

---

# MULTI-VIEW KERNEL PCA FOR TIME SERIES FORECASTING

---

A PREPRINT

**Arun Pandey**

Department of Electrical Engineering  
ESAT-STADIUS, KU Leuven  
Kasteelpark Arenberg 10, B-3001 Leuven, Belgium  
arun.pandey@esat.kuleuven.be

**Hannes De Meulemeester**

Department of Electrical Engineering  
ESAT-STADIUS, KU Leuven  
Kasteelpark Arenberg 10, B-3001 Leuven, Belgium  
hannes.demeulemeester@esat.kuleuven.be

**Bart De Moor**

Department of Electrical Engineering  
ESAT-STADIUS, KU Leuven  
Kasteelpark Arenberg 10, B-3001 Leuven, Belgium  
bart.demoor@kuleuven.be

**Johan A.K. Suykens**

Department of Electrical Engineering  
ESAT-STADIUS, KU Leuven  
Kasteelpark Arenberg 10, B-3001 Leuven, Belgium  
johan.suykens@kuleuven.be

January 25, 2023

## ABSTRACT

In this paper, we propose a kernel principal component analysis model for multi-variate time series forecasting, where the training and prediction schemes are derived from the multi-view formulation of Restricted Kernel Machines. The training problem is simply an eigenvalue decomposition of the summation of two kernel matrices corresponding to the views of the input and output data. When a linear kernel is used for the output view, it is shown that the forecasting equation takes the form of kernel ridge regression. When that kernel is non-linear, a pre-image problem has to be solved to forecast a point in the input space. We evaluate the model on several standard time series datasets, perform ablation studies, benchmark with closely related models and discuss its results.

## 1 Introduction

Kernel methods have seen great success in many applications with very-high dimensional data but low number of samples, and are therefore one of the most popular non-parametric models. In critical machine learning applications, kernel methods are preferred due to their strong theoretical foundation in learning theory [1, 2, 3, 4]. Kernel methods map the data into a high-dimensional (possibly infinite) feature space by using the kernel trick. This kernel trick allows for natural, non-linear extensions to the traditional linear methods in terms of a dual representation using a suitable kernel function. This led to numerous popular methods such as kernel principal component analysis [5], kernel fisher discriminant analysis [6] and the least-squares support vector machine [1].

However, when it comes to learning large-scale problems, kernel methods fall behind deep learning techniques due to their time and memory complexity. This also holds in the time series analysis and forecasting domain, which has recently been dominated by specialized deep neural network models [7, 8, 9, 10].

Attempts have been made to combine kernel and deep learning methods [11, 12], especially for specific cases such as with deep gaussian processes [13] and multi-layer support vector machines [14]. Recently, a new unifying framework, named Restricted Kernel Machines (RKM), was proposed [15] that attempts to bridge kernel methods with deep learning. The Lagrangian function of the Least-Squares Support Vector Machine (LS-SVM) is similar to the energy function of Restricted Boltzmann Machine (RBMs), thereby drawing link between kernel methods and RBMs; hence the name Restricted Kernel Machines.

**Contribution:** In this work, we propose a novel kernel autoregressive time series forecasting model based on the RKM framework, where the training problem is the eigen-decomposition of two kernel matrices. Additionally, we use the same objective function to derive a novel prediction scheme to recursively forecast several steps ahead in the future.

Further, we experimentally compare the model with several closely related models on multiple publicly available time series datasets.

## 2 Related Works

**Energy-based methods:** Energy-based models encode dependencies between visible and latent variables using the so-called energy function. Such a function is used to model joint-probability distribution over visible and latent variables of the system. Such a model is trained by finding a configuration with the lowest total energy. Energy based models have the advantage over probabilistic methods that they do not require proper normalization, which in turn allows for more freedom when designing a model. One of the most well known energy-based models is the Restricted Boltzmann Machine [16]. Among other applications, this model has been extended for time series modeling [17, 18]. For an excellent review of Energy-based Models and Restricted Boltzmann Machines for time series modeling, see [19, 20].

