$$\mu_{X|Y}^\pi = C_{X|Y}^\pi \phi(\mathbf{y}). \quad (14)$$

This prior-modified conditional operator is constructed from two prior-modified covariance operators C_{YY}^π and C_{YX}^π obtained from the kernel sum and the kernel chain rule, respectively, using the relation

$$C_{X|Y}^\pi = C_{XY}^\pi (C_{YY}^\pi)^{-1}. \quad (15)$$

In the first version, which we denote by KBR(b) following the notation of Song et al (2013), Fukumizu et al (2013) derived the kernel Bayes' rule using the tensor product conditional operator in the kernel chain rule (c.f. Equation 13) and arrived at

$$\hat{\mu}_{X|y}^\pi = \Upsilon_x \mathbf{D} \mathbf{G} \left((\mathbf{D} \mathbf{G})^2 + \kappa \mathbf{I}_m \right)^{-1} \mathbf{D} \mathbf{g}_y, \quad (16)$$

with the diagonal $\mathbf{D} := \text{diag}((\mathbf{K}_{xx} + \lambda \mathbf{I}_m)^{-1} \mathbf{K} \alpha)$, the gram matrix $\mathbf{G} = \Phi_y^\top \Phi_y$, the kernel vector \mathbf{g}_y and κ as regularization parameter. Song et al (2013) derived the KBR using the first formulation of the kernel chain rule shown in Equation 3.2 which results in

$$\hat{\mu}_{X|y}^\pi = \Upsilon_x \mathbf{\Lambda}^\top \left((\mathbf{D} \mathbf{G})^2 + \kappa \mathbf{I}_m \right)^{-1} \mathbf{G} \mathbf{D} \mathbf{g}_y, \quad (17)$$

with $\mathbf{\Lambda} := (\mathbf{K} + \lambda \mathbf{I}_m)^{-1} \mathbf{K} \text{diag}(\alpha)$. This second version of the kernel Bayes' rule is denoted by KBR(a). As the matrix $\mathbf{D} \mathbf{G}$ is typically not invertible, both of these versions of the KBR use a form of the Tikhonov regularization in which the matrix in the inverse is squared. Boots et al (2013) use a third form of the KBR which is derived analogously to the first version but does not use the squared form of Tikhonov regularization, i.e.,

$$\hat{\mu}_{X|y}^\pi = \Upsilon_x (\mathbf{D} \mathbf{G} + \kappa \mathbf{I}_m)^{-1} \mathbf{D} \mathbf{g}_y. \quad (18)$$

Since the product $\mathbf{D} \mathbf{G}$ is often not positive definite, a strong regularization parameter is required to make the matrix invertible. We denote this third version of the kernel Bayes' rule consequently by KBR(c).

All three versions of the kernel Bayes' rule presented above have drawbacks. First, due to the approximation of the prior modified covariance operators, these operators are not guaranteed to be positive definite and, thus, their inversion requires either a harsh form of the Tikhonov regularization or a strong regularization factor and are often still numerically instable. Furthermore, the inverse is dependent on the embedding of the prior distribution and, hence, needs to be recomputed for every Bayesian update of the mean map. This recomputation significantly increases the computational costs, for example, if we want to apply hyper-parameter optimization techniques.

2.2 The Kalman Filter

The Kalman filter is a well known technique for state estimation, prediction, and smoothing in environments with linear system dynamics that are subject to zero-mean Gaussian noise with known covariances (Kalman 1960). The system equations can be formulated as

$$\mathbf{x}_{t+1} = \mathbf{F} \mathbf{x}_t + \mathbf{v}_t, \quad \mathbf{y}_t = \mathbf{H} \mathbf{x}_t + \mathbf{w}_t, \quad (19)$$

where \mathbf{x}_t is the latent state of the system at time t and \mathbf{y}_t is the corresponding observation. The linear Gaussian model is defined by the system matrix \mathbf{F} , the observation matrix \mathbf{H} , and noise vectors \mathbf{v}_t and \mathbf{w}_t which are sampled from $\mathcal{N}(\mathbf{0}, \mathbf{P})$ and $\mathcal{N}(\mathbf{0}, \mathbf{R})$, respectively.

From the assumption of Gaussian transition noise and Gaussian observation noise, it follows naturally that the belief state over the latent state \mathbf{x}_t is as well a Gaussian distribution with mean $\boldsymbol{\mu}_{\mathbf{x},t}$ and covariance $\boldsymbol{\Sigma}_{\mathbf{x},t}$. The Kalman filter applies iteratively two update procedures to the belief state to which we will refer to as *prediction* and *innovation* update. During the prediction update the Kalman filter propagates the belief state in time by applying the transition model, i.e.,

$$\boldsymbol{\mu}_{\mathbf{x},t+1}^- = \mathbf{F}\boldsymbol{\mu}_{\mathbf{x},t}^+, \quad \boldsymbol{\Sigma}_{\mathbf{x},t+1}^- = \mathbf{F}\boldsymbol{\Sigma}_{\mathbf{x},t}^+\mathbf{F}^\top + \mathbf{P}. \quad (20)$$

On new observations \mathbf{y}_t , the innovation update applies Bayes' theorem to the *a-priori* belief state $\{\boldsymbol{\mu}_{\mathbf{x},t}^-, \boldsymbol{\Sigma}_{\mathbf{x},t}^-\}$ to obtain the *a-posteriori* mean and covariance as

$$\boldsymbol{\mu}_{\mathbf{x},t}^+ = \boldsymbol{\mu}_{\mathbf{x},t}^- + \mathbf{Q}_t(\mathbf{y}_t - \mathbf{H}\boldsymbol{\mu}_{\mathbf{x},t}^-), \quad (21)$$

$$\boldsymbol{\Sigma}_{\mathbf{x},t}^+ = \boldsymbol{\Sigma}_{\mathbf{x},t}^- - \mathbf{Q}_t\mathbf{H}\boldsymbol{\Sigma}_{\mathbf{x},t}^-, \quad (22)$$

with the Kalman gain matrix

$$\mathbf{Q}_t = \boldsymbol{\Sigma}_{\mathbf{x},t}^- \mathbf{H}^\top (\mathbf{H}\boldsymbol{\Sigma}_{\mathbf{x},t}^- \mathbf{H}^\top + \mathbf{R})^{-1}.$$

Another approach to derive the Kalman filter equations follows from a least-squares objective between the state estimator a-priori and a-posteriori to the observation Simon (2006). This second approach does not make the explicit assumption that the belief state can be represented as a Gaussian random variable. Rather this representation follows from the objective to minimize the variance of the error between the a-priori and a-posteriori estimators. We will take this second approach as inspiration to derive the kernel Kalman rule in Section 4.

3 Efficient Nonparametric Inference in a Subspace

A general drawback of kernel methods is that the complexities of the algorithms scale poorly with the number of samples in the kernel matrices. As the conditional embedding operator and the kernel inference rules require the inversion of a kernel matrix, the complexity scales cubically with the number of data points. To overcome this drawback, several approaches exist that aim to find a good trade-off between a compact representation and leveraging from a large data set. Examples are the sparse Gaussian processes that use pseudo-inputs (Snelson and Ghahramani 2006; Csató and Oppé 2002), or a sparse subset of the data which is selected by maximizing the posterior probability (Smola and Bartlett 2001). Other techniques are based on approximating the kernel matrices using the Nyström method (Williams and Seeger 2000) or random Fourier features (Rahimi and Recht 2007). We approach this problem by proposing the subspace conditional embedding operators (Gebhardt et al 2015). The basic idea is to use only a subset of the available training data as representation for the embeddings but the full data set to learn the conditional operators. In the following sections, we will recapitulate this approach and show how it can be applied to the framework for nonparametric inference.

Given the feature matrices $\boldsymbol{\Phi}_y := [\boldsymbol{\phi}(\mathbf{y}_1), \dots, \boldsymbol{\phi}(\mathbf{y}_n)]$ and $\boldsymbol{\Upsilon}_x := [\boldsymbol{\varphi}(\mathbf{x}_1), \dots, \boldsymbol{\varphi}(\mathbf{x}_n)]$, we can define the respective subsets $\boldsymbol{\Psi}_y \subset \boldsymbol{\Phi}_y$ and $\boldsymbol{\Gamma}_x \subset \boldsymbol{\Upsilon}_x$, where $|\boldsymbol{\Psi}_y| = |\boldsymbol{\Gamma}_x| = m \ll n$. We assume that the subsets are representative for the embedded distributions. Similar to the conditional operator discussed in Section 2.1.2, we define the subspace conditional operator $C_{Y|X}^S$ as the mapping from an embedding $\boldsymbol{\varphi}(\mathbf{x}) \in \mathcal{H}_X$ to the mean embedding $\mu_{Y|x} \in \mathcal{H}_Y$ of the conditional distribution $P(Y|\mathbf{x})$ conditioned on a certain variate \mathbf{x} . To obtain this subspace conditional operator, we first introduce an auxiliary conditional operator $C_{Y|X}^{\text{aux}}$ which maps

from the subspace projection of the embedding $\mathbf{\Gamma}_x^\top \boldsymbol{\varphi}(\mathbf{x})$ to the mean map of the conditional distribution, i.e.,

$$\mu_{Y|x} = C_{Y|X}^{\text{aux}} \mathbf{\Gamma}_x^\top \boldsymbol{\varphi}(\mathbf{x}). \quad (23)$$

We can then derive this auxiliary conditional operator by minimizing the squared error on the full data set

$$C_{Y|X}^{\text{aux}} = \arg \min_C \|\Phi_y - C \mathbf{\Gamma}_x^\top \Upsilon_x\|_2 \quad (24)$$

$$= \Phi_y \Upsilon_x^\top \mathbf{\Gamma}_x (\mathbf{\Gamma}_x^\top \Upsilon_x \Upsilon_x^\top \mathbf{\Gamma}_x + \lambda \mathbf{I})^{-1}. \quad (25)$$

Substituting this result for the auxiliary conditional operator in Equation 23 gives us the subspace conditional operator as

$$\begin{aligned} C_{Y|X}^S &= C_{Y|X}^{\text{aux}} \mathbf{\Gamma}_x^\top \\ &= \Phi_y \bar{\mathbf{K}} \left(\bar{\mathbf{K}}^\top \bar{\mathbf{K}} + \lambda \mathbf{I} \right)^{-1} \mathbf{\Gamma}_x^\top, \end{aligned} \quad (26)$$

where $\bar{\mathbf{K}} = \Upsilon_x^\top \mathbf{\Gamma}_x \in \mathbb{R}^{n \times m}$ is the kernel matrix of the sample feature set Υ_x and its subset $\mathbf{\Gamma}_x$. Since we assume that $m \ll n$, the inverse in the subspace conditional operator is in $\mathbb{R}^{m \times m}$ and, thus, of a much smaller size than the inverse in the standard conditional operator shown in Equation 6. Additionally, we can make use the feature matrix $\mathbf{\Gamma}_x^\top$ on the right hand side of the subspace conditional operator and represent the mean embedding always in the subspace spanned by the features $\mathbf{\Gamma}_x$. This allows to avoid representations and computations in the high-dimensional space spanned by the features of the full sample set. In the following sections we will rederive the kernel sum rule, the kernel chain rule and the kernel Bayes rule based on the subspace conditional embedding operator analogously to Song et al (2013).

3.1 Relation to Other Sparsification Approaches

Many other sparsification techniques for kernel methods exist. The two most important techniques among these are probably the Nyström method (Williams and Seeger 2000; Drineas and Mahoney 2005) and the random Fourier features (Rahimi and Recht 2007). Both methods are closely related to our approach. However, a major difference is that the subspace conditional operator still is an operator in the RKHS defined by the kernel function. The Nyström method as well as the random Fourier features use finite dimensional approximation of the feature function intrinsic to the kernel function and thus the resulting operators are finite matrices that operate on the finite dimensional weight vectors of the mean embedding. See Appendix A.2 on how these sparsification approaches can be applied to the conditional operator.

3.2 The Subspace Kernel Sum Rule

Analogously to Song et al (2013), the subspace kernel sum rule is simply the subspace conditional operator applied to the embedding of a distribution $\pi(X)$, i.e.,

$$\mu_Y^\pi = C_{Y|X}^S \mu_X^\pi = \Phi_y \bar{\mathbf{K}} \left(\bar{\mathbf{K}}^\top \bar{\mathbf{K}} + \lambda \mathbf{I} \right)^{-1} \mathbf{\Gamma}_x^\top \Upsilon_x \alpha, \quad (27)$$

where $\mu_X^\pi = \Upsilon_x \alpha$ is the embedding of the prior distribution $\pi(X)$. We construct the subspace kernel sum rule for tensor product features differently than Song et al (2013). Instead of applying the conditional operator to the mean embedding and then approximating the covariance operator with the resulting weights (c.f. Equation 11), we first approximate the covariance operator C_{XX}^π from the weights α of the mean map μ_X^π and then multiply the subspace conditional operator to both sides of it

$$\begin{aligned} C_{YY}^{S,\pi} &= C_{Y|X}^S C_{XX}^\pi \left(C_{Y|X}^S \right)^\top = C_{Y|X}^S \Upsilon_x \text{diag}(\alpha) \Upsilon_x^\top \left(C_{Y|X}^S \right)^\top \\ &= \Phi_y \bar{K} L_{\bar{K}} \text{diag}(\alpha) L_{\bar{K}}^\top \bar{K}^\top \Phi_y^\top. \end{aligned} \quad (28)$$

Here, we denote $L_{\bar{K}} := \left(\bar{K}^\top \bar{K} + \lambda I \right)^{-1} \bar{K}^\top \in \mathbb{R}^{m \times n}$ to keep the notation uncluttered. This definition of the subspace kernel sum rule follows from the kernel chain rule for tensor product features, where the conditional operator $C_{Y|X}$ is applied to the covariance embedding C_{XX}^π to obtain the covariance embedding C_{YX}^π (c.f. Equation). The subspace kernel sum rule follows from applying the transpose of the condition operator a second time from the right-hand side.

3.3 The Subspace Kernel Chain Rule

The subspace kernel chain rule is a straight forward modification of the kernel chain rule by Song et al (2013). We simply apply the subspace conditional operator $C_{Y|X}^S$ from the left side to a covariance operator $C_{XX}^\pi = \Upsilon_x \text{diag}(\alpha) \Upsilon_x^\top$ approximated from the weights α of the prior mean map μ_X^π

$$C_{YX}^{S,\pi} = C_{Y|X}^S C_{XX}^\pi = \Phi_y \bar{K} L_{\bar{K}} \text{diag}(\alpha) \Upsilon_x^\top. \quad (29)$$

With the subspace kernel sum rule and the subspace kernel chain rule we can now construct the subspace kernel Bayes' rule.

3.4 The Subspace Kernel Bayes' Rule

The Bayes' rule computes a posterior distribution $P(X|Y)$ from a prior distribution $\pi(X)$ and a likelihood function $P(Y|X)$. Fukumizu et al (2013) derive a conditional operator $C_{X|Y}^\pi$ from the prior modified covariance operators C_{XY}^π and C_{YY}^π . We follow this approach and construct the subspace kernel Bayes' rule (subKBR) from the prior modified covariance operators $C_{XY}^{S,\pi}$ and $C_{YY}^{S,\pi}$ which we obtain from the subspace kernel chain rule and the subspace kernel sum rule for tensor product features, respectively. When applied to the embedding of a variate y , the subspace kernel Bayes' rule returns the mean embedding of the conditional distribution $P(X|y)$ as

$$\mu_{X|y}^\pi = C_{X|Y}^{S,\pi} \phi(y) \quad (30)$$

$$\mu_{X|y}^\pi = C_{XY}^{S,\pi} \left(C_{YY}^{S,\pi} \right)^{-1} \phi(y). \quad (31)$$

From the subspace kernel chain rule, we obtain

$$C_{XY}^{S,\pi} = \left(C_{Y|X}^S C_{XX}^\pi \right)^\top \quad (32)$$

$$= \Upsilon_x \text{diag}(\alpha) L_{\bar{K}}^\top \bar{K}^\top \Phi_y^\top \quad (33)$$

and from the subspace kernel sum rule

$$C_{YY}^{S,\pi} = C_{Y|X}^S C_{XX}^\pi \left(C_{Y|X}^S \right)^\top \quad (34)$$

$$= \Phi_y \bar{K} L_{\bar{K}} \text{diag}(\alpha) L_{\bar{K}}^\top \bar{K}^\top \Phi_y^\top. \quad (35)$$

To keep the notation of the subspace kernel Kalman rule uncluttered, we define the following matrices

$$\bar{\Lambda} := \text{diag}(\alpha) L_{\bar{K}}^\top \quad (36)$$

$$\bar{D} := L_{\bar{K}} \text{diag}(\alpha) L_{\bar{K}}^\top \in \mathbb{R}^{m \times m}, \quad (37)$$

$$E := \bar{K}^\top G \bar{K} \in \mathbb{R}^{m \times m}, \quad (38)$$

where $G = \Phi_y^\top \Phi_y$ is the kernel matrix of the samples y_i . Using the same form of the Tikhonov regularization as the kernel Bayes' rule in Fukumizu et al (2011) and substituting the prior modified subspace covariance operators from Equations 33 and 35 results in

$$\mu_{X|y}^\pi = C_{XY}^{S,\pi} \left[\left(C_{YY}^{S,\pi} \right)^2 + \gamma I \right]^{-1} C_{YY}^{S,\pi} \phi(y) \quad (39)$$

$$= \Upsilon_x \bar{\Lambda} \bar{K}^\top \Phi_y^\top \left[\left(\Phi_y \bar{K} \bar{D} \bar{K}^\top \Phi_y^\top \right) \Phi_y \bar{K} \bar{D} \bar{K}^\top \Phi_y^\top + \gamma I \right]^{-1} \Phi_y \bar{K} \bar{D} \bar{K}^\top \Phi_y^\top \phi(y) \quad (40)$$

$$= \Upsilon_x \bar{\Lambda} E \left[(\bar{D} E)^2 + \gamma I \right]^{-1} \bar{D} \bar{K}^\top g_y, \quad (41)$$

where we apply the matrix identity $A(BA + \lambda I)^{-1} = (AB + \lambda I)^{-1}A$ with $A = \Phi_y \bar{K}$ to obtain a finite matrix in the inverse. The term $g_y = \Phi_y^\top \phi(y)$ denotes the kernel vector of the new observation y . Since E and \bar{D} are both in $\mathbb{R}^{m \times m}$, the matrix inversion is only in $O(m^3)$ instead of $O(n^3)$. The whole kernel Bayes' rule is in $O(nm^2)$ and, thus, scales linearly with the number of sample points (given a fixed reference set) instead of cubically as for the original kernel Bayes' rule.