**Kernel methods for time series:** Kernel methods have long been a staple in the time series analysis toolkit. Thanks to Mercer’s theorem [21], these methods can implicitly work in very high dimensional spaces and can often allow for complicated non-linear problems to be turned into convex optimization problems.

In the case of time series forecasting, one can change the traditional kernel ridge regression formulation into an online variant where only a number of data points are available at each time step. The solution to this problem can then be iteratively updated for each new time step. This technique is known as Kernel Recursive Least-Squares [22] (KRLS). One disadvantage to this method is that the model increases in size for each new data point that gets introduced. Various solutions for this have been proposed that attempt to limit the final model size [23, 24]. Alternatively, a bayesian formulation for KRLS exists which also includes a forgetting mechanism [25].

Support vector machines [26, 3] are a particular popular type of kernel machine since training an SVM corresponds with solving a convex optimization problem with sparse solutions. The standard, static, regression variants of support vector machines can be used in feedforward fashion for time series forecasting. Alternatively, explicitly recursive extensions have been formulated [27]. The disadvantage is that the resulting optimization problems become non-convex.

Gaussian processes [28] are the extension of kernel methods to the probabilistic setting. Unlike support vector machines, gaussian processes are not sparse, which can be problematic for large scale problems. Various solutions have been proposed to reduce the effective number of data points and train on large data sets, such as by using inducing variables [29] or by performing the computations on subsets of the data [30]. [31] showed that a gaussian process could be formulated as a latent variable model. This allowed for hierarchical, deep, gaussian processes [13].

### 2.1 Restricted Kernel Machines

Restricted Kernel Machines (RKM) [15] is a general kernel methods framework characterized by visible and hidden units. Using the Fenchel-Young conjugate for quadratic functions, RKMs make the connection between Least-Squares Support Vector Machines [3] and Restricted Boltzmann Machines [16]. The workhorse of restricted kernel machines is kernel PCA, of which the energy function can be defined in the RKM framework using visible ( $v_i$ ) and latent ( $h_i$ ) units as follows:

$$J(\mathbf{W}, \mathbf{h}_i) = \sum_{i=1}^n \langle v_i, \mathbf{W}h_i \rangle_{\mathcal{H}} + \frac{1}{2} \mathbf{h}_i^\top \mathbf{\Lambda} \mathbf{h}_i + \frac{1}{2} \text{Tr}(\mathbf{W}^\top \mathbf{W}) \quad (1)$$

The stationary points of the energy function are given by:

$$\begin{aligned} \frac{\partial J}{\partial \mathbf{h}_i} = 0 &\implies \mathbf{W}^\top v_i = \mathbf{\Lambda} \mathbf{h}_i, \quad \forall i \\ \frac{\partial J}{\partial \mathbf{W}} = 0 &\implies \mathbf{W} = \sum_{i=1}^n v_i \mathbf{h}_i^\top. \end{aligned}$$

Eliminating the  $\mathbf{W}$  from above yields the following eigenvalue problem:

$$\mathbf{K} \mathbf{H}^\top = \mathbf{H}^\top \mathbf{\Lambda} \quad (2)$$

with  $\mathbf{H} = [\mathbf{h}_1 \dots \mathbf{h}_n] \in \mathbb{R}^{s \times N}$  and  $\mathbf{\Lambda} = \text{diag}(\lambda_i) \in \mathbb{R}^{s \times s}$  where generally  $s \leq N$ , indicating the number of components. Training an RKM model consists of solving an eigenvalue problem. This is in contrast with restricted boltzmann machines where the model is trained by approximating the gradient of the log-likelihood of  $\mathbf{W}$  by using methods such as contrastive divergence [32].

The objective function  $J(\mathbf{W}, \mathbf{h}_i)$  in Eq. (1) forms the basis for more specific RKM models which can be used in downstream tasks. Additionally, a key advantage of characterizing kernel methods by visible and hidden units is that deeper, multi-layer, models can be developed similarly to deep boltzmann machines [33].

Classification and regression can be performed by introducing the primal formulation of LS-SVM into the objective function [15]. This is motivated by the fact that kernel PCA corresponds to a one-class LS-SVM problem with zero target value around which the variance is maximized [34]. Classification using RKMs has been extended to multi-view classification using a tensor-based representation of the methods [35]. RKMs have been shown to be powerful feature extractors and can be used for disentangled representation learning [36, 37]. This, in turn, allows them to be used in outlier and out-of-distribution detection tasks [38]. Finally, the latent space of an RKM can be sampled from to generate new samples from the learned distribution [39, 40, 41, 42].

One aspect that was missing, until now, was a general time series model situated in the RKM setting. This work attempts to fill this gap by introducing a framework with autoregressive RKMs. This idea was briefly introduced in [43] where an RKM model with autoregression in the latent space was discussed. In the next section, we generalize this method and introduce a complete framework for time series analysis using restricted kernel machines.

### 3 Proposed Model

Our objective is to capture the dynamics of the multi-variate time series  $\{\mathbf{x}_i\}_{i=1}^n, \mathbf{x}_i \in \mathbb{R}^d$  containing  $n$  time steps. Given the lag  $p \in \mathbb{Z}^+$ , we transform the given time series such that the input vector is a vector of lagged values and the corresponding output vector is one-step ahead data-point. Let the input data-vectors be  $\{[\mathbf{x}_i^{(1)}, \dots, \mathbf{x}_{i-p}^{(1)}, \mathbf{x}_i^{(d)}, \dots, \mathbf{x}_{i-p}^{(d)}]\}_{i=p+1}^{n-1} \subset \mathcal{X}$  and the corresponding target values be  $\{\mathbf{x}_{i+1}\}_{i=p+1}^{n-1} \subset \mathcal{Y}$ . This model structure is employed to model the dynamics of the system which enables us to use static regression models for training. However, for prediction, we employ the *recursive* model such that the predictions are used as input for further predictions  $\hat{\mathbf{y}} = f(\hat{\mathbf{y}}_{i-1}, \dots, \hat{\mathbf{y}}_{i-p})$ .

We define feature maps  $\phi : \mathcal{X} \rightarrow \mathcal{H}_x, \psi : \mathcal{Y} \rightarrow \mathcal{H}_y$  with  $\mathcal{H}_x, \mathcal{H}_y$  indicating the Reproducing Kernel Hilbert Space<sup>1</sup> (RKHS; see [44] for more details). Such a feature map could be constructed explicitly or implicitly via a kernel function  $k(\mathbf{x}_i, \mathbf{x}_j) : \mathcal{X}^2 \rightarrow \mathbb{R} : (\mathbf{x}_i, \mathbf{x}_j) \mapsto \langle \phi(\mathbf{x}_i), \phi(\mathbf{x}_j) \rangle_{\mathcal{H}_x}$ . Further we define two linear operators<sup>2</sup>  $\mathbf{U}_x : \mathbb{R}^s \rightarrow \mathcal{H}_x$  and  $\mathbf{U}_y : \mathbb{R}^s \rightarrow \mathcal{H}_y, s \leq n$  with their ad-joints  $\mathbf{U}_x^\top : \dim(\mathcal{H}_x) \rightarrow \mathbb{R}^s, \mathbf{U}_y^\top : \dim(\mathcal{H}_y) \rightarrow \mathbb{R}^s$ . Data points  $\mathbf{x}_i$  and  $\mathbf{y}_i$  will be associated to a latent variable  $\mathbf{h}_t \in \mathbb{R}^s$  through the pairing terms  $\langle \phi(\mathbf{x}_i), \mathbf{U}_x \mathbf{h}_i \rangle_{\mathcal{H}_x}$  and  $\langle \psi(\mathbf{y}_i), \mathbf{U}_y \mathbf{h}_i \rangle_{\mathcal{H}_y}$ .