3.5 Experimental Evaluation

We compare the performance, learning time and run time of the subspace kernel Bayes' filter in comparison to the standard kernel Bayes' filter with a simple toy task. We simulate a pendulum which we randomly initialize in the ranges $[0.1\pi, 0.4\pi]$ and $[-0.5\pi, 0.5\pi]$ for the angle θ and the angular velocity $\dot{\theta}$, respectively. The pendulum has a mass of 5kg and a friction coefficient of 1. We apply Gaussian white noise to the system with a variance of 1, and to the observations with a variance of 0.1. Additionally, the observed angles are randomly perturbed by an offset of $\pi/4$. These random perturbations occur with a probability of 0.1 in every time step. Each episode consists of 30 time steps with $\Delta t = 0.1$.

Figure 1 shows that the subspace KBF has a slightly better performance when the training set equals the subspace set and maintains the performance of the standard KBF with an increasing number of training samples while the subspace set is fixed to 100 samples.




Fig. 1: The subspace variant of the kernel Bayes' filter outperforms the standard kernel Bayes' filter in both, the training time depicted in the left plot and the run time depicted in the middle plot, while maintaining a similar performance to the standard kernel Bayes' rule as depicted in the right plot. The size of the subspace is fixed to 100 samples.

However, at the same time the learning time of the subspace KBF increases at a much lower magnitude and the run time is nearly constant while the learning and run time of the standard KBR grow cubically. The samples for the subspace kernel Bayes rule are drawn uniformly without replacement from the full sample set.

4 The Kernel Kalman Rule

In this section we will present the *kernel Kalman rule* (KKR) as an alternative to the kernel Bayes' rule (Fukumizu et al 2011). We assume a prior belief state over a variable X embedded in a Hilbert space \mathcal{H}_k as

$$\mu_{X,t}^- = \mathbb{E}_{X_t | y_{1:t-1}}[\varphi(X)]$$

and new measurement y_t embedded in a Hilbert space \mathcal{H}_g as $\phi(y_t)$. With the kernel Kalman rule we want to infer the embedding of the posterior belief state

$$\mu_{X,t}^+ = \mathbb{E}_{X_t | y_{1:t}}[\varphi(X)] \in \mathcal{H}_k$$

from the prior belief and the new measurement. The derivations for the KKR are inspired by the ansatz from recursive least squares (Gauss 1823; Sorenson 1970; Simon 2006), and thus the resulting update equations follow from a clear optimality criterion.

4.1 Derivation using Recursive Least Squares

Let $C_{Y|X}$ be a conditional embedding operator of the observation model $P(Y|X)$ that yields for a given belief state embedded in the Hilbert space \mathcal{H}_k the distribution over possible observations embedded into the Hilbert space \mathcal{H}_g . We call this conditional embedding operator also *observation operator*. The objective of the KKR is then to find the mean embedding μ_X that minimizes the squared error

$$\sum_t \left(\phi(y_t) - C_{Y|X} \mu_X \right)^T \mathcal{R}^{-1} \left(\phi(y_t) - C_{Y|X} \mu_X \right), \quad (42)$$

with some metric \mathcal{R} , in an iterative fashion. In each iteration, we want to update the prior mean map $\mu_{X,t}^-$ based on the measurement y_t to obtain the posterior mean map $\mu_{X,t}^+$.

From the recursive least squares solution, we know that the update rule for obtaining the posterior mean map $\mu_{X,t}^+$ is

$$\mu_{X,t}^+ = \mu_{X,t}^- + Q_t \left(\phi(y_t) - C_{Y|X} \mu_{X,t}^- \right), \quad (43)$$

where Q_t is the Hilbert space *Kalman gain operator* that is applied to the correction term $\delta_t = \phi(y_t) - C_{Y|X} \mu_{X,t}^-$. We call this rule the *kernel Kalman rule* (KKR). It remains to show how to obtain the Kalman gain operator. As we will show in the next paragraphs that this rule is an unbiased estimator of the posterior mean map, we cannot obtain the optimal Q_t by minimizing the error directly. Instead, we will show in Section 4.1.2, how we can find an optimal Q_t by minimizing the covariance of the error instead.

4.1.1 The Kernel Kalman Update is an Unbiased Estimator of the Posterior Mean Map

The embedding of the observation $\phi(y_t)$ in the correction term δ_t is a single-sample estimator of the embedding of the true distribution over observations $\mu_{Y,t}^-$. We make the assumption that the error of this single-sample estimator is independent from the state x_t , as we can obtain the embedding of the distribution over the observations from the observation operator and the prior distribution $\mu_{X,t}^-$ as

$$\phi(y_t) = \mu_{Y|X,t}^- + \zeta_t = C_{Y|X} \mu_{X,t}^- + \zeta_t, \quad (44)$$

where ζ_t denotes the error of the single sample estimator to the embedding of the true distribution. By taking the expectation it is easy to show that the error of the single-sample estimator is zero-mean and thus $\phi(y_t)$ is an unbiased estimator for $\mu_{Y|X,t}^-$. We refer to Appendix B.1 for a more detailed derivation. Similarly, the error of the a-posteriori estimator to the embedding of the true state is given as

$$\varepsilon_t^+ = \varphi(x_t) - \mu_{X,t}^+ \quad (45)$$

$$= \varphi(x_t) - \mu_{X,t}^- - Q_t (\phi(y_t) - C_{Y|X} \mu_{X,t}^-), \quad (46)$$

where we use Equation 43 to substitute the embedding of the posterior belief. By substituting $\phi(y_t)$ with Equation 44 and defining the error of the a-priori estimator analogously as $\varepsilon_t^- = \varphi(x_t) - \mu_{X,t}^-$, we arrive at

$$\varepsilon_t^+ = \varphi(x_t) - \mu_{X,t}^- - Q_t (C_{Y|X} \varphi(x_t) + \zeta_t - C_{Y|X} \mu_{X,t}^-) \quad (47)$$

$$= \left(I - Q_t C_{Y|X} \right) (\varphi(x_t) - \mu_{X,t}^-) - Q_t \zeta_t \quad (48)$$

$$= \left(I - Q_t C_{Y|X} \right) \varepsilon_t^- - Q_t \zeta_t \quad (49)$$

with identity operator I . We can now apply the expectation operator and exploit its linearity to obtain

$$\begin{aligned} \mathbb{E} [\varepsilon_t^+] &= \mathbb{E} \left[\left(I - Q_t C_{Y|X} \right) \varepsilon_t^- - Q_t \zeta_t \right] \\ &= \left(I - Q_t C_{Y|X} \right) \mathbb{E} [\varepsilon_t^-] - Q_t \mathbb{E} [\zeta_t]. \end{aligned} \quad (50)$$

Since the residual of the observation operator is zero mean ($\mathbb{E}[\zeta_t] = 0$), we see that, given an unbiased a-priori estimator ($\mathbb{E}[\varepsilon_t^-] = 0$), the a-posteriori estimator obtained from the kernel Kalman update is unbiased ($\mathbb{E}[\varepsilon_t^+] = 0$) independent of the choice of Q . Thus, we cannot use the expected error of the estimator as an optimality criterion for the Kalman gain operator.

4.1.2 Finding the Optimal Kernel Kalman Gain Operator

If the expected error—or the first moment of the error distribution—is already zero, taking the covariance of the error—or the second moment—is a consequent choice. Hence, we chose the kernel Kalman gain operator \mathbf{Q}_t which minimizes the expected squared loss $\mathbb{E}[(\boldsymbol{\varepsilon}_t^+)^T \boldsymbol{\varepsilon}_t^+]$ or equivalently the variance of the estimator. The objective for minimizing the variance can also be reformulated as minimizing the trace of the a-posteriori covariance operator $C_{XX,t}^+$ of the state \mathbf{x}_t at time t , i.e., $\min_{\mathbf{Q}_t} \mathbb{E}[(\boldsymbol{\varepsilon}_t^+)^T \boldsymbol{\varepsilon}_t^+] = \min_{\mathbf{Q}_t} \text{Tr } C_{XX,t}^+$. Using the formulation of the posterior error from Equation 49 and the independence assumption of $\boldsymbol{\zeta}_t$ and $\boldsymbol{\varepsilon}_t$ allows us to reformulate the a-posteriori covariance operator as

$$C_{XX,t}^+ = (\mathbf{I} - \mathbf{Q}_t C_{Y|X}) C_{XX,t}^- (\mathbf{I} - \mathbf{Q}_t C_{Y|X})^T + \mathbf{Q}_t \mathcal{R} \mathbf{Q}_t^T, \quad (51)$$

where $\mathcal{R} = \mathbb{E}[\boldsymbol{\zeta}_t \boldsymbol{\zeta}_t^T]$ is the covariance of the residual of the observation operator. Taking the derivative of the trace of the covariance operator and setting it to zero leads to the solution for the kernel Kalman gain operator

$$\mathbf{Q}_t = C_{XX,t}^- C_{Y|X}^T (C_{Y|X} C_{XX,t}^- C_{Y|X}^T + \mathcal{R})^{-1}. \quad (52)$$

We provide a detailed derivation of the optimal Kalman gain operator in Appendix B.2. From Equations 51 and 52, we can also see that it is possible to recursively estimate the covariance embedding operator independently of the mean map or of the observations. This property will allow us later to precompute the covariance embedding operator as well as the kernel Kalman gain operator to further improve the computational complexity of our algorithm. Following Simon (2006), the update of the covariance operator can be further simplified to

$$C_{XX,t}^+ = C_{XX,t}^- - \mathbf{Q}_t C_{Y|X} C_{XX,t}^-. \quad (53)$$

The derivations of this simplification can be found in Appendix B.3. In the following section we will show how to obtain the empirical Kalman update rule from a finite data set.

4.2 Empirical Kernel Kalman Rule

The equations for the kernel Kalman rule that we derived in the previous section are based on embeddings in infinite dimensional Hilbert spaces and operators that map between these spaces. In practice, these embeddings and operators are estimated from a finite set of samples $\mathcal{D} = [(\mathbf{x}_1, \mathbf{y}_1), \dots, (\mathbf{x}_n, \mathbf{y}_n)]$. In this section we will show how we can reformulate the kernel Kalman rule with manipulations of finite matrices by applying the kernel trick, i.e., matrix identities. Based on the data set \mathcal{D} and the corresponding feature matrices \mathbf{Y}_x and $\mathbf{\Phi}_y$, the finite sample estimators of the prior mean embedding and the prior covariance operator are given as

$$\hat{\boldsymbol{\mu}}_{X,t}^- = \mathbf{Y}_x \mathbf{m}_t^- \quad \text{and} \quad \hat{C}_{XX,t}^- = \mathbf{Y}_x \mathbf{S}_t^- \mathbf{Y}_x^T, \quad (54)$$

respectively, with weight vector \mathbf{m}_t^- and positive definite weight matrix \mathbf{S}_t^- . Using this finite sample estimator of the covariance operator, the finite sample estimator of the conditional operator from Equation 6, and by approximating the covariance of the residual of the observation operator with a small diagonal $\mathcal{R} = \kappa \mathbf{I}$, we can rewrite the kernel Kalman gain operator as

$$\mathbf{Q}_t = \mathbf{Y}_x \mathbf{S}_t^- \mathbf{O}^T \mathbf{\Phi}_y^T (\mathbf{\Phi}_y \mathbf{O} \mathbf{S}_t^- \mathbf{O}^T \mathbf{\Phi}_y^T + \kappa \mathbf{I})^{-1}, \quad (55)$$

with the observation matrix $\mathbf{O} = (\mathbf{K} + \lambda\mathbf{I})^{-1}\mathbf{K}$. Here, the approximation of $\mathcal{R} = \kappa\mathcal{I}$ also acts as a small regularization in the inverse to ensure its positive definiteness and to improve the numerical stability of the kernel Kalman rule. However, \mathbf{Q}_t still requires the inversion of an infinite dimensional matrix. Using matrix identities, we can solve this problem and arrive at

$$\mathbf{Q}_t = \Upsilon_x \underbrace{\mathbf{S}_t^- \mathbf{O}^\top (\mathbf{GOS}_t^- \mathbf{O}^\top + \kappa\mathbf{I})^{-1}}_{\mathbf{Q}_t} \Phi_y^\top, \quad (56)$$

where we defined $\mathbf{Q}_t = \mathbf{S}_t^- \mathbf{O}^\top (\mathbf{GOS}_t^- \mathbf{O}^\top + \kappa\mathbf{I})^{-1} \in \mathbb{R}^{n \times n}$ with the Gram matrix of the observations $\mathbf{G} = \Phi_y^\top \Phi_y$. Based on this reformulation of the kernel Kalman gain operator, we can obtain finite vector/matrix representations of the update equations for the estimator of the mean embedding (Equation 43) and the estimator of the covariance operator (Equation 53). For the weight vector \mathbf{m}_t , we arrive at

$$\mathbf{m}_t^+ = \mathbf{m}_t^- + \mathbf{Q}_t (\mathbf{g}_{y_t} - \mathbf{G}\mathbf{O}\mathbf{m}_t^-), \quad (57)$$

where $\mathbf{g}_{y_t} = \Phi_y^\top \phi(y_t)$ is the embedding of the measurement at time t . Similarly, we can also obtain the update equation for the weight matrix \mathbf{S}_t as

$$\mathbf{S}_t = \mathbf{S}_t^- - \mathbf{Q}_t \mathbf{GOS}_t^-. \quad (58)$$

The algorithm requires the inversion of a $m \times m$ matrix in every iteration for computing the kernel Kalman gain matrix \mathbf{Q}_t . Hence, similar to the kernel Bayes' rule, the computational complexity of a straightforward implementation would scale cubically with the number of data points m . However, in contrast to the KBR, the inverse in \mathbf{Q}_t is only dependent on time and not on the estimate of the mean map. Thus, the kernel Kalman gain matrix can be precomputed since it is identical for multiple parallel runs of the algorithm. Furthermore, if the stream of incoming measurements is reliable (no time steps without incoming measurement), \mathbf{S}_t will converge to a stationary matrix and by that \mathbf{Q}_t will become stationary as well. While many applications do not require to perform state estimations in parallel, it is a huge advantage for hyper-parameter optimization as we can evaluate multiple trajectories from a validation set simultaneously. As for most kernel-based methods, hyper-parameter optimization is crucial for scaling the approach to complex tasks. So far, the hyper-parameters of the kernels for the KBF have typically been set by heuristics as optimization would be too expensive.

4.3 The Subspace Kernel Kalman Rule

In Section 3 we have already shown how we can apply the subspace conditional embedding operator to leverage from large data sets but at the same time maintain the computational tractability of the learned models. In this section, we will now show how we can apply this technique to the kernel Kalman rule to obtain the *subspace kernel Kalman rule* (subKBR).