#### 3.1 Training

$$\begin{aligned}
J_n(\mathbf{U}_x, \mathbf{U}_y, \mathcal{H}_n) = & \sum_{i=1}^n \left[ \underbrace{-\langle \phi(\mathbf{x}_i), \mathbf{U}_x \mathbf{h}_i \rangle_{\mathcal{H}_x}}_{\text{1st view}} - \underbrace{\langle \psi(\mathbf{y}_i), \mathbf{U}_y \mathbf{h}_i \rangle_{\mathcal{H}_y}}_{\text{2nd view}} + \frac{1}{2} \mathbf{h}_i^\top \mathbf{\Lambda} \mathbf{h}_i \right. \\
& \left. + \underbrace{\|\phi(\mathbf{x}_i)\|_{\mathcal{H}_x}^2 + \|\psi(\mathbf{y}_i)\|_{\mathcal{H}_y}^2}_{\text{regularization}} + \frac{1}{2} [\|\mathbf{U}_x\|_{\text{F}}^2 + \|\mathbf{U}_y\|_{\text{F}}^2] \right] \quad (3)
\end{aligned}$$

The first two terms in the objective maximizes the pairing between the visible variables in the feature-space  $\{\phi(\mathbf{x}) : \mathbf{x} \in \mathcal{X}\}, \{\psi(\mathbf{y}) : \mathbf{y} \in \mathcal{Y}\}$  and latent variables  $\{\mathbf{h} \in \mathbb{R}^s\}$ . In this way we map the points from the two data-sources to a single point in the latent space. The regularization terms and constraints are meant to bound the objective.

<sup>1</sup>Throughout our discussion, we assume that the feature vectors are centered in the feature-space *i.e.* (possibly infinite dimensional)  $\bar{\phi}(\mathbf{x}) = \phi(\mathbf{x}) - \mu_\phi$  with  $\mu_\phi = \mathbb{E}_{\xi \sim \mathcal{X}} [\phi(\xi)]$ . Using an implicit formulation, it suffices to notice that  $\langle \bar{\phi}(\mathbf{x}), \bar{\phi}(\mathbf{y}) \rangle_{\mathcal{H}} = \langle \phi(\mathbf{x}) - \mu_\phi, \phi(\mathbf{y}) - \mu_\phi \rangle_{\mathcal{H}} = \langle \phi(\mathbf{x}), \phi(\mathbf{y}) \rangle_{\mathcal{H}} - \langle \mu_\phi, \phi(\mathbf{y}) \rangle_{\mathcal{H}} - \langle \phi(\mathbf{x}), \mu_\phi \rangle_{\mathcal{H}} + \langle \mu_\phi, \mu_\phi \rangle_{\mathcal{H}} = k(\mathbf{x}, \mathbf{y}) - \mathbb{E}_{\xi \sim \mathcal{X}} [k(\xi, \mathbf{y})] - \mathbb{E}_{\xi \sim \mathcal{X}} [k(\mathbf{x}, \xi)] + \mathbb{E}_{\xi, \zeta \sim \mathcal{X}} [k(\xi, \zeta)]$ . In practice, we can compute statistics on  $\mathcal{X}$ .

<sup>2</sup>The linear operator  $\mathbf{U}_x$  is often referred to as a *matrix* as it only exists explicitly in the case of finite dimensional Hilbert spaces  $\mathcal{H}$ . It then takes the form  $\mathbf{U}_x \in \mathbb{R}^{\dim(\mathcal{H}_x) \times s}$ .

**Solving the objective** Given the visible variables, characterizing the stationary points of  $J(\mathbf{U}_x, \mathbf{U}_y, \mathcal{H}_T | \mathcal{X}_T)$  in the pairing linear operators and the latent variables leads to the following equations for  $1 \leq i \leq n$ :

$$\begin{cases} \frac{\partial J}{\partial \mathbf{h}_i} = 0 \Rightarrow \mathbf{\Lambda} \mathbf{h}_i = \mathbf{U}_x^\top \phi(\mathbf{x}_i) + \mathbf{U}_y^\top \psi(\mathbf{y}_i) \end{cases} \quad (4)$$

$$\begin{cases} \frac{\partial J}{\partial \mathbf{U}_x} = 0 \Rightarrow \mathbf{U}_x = \sum_{i=1}^n \phi(\mathbf{x}_i) \mathbf{h}_i^\top \end{cases} \quad (5)$$

$$\begin{cases} \frac{\partial J}{\partial \mathbf{U}_y} = 0 \Rightarrow \mathbf{U}_y = \sum_{i=1}^n \psi(\mathbf{y}_i) \mathbf{h}_i^\top. \end{cases} \quad (6)$$