A core difference between the KBR and the subKBR is the representation of the embedded distributions. While we represent the embeddings for the kernel Kalman rule as weight vector \mathbf{m}_t^- and weight matrix \mathbf{S}_t^- , we use the projections into a subspace

$$\mathbf{n}_t = \Gamma_x^\top \mu_t = \Gamma_x^\top \Upsilon_x \mathbf{m}_t = \bar{\mathbf{K}}^\top \mathbf{m}_t, \quad (59)$$

$$\mathbf{P}_t = \Gamma_x^\top C_{XX,t} \Gamma_x = \Gamma_x^\top \Upsilon_x \mathbf{S}_t \Upsilon_x^\top \Gamma_x = \bar{\mathbf{K}}^\top \mathbf{S}_t \bar{\mathbf{K}}. \quad (60)$$

to represent the distribution for the subspace kernel Kalman rule. These projections will later allow us to express all operations in the lower dimensional subspace instead of the

space spanned by the full data set. Matrix manipulations with the full data set are then only necessary during the learning phase of the KKR not while performing inference.

We use a slightly modified version of the kernel Kalman gain from Equation 52 where we approximate the covariance operator \mathcal{R} with a diagonal matrix $\kappa\mathcal{I}$. With the subspace conditional embedding operator of the distribution $P(Y|X)$, as derived in Section 2.1.2

$$C_{Y|X}^S = \Phi_y \bar{\mathbf{K}} \left(\bar{\mathbf{K}}^\top \bar{\mathbf{K}} + \lambda \mathbf{I} \right)^{-1} \Gamma_x^\top, \quad (61)$$

we obtain the subspace kernel Kalman gain operator as

$$\mathbf{Q}_t^S = C_{XX,t}^- \left(C_{Y|X}^S \right)^\top \left(C_{Y|X}^S C_{XX,t}^- \left(C_{Y|X}^S \right)^\top + \kappa \mathcal{I} \right)^{-1} \quad (62)$$

We can further derive a finite matrix representation of the operator using matrix identities and the projection into the subspace spanned by the features Γ_x^\top as

$$\Gamma_x^\top \mathbf{Q}_t^S = \underbrace{P_t^- (\mathbf{O}^S)^\top \left(\bar{\mathbf{K}}^\top \mathbf{G} \bar{\mathbf{K}} \mathbf{O}^S P_t^- (\mathbf{O}^S)^\top + \kappa \mathbf{I} \right)^{-1} \bar{\mathbf{K}}^\top \Phi_y^\top}_{\mathbf{Q}_t^S}, \quad (63)$$

where we define the subspace kernel Kalman gain matrix \mathbf{Q}_t^S using the short hand $\mathbf{O}^S := (\bar{\mathbf{K}}^\top \bar{\mathbf{K}} + \lambda \mathbf{I})^{-1}$. A detailed derivation can be found in Appendix B.5. Note that $\mathbf{Q}_t^S \in \mathbb{R}^{m \times n}$ and not $\mathbb{R}^{m \times m}$, however when applying the subspace KKR in an inference algorithm, we can use the matrix $\bar{\mathbf{K}}^\top$ on the right side as a projection for the high-dimensional embedding of the distribution over the variable Y to which the gain is applied (see Algorithm 2 for an example). From here, the update equation for the projection of the mean map becomes

$$\mathbf{n}_{X,t}^+ = \mathbf{n}_{X,t}^- + \mathbf{Q}_t^S \left(\mathbf{g}_{:y_t} - \mathbf{G} \bar{\mathbf{K}} \mathbf{O}^S \mathbf{n}_t^- \right), \quad (64)$$

And similarly, we can derive the update equation for the covariance embedding as

$$\mathbf{P}_t^+ = \mathbf{P}_t^- - \mathbf{Q}_t^S \mathbf{G} \bar{\mathbf{K}} \mathbf{O}^S \mathbf{P}_t^- \quad (65)$$

In contrast to the kernel Kalman gain presented in the previous section, but also in contrast to the kernel Bayes rules discussed in Section 2.1.5, the subspace kernel Kalman gain requires only the inversion of an $m \times m$ matrix instead of an $n \times n$ matrix, where $m \ll n$. Still, the full data set of n samples can be used to learn the Kalman gain operator.

4.4 Experimental Comparison of (sub)KKR and (sub)KBR

We compare the performance of the (subspace) kernel Kalman rule to the performance of the (subspace) kernel Bayes rule on a simple stationary filtering task for estimating the expectation of a Gaussian distribution. We sample 500 context variables c_i uniformly from the interval $[-5, 5]$ as the mean of the Gaussian distributions. Afterwards we draw one single sample s_i for each context from $\mathcal{N}(c_i, \frac{1}{3})$ and learn the kernel Kalman rule and the different versions of the kernel Bayes rule with the context variables as states and the samples as observations. For the performance comparison (Figure 2), the KKR and the KBR are learned with a kernel size of 200 samples, subKKR and subKBR are both learned with 200 samples to span the subspace and the full set of 500 samples to learn the operators. In the comparison of time efficiency, the respective kernel size and subspace size is denoted as column header,




Fig. 2: Performance of the KKR updates vs the KBR updates for estimating the mean of a Gaussian distribution with 1-10 seen samples. The ML estimate serves as a baseline. Depicted are the median and the (0.15, 0.85)-quantiles of the MSE to the true mean over 20 runs.

#	200	300	400	500
KKR	0.1110	0.3360	0.7155	1.3965
subKKR	0.1820	0.4655	0.9130	1.6075
KBR(a)	2.9395	9.2395	21.5035	41.6955
KBR(b)	0.8695	2.5840	5.7580	10.9900
KBR(c)	0.5455	1.5575	3.3695	6.4465
subKBR	2.3530	4.7625	7.6540	12.7960

Table 1: Time consumptions of the KKR and KBR update methods for different kernel sizes. The update was performed on 10 samples from 10 different distributions. The subspace KKR and KBR updates are trained with 500 samples in the full data set. Both KKR methods outperform the KBR methods clearly as they are able to process the update on the 10 different distributions in parallel.

while the subKKR and subKBR have always been learned with the full data set of 500 samples. All data were drawn uniformly without replacement from the full data set.

To evaluate the methods, we generate a data set with ten context variables from the same uniform distribution. Now, we draw ten samples from the Gaussian around each context and update each method iteratively with these ten samples. For each update we compute the squared error to the true context and take the mean over all ten context variables. We use squared exponentials as kernel functions and optimize their bandwidths as well as the regularization parameters using CMA-ES (Hansen 2006). Figure 2 shows the median and the (0.15, 0.85)-quantiles of the MSE to the true context over the number of seen samples. As a baseline we depict the maximum-likelihood (ML) estimate of the expectation. We see that while in the beginning all methods perform similar to the ML estimate, with more seen samples KKR and subKKR outperform all variants of the KBR. Moreover the choice to depict the median and (0.15, 0.85)-quantiles over mean and standard deviation is due to the instable optimization behavior of the KBR which produced a lot of outliers. In Table 1 we state the time consumed to perform ten KKR/KBR updates on 10 estimation tasks for different kernel sizes. Here, the KKR/subKKR methods benefit from their ability to process the updates for all 10 estimation tasks in parallel. Yet, the ability to precompute $\mathbf{Q}_t/\mathbf{Q}_t^S$ and $\mathbf{S}_t/\mathbf{P}_t$ is not even exploited.

5 Applications of the Kernel Kalman Rule

In Section 4, we have shown how we can derive the KKR as an alternative operator for Bayesian updates in the framework for nonparametric inference. In this section we will present two applications of the kernel Kalman rule. In Section 5.1 we will first present the *kernel Kalman filter* (KKF) and discuss details about the implementation. A subspace variate of the KKF is presented in Section 5.2 and experimental results of both are shown in Section 5.3. The *kernel forward backward smoother* (KFBS) is presented as another application of the KKR in Section 5.4 and a subspace variate is discussed in Section 5.5. We finally show experimental evaluations of the KFBS and the subKFBS in Section 5.6.

5.1 The Kernel Kalman Filter

Similar to the kernel Bayes' filter (Fukumizu et al 2013; Song et al 2013), we can combine the kernel Kalman rule with the kernel sum rule to formulate the *kernel Kalman filter* (KKF). To learn the models of the KKF, we assume a data set $\mathcal{D} = \{(\tilde{x}_1, x_1, y_1), \dots, (\tilde{x}_m, x_m, y_m)\}$ consisting of triples with preceding state \tilde{x}_i , state x_i , and measurement y_i as given. We further assume the states to be Markov, i.e., the state x_i is only dependent on its predecessor \tilde{x}_i . Based on this data set we define the feature matrices $\mathbf{Y}_x = [\boldsymbol{\varphi}(x_1), \dots, \boldsymbol{\varphi}(x_m)]$, $\mathbf{Y}_{\tilde{x}} = [\boldsymbol{\varphi}(\tilde{x}_1), \dots, \boldsymbol{\varphi}(\tilde{x}_m)]$, and $\mathbf{\Phi}_y = [\boldsymbol{\phi}(y_1), \dots, \boldsymbol{\phi}(y_m)]$. In contrast to the KBF, we represent the belief state as mean map $\mu_{X,t} = \mathbf{Y}_x \mathbf{m}_t$ and as covariance operator $C_{XX,t} = \mathbf{Y}_x \mathbf{S}_t \mathbf{Y}_x^\top$.

The forward model $P(X|\tilde{X})$ that propagates the posterior belief state at time t to the prior belief state at time $t + 1$ can then be learned as conditional embedding operator

$$C_{X|\tilde{X}} = \mathbf{Y}_x (\mathbf{K}_{\tilde{x}\tilde{x}} + \lambda \mathbf{I})^{-1} \mathbf{Y}_{\tilde{x}}^\top, \quad (66)$$

which we also call transition operator. Here, $\mathbf{K}_{\tilde{x}\tilde{x}}$ is the Gram matrix of the features of the preceding states $\mathbf{Y}_{\tilde{x}}$. The posterior belief state at time t is then propagated to the prior belief state at time $t + 1$ time by applying the kernel sum rule. That is, we apply the transition operator to the posterior mean map and the posterior covariance embedding at time t and obtain prior mean map and prior covariance embedding at time $t + 1$, i.e.,

$$\mu_{X,t+1} = C_{X|\tilde{X}} \mu_{X,t}^+ = \mathbf{Y}_x \mathbf{T} \mathbf{m}_t^+, \quad \Leftrightarrow \quad \mathbf{m}_{t+1}^- = \mathbf{T} \mathbf{m}_t^+ \quad (67)$$

$$C_{XX,t+1}^- = C_{X|\tilde{X}} C_{XX,t}^+ C_{X|\tilde{X}}^\top + \mathcal{V} \quad \Leftrightarrow \quad \mathbf{S}_{t+1}^- = \mathbf{T} \mathbf{S}_t^+ \mathbf{T}^\top + \mathbf{V}. \quad (68)$$

Note that the propagation of the covariance embedding is slightly different to the kernel sum rule by Song et al (2013), however this formulation follows directly from the kernel chain rule (c.f. Equations 28 and 11). Analog to the observation matrix \mathbf{O} (c.f. Section 4.2), we denote the transition matrix $\mathbf{T} = (\mathbf{K}_{\tilde{x}\tilde{x}} + \lambda \mathbf{I})^{-1} \mathbf{K}_{\tilde{x}x}$, where $\mathbf{K}_{\tilde{x}x} = \mathbf{Y}_{\tilde{x}}^\top \mathbf{Y}_x$ is the kernel matrix of the preceding states and the states. The covariance of the transition residual \mathcal{V} and its finite matrix representation \mathbf{V} can be obtained as

$$\mathcal{V} = \frac{1}{m} \left(C_{X|\tilde{X}} \mathbf{Y}_x - \mathbf{Y}_{\tilde{x}} \right) \left(C_{X|\tilde{X}} \mathbf{Y}_x - \mathbf{Y}_{\tilde{x}} \right)^\top \quad (69)$$

$$= \frac{1}{m} \left(\mathbf{Y}_x (\mathbf{K}_{\tilde{x}\tilde{x}} + \lambda \mathbf{I})^{-1} \mathbf{Y}_{\tilde{x}}^\top \mathbf{Y}_x - \mathbf{Y}_{\tilde{x}} \right) \left(\mathbf{Y}_x (\mathbf{K}_{\tilde{x}\tilde{x}} + \lambda \mathbf{I})^{-1} \mathbf{Y}_{\tilde{x}}^\top \mathbf{Y}_x - \mathbf{Y}_{\tilde{x}} \right)^\top \quad (70)$$

$$= \mathbf{Y}_{\tilde{x}} \frac{1}{m} \underbrace{\left((\mathbf{K} + \lambda \mathbf{I})^{-1} \mathbf{K} - \mathbf{I} \right) \left((\mathbf{K} + \lambda \mathbf{I})^{-1} \mathbf{K} - \mathbf{I} \right)^\top}_{\mathbf{V}} \mathbf{Y}_{\tilde{x}}^\top \quad (71)$$

On the new prior belief state that we obtain from the transition update, we can afterwards apply the kernel Kalman rule as observation update. Before we give a condensed summary of the kernel Kalman filter in Algorithm 1, we will discuss how we obtain the embedding of the distribution over the initial states in the next section. To extract some meaningful information from the RKHS-embedded distributions, we furthermore need to find mapping of the embedded distribution back into the state space. In Section 5.1.2, show how we approached the so-called *preimage problem* and shortly discuss other solutions.

5.1.1 Embedding the Initial State Distribution

Before running the filter on incoming measurements y_t , we need to initialize the belief state with an initial mean map $\mu_{X,0}$ and an initial covariance operator $C_{XX,0}$. We can obtain these initial embeddings from a data set $\mathcal{D}_0 = \{x_1^0, \dots, x_N^0\}$ which consists in general of samples from the initial distribution of the system. Practically, we can obtain this data set by taking the initial states from multiple training episodes or—if we assume a stationary distribution—we can also take all training samples for the initialization. We can obtain the initial mean map by first embedding a uniform distribution into the RKHS spanned by the features of the initial states $\Upsilon_{x,0}$. Afterwards, we apply a conditional operator to map this distribution into the Hilbert space spanned by the features Υ_x as

$$\mu_{X,0} = \Upsilon_x \mathbf{m}_0 = \Upsilon_x (\mathbf{K} + \lambda \mathbf{I})^{-1} \Upsilon_x^\top \Upsilon_{x,0} \mathbf{1}_N \frac{1}{N} \quad (72)$$

$$\Leftrightarrow \mathbf{m}_0 = (\mathbf{K} + \lambda \mathbf{I})^{-1} \mathbf{K}_0 \mathbf{1}_N \frac{1}{N}. \quad (73)$$

where $\mathbf{K}_0 = \Upsilon_x^\top \Upsilon_{x,0}$ is the kernel matrix of the training samples and the samples in \mathcal{D}_0 , and $\mathbf{1}_N$ denotes the N -dimensional all-ones vector. Similarly, we can obtain the initial covariance embedding operator as

$$C_{XX,0} = \Upsilon_x \mathbf{S}_0 \Upsilon_x^\top = \frac{1}{N} \Upsilon_x (\mathbf{K} + \lambda \mathbf{I})^{-1} \mathbf{K}_0 \mathbf{K}_0^\top (\mathbf{K} + \lambda \mathbf{I})^{-1} \Upsilon_x^\top - \Upsilon_x \mathbf{m}_0 \mathbf{m}_0^\top \Upsilon_x^\top. \quad (74)$$

Hence, we can obtain the initial weight vector \mathbf{m}_0 and the initial weight matrix \mathbf{S}_0 by computing the mean and the covariance over the columns of the matrix $\mathbf{C}_0 = (\mathbf{K} + \lambda \mathbf{I})^{-1} \mathbf{K}_0$.

5.1.2 The Pre-image Problem / Recovering the State-Space Distribution

Recovering a distribution in the state space that is a pre-image of a given mean map is still a topic of ongoing research. There are several approaches to this problem, such as fitting a Gaussian mixture model (Mccalman et al 2013), or sampling from the embedded distribution by optimization (Chen et al 2012). In the experiments conducted for this paper, we approach the preimage problem by matching a Gaussian distribution, which is a reasonable choice if the recovered distribution is unimodal. Since we embed the belief state for the kernel Kalman rule as a mean map and as a covariance operator, obtaining the mean and covariance of a Gaussian approximation is done by simple matrix manipulations. The space of the samples \mathbb{R}^d together with the linear kernel $k(\mathbf{x}_1, \mathbf{x}_2) = \langle \mathbf{x}_1, \mathbf{x}_2 \rangle = \mathbf{x}_1^\top \mathbf{x}_2$ forms an RKHS as well. Therefore, we can simply define a conditional embedding operator that maps from the Hilbert space of the feature vectors to the Hilbert space of the samples as

$$C_{\text{pre}} = \mathbf{X} (\mathbf{K} + \lambda \mathbf{I})^{-1} \Upsilon_x^\top. \quad (75)$$

By applying this conditional operator now to the belief state, we obtain the mean of the embedded distribution in the sample space

$$\boldsymbol{\eta}_t = C_{\text{pre}} \boldsymbol{\mu}_{X,t} = C_{\text{pre}} \mathbb{E}_{b_t} [\boldsymbol{\varphi}(X)] = \mathbb{E}_{b_t} [C_{\text{pre}} \boldsymbol{\varphi}(X)] = \mathbb{E}_{b_t} [X]. \quad (76)$$

Similarly, we can also apply this operator to the covariance embedding to obtain the covariance of the belief state in the sample space

$$\begin{aligned}
\Sigma_t &= C_{\text{pre}} C_{XX,t} C_{\text{pre}}^\top \\
&= C_{\text{pre}} (\mathbb{E}_{b_t} [\boldsymbol{\varphi}(X) \otimes \boldsymbol{\varphi}(X)] - \boldsymbol{\mu}_{X,t} \otimes \boldsymbol{\mu}_{X,t}) C_{\text{pre}}^\top \\
&= \mathbb{E}_{b_t} [C_{\text{pre}} \boldsymbol{\varphi}(X) \otimes \boldsymbol{\varphi}(X) C_{\text{pre}}^\top] - C_{\text{pre}} \boldsymbol{\mu}_{X,t} \otimes \boldsymbol{\mu}_{X,t} C_{\text{pre}}^\top \\
&= \mathbb{E}_{b_t} [X \otimes X] - \boldsymbol{\eta}_t \otimes \boldsymbol{\eta}_t^\top
\end{aligned} \tag{77}$$

However, also any other approach from the literature can be used in the kernel Kalman filter algorithm.