Eliminating  $\mathbf{U}_x, \mathbf{U}_y$  from the above equations and denoting  $\mathbf{H} = [\mathbf{h}_1 \dots \mathbf{h}_n] \in \mathbb{R}^{s \times n}$  and  $\mathbf{\Lambda} = \text{diag}(\lambda_i) \in \mathbb{R}^{s \times s}$ , we get the following eigen-value problem:

$$[\mathbf{K}(\mathcal{X}) + \mathbf{K}(\mathcal{Y})] \mathbf{H}^\top = \mathbf{H}^\top \mathbf{\Lambda}, \quad (7)$$

where the elements of  $\mathbf{K}(\mathcal{X})_{i,j} = k_x(\mathbf{x}_i, \mathbf{x}_j)$  and  $\mathbf{K}(\mathcal{Y})_{i,j} = k_y(\mathbf{y}_i, \mathbf{y}_j)$ . Further,  $\mathbf{K}(\mathcal{X})$  and  $\mathbf{K}(\mathcal{Y})$  are Hermitian matrices and their sum is also Hermitian. *Spectral theorem* guarantees the existence of real eigenvalues when diagonalized by a unitary matrix. Moreover these eigenvalues are non-negative, since these matrices are positive semi-definite.

**Remarks** The proposed model is inspired by [43], however there are three important differences. (i) The data in the proposed model is in the auto-regressive format. This enables us to formulate the problem in the well-studied Non-linear Auto-Regressive (NAR) model setting. (ii) The second-view  $\mathbf{K}(\mathcal{Y})$  is now a kernel matrix i.e. a positive semi-definite matrix. This was not a constraint with the *time-kernel* in [43]. Moreover,  $\mathbf{K}(\mathcal{Y})$  implicitly captures the weights on the past latent-variables via a kernel function instead of hand-crafting it. (iii) The prediction scheme in [43] used past latent variables to predict the next latent variable. This was sufficient to simulate the dynamics forward after the training-set, however there was no way to start with *any* given initial state vector. In contrast, here we require the initial state vector from the input space to predict the dynamics forward.

### 3.2 Prediction

The main goal is to forecast ahead of the training set i.e. find  $\{\mathbf{x}_{n+1}, \dots, \mathbf{x}_{n'}\}$ ,  $n' > n$ . To do so, we now work in  $\mathcal{X}_{n'} = \mathcal{X}_n \cup \{\mathbf{x}_{n+1}, \dots, \mathbf{x}_{n+n'}\}$ ,  $\mathcal{H}_{n'} = \mathcal{H}_n \cup \{\mathbf{h}_{n+1}, \dots, \mathbf{h}_{n+n'}\}$  and consider  $\mathcal{A}_{n'}$ . This gives the following objective

$$\begin{aligned} J_{n'}(\mathbf{U}_x, \mathbf{U}_y, \mathcal{H}_{n'}) &= \sum_{i=1}^{n'} [-\langle \phi(\mathbf{x}_i), \mathbf{U}_x \mathbf{h}_i \rangle_{\mathcal{H}_x} - \langle \psi(\mathbf{y}_i), \mathbf{U}_y \mathbf{h}_i \rangle_{\mathcal{H}_y} + \frac{1}{2} \mathbf{h}_i^\top \mathbf{\Lambda} \mathbf{h}_i \\ &\quad + \|\phi(\mathbf{x}_i)\|_{\mathcal{H}_x}^2 + \|\psi(\mathbf{y}_i)\|_{\mathcal{H}_y}^2] + \frac{1}{2} [\|\mathbf{U}_x\|_{\mathbb{F}}^2 + \|\mathbf{U}_y\|_{\mathbb{F}}^2] \end{aligned} \quad (8)$$

Given the learned  $\mathbf{U}_x, \mathbf{U}_y$  from the training, characterizing the stationary points of  $J_{n'}(\mathcal{X}_{n'}, \mathcal{H}_{n'} | \mathbf{U}_x, \mathbf{U}_y)$  in terms of visible and latent variables gives for  $1 \leq i \leq n'$ :

$$\begin{cases} \frac{\partial J}{\partial \phi(\mathbf{x}_i)} = 0 \Rightarrow \phi(\mathbf{x}_i) = \mathbf{U}_x \mathbf{h}_i \end{cases} \quad (9)$$

$$\begin{cases} \frac{\partial J}{\partial \psi(\mathbf{y}_i)} = 0 \Rightarrow \psi(\mathbf{y}_i) = \mathbf{U}_y \mathbf{h}_i \end{cases} \quad (10)$$

$$\begin{cases} \frac{\partial J}{\partial \mathbf{h}_i} = 0 \Rightarrow \mathbf{\Lambda} \mathbf{h}_i - \mathbf{U}_y^\top \psi(\mathbf{y}_i) = \mathbf{U}_x^\top \phi(\mathbf{x}_i). \end{cases} \quad (11)$$

Using the above equations, we can obtain an expression for  $\mathbf{h}_i$  with  $\phi(\mathbf{x}_i)$ :

$$\mathbf{h}_i = (\mathbf{\Lambda} - \mathbf{U}_y^\top \mathbf{U}_y)^{-1} \mathbf{U}_x^\top \phi(\mathbf{x}_i). \quad (12)$$

Substituting it back in Eq. (10), we obtain an expression for output in the feature-space:

$$\psi(\mathbf{y}_i) = \mathbf{U}_y (\mathbf{\Lambda} - \mathbf{U}_y^\top \mathbf{U}_y)^{-1} \mathbf{U}_x^\top \phi(\mathbf{x}_i). \quad (13)$$

**Case I: Non-linear  $k_x$ , linear  $k_y$ :** In this case, Eq. (13) reduces to the following variant of kernel ridge regression  $\mathbf{y}_i = (\sum_{j=1}^n \tilde{\mathbf{h}}_j k_x(\mathbf{x}_j, \mathbf{x}_i))$ , where  $\tilde{\mathbf{h}} = \mathbf{U}_y (\mathbf{\Lambda} - \mathbf{U}_y^\top \mathbf{U}_y)^{-1} \mathbf{h}_j$ . The way in which those coefficients are obtained, differs between our method and kernel ridge regression.

**Case II: Non-linear  $k_x$ , non-linear  $k_y$ :** If the feature-map  $\psi(\cdot)$  is explicitly known and *invertible*, then after the inversion, we have the desired output  $\mathbf{y}_i$ . However typically in kernel methods, for instance when using Gaussian kernel,  $\psi(\cdot)$  is unknown. As a consequence, the closed-form expression in Eq. (13) cannot be computed. In such situations, we leverage the *kernel-trick* to obtain the distances between the implicit feature-vector and the feature-vectors of the training-set.

**Kernel-trick:** Left-multiplying Eq. (13) with  $\phi(\mathbf{y}_j)^\top$  and substituting for  $\mathbf{U}_y$  from Eq. (6), we obtain  $\forall j \in \{1, \dots, n\}$ :

$$k_y(\mathbf{y}_j, \mathbf{y}_i) = \left( \sum_{k=1}^n k_y(\mathbf{y}_j, \mathbf{y}_k) \mathbf{h}_k^\top \right) \left( \mathbf{\Lambda} - \sum_{l,m=1}^n \mathbf{h}_l k_y(\mathbf{y}_l, \mathbf{y}_m) \mathbf{h}_m^\top \right)^{-1} \left( \sum_{k=1}^n \mathbf{h}_k k(\mathbf{x}_k, \mathbf{x}_i) \right).$$