Algorithm 1: The Kernel Kalman Filter

input: triples $\{(\tilde{x}_1, x_1, y_1), \dots, (\tilde{x}_m, x_m, y_m)\}$,
kernel functions k and g , regularization parameters λ and κ ,
let \mathbf{X} be the matrix of all x_i as columns, $\tilde{\mathbf{X}}$ and \mathbf{Y} analogously,
let \mathbf{X}_0 be the matrix of all data points used to compute the initial embeddings

compute kernel matrices
 $\mathbf{K} = k(\mathbf{X}, \mathbf{X})$, $\mathbf{K}_{\tilde{x}\tilde{x}} = k(\tilde{\mathbf{X}}, \tilde{\mathbf{X}})$, $\mathbf{K}_{\tilde{x}\mathbf{X}} = k(\tilde{\mathbf{X}}, \mathbf{X})$, and $\mathbf{G} = g(\mathbf{Y}, \mathbf{Y})$

compute model matrices
 $\mathbf{T} = (\mathbf{K}_{\tilde{x}\tilde{x}} + \lambda \mathbf{I})^{-1} \mathbf{K}_{\tilde{x}\mathbf{X}}$ and $\mathbf{O} = (\mathbf{K} + \lambda \mathbf{I})^{-1} \mathbf{K}$

compute initial embeddings
kernel matrix with samples of the initial distribution: $\mathbf{K}_0 = k(\mathbf{X}, \mathbf{X}_0)$, $\mathbf{C}_0 = (\mathbf{K} + \lambda \mathbf{I})^{-1} \mathbf{K}_0$,
compute mean and variance over the columns: $\mathbf{m}_0 = \text{mean}(\mathbf{C}_0)$, $\mathbf{S}_0 = \text{var}(\mathbf{C}_0)$

loop

- if** new observation \mathbf{y}_t available **then**
- compute kernel Kalman gain*
- $\mathbf{Q}_t = \mathbf{S}_t^- \mathbf{O}^\top (\mathbf{G} \mathbf{O} \mathbf{S}_t^- \mathbf{O}^\top + \kappa \mathbf{I})^{-1}$
- innovation update*
- $\mathbf{m}_t^+ = \mathbf{m}_t^- + \mathbf{Q}_t (\mathbf{g}_{:\mathbf{y}_t} - \mathbf{G} \mathbf{O} \mathbf{m}_t^-)$
- $\mathbf{S}_t^+ = \mathbf{S}_t^- - \mathbf{Q}_t \mathbf{G} \mathbf{O} \mathbf{S}_t^-$
- transition update*
- $\mathbf{m}_{t+1}^- = \mathbf{T} \mathbf{m}_t^+$, $\mathbf{S}_{t+1}^- = \mathbf{T} \mathbf{S}_t^+ \mathbf{T}^\top + \mathbf{V}$
- project into state space*
- $\boldsymbol{\eta}_t = \mathbf{X} \mathbf{O} \mathbf{m}_t$, $\boldsymbol{\Sigma}_t = \mathbf{X} \mathbf{O} \mathbf{S}_t \mathbf{O}^\top \mathbf{X}^\top$

5.1.3 Embedding Observation Windows

So far, we assumed that we have access to the latent states x_i in our training set. However, in many setups we only have access to the partial observations y_i which do not have the Markov property. Yet, we can still learn a KKR model from the provided data by embedding time windows $\mathbf{y}_{t-k+1:t}$ of size k as internal state representation. Similar approaches have been used by auto-regressive HMMs (Shannon et al 2013). With longer data windows, the transitions become more and more Markov. How many observation each data window has to contain depends on two factors: on the dimensionality of the underlying system and on the signal-to-noise ratio of the measurements y_i .

5.2 The Subspace Kernel Kalman Filter

The subspace kernel Kalman filter (subKKF) is a natural extension of the KKF that leverages the subspace formulation of the conditional embedding operator presented in Equation 26 and its application to the framework for nonparametric inference as well as the subspace formulation of the kernel Kalman rule derived in Section 4.3.

In contrast to the KKF, we assume for the subKKF a data set of triples $\{(x_1, x'_1, y_1), \dots, (x_m, x'_m, y_m)\}$, where x'_i is the successor state to x_i . The representation of the belief state changes from weight vector \mathbf{m}_t and weight matrix \mathbf{S}_t to the subspace projections of the embeddings $\mathbf{n}_t = \mathbf{\Gamma}_x^\top \mathbf{Y}_x \mathbf{m}_t = \bar{\mathbf{K}}^\top \mathbf{m}_t$ and $\mathbf{P}_t = \mathbf{\Gamma}_x^\top \mathbf{Y}_x \mathbf{S}_t \mathbf{Y}_x^\top \mathbf{\Gamma}_x = \bar{\mathbf{K}}^\top \mathbf{S}_t \bar{\mathbf{K}}$, respectively. Additionally, both update procedures of the kernel Kalman filter, the transition update and the innovation update, have to be substituted by their subspace counterparts. The transition update is realized by the subspace kernel sum rule and the innovation update by the subspace kernel Kalman rule. The equations are depicted in Algorithm 2.

Since we represent the belief state as a projection into the subspace defined by $\mathbf{\Gamma}_x$, we can directly obtain the initial belief state by projecting the uniform embedding in the RKHS spanned by the samples from the initial distribution as

$$\mathbf{n}_0 = \mathbf{\Gamma}_x \mathbf{Y}_{x,0} \mathbf{1}_N \frac{1}{N} = \mathbf{K}_0^S \mathbf{1}_N \frac{1}{N}, \quad (78)$$

$$\mathbf{P}_0 = \frac{1}{N} \mathbf{\Gamma}_x \mathbf{Y}_{x,0} \mathbf{Y}_{x,0}^\top \mathbf{\Gamma}_x^\top = \frac{1}{N} \mathbf{K}_0^S (\mathbf{K}_0^S)^\top. \quad (79)$$

Here, \mathbf{K}_0^S is the feature matrix of the reduced sample set and the samples from the initial state distribution. For the mapping back into the state space, we can similarly to the KKF define a subspace conditional operator as

$$\mathbf{C}_{\text{pre}}^S := \mathbf{X} \bar{\mathbf{K}} (\bar{\mathbf{K}}^\top \bar{\mathbf{K}} + \lambda \mathbf{I})^{-1} \mathbf{\Gamma}_x^\top, \quad (80)$$

By applying this operator to mean map and covariance embedding, we obtain the mean and variance in state space from the subspace projections as

$$\boldsymbol{\eta}_t = \mathbf{X} \bar{\mathbf{K}} (\bar{\mathbf{K}}^\top \bar{\mathbf{K}} + \lambda \mathbf{I})^{-1} \mathbf{n}_t, \quad (81)$$

$$\boldsymbol{\Sigma}_t = \mathbf{X} \bar{\mathbf{K}} (\bar{\mathbf{K}}^\top \bar{\mathbf{K}} + \lambda \mathbf{I})^{-1} \mathbf{P}_t (\bar{\mathbf{K}}^\top \bar{\mathbf{K}} + \lambda \mathbf{I})^{-1} \bar{\mathbf{K}}^\top \mathbf{X}^\top. \quad (82)$$

A concise description of the subspace kernel Kalman filter can be found in Algorithm 2.

5.2.1 Selecting the Sample Set to Span the Subspace

To learn the subspace kernel Kalman filter models, we assume that we have a large data set of size m but we represent the belief states in a subspace spanned by only n samples, where $n \ll m$. A crucial problem is though how to select the subset of n samples from the full data set of m samples. We propose two approaches to address this problem which aim at different characteristics of the subset.

The first approach simply samples uniformly without replacement from the full data set. The result is a subset that resembles statistically the full data set, i.e., regions that have a high density in the full data set will have a high density in the subset and vice versa.

The second approach is a bit more sophisticated in its selection strategy. The goal is here to get an optimal coverage of the sample space in the subset according to a certain

measure. To do so, we select the first sample randomly from the full data set into the subset. Afterwards, we iteratively extend the subset by adding samples according to the following criterion: we compute the maximum kernel activation for each sample in the full data set with the samples in the current subset, then we extend the current subset by taking the sample from the full data set with the minimal maximum activation. We call this second strategy the *kernel activation heuristic*.

Algorithm 2: The Subspace Kernel Kalman Filter

input: triples $\{(x_1, x'_1, y_1), \dots, (x_m, x'_m, y_m)\}$,
 kernel functions k and g , regularization parameters λ and κ ,
 let \mathbf{X} be the matrix of all x_i as columns, \mathbf{X}' and \mathbf{Y} analogously,
 let \mathbf{X}_0 be the matrix of all data points used to compute the initial embeddings

select subset
 $\mathbf{X}_S \sim$ random strategy, or $\mathbf{X}_S \sim$ min-max-activation strategy

compute kernel matrices
 $\bar{\mathbf{K}} = k(\mathbf{X}, \mathbf{X}_S)$, $\bar{\mathbf{K}}' = k(\mathbf{X}', \mathbf{X}_S)$, and $\mathbf{G} = g(\mathbf{Y}, \mathbf{Y})$

compute model matrices
 $\mathbf{T}^S = \bar{\mathbf{K}}'^T \bar{\mathbf{K}} (\bar{\mathbf{K}}^T \bar{\mathbf{K}} + \lambda \mathbf{I})^{-1}$ and $\mathbf{O}^S := (\bar{\mathbf{K}}^T \bar{\mathbf{K}} + \lambda \mathbf{I})^{-1}$

compute initial embeddings
 kernel matrix with samples of the initial distribution: $\mathbf{K}_0^S = k(\mathbf{X}_S, \mathbf{X}_0)$
 compute mean and variance over the columns: $\mathbf{m}_0 = \text{mean}(\mathbf{K}_0^S)$, $\mathbf{S}_0 = \text{var}(\mathbf{K}_0^S)$

loop

if new observation \mathbf{y}_t available **then**

compute subspace kernel Kalman gain
 $\mathbf{Q}_t^S = \mathbf{P}_t^- (\mathbf{O}^S)^T (\bar{\mathbf{K}}^T \mathbf{G} \bar{\mathbf{K}} \mathbf{O}^S \mathbf{P}_t^- (\mathbf{O}^S)^T + \kappa \mathbf{I})^{-1} \bar{\mathbf{K}}^T$
*note that you can apply the matrix $\bar{\mathbf{K}}^T$ to the kernel matrix \mathbf{G} in the innovation update
 already at learning time to increase computational efficiency.*

innovation update
 $\mathbf{n}_t^+ = \mathbf{n}_t^- + \mathbf{Q}_t^S (\mathbf{g}_{\cdot \mathbf{y}_t} - \mathbf{G} \bar{\mathbf{K}} \mathbf{O}^S \mathbf{n}_t^-)$
 $\mathbf{P}_t^+ = \mathbf{P}_t^- - \mathbf{Q}_t^S \mathbf{G} \bar{\mathbf{K}} \mathbf{O}^S \mathbf{P}_t^-$

transition update
 $\mathbf{n}_{t+1}^- = \mathbf{T}^S \mathbf{n}_t^+$, $\mathbf{P}_{t+1}^- = \mathbf{T}^S \mathbf{P}_t^+ (\mathbf{T}^S)^T + \mathbf{V}_S$

project into state space
 $\boldsymbol{\eta}_t = \mathbf{X} \bar{\mathbf{K}} \mathbf{O}^S \mathbf{n}_t$, $\boldsymbol{\Sigma}_t = \mathbf{X} \bar{\mathbf{K}} \mathbf{O}^S \mathbf{P}_t \mathbf{O}^S \bar{\mathbf{K}}^T \mathbf{X}^T$

5.3 Experimental Evaluation of the Kernel Kalman Filter

We evaluate the performance of the KKF and the subKKF on two experiments on simulated environments, a pendulum and a quad-link, and one experiment on real-world data from a human motion tracking data set (Wojtusich and von Stryk 2015). For all kernel based methods, we use the squared exponential kernel, where we choose the kernel bandwidths according to the *median trick* (Jaakkola et al 1999) and scale the median distances with a single optimized parameter.

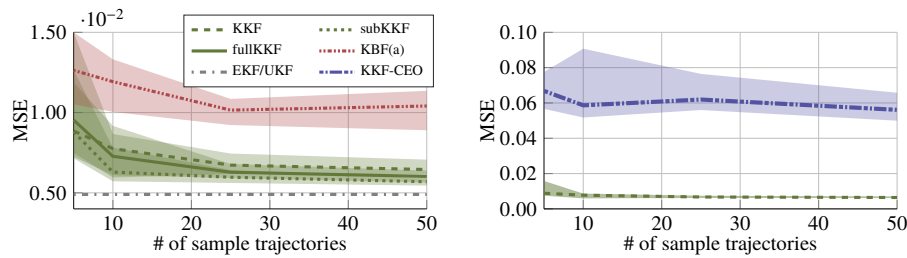


Fig. 3: Comparison of KKF to KBF(b), KKF-CEO, EKF and UKF. All kernel methods (except fullKKF) use kernel matrices of 100 samples. The subKKF method uses a subset of 100 samples and the whole dataset to learn the conditional operators. Depicted is the median MSE to the ground-truth of 20 trials with the [0.25 0.75] quantiles.

5.3.1 Pendulum

In this experiment, we used a simulated pendulum as system dynamics. It is uniformly initialized in the range $[0.1\pi, 0.4\pi]$ with a velocity sampled from the range $[-0.5\frac{\pi}{s}, 0.5\frac{\pi}{s}]$. The observations of the filters were the joint positions with additive Gaussian noise sampled from $\mathcal{N}(0, 0.01)$. We compared the KKF, the subspace KKF (subKKF) and the KKF learned with the full dataset (fullKKF) to version (a) of the kernel Bayes filter (KBF(a)) (Song et al 2013) (the other versions, KBF(b) and KBF(c), yielded worse results in this experiment) and the kernel Kalman filter with covariance embedding operator (KKF-CEO) (Zhu et al 2014), as well as to standard filtering approaches such as the EKF (Julier and Uhlmann 1997) and the UKF (Wan and Van Der Merwe 2000) (which require a model of the system dynamics). To learn the models, we simulated episodes with a length of 30 steps (3 sec). For the KKF and all KBF models, we use a kernel size of 100 samples, for the fullKKF, we use all available training samples and for the subKKF we use a set of 100 samples to span the subspace and the full data set to learn the operators. The samples are selected from the full data set using the kernel activation heuristic. The results are shown in Figure 3. The KKF and subKKF show clearly better results than all other non-parametric filtering methods and reach a performance level close to the EKF and UKF.

5.3.2 Quad-Link

In this experiment, we used a simulated 4-link pendulum where we observe the 2-D end-effector positions. The state of the pendulum consists of the four joint angles and joint velocities. We evaluate the prediction performance of the subKKF in comparison to the KKF-CEO, the EKF and the UKF. All other non-parametric filtering methods could not achieve a good performance or showed to be not feasible down due to the very high computation times. As the subKKF outperformed the KKF in the previous experiments and is also computationally much cheaper, we skip the comparison to the standard KKF in this and the subsequent experiments.

In a first qualitative evaluation, we compare the long-term prediction performance of the subKKF in comparison to the UKF, the EKF and the Monte-Carlo filter (MCF) as a baseline. This evaluation can be seen in Figure 4. The first five steps of of the end-effector trajectories were observed by the filters, the following 30 steps were predicted. The UKF is not able to

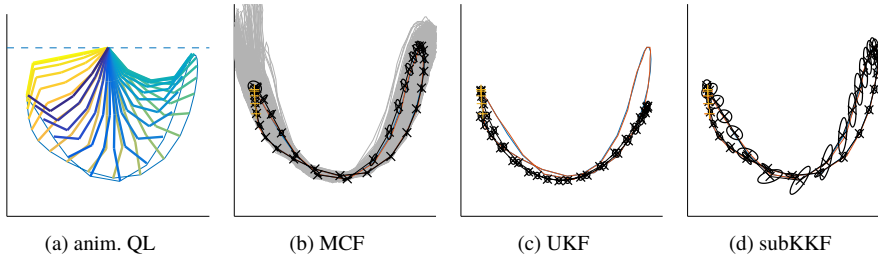


Fig. 4: Example trajectory of the quad-link end-effector. The filter outputs in black, where the ellipses enclose 90% of the probability mass. All filters were updated with the first five measurements (yellow marks) and predicted the following 30 steps. Figure (a) is an animation of the trajectory.




Fig. 5: 1, 2 and 3 step prediction performances in mean euclidean distances (MED) to the true end-effector positions of the quad-link.

predict the movements of the quad-link end-effector due to the high non-linearity, while the subKKF is able to predict the whole trajectory.

We also compared the 1, 2 and 3-step prediction performance of the subKKF to the KKF-CEO, EKF and UKF (Fig. 5). The KKF-CEO provides poor results already for the filtering task. The EKF performs equally bad, since the observation model is highly non-linear. The UKF already yields a much better performance as it does not suffer from the linearization of the system dynamics. The subKKF outperformed the UKF.

5.3.3 Human Motion Data

The human motion dynamics (HuMoD) database by Wojtusich and von Stryk (2015) consists of the datasets of several motions executed by two subjects. All datasets contain the recordings from a motion capture system with 36 markers as well as the recordings of the electrical activity of 14 muscles in the legs. Additionally, data from the treadmill such as ground reaction forces and velocities are available. The x -, y -, and z -locations of the markers were recorded at 500Hz, the muscle activities at 2000Hz, and the data from the treadmill at 1000Hz. Furthermore, the database contains joint positions and joint trajectories derived from the marker positions via a kinematic model of the human body. In our experiments, we use the marker locations, the derived locations of the joints and the muscle activities. We subsample all data to a common frame rate of 50Hz and transpose the x - and z -position of all markers such that the T12-marker (marker at the 12th thoracic vertebra) has $(x = 0, z = 0)$ in all frames. Note that in the HuMoD database, the x -axis points in the motion direction (i.e.,

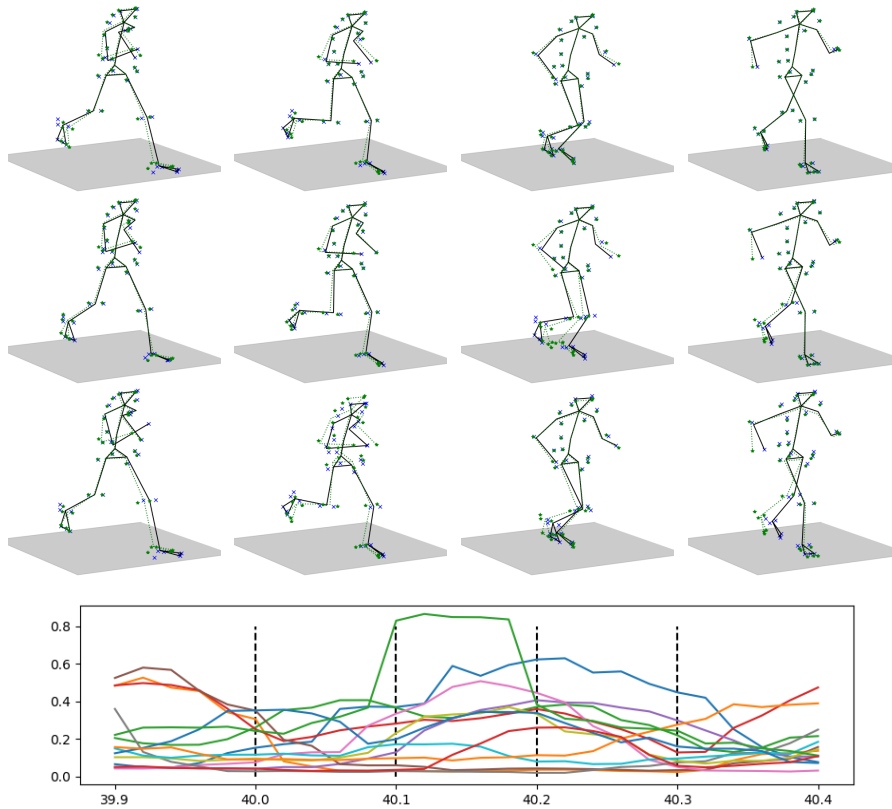


Fig. 6: Example sequence of 4 postures and the measured muscle activities. The marker and skeleton in green depict the ground-truth, the estimates from the models are depicted in black/blue. The learned models estimate the marker and joint positions from the muscle activities. The first row shows the estimated positions from the subKKF, the second row shows the estimated positions from the subKBF and the third row shows the estimated positions from a sparse GP. For all three models, we used a sample set of 2000 samples and a sparse subset of 500 samples.

along the treadmill), the y -axis points upwards and the z -axis forms a right-hand coordinate system towards the right side of the treadmill. We used walking motions at 1.0m/s, 1.5m/s, 2.0m/s, and running motions at 2.0m/s, 3.0m/s, and 4m/s, captured from one subject. For evaluating the trained model, we used a test data-set in which the subject transitions linearly from 0m/s up to 4m/s and back to 0m/s.

In the experiment, we compare the performance of subKKF, subKBF, and sparse Gaussian process (SGP) in restoring the marker and joint positions from the muscle activities. We learn all three models using the marker and joint positions as state variables (or outputs) x_i and the muscle activities as observations (or inputs) y_i . We use a set of 2000 samples to learn the kernel matrices and a subset of 500 samples to define the subspace (or as inducing inputs). For subKKF and subKBF, we used a window size of 2. While we could easily carry out the optimization of the parameters for the subKKF and for the SGP, the optimization of the parameters for the subKBF was not feasible in a considerable amount of time.




Fig. 7: Performance of the subKKF and the subKBF on the HuMoD transition data for different sizes of the subspace.




Fig. 8: Performance of the subKKF on HuMoD test sequences after 0, 10, 20, 30 and 40 iterations of the CMA-ES optimizer.

#	200	400	600	800	1000
subKKF	6.66	5.91	5.94	5.87	5.92
subKBF	7.48	6.73	6.63	6.55	6.70
SGP	76.33	70.15	68.46	63.13	58.97

subKKF	0.61±0.10 s	1.71±0.17 s	3.84±0.40 s	7.21±0.85 s	11.72±0.41 s
subKBF	59±1.3 s	170±11.3 s	352±15 s	620±15 s	927±174 s

Table 2: Top: performance of subKKF, subKBF, and SGP for the HuMoD transition data for different sizes of the subspace. Bottom: time consumptions of subKKF and subKBF for filtering 100 test sequences of length 50 from the HuMoD data set.

Figure 6 depicts marker and joint positions of four exemplary postures together with the muscle activities during that period of time. The locations of the exemplary postures in the time linen are depicted by vertical lines in the plot of the muscle activities. While this is only a qualitative example, it depicts how the subKKF outperforms the subKBF and the SGP in restoring the positions of the markers and the joints.

We compare the performance of subKKF, subKBF and SGP for different sizes of the subspace (inducing inputs). Figure 7 depicts the performance of subKKF and subKBF. The results of the SGP can be seen in Table 2 which are clearly worse in comparison to the filtering approaches which take the temporal correlation of the data into account. Furthermore, Table 2 also depicts the time consumption of subKKF and subKBF for the evaluation of the transition data in the HuMoD data base. The gain in efficiency of the subKKF over the subKBF which is around the factor 100 can be seen clearly. In Figure 8, we depict the performance gain of the subKKF over the number of iterations of the CMA-ES optimizer. We see that in this case, the first 10 iterations yield a bigger jump in performance than the following 40 steps. However, from our experience, this is very specific to the problem and the initial setting of the parameters, which were in this experiment already very close to the optimal parameters.

5.4 The Kernel Forward-Backward Smoother

Smoothing is in contrast to filtering a post-processing routine. While filtering refers to a routine where the current state is estimated recursively from all past observations, smoothing computes the best state estimates given all available observations from the past and the

future. Hence, for a given time series of observations $[y_1, \dots, y_T]$, we want to obtain the belief $p(x_t | y_1, \dots, y_T)$ for all $1 \leq t \leq T$.

One well-known and simple approach to smoothing is the forward-backward smoother. During a forward pass the standard filtering algorithm is applied to the observations. Afterwards, on the backward pass, an inverse filter is applied to the same time series of observations. Finally the filter estimates of forward and backward pass are combined into the smoothed estimates. Since the information from the observation should be incorporated only once into the smoothed estimate, we need to combine the posterior estimate of the forward pass with the prior estimate of the backward pass (or vice versa). The problems of this routine are usually that it requires an inverse model of the underlying system for the backward pass and an initialization of the final state. If we assume that the system has a finite state space, these issues are however not applicable for the KKF as we can learn the models in the feature space from data.

5.4.1 Computing the Smoothed Belief State as a Weighted Average

Assuming that we have the a-posteriori belief states from the forward pass and the a-priori belief states from the backward pass as

$$\left\{ \left(\mu_{f,1}^+, C_{f,1}^+ \right), \dots, \left(\mu_{f,T}^+, C_{f,T}^+ \right) \right\} \quad \text{and} \quad \left\{ \left(\mu_{b,1}^-, C_{b,1}^- \right), \dots, \left(\mu_{b,T}^-, C_{b,T}^- \right) \right\}, \quad (83)$$

respectively, we can combine the mean maps into a smoothed belief state as the weighted average

$$\mu_{s,t} = \mathcal{Z}_{f,t} \mu_{f,t}^+ + \mathcal{Z}_{b,t} \mu_{b,t}^-. \quad (84)$$

Since both, the estimator from the forward pass and the estimator from the backward pass, are unbiased, the weighting operators $\mathcal{Z}_{f,t}$ and $\mathcal{Z}_{b,t}$ need to satisfy $\mathcal{I} = \mathcal{Z}_{f,t} + \mathcal{Z}_{b,t}$, i.e.,

$$\mathbb{E} \left[\varphi(X_t) - \mu_t^s \right] \stackrel{!}{=} 0 \quad (85)$$

$$\mathbb{E} \left[\varphi(X_t) - \mathcal{Z}_{f,t} \mu_t^{f+} + \mathcal{Z}_{b,t} \mu_t^{b-} \right] \stackrel{!}{=} 0 \quad (86)$$

$$\mathbb{E} \left[\varphi(X_t) \right] - \mathcal{Z}_{f,t} \mathbb{E} \left[\mu_t^{f+} \right] + \mathcal{Z}_{b,t} \mathbb{E} \left[\mu_t^{b-} \right] \stackrel{!}{=} 0 \quad (87)$$

$$\mu_t - (\mathcal{Z}_{f,t} + \mathcal{Z}_{b,t}) \mu_t \stackrel{!}{=} 0 \quad (88)$$

$$\Rightarrow \mathcal{Z}_{f,t} + \mathcal{Z}_{b,t} = \mathcal{I} \quad (89)$$

Thus, the weighting operators can be expressed by each other as $\mathcal{Z}_{b,t} = \mathcal{I} - \mathcal{Z}_{f,t}$ and vice versa. We substitute this representation back into the smoothing update in Equation 84 to obtain

$$\mu_{s,t} = \mathcal{Z}_{f,t} \mu_{f,t}^+ + (\mathcal{I} - \mathcal{Z}_{f,t}) \mu_{b,t}^-. \quad (90)$$

5.4.2 Finding the Optimal Weighting Operators

We obtain the optimal weighting operators by minimizing the squared error of the smoothing mean map which is equivalent to minimizing the trace of the smoothed covariance embedding operator $C_{s,t}$, i.e.,

$$\mathbb{E} \left[(\varphi(X_t) - \mu_t^s)^\top (\varphi(X_t) - \mu_t^s) \right] = \mathbb{E} \left[\text{Tr} \left((\varphi(X_t) - \mu_t^s) (\varphi(X_t) - \mu_t^s)^\top \right) \right] = \text{Tr} C_{s,t}. \quad (91)$$

We can then use Equation 90 to rewrite the covariance operator as

$$C_{s,t} = \mathbb{E} \left[\left(\boldsymbol{\varphi}(X_t) - \mathcal{Z}_{f,t} \mu_{f,t}^+ + (I - \mathcal{Z}_{f,t}) \mu_{b,t}^- \right) \left(\dots \right)^\top \right] \quad (92)$$

$$= \mathbb{E} \left[\left(\mathcal{Z}_{f,t} (\epsilon_f - \epsilon_b) + \epsilon_b \right) \left(\dots \right)^\top \right], \quad (93)$$

where $\epsilon_f = \boldsymbol{\varphi}(x_t) - \mu_{f,t}^+$ is the error of the a-posteriori estimate of the forward pass and $\epsilon_b = \boldsymbol{\varphi}(x_t) - \mu_{b,t}^-$ is the error of the a-priori estimate of the backward pass and where we used the relation $\boldsymbol{\varphi}(X_t) = I \boldsymbol{\varphi}(X_t) = (\mathcal{Z}_{f,t} + \mathcal{Z}_{b,t}) \boldsymbol{\varphi}(X_t) = (\mathcal{Z}_{f,t} + (I - \mathcal{Z}_{f,t})) \boldsymbol{\varphi}(X_t)$. By expanding the square and since the cross-covariances of the errors from forward and the backward pass are zero (i.e., $\mathbb{E}[\epsilon_f \epsilon_b^\top] = 0$), we arrive at

$$C_{s,t} = \mathbb{E} \left[\mathcal{Z}_{f,t} \left(\epsilon_f \epsilon_f^\top + \epsilon_b \epsilon_b^\top \right) \mathcal{Z}_{f,t}^\top - \mathcal{Z}_{f,t} \epsilon_b \epsilon_b^\top - \epsilon_b \epsilon_b^\top \mathcal{Z}_{f,t}^\top + \epsilon_b \epsilon_b^\top \right]. \quad (94)$$

Lastly, we can take the derivative and set it to zero to obtain the optimal $\mathcal{Z}_{f,t}$ as

$$0 \stackrel{!}{=} \frac{\partial \text{Tr } C_{s,t}}{\partial \mathcal{Z}_{f,t}} = 2 \mathbb{E} \left[\mathcal{Z}_{f,t} \left(\epsilon_f \epsilon_f^\top + \epsilon_b \epsilon_b^\top \right) - \epsilon_b \epsilon_b^\top \right] \quad (95)$$

$$0 \stackrel{!}{=} \mathcal{Z}_{f,t} \left(\mathbb{E} \left[\epsilon_f \epsilon_f^\top \right] + \mathbb{E} \left[\epsilon_b \epsilon_b^\top \right] \right) - \mathbb{E} \left[\epsilon_b \epsilon_b^\top \right] \quad (96)$$

$$0 \stackrel{!}{=} \mathcal{Z}_{f,t} \left(C_{f,t}^+ + C_{b,t}^- \right) - C_{b,t}^- \quad (97)$$

$$\mathcal{Z}_{f,t} = C_{b,t}^- \left(C_{f,t}^+ + C_{b,t}^- \right)^{-1}. \quad (98)$$

From the condition on the weighting operators stated in Equation 89, it furthermore follows that $\mathcal{Z}_{b,t} = C_{f,t}^- \left(C_{f,t}^+ + C_{b,t}^- \right)^{-1}$.

5.4.3 Smoothing the Covariance Embedding Operator

Taking the representation of the smoothed covariance in Equation 94 and substituting the covariance operators and the optimal weighting operator $\mathcal{Z}_{f,t}$ gives us the following smoothed covariance operator

$$C_{s,t} = C_{b,t}^- - C_{b,t}^- \left(C_{f,t}^+ + C_{b,t}^- \right)^{-1} C_{b,t}^- \quad (99)$$

$$= C_{b,t}^- - \left(I - C_{f,t}^+ \left(C_{f,t}^+ + C_{b,t}^- \right)^{-1} \right) C_{b,t}^- \quad (100)$$

$$= C_{f,t}^+ \left(C_{f,t}^+ + C_{b,t}^- \right)^{-1} C_{b,t}^- \quad (101)$$

From the optimal solution of the weighting operator, we can now see that the smoothing update of the covariance embedding operator can be expressed as

$$C_{s,t} = \mathcal{Z}_{b,t} C_{b,t}^- = \mathcal{Z}_{f,t} C_{f,t}^+ \quad (102)$$

In the following section, we will show how the smoothing update can be expressed with finite samples using vector/matrix operations.

5.4.4 The Empirical Kernel Forward-Backward Smoother

We assume that we are given the weight vectors and weight matrices from the forward and the backward pass as $\{(\mathbf{m}_{f,1}^+, \mathbf{S}_{f,1}^+), \dots, (\mathbf{m}_{f,T}^+, \mathbf{S}_{f,T}^+)\}$ and $\{(\mathbf{m}_{b,1}^-, \mathbf{S}_{b,1}^-), \dots, (\mathbf{m}_{b,T}^-, \mathbf{S}_{b,T}^-)\}$, respectively. Since the weighting operator $\mathcal{Z}_{b,t}$ can be expressed by $\mathcal{Z}_{f,t}$ and vice versa, we only need to compute one of the weighting operators which we choose to be $\mathcal{Z}_{f,t}$. We add the identity operator with a small scalar γ in the inverse to improve the numerical stability and obtain

$$\mathcal{Z}_{f,t} = C_{b,t}^- \left(C_{f,t}^+ + C_{b,t}^- + \gamma \mathbf{I} \right)^{-1} \quad (103)$$

$$= \Upsilon_x \mathbf{S}_{b,t}^- \Upsilon_x^\top \left(\Upsilon_x \left(\mathbf{S}_{f,t}^+ + \mathbf{S}_{b,t}^- \right) \Upsilon_x^\top + \gamma \mathbf{I} \right)^{-1} \quad (104)$$

$$= \Upsilon_x \mathbf{S}_{b,t}^- \left(\Upsilon_x^\top \Upsilon_x \left(\mathbf{S}_{f,t}^+ + \mathbf{S}_{b,t}^- \right) + \gamma \mathbf{I} \right)^{-1} \Upsilon_x^\top \quad (105)$$

$$= \Upsilon_x \mathbf{S}_{b,t}^- \underbrace{\left(\mathbf{K} \left(\mathbf{S}_{f,t}^+ + \mathbf{S}_{b,t}^- \right) + \gamma \mathbf{I} \right)^{-1}}_{\mathbf{Z}_{f,t}} \Upsilon_x^\top, \quad (106)$$

where we used again the matrix identity $\mathbf{A}(\mathbf{B}\mathbf{A} + \mathbf{I})^{-1} = (\mathbf{A}\mathbf{B} + \mathbf{I})^{-1}\mathbf{A}$ and defined the finite weighting matrix $\mathbf{Z}_{f,t}$. With this weighting matrix we can now combine the mean maps of the forward and the backward pass as

$$\mu_{s,t} = \mathcal{Z}_{f,t} \mu_{f,t}^+ + (\mathbf{I} - \mathcal{Z}_{f,t}) \mu_{b,t}^- \quad (107)$$

$$= \Upsilon_x \mathbf{Z}_{f,t} \Upsilon_x^\top \mu_{f,t}^+ + (\mathbf{I} - \Upsilon_x \mathbf{Z}_{f,t} \Upsilon_x^\top) \mu_{b,t}^- \quad (108)$$

$$= \Upsilon_x \mathbf{Z}_{f,t} \mathbf{K} \mathbf{m}_{f,t}^+ + \Upsilon_x \mathbf{m}_{b,t}^- - \Upsilon_x \mathbf{Z}_{f,t} \mathbf{K} \mathbf{m}_{b,t}^- \quad (109)$$

$$\Upsilon_x \mathbf{m}_t^s = \Upsilon_x \left(\mathbf{m}_{b,t}^- + \mathbf{Z}_{f,t} \mathbf{K} \left(\mathbf{m}_{f,t}^+ - \mathbf{m}_{b,t}^- \right) \right). \quad (110)$$

And similarly we can obtain the smoothed estimate of the covariance operator as

$$C_t^s = \mathcal{Z}_{f,t} C_{f,t}^+ \quad (111)$$

$$= \Upsilon_x \mathbf{Z}_{f,t} \Upsilon_x^\top \Upsilon_x \mathbf{S}_{f,t}^+ \Upsilon_x^\top \quad (112)$$

$$\Upsilon_x \mathbf{S}_t^s \Upsilon_x^\top = \Upsilon_x \mathbf{Z}_{f,t} \mathbf{K} \mathbf{S}_{f,t}^+ \Upsilon_x^\top \quad (113)$$

A concise description of the kernel forward-backward smoothing algorithm can be found in Algorithm 3.

5.4.5 Initialization of the Backward Kernel Kalman Filter

A critical aspect of the classical forward-backward smoothing algorithm is the initialization of the belief state for the backward pass. Often the distribution over the initial state is well known but not a distribution over the terminal state, where it is often not even clear how a terminal state is defined. For the backward kernel Kalman filter, two approaches can be used to initialize the belief state. The first approach assumes that we have multiple episodes in the training data, where each episode terminates in a terminal state of the system. We can then compute the initialization for the backward pass analogously to the initialization for the KKF described in Section 5.1.1. The second approach simply assumes that the system has a stationary distribution which is covered by the training data. The initialization is then the embedding of the distribution over all the samples in the training set.

Algorithm 3: The Kernel Forward-Backward Smoother

input: triples $\{(\tilde{x}_1, x_1, x'_1, y_1), \dots, (\tilde{x}_m, x_m, x'_m, y_m)\}$,
kernel functions k and g , regularization parameters λ and κ ,
let \mathbf{X} be the matrix of all x_i as columns, $\tilde{\mathbf{X}}$, \mathbf{X}' and \mathbf{Y} analogously,
let \mathbf{X}_0 be the matrix of all data points used to initialize the forward pass
let \mathbf{X}_T be the matrix of all data points used to initialize the backward pass

learn the forward filter
See Algorithm 1 with data matrices \mathbf{X}_0 , $\tilde{\mathbf{X}}$, \mathbf{X} , and \mathbf{Y} .

learn the backward filter
See Algorithm 1 with data matrices \mathbf{X}_T , \mathbf{X} , \mathbf{X}' , and \mathbf{Y} , note that you need to learn the transition model from \mathbf{X}' to \mathbf{X} .

apply forward filter
Compute filtered estimates $\left\{ \left(\mathbf{m}_{f,1}^+, \mathbf{S}_{f,1}^+ \right), \dots, \left(\mathbf{m}_{f,T}^+, \mathbf{S}_{f,T}^+ \right) \right\}$

apply backward filter and compute smoothed estimates

loop

- if** observation \mathbf{y}_t available **then**
 - compute kernel Kalman gain
$$\mathbf{Q}_{b,t} = \mathbf{S}_{b,t}^- \mathbf{O}^\top \left(\mathbf{G} \left(\mathbf{O} \mathbf{S}_{b,t}^- \mathbf{O}^\top + \mathbf{R}_b \right) + \kappa \mathbf{I} \right)^{-1}$$
 - innovation update
$$\begin{aligned} \mathbf{m}_{b,t}^+ &= \mathbf{m}_{b,t}^- + \mathbf{Q}_{b,t} (\mathbf{g}_{:y_t} - \mathbf{G} \mathbf{O} \mathbf{m}_{b,t}^-) \\ \mathbf{S}_{b,t}^+ &= \mathbf{S}_{b,t}^- - \mathbf{Q}_{b,t} \mathbf{G} \mathbf{O} \mathbf{S}_{b,t}^- \end{aligned}$$
- transition update
$$\mathbf{m}_{b,t-1}^- = \mathbf{T}_b \mathbf{m}_{b,t}^+, \quad \mathbf{S}_{b,t-1}^- = \mathbf{T}_b \mathbf{S}_{b,t}^+ \mathbf{T}_b^\top + \mathbf{V}_b$$
- compute smoothed estimate
$$\mathbf{Z}_{f,t} = \mathbf{S}_{b,t}^- \left(\mathbf{K} \left(\mathbf{S}_{f,t}^+ + \mathbf{S}_{b,t}^- \right) + \gamma \mathbf{I} \right)^{-1}$$

$$\begin{aligned} \mathbf{m}_t^s &= \mathbf{m}_{b,t}^- + \mathbf{Z}_{f,t} \mathbf{K} \left(\mathbf{m}_{f,t}^+ - \mathbf{m}_{b,t}^- \right) \\ \mathbf{S}_t^s &= \mathbf{Z}_{f,t} \mathbf{K} \mathbf{S}_{f,t}^+ \end{aligned}$$
- project into state space
$$\boldsymbol{\eta}_t^s = \mathbf{X} \mathbf{O} \mathbf{m}_t^s, \quad \boldsymbol{\Sigma}_t^s = \mathbf{X} \mathbf{O} \mathbf{S}_t^s \mathbf{O}^\top \mathbf{X}^\top$$

5.5 The Subspace Kernel Forward-Backward Smoother

If we use the subspace kernel Kalman filter to perform the forward and the backward pass, we obtain as outcome the subspace projections \mathbf{n}_t of the mean map and \mathbf{P}_t covariance embedding instead of the weight vectors \mathbf{m}_t and weight matrices \mathbf{S}_t , respectively. To perform smoothing on these subspace projections, we need to find the weighting matrices for the smoothing update analog to Equation 90. Though, as the representation is already in a finite domain, we can directly apply the optimal solution found in Equation 98 to the subspace projections of the covariance operator. Hence, the weighting matrix for the subspace kernel forward-backward smoother (subKFBS) becomes

$$\mathbf{Z}_{f,t}^s = \mathbf{P}_{b,t}^- \left(\mathbf{P}_{f,t}^+ + \mathbf{P}_{b,t}^- \right)^{-1} \quad (114)$$

From here, we can obtain the equations for the smoothing update of the subspace kernel forward-backward smoother easily by

$$\mathbf{n}_{s,t} = \mathbf{Z}_{f,t}^S \mathbf{n}_{f,t}^+ + (\mathbf{I} - \mathbf{Z}_{f,t}^S) \mathbf{n}_{b,t}^-, \quad \text{and} \quad (115)$$

$$\mathbf{P}_{s,t} = \mathbf{Z}_{f,t}^S \mathbf{P}_{f,t}^+. \quad (116)$$

Algorithm 4 gives a compact description of the subKFBS.

Algorithm 4: The Subspace Kernel Forward-Backward Smoother

input: triples $\{(\tilde{x}_1, x_1, x'_1, y_1), \dots, (\tilde{x}_m, x_m, x'_m, y_m)\}$,
kernel functions k and g , regularization parameters λ and κ ,
let \mathbf{X} be the matrix of all x_i as columns, $\tilde{\mathbf{X}}$, \mathbf{X}' and \mathbf{Y} analogously,
let \mathbf{X}_0 be the matrix of all data points used to initialize the forward pass
let \mathbf{X}_T be the matrix of all data points used to initialize the backward pass

learn forward and backward filter

See Algorithm 1 with data matrices \mathbf{X}_0 , \mathbf{X} , \mathbf{X}' , and \mathbf{Y} for the forward filter and data matrices $\tilde{\mathbf{X}}$, \mathbf{X} , and \mathbf{Y} for the backward filter. Note that you need to learn the transition model for the backward filter from \mathbf{X} to $\tilde{\mathbf{X}}$.

apply forward filter

Compute filtered estimates $\left\{ \left(\mathbf{n}_{f,1}^+, \mathbf{P}_{f,1}^+ \right), \dots, \left(\mathbf{n}_{f,T}^+, \mathbf{P}_{f,T}^+ \right) \right\}$

apply backward filter and compute smoothed estimates

loop

if observation \mathbf{y}_t available **then**

compute subspace kernel Kalman gain

$$\mathbf{Q}_t^S = \mathbf{P}_{b,t}^- (\mathbf{O}^S)^\top (\tilde{\mathbf{K}}^\top \mathbf{G} \tilde{\mathbf{K}} \mathbf{O}^S \mathbf{P}_{b,t}^- (\mathbf{O}^S)^\top + \kappa \mathbf{I})^{-1} \tilde{\mathbf{K}}^\top$$

innovation update

$$\mathbf{n}_{b,t}^+ = \mathbf{n}_{b,t}^- + \mathbf{Q}_t^S (\mathbf{g}_{\cdot \mathbf{y}_t} - \mathbf{G} \tilde{\mathbf{K}} \mathbf{O}^S \mathbf{n}_{b,t}^-), \quad \mathbf{P}_{b,t}^+ = \mathbf{P}_{b,t}^- - \mathbf{Q}_t^S \mathbf{G} \tilde{\mathbf{K}} \mathbf{O}^S \mathbf{P}_{b,t}^-$$

transition update

$$\mathbf{n}_{b,t-1}^- = \mathbf{T}^S \mathbf{n}_{b,t}^+, \quad \mathbf{P}_{b,t-1}^- = \mathbf{T}^S \mathbf{P}_{b,t}^+ (\mathbf{T}^S)^\top + \mathbf{V}^S$$

compute smoothed estimate

$$\mathbf{Z}_{f,t}^S = \mathbf{P}_{b,t}^- (\mathbf{P}_{f,t}^+ + \mathbf{P}_{b,t}^-)^{-1}$$

$$\mathbf{n}_{s,t} = \mathbf{Z}_{f,t}^S \mathbf{n}_{f,t}^+ + (\mathbf{I} - \mathbf{Z}_{f,t}^S) \mathbf{n}_{b,t}^-, \quad \mathbf{P}_{s,t} = \mathbf{Z}_{f,t}^S \mathbf{P}_{f,t}^+$$

project into state space

$$\boldsymbol{\eta}_{s,t} = \mathbf{X} \tilde{\mathbf{K}} \mathbf{O}^S \mathbf{n}_{s,t}, \quad \boldsymbol{\Sigma}_{s,t} = \mathbf{X} \tilde{\mathbf{K}} \mathbf{O}^S \mathbf{P}_{s,t} \mathbf{O}^S \tilde{\mathbf{K}}^\top \mathbf{X}^\top$$

5.6 Experimental Evaluation of the Kernel Forward Backward Smoother

We evaluate the kernel forward backward smoother with two experiments. In the first experiment on data from a simulated pendulum, we show the performance gain of the KFBS over the KKF. In the second experiment, we apply the KFBS on data of a table tennis ball recorded with a camera-based tracking system and show how the KFBS and subKFBS are able to restore the full trajectory of the ball while only having observations at the first four and at the last time step.




Fig. 9: Qualitative comparison of the forward and the backward pass to the smoothed estimates of the KFBS on a simulated pendulum. The upper plots show the mean and variance output of the filter/smoothen, the lower plots show the profiles of the standard deviation. While the forward pass already yields good estimates in the first half of the time series, the smoother incorporates the good estimates from the backward pass in the second half and outperforms the filters. In addition, the smoother yields a more confident about its estimate.

5.6.1 Pendulum

We simulate a pendulum similar to the one from Section 5.3.1, however we initialize the pendulum in the range $[-0.25\pi, 0.25\pi]$ and with a angular velocity sampled from the range $[-2\frac{\pi}{s}, 2\frac{\pi}{s}]$. During the simulation, we apply Gaussian process noise with $\sigma = 0.01$ and as observations we use the angular displacement with Gaussian observation noise with $\sigma = 0.2$. To learn the KFBS models, we sample 100 episodes with each 30 steps where a step corresponds to 0.1 second. The training samples are 200 windows of four observations, which we select by the kernel activation heuristic explained in Section 5.2.1. To find the optimal parameters, we apply CMA-ES (Hansen 2006) where we use the negative log-likelihood of the ground-truth to the smoothed estimate as optimality criterion. During the optimization, we use a test-dataset of 10 episodes, where we still observe at each time step. Later, we evaluate the smoothing performance on an evaluation-dataset where we do not observe at each time step but only at $t = [1 - 4, 6, 11, 16, 21, 27 - 30]$. This optimization procedure yielded better results than directly optimizing with only partial observations.

In Figure 9, we show a qualitative comparison of the forward and the backward pass to the smoother. The results are as expected: the forward pass yields better results in the first half of the episode, and the backward pass yields better results in the second half. The smoother combines both estimates and outperforms the filter results. The smoothing can also be observed in the profiles of the standard deviation. While the variance from the filters increases at each time step without observation until the next measurement, the variance of the smoother is much smaller and only rises slightly between the observations.

In Figure 10, we compare the performance of a standard KKF to the KFBS and the subKFBS for different kernel sizes on the same state estimation task for a simulated pendulum. The subKFBS has been learned with 300 samples in the full data set. Depicted are the median and the $[0.15, 0.85]$ -quantiles of the MSE over 20 repetitions. The KFBS and the subKFBS clearly outperform the KKF for small kernel sizes (50, 100) and also yield better results for larger kernel sizes (150, 200). The subKFBS yields slightly better results than the KFBS. In




Fig. 10: The KFBS and the subKFBS outperform the KKF clearly for small kernel sizes but also with more samples in the gram matrices. The task was to estimate the state of a pendulum from noisy partial observations. Depicted are the median and the $[0.15, 0.85]$ -quantiles of the MSE over 20 repetitions.

addition, we see from the quantiles of the MSE that the KFBS and the subKFBS have a more stable behaviour in the optimization process than the KKF.

5.6.2 Tabletennis

In a second experiment, we perform smoothing on observations of a table tennis ball (Gomez-Gonzalez et al 2016). The data set contains 54 trajectories of a table tennis ball tracked with a camera system, where each trajectory contains 51 observations which are recorded with a frequency of 100 Hz. We train the subKFBS with 34 trajectories, use 10 trajectories for optimizing the parameters using CMA-ES (Hansen 2006). The remaining 10 trajectories are used for evaluating the results. For the smoothing task, the ball has been observed at the first five time steps and then again at the last time step.

Figure 11 shows qualitative examples of smoothed trajectories using the subKFBS in comparison the output of the subKKF. Here, we used data windows of size 4 and learned the models with 300 samples in the training data set and 100 samples in the subset. We optimized the regularization parameters and all bandwidths with CMA-ES (Hansen 2006). The plot shows how the subKFBS can estimate accurately the path of the ball only from observations at the beginning and at the end of the trajectory, while the subKKF diverges from the actual trajectory over time. In addition, the plot also depicts how the subKFBS, which is only learned from noisy observations, returns a trajectory that is more smooth than the noisy data recorded by a camera-based tracking system.

We also compare the KFBS to the subKFBS for different kernel sizes with the same smoothing task on recorded table tennis ball data. Figure 12 shows a comparison of the MSE, depicting the median and the $[0.05, 0.95]$ -quantiles over 20 repetitions. The KFBS has been learned with a varying kernel size of 50, 100, 150, and 200 samples. The subKFBS uses the same number of samples to span the subspace but learns the models always with 400 samples in the full training set. While the subKFBS outperforms the KFBS for all kernel sizes, the KFBS achieves a similar performance to the KFBS when learned with 200 samples.

6 Conclusion & Future Work

In this paper, we have presented the kernel Kalman rule (KKR) as an alternative to the kernel Bayes' rule (KBR) in the framework for nonparametric inference Song et al (2013). In contrast to the KBR, the KKR is computationally more efficient, numerically more stable and follows from a clear optimization objective. We have further combined the KKR as Bayesian update with the kernel sum rule to formulate the kernel Kalman filter (KKF). The kernel




Fig. 11: In comparison to the KKF, the KFBS is able to reconstruct the trajectories of a table tennis ball from observations at the first five and at the last two time steps. The plot shows the z-coordinate of two trajectories of a table tennis ball recorded with a camera-based tracking system. We learned the subKFBS/subKKF with 100 samples in the subspace and 300 points in the training set.




Fig. 12: The subKFBS performs better than the KFBS in the table tennis ball smoothing task. This difference in the MSE between the estimates and the noisy recorded data is more prevalent for small kernel sizes and decreases with the number of samples in the Gram matrices. Depicted is the median and the $[0.05, 0.95]$ -quantiles of the MSE over 20 repetitions.

Kalman filter can be applied to nonlinear state estimation tasks as it learns the probabilistic transition and observation dynamics as linear functions on embeddings of the belief state in high-dimensional Hilbert spaces from data. In difference to existing kernel Kalman filter formulations, the KKF also provides a more general formulation that is much closer to the original Kalman filter equations and can also be applied to partially observable systems.

While the KKF can be applied to state estimation and prediction based on past observations, we extend this work by introducing the kernel forward backward smoother (KFBS) which infers the belief state from current, past, and future information. We have shown in an experimental evaluation how this additional information leads to a performance gain of the KFBS over the KKF. As kernel methods typically scale poorly with the number of data points in the kernel matrices, we have introduced a sparsification technique that leverages from the full training set while representing the embeddings only with a small subset of the data. This technique leads to significant gains of the computational efficiency while yielding similar or even slightly better results than without the sparsification.

We have shown that it is possible to learn the kernel Kalman rule and other kernelized inference methods also from partial observations if sliding windows of the time series provide sufficient statistics. However in future work, we want to concentrate on learning the transition dynamics in the RKHS with an expectation-maximization algorithm in case of missing information about the latent state as we think that this leads to better models of the dynamics and also improves the accuracy of the estimated variance.

References

- Aronszajn N (1950) Theory of reproducing kernels. *Transactions of the American Mathematical Society* 68(3):337–404
- Boots B, Gretton A, Gordon GJ (2013) Hilbert space embeddings of predictive state representations. In: *Proceedings of the 29th International Conference on Uncertainty in Artificial Intelligence (UAI)*

- Chen Y, Welling M, Smola A (2012) Super-samples from kernel herding
- Csató L, Opper M (2002) Sparse on-line gaussian processes. *Neural computation* 14(3):641–668
- Drineas P, Mahoney MW (2005) On the nyström method for approximating a gram matrix for improved kernel-based learning. *Journal of Machine Learning Research* 6:23
- Fukumizu K, Song L, Gretton A (2011) Kernel bayes' rule. In: *Advances in Neural Information Processing Systems 24*, Curran Associates, Inc., pp 1737–1745
- Fukumizu K, Song L, Gretton A (2013) Kernel bayes' rule: Bayesian inference with positive definite kernels. *Journal of Machine Learning Research* 14(1):3683–3719
- Gauss CF (1823) *Theoria combinationis observationum erroribus minimis obnoxiae*. H. Dieterich
- Gebhardt GHW, Kupcsik A, Neumann G (2015) Learning subspace conditional embedding operators. *Workshop on Large-Scale Kernel Learning at ICML 2015*
- Gebhardt GHW, Kupcsik A, Neumann G (2017) The kernel kalman rule - efficient nonparametric inference with recursive least squares. In: *AAAI Conference on Artificial Intelligence*
- Gomez-Gonzalez S, Neumann G, Schölkopf B, Peters J (2016) Using probabilistic movement primitives for striking movements. *IEEE-RAS International Conference on Humanoid Robots* pp 502–508
- Grünewälder S, Lever G, Baldassarre L, Patterson S, Gretton A, Pontil M (2012) Conditional mean embeddings as regressors. In: *Proceedings of the 29th International Conference on Machine Learning*
- Hansen N (2006) The cma evolution strategy: A comparing review. In: *Towards a New Evolutionary Computation. Advances on estimation of distribution algorithms*, Springer, Berlin/Heidelberg, pp 75–102
- Hsu D, Kakade SM, Zhang T (2012) A spectral algorithm for learning hidden markov models. *Journal of Computer and System Sciences* 78(5):1460–1480
- Jaakkola T, Diekhans M, Haussler D (1999) Using the fisher kernel method to detect remote protein homologies. *Proceedings of the International Conference on Intelligent Systems for Molecular Biology; ISMB*
- Jaeger H (2000) Observable operator models for discrete stochastic time series. *Neural Computation* 12(6):1371–1398
- Julier SJ, Uhlmann JK (1997) A new extension of the kalman filter to nonlinear systems. In: *Int. symp. aerospace/defense sensing, simul. and controls*, vol 3, pp 3–2
- Kalman RE (1960) A new approach to linear filtering and prediction problems. *Journal of Fluids Engineering* 82(1):35–45
- Kawahara Y, Yairi T, Machida K (2007) A kernel subspace method by stochastic realization for learning nonlinear dynamical systems. *Advances in Neural Information Processing Systems 19*:665–672
- Mccalman L, O'callaghan S, Ramos F (2013) Multi-modal estimation with kernel embeddings for learning motion models. In: *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, IEEE, Karlsruhe, pp 2845–2852
- McElhove BA (1966) An assessment of the navigation and course corrections for a manned flyby of mars or venus. *IEEE Transactions on Aerospace and Electronic Systems AES-2*
- Nishiyama Y, Afsharinejad A, Naruse S, Boots B, Song L (2016) The nonparametric kernel bayes smoother. In: *International Conference on Artificial Intelligence and Statistics*, vol 41, pp 547–555
- Rahimi A, Recht B (2007) Random features for large-scale kernel machines. *Advances in Neural Information Processing Systems (1)*:1–8
- Ralaivola L, d'Alche Buc F (2005) Time series filtering, smoothing and learning using the kernel kalman filter. *Proceedings 2005 IEEE International Joint Conference on Neural Networks*, 2005 3:1449–1454
- Schölkopf B, Herbrich R, Smola AJ (2001) A generalized representer theorem
- Shannon M, Zen H, Byrne W (2013) Autoregressive models for statistical parametric speech synthesis. *IEEE Transactions on Audio, Speech and Language Processing* 21(3):587–597
- Simon D (2006) *Optimal State Estimation*. John Wiley & Sons, Inc., Hoboken, NJ, USA
- Smith GL, Schmidt SF, McGee LA (1962) Application of statistical filter theory to the optimal estimation of position and velocity on board a circumlunar vehicle. *National Aeronautics and Space Administration*
- Smola AJ, Bartlett PP (2001) Sparse greedy gaussian process regression. *Advances in Neural Information Processing Systems 13* 13:619–625
- Smola AJ, Gretton A, Song L, Schölkopf B (2007) A hilbert space embedding for distributions. In: *Proceedings of the 18th international conference on Algorithmic Learning Theory*, Springer Berlin Heidelberg, Berlin, Heidelberg, Lecture Notes in Computer Science, vol 4754, pp 13–31
- Snelson E, Ghahramani Z (2006) Sparse gaussian processes using pseudo-inputs. In: *Advances in Neural Information Processing Systems 18*, MIT Press, pp 1257–1264
- Song L, Huang J, Smola A, Fukumizu K (2009) Hilbert space embeddings of conditional distributions with applications to dynamical systems. In: *Proceedings of the 26th Annual International Conference on Machine Learning - ICML '09*, ACM Press, New York, New York, USA, pp 1–8

- Song L, Boots B, Siddiqi SM, Gordon GJ, Smola AJ (2010) Hilbert space embeddings of hidden markov models. In: Proceedings of the 27th international conference on machine learning (ICML-10), pp 991–998
- Song L, Fukumizu K, Gretton A (2013) Kernel embeddings of conditional distributions: A unified kernel framework for nonparametric inference in graphical models. *IEEE Signal Processing Magazine* 30(4):98–111
- Sorenson HW (1970) Least-squares estimation: from gauss to kalman. *Spectrum, IEEE* 7(7):63–68
- Sun W, Capobianco R, Gordon GJ, Bagnell JA, Boots B (2016a) Learning to smooth with bidirectional predictive state inference machines. In: *The Conference on Uncertainty in Artificial Intelligence (UAI 2016)*
- Sun W, Venkatraman A, Boots B, Bagnell JA (2016b) Learning to filter with predictive state inference machines. In: *International Conference on Machine Learning*, pp 1197–1205
- Wan EAE, Van Der Merwe R (2000) The unscented kalman filter for nonlinear estimation. In: *Proceedings of the IEEE 2000 Adaptive Systems for Signal Processing, Communications, and Control Symposium*, IEEE, pp 153–158
- Williams CKI, Seeger M (2000) Using the nyström method to speed up kernel machines. In: *Proceedings of the 13th International Conference on Neural Information*, pp 661–667
- Wojtusich J, von Stryk O (2015) Humod - a versatile and open database for the investigation, modeling and simulation of human motion dynamics on actuation level. In: *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, IEEE
- Zhu P, Chen B, Principe JC (2014) Learning nonlinear generative models of time series with a kalman filter in rkhs. *IEEE Transactions on Signal Processing* 62(1):141–155

A Derivations for the Subspace Conditional Operator

A.1 Derivation of the Subspace Conditional Operator

We define the subspace conditional operator $C_{Y|X}^S$ as the mapping from an embedding $\varphi(\mathbf{x}) \in \mathcal{H}_X$ to the mean embedding $\mu_{Y|\mathbf{x}} \in \mathcal{H}_Y$ of the conditional distribution $P(Y|\mathbf{x})$ conditioned on a certain variate \mathbf{x} . To obtain this subspace conditional operator, we first introduce an auxiliary conditional operator $C_{Y|X}^{\text{aux}}$ which maps from the subspace projection of the embedding $\Gamma_x^\top \varphi(\mathbf{x})$ to the mean map of the conditional distribution, i.e.,

$$\mu_{Y|\mathbf{x}} = C_{Y|X}^{\text{aux}} \Gamma_x^\top \varphi(\mathbf{x}). \quad (117)$$

We can then derive this auxiliary conditional operator by minimizing the squared error on the full data set

$$C_{Y|X}^{\text{aux}} = \arg \min_C \|\Phi_y - C \Gamma_x^\top \Upsilon_x\|_2 \quad (118)$$

$$0 = \frac{\partial}{\partial C_{Y|X}^{\text{aux}}} \|\Phi_y - C_{Y|X}^{\text{aux}} \Gamma_x^\top \Upsilon_x\|_2 \quad (119)$$

$$0 = -2 \left(\Phi_y - C_{Y|X}^{\text{aux}} \Gamma_x^\top \Upsilon_x \right) \Upsilon_x^\top \Gamma_x \quad (120)$$

$$C_{Y|X}^{\text{aux}} \Gamma_x^\top \Upsilon_x \Upsilon_x^\top \Gamma_x = \Phi_y \Upsilon_x^\top \Gamma_x \quad (121)$$

$$C_{Y|X}^{\text{aux}} = \Phi_y \Upsilon_x^\top \Gamma_x \left(\Gamma_x^\top \Upsilon_x \Upsilon_x^\top \Gamma_x + \lambda \mathbf{I} \right)^{-1}. \quad (122)$$

We can then substitute this result for the auxiliary conditional operator in Equation 23 and obtain the subspace conditional operator as

$$C_{Y|X}^S = C_{Y|X}^{\text{aux}} \Gamma_x^\top \quad (123)$$

$$= \Phi_y \Upsilon_x^\top \Gamma_x \left(\Gamma_x^\top \Upsilon_x \Upsilon_x^\top \Gamma_x + \lambda \mathbf{I} \right)^{-1} \Gamma_x^\top \quad (124)$$

$$= \Phi_y \bar{\mathbf{K}} \left(\bar{\mathbf{K}}^\top \bar{\mathbf{K}} + \lambda \mathbf{I} \right)^{-1} \Gamma_x^\top, \quad (125)$$

where $\bar{\mathbf{K}} = \Upsilon_x^\top \Gamma_x \in \mathbb{R}^{n \times m}$ is the kernel matrix of the sample feature set Υ_x and its subset Γ_x .

A.2 Comparison to Other Sparsification Approaches

The Nystrom method replaces a $n \times n$ Gram matrix \mathbf{K} with the low-rank approximation $\tilde{\mathbf{K}} = \mathbf{C}\mathbf{W}_k^+\mathbf{C}^\top$, where $\mathbf{C} \in \mathbb{R}^{n \times m}$ is a matrix of m sampled columns from \mathbf{G} and \mathbf{W}_k is “the best rank- k approximation to the matrix \mathbf{W} , which is a matrix whose elements consist of those elements in \mathbf{G} in the intersection of the chosen columns and the corresponding rows” (Drineas and Mahoney 2005). Applying this technique to the conditional embedding operator applied to an embedding $\boldsymbol{\mu}_X = \Upsilon_X \mathbf{m}_X$ would result in

$$\hat{C}_{Y|X} \boldsymbol{\mu}_X = \Phi_Y (\mathbf{C}\mathbf{W}_k^+\mathbf{C}^\top + \lambda \mathbf{I}_n)^{-1} \mathbf{C}\mathbf{W}_k^+\mathbf{C}^\top \mathbf{m}_X \quad (126)$$

$$= \Phi_Y \mathbf{C}\mathbf{W}_k^+ (\mathbf{C}^\top \mathbf{C}\mathbf{W}_k^+ + \lambda \mathbf{I}_m)^{-1} \mathbf{C}^\top \mathbf{m}_X. \quad (127)$$

Thus, the subspace conditional operator implements the Nystrom method with $\mathbf{C} = \tilde{\mathbf{K}}$ and $\mathbf{W} = \mathbf{I}$.

The idea of the random Fourier features is to compute the Fourier transform p of the kernel method k . Random samples are drawn from the distribution over frequencies p which are then used to construct a feature function $\mathbf{z}(\mathbf{x})$ which approximates the feature function intrinsic to the kernel function, i.e., $\mathbf{z}(\mathbf{x})^\top \mathbf{z}(\mathbf{y}) \approx k(\mathbf{x}, \mathbf{y})$. Let $\mathbf{Z} \in \mathbb{R}^{n \times m}$ be the feature matrix of the n data points with m random features. We could approximate the conditional embedding operator as

$$\hat{C}_{Y|X} \boldsymbol{\mu}_X = \Phi_Y (\mathbf{Z}\mathbf{Z}^\top + \lambda \mathbf{I}_n)^{-1} \mathbf{Z}\mathbf{Z}^\top \mathbf{m}_X \quad (128)$$

$$= \Phi_Y \mathbf{Z} (\mathbf{Z}^\top \mathbf{Z} + \lambda \mathbf{I}_m)^{-1} \mathbf{Z}^\top \mathbf{m}_X. \quad (129)$$

Again, it is easy to observe the similarity to our approach if we replace \mathbf{Z} by $\tilde{\mathbf{K}}$.

B Derivations for the Kernel Kalman Rule and its Applications

B.1 The Residual of the Observation Operator is Unbiased

We can easily show that the residual of the observation operator is unbiased by taking the expectation

$$\mathbb{E}_Y[\zeta_t] = \mathbb{E}_Y[\boldsymbol{\phi}(\mathbf{y}_t) - C_{Y|X} \boldsymbol{\varphi}(2x_t)] \quad (130)$$

$$= \mathbb{E}_{Y,X}[\boldsymbol{\phi}(\mathbf{y}_t) - C_{Y|X} \boldsymbol{\varphi}(x_t)] \quad (131)$$

$$= \mathbb{E}_{Y,X}[\boldsymbol{\phi}(\mathbf{y}_t)] - C_{Y|X} \boldsymbol{\mu}_{X,t} \quad (132)$$

$$= \mathbb{E}_X[\mathbb{E}_{Y|X}[\boldsymbol{\phi}(\mathbf{y}_t)]] - \mathbb{E}_X[\mathbb{E}_{Y|X}[\boldsymbol{\phi}(\mathbf{y}_t)]] = 0. \quad (133)$$

Here, we used the definition of the conditional embedding operator found in Song et al (2013).

B.2 Derivation of the Optimal Kalman Gain Operator

We want to find the kernel Kalman gain operator Q_t which minimizes the expected squared loss $\mathbb{E}[(\boldsymbol{\varepsilon}_t^+)^\top \boldsymbol{\varepsilon}_t^+]$ or equivalently the variance of the estimator. The objective for minimizing the variance can also be reformulated as minimizing the trace of the a-posteriori covariance operator $C_{XX,t}^+$ of the state \mathbf{x}_t at time t , i.e.,

$$\min_{Q_t} \mathbb{E}[(\boldsymbol{\varepsilon}_t^+)^\top \boldsymbol{\varepsilon}_t^+] = \min_{Q_t} \text{Tr} \mathbb{E}[\boldsymbol{\varepsilon}_t^+ (\boldsymbol{\varepsilon}_t^+)^\top] \quad (134)$$

$$= \min_{Q_t} \text{Tr} \mathbb{E}\left[\left(\boldsymbol{\varphi}(x_t) - \boldsymbol{\mu}_{X,t}^+\right) \left(\boldsymbol{\varphi}(x_t) - \boldsymbol{\mu}_{X,t}^+\right)^\top\right] \quad (135)$$

$$= \min_{Q_t} \text{Tr} C_{XX,t}^+. \quad (136)$$

By substituting the posterior error with $\boldsymbol{\varepsilon}_t^+ = (I - Q_t C_{Y|X})\boldsymbol{\varepsilon}_t^- - Q_t \boldsymbol{\zeta}_t$, we can now rewrite this a-posteriori covariance operator as

$$C_{XX,t}^+ = \mathbb{E} [\boldsymbol{\varepsilon}_t^+ (\boldsymbol{\varepsilon}_t^+)^T] \quad (137)$$

$$= \mathbb{E} \left[\left((I - Q_t C_{Y|X}) \boldsymbol{\varepsilon}_t^- - Q_t \boldsymbol{\zeta}_t \right) \left((I - Q_t C_{Y|X}) \boldsymbol{\varepsilon}_t^- - Q_t \boldsymbol{\zeta}_t \right)^T \right] \quad (138)$$

$$= (I - Q_t C_{Y|X}) \mathbb{E} [\boldsymbol{\varepsilon}_t^- (\boldsymbol{\varepsilon}_t^-)^T] (I - Q_t C_{Y|X})^T - \quad (139)$$

$$Q_t \mathbb{E} [\boldsymbol{\zeta}_t (\boldsymbol{\varepsilon}_t^-)^T] (I - Q_t C_{Y|X})^T - (I - Q_t C_{Y|X}) \mathbb{E} [\boldsymbol{\varepsilon}_t^- \boldsymbol{\zeta}_t^T] Q_t^T + \quad (140)$$

$$Q_t \mathbb{E} [\boldsymbol{\zeta}_t \boldsymbol{\zeta}_t^T] Q_t^T \quad (141)$$

Since the residual of the observation operator $\boldsymbol{\zeta}_t$ is assumed to be independent from the estimation error and since we assume the a-priori estimate to be zero mean we get

$$\mathbb{E} [\boldsymbol{\zeta}_t (\boldsymbol{\varepsilon}_t^-)^T] = \mathbb{E} [\boldsymbol{\zeta}_t] \mathbb{E} [(\boldsymbol{\varepsilon}_t^-)^T] = \mathbb{E} [(\boldsymbol{\varepsilon}_t^-)^T] \mathbb{E} [\boldsymbol{\zeta}_t] = \mathbb{E} [\boldsymbol{\varepsilon}_t^- \boldsymbol{\zeta}_t^T] = 0. \quad (142)$$

With this insight the posterior covariance operator can be reformulated as

$$C_{XX,t}^+ = (I - Q_t C_{Y|X}) C_{XX,t}^- (I - Q_t C_{Y|X})^T + Q_t \mathcal{R} Q_t^T, \quad (143)$$

where $\mathcal{R} = \mathbb{E}[\boldsymbol{\zeta}_t \boldsymbol{\zeta}_t^T]$ is the covariance of the residual of the observation operator. Taking the derivative of the trace of the covariance operator and setting it to zero leads to the solution for the kernel Kalman gain operator as

$$0 = 2 (I - Q_t C_{Y|X}) C_{XX,t}^- (-C_{Y|X}^T) + 2 Q_t \mathcal{R} \quad (144)$$

$$Q_t C_{Y|X} C_{XX,t}^- C_{Y|X}^T + Q_t \mathcal{R} = C_{XX,t}^- C_{Y|X}^T \quad (145)$$

$$Q_t (C_{Y|X} C_{XX,t}^- C_{Y|X}^T + \mathcal{R}) = C_{XX,t}^- C_{Y|X}^T \quad (146)$$

$$Q_t = C_{XX,t}^- C_{Y|X}^T (C_{Y|X} C_{XX,t}^- C_{Y|X}^T + \mathcal{R})^{-1}. \quad (147)$$

B.3 Simplifying the Update of the Covariance Operator

Following the derivations in (Simon 2006), we can find a simpler formulation for the update of the covariance operator in Equation 51. First, we substitute the kernel Kalman gain from Equation 52 using the notation $\mathcal{U} = (C_{Y|X} C_{XX,t}^- C_{Y|X}^T + \mathcal{R})$ which yields

$$C_{XX,t}^+ = (I - C_{XX,t}^- C_{Y|X}^T \mathcal{U}^{-1} C_{Y|X}) C_{XX,t}^- (\dots)^T + C_{XX,t}^- C_{Y|X}^T \mathcal{U}^{-1} \mathcal{R} (\dots)^T \quad (148)$$

By expanding both terms we obtain

$$C_{XX,t}^+ = C_{XX,t}^- - C_{XX,t}^- C_{Y|X}^T \mathcal{U}^{-1} C_{Y|X} C_{XX,t}^- - C_{XX,t}^- C_{Y|X}^T \mathcal{U}^{-1} C_{Y|X} C_{XX,t}^- + \quad (149)$$

$$C_{XX,t}^- C_{Y|X}^T \mathcal{U}^{-1} C_{Y|X} C_{XX,t}^- C_{Y|X}^T \mathcal{U}^{-1} C_{Y|X} C_{XX,t}^- + \quad (150)$$

$$C_{XX,t}^- C_{Y|X}^T \mathcal{U}^{-1} \mathcal{R} \mathcal{U}^{-1} C_{Y|X} C_{XX,t}^- \quad (151)$$

If we now combine the second and third as well as the last two terms in this equation, we arrive at

$$C_{XX,t}^+ = C_{XX,t}^- - 2 C_{XX,t}^- C_{Y|X}^T \mathcal{U}^{-1} C_{Y|X} C_{XX,t}^- + \quad (152)$$

$$C_{XX,t}^- C_{Y|X}^T \mathcal{U}^{-1} (C_{Y|X} C_{XX,t}^- C_{Y|X}^T + \mathcal{R}) \mathcal{U}^{-1} C_{Y|X} C_{XX,t}^- \quad (153)$$

$$= C_{XX,t}^- - 2 C_{XX,t}^- C_{Y|X}^T \mathcal{U}^{-1} C_{Y|X} C_{XX,t}^- + C_{XX,t}^- C_{Y|X}^T \mathcal{U}^{-1} \mathcal{U} \mathcal{U}^{-1} C_{Y|X} C_{XX,t}^- \quad (154)$$

$$= C_{XX,t}^- - C_{XX,t}^- C_{Y|X}^T \mathcal{U}^{-1} C_{Y|X} C_{XX,t}^- \quad (155)$$

$$= C_{XX,t}^- - Q_t C_{Y|X} C_{XX,t}^- \quad (156)$$

With this update equation for the estimator covariance, which is more concise and more similar to the original equations of the Kalman filter, we can summarize the kernel Kalman rule as follows: I. compute the kernel Kalman gain operator Q_t (Equation 52), II. update the estimator of the mean embedding $\mu_{X,t}^+$ (Equation 43), III. update the estimator of the covariance embedding $C_{XX,t}^+$ (Equation 53).

B.4 Derivations for the Sample-based Kernel Kalman Rule

The Kalman gain operator can be rewritten with the sample estimators of mean embedding, covariance embedding and conditional embedding operator as

$$Q_t = \Upsilon_x S_t^- O^\top \Phi_y^\top \left(\Phi_y O S_t^- O^\top \Phi_y^\top + \kappa I \right)^{-1}, \quad (157)$$

However, this formulation still has the inversion of a potentially infinite dimensional operator. To this end, we can apply the matrix identity $A(BA + I)^{-1} = (AB + I)^{-1}A$ to the kernel Kalman gain matrix as follows

$$Q_t = \Upsilon_x S_t^- O^\top \underbrace{\Phi_y^\top}_A \left[\underbrace{\Phi_y O S_t^- O^\top \Phi_y^\top}_B + \underbrace{\kappa I}_A \right]^{-1} \quad (158)$$

$$= \Upsilon_x S_t^- O^\top \left(\underbrace{\Phi_y^\top}_A \underbrace{\Phi_y O S_t^- O^\top}_B + \underbrace{\kappa I}_A \right)^{-1} \underbrace{\Phi_y^\top}_A \quad (159)$$

$$= \Upsilon_x \underbrace{S_t^- O^\top (G O S_t^- O^\top + \kappa I)^{-1}}_{Q_t} \Phi_y^\top, \quad (160)$$

where we defined $Q_t = S_t^- O^\top (G O S_t^- O^\top + \kappa I)^{-1} \in \mathbb{R}^{n \times n}$ with the Gram matrix of the observations $G = \Phi_y^\top \Phi_y$. Based on this formulation of the kernel Kalman gain operator, we can now derive finite vector/matrix formulations for the update equations of the estimator of the mean embedding (Equation 43) and the estimator of the covariance operator (Equation 53). For the estimator of the mean embedding, we obtain

$$\mu_{X,t}^+ = \mu_{X,t}^- + Q_t \left(\phi(y_t) - C_{Y|X} \mu_{X,t}^- \right) \quad (161)$$

$$\Upsilon_x m_t^+ = \Upsilon_x m_t^- + \Upsilon_x Q_t \Phi_y^\top \left(\phi(y_t) - \Phi_y (K + \lambda I_m)^{-1} \Upsilon_x^\top \Upsilon_x m_t^- \right), \quad (162)$$

$$= \Upsilon_x m_t^- + \Upsilon_x Q_t \left(\Phi_y^\top \phi(y_t) - \Phi_y^\top \Phi_y (K + \lambda I_m)^{-1} \Upsilon_x^\top \Upsilon_x m_t^- \right), \quad (163)$$

$$m_t^+ = m_t^- + Q_t \left(g_{y_t} - G O m_t^- \right), \quad (164)$$

where $g_{y_t} = \Phi_y^\top \phi(y_t)$ is the embedding of the measurement at time t . And the estimator for the covariance operator gets

$$\hat{C}_{XX,t}^+ = \hat{C}_{XX,t}^- - \Upsilon_x Q_t \Phi_y^\top C_{Y|X} \hat{C}_{XX,t}^- \quad (165)$$

$$\Upsilon_x S_t \Upsilon_x^\top = \Upsilon_x S_t^- \Upsilon_x^\top - \Upsilon_x Q_t \Phi_y^\top \Phi_y (K + \lambda I_m)^{-1} \Upsilon_x^\top \Upsilon_x S_t^- \Upsilon_x^\top \quad (166)$$

$$= \Upsilon_x S_t^- \Upsilon_x^\top - \Upsilon_x Q_t G (K + \lambda I_m)^{-1} K S_t^- \Upsilon_x^\top \quad (167)$$

$$S_t = S_t^- - Q_t G O S_t^-. \quad (168)$$

B.5 Derivation of the Subspace Kernel Kalman Gain Operator/Matrix

To derive the gain operator of the subspace kernel Kalman rule, we apply the subspace conditional operator

$$C_{Y|X}^S = \Phi_y \bar{K} \left(\bar{K}^\top \bar{K} + \lambda I \right)^{-1} \Gamma_x^\top, \quad (169)$$

to the kernel Kalman gain from Equation 52 and obtain

$$Q_t^S = C_{XX,t}^- \left(C_{Y|X}^S \right)^\top \left(C_{Y|X}^S C_{XX,t}^- \left(C_{Y|X}^S \right)^\top + \kappa I \right)^{-1} \quad (170)$$

$$= \Upsilon_x S_t^- \bar{K} (\mathcal{O}^S)^\top \underbrace{\bar{K}^\top \Phi_y^\top}_A \left(\underbrace{\Phi_y \bar{K} \mathcal{O}^S \bar{K}^\top S_t^- \bar{K} (\mathcal{O}^S)^\top}_B \underbrace{\bar{K}^\top \Phi_y^\top + \kappa I}_A \right)^{-1}. \quad (171)$$

Here, we denote $\mathcal{O}^S := \left(\bar{K}^\top \bar{K} + \lambda I \right)^{-1}$. To obtain a finite dimensional matrix in the inverse, we apply again the matrix identity $A(BA + I)^{-1} = (AB + I)^{-1}A$ and arrive at

$$Q_t^S = \Upsilon_x S_t^- \bar{K} (\mathcal{O}^S)^\top \left(\underbrace{\bar{K}^\top \Phi_y^\top}_A \underbrace{\Phi_y \bar{K} \mathcal{O}^S \bar{K}^\top S_t^- \bar{K} (\mathcal{O}^S)^\top}_B + \kappa I \right)^{-1} \underbrace{\bar{K}^\top \Phi_y^\top}_A \quad (172)$$

$$= \Upsilon_x S_t^- \bar{K} (\mathcal{O}^S)^\top \left(\bar{K}^\top G \bar{K} \mathcal{O}^S \bar{K}^\top S_t^- \bar{K} (\mathcal{O}^S)^\top + \kappa I \right)^{-1} \bar{K}^\top \Phi_y^\top. \quad (173)$$

Using the projecting the subspace kernel Kalman gain into the subspace spanned by the features Γ_x^\top leads to

$$\Gamma_x^\top Q_t^S = \Gamma_x^\top \Upsilon_x S_t^- \bar{K} (\mathcal{O}^S)^\top \left(\bar{K}^\top G \bar{K} \mathcal{O}^S \bar{K}^\top S_t^- \bar{K} (\mathcal{O}^S)^\top + \kappa I \right)^{-1} \bar{K}^\top \Phi_y^\top \quad (174)$$

$$= \underbrace{P_t^- (\mathcal{O}^S)^\top \left(\bar{K}^\top G \bar{K} \mathcal{O}^S P_t^- (\mathcal{O}^S)^\top + \kappa I \right)^{-1} \bar{K}^\top \Phi_y^\top}_{Q_t^S}, \quad (175)$$

where we define the subspace kernel Kalman gain matrix Q_t^S . Using this representation, we can obtain the update equations for the subspace projections of mean embedding as

$$\Gamma_x^\top \mu_{X,t}^+ = \Gamma_x^\top \left(\mu_{X,t}^- + Q_t^S \left(\phi(y_t) - C_{Y|X}^S \mu_{X,t}^- \right) \right) \quad (176)$$

$$n_{X,t}^+ = n_{X,t}^- + Q_t^S \left(g_{y_t} - G \bar{K} \mathcal{O}^S n_t^- \right), \quad (177)$$

and similarly for the covariance operator as

$$\Gamma_x^\top \hat{C}_{XX,t}^+ \Gamma_x = \Gamma_x^\top \left(\hat{C}_{XX,t}^- - Q_t^S C_{Y|X}^S \hat{C}_{XX,t}^- \right) \Gamma_x \quad (178)$$

$$P_t^+ = P_t^- - Q_t^S \Phi_y^\top \Phi_y \bar{K} \mathcal{O}^S \Gamma_x^\top \hat{C}_{XX,t}^- \Gamma_x \quad (179)$$

$$= P_t^- - Q_t^S G \bar{K} \mathcal{O}^S P_t^- \quad (180)$$