The above could be reduced in the matrix form as follows:

$$\mathbf{k}_y(\mathbf{y}_i) = \mathbf{K}(\mathcal{Y}) \mathbf{H}^\top \left( \mathbf{\Lambda} - \mathbf{H} \mathbf{K}(\mathcal{Y}) \mathbf{H}^\top \right)^{-1} \left( \mathbf{H} \mathbf{k}_x(\mathbf{x}_i) \right), \quad (14)$$

where  $\mathbf{k}_y(\mathbf{y}_i) = [k_y(\mathbf{y}_1, \mathbf{y}_i), \dots, k_y(\mathbf{y}_n, \mathbf{y}_i)]^\top$ . These distances in the feature space can be used in various *pre-image methods* to obtain the output  $\mathbf{y}_i$ . In the next section, we briefly discuss the pre-image problem and some common methods.

### 3.2.1 Pre-image Problem

The advantage of using a kernel,  $k(\mathbf{y}_t, \mathbf{y}_j) = \langle \phi(\mathbf{y}_t), \phi(\mathbf{y}_j) \rangle_{\mathcal{H}_y}$ , is that all computations can be implicitly performed in feature space and the exact mapping  $\psi: \mathcal{Y} \rightarrow \mathcal{H}_y$  is not required. However, this gives rise to the pre-image problem: given a point  $\Psi \in \mathcal{H}$ , find  $\mathbf{y} \in \mathcal{Y}$  such that  $\Psi = \psi(\mathbf{y})$ . This pre-image problem is known to be ill-posed [45] as the pre-image might not exist [2] or, if it exists, it may not be unique. Instead, an optimization problem is formulated to find the approximate pre-image  $\mathbf{y}^*$ :

$$\operatorname{argmin}_{\mathbf{y}^*} \|\Psi - \psi(\mathbf{y}^*)\|^2 \quad (15)$$

While various techniques exist to approximately solve this problem [46, 47, 48, 49], we discuss two particular methods that can conveniently be incorporated into the RKM framework.

**Kernel Smoother:** A kernel smoother [39] is essentially a weighted-average of nearby data-points in the input space where both the weights and the nearby-points are given by the kernels in Eq. (14). Let  $n_r < N$  be the number of nearest-neighbors, then the approximate pre-image for time-step  $i$  is given by:

$$\tilde{\mathbf{y}}_i = \sum_{j=1}^{n_r} \frac{k_y(\mathbf{y}_j, \mathbf{y}_i) \mathbf{y}_j}{\sum_{l=1}^{n_r} k_y(\mathbf{y}_l, \mathbf{y}_i)}.$$

With this method, the implicit assumption is that the forecasted point lives within the simplex of training data-points. Further, in the limit-case  $n_r = 1$ , only the nearest training point is chosen as the prediction. In our experiments, we treat  $n_r$  as the hyperparameter.

**Kernel Ridge Regression:** Alternatively, one can *learn* the pre-image map using kernel ridge regression [49]. Consider a specific case of Eq. (15) where the denoised point  $\Psi = \mathbf{U}_y \mathbf{h}^*$  (see Eq. (10)) i.e.  $\operatorname{argmin}_{\mathbf{x}^*} \|\mathbf{U}_y \mathbf{h}^* - \psi(\mathbf{y}^*)\|^2$ . Then one could formulate the following problem to learn the pre-image map  $\Gamma: \mathcal{H} \mapsto \mathcal{Y}$  as follows:

$$\operatorname{argmin}_{\Gamma} \sum_{i=1}^n \ell(\mathbf{y}_i, \mathbf{U}_y \mathbf{h}_i) + \lambda \Omega(\Gamma).$$

Since  $\mathcal{Y}$  is a vector space  $\mathbb{R}^d$ , we can consider the above as kernel ridge regression problem with kernel  $k$ . This gives the mapping  $\Gamma_j(\mathbf{U}_y \mathbf{h}_i) = \sum_{r=1}^n \alpha_r^j k(\mathbf{U}_y \mathbf{h}_r, \mathbf{U}_y \mathbf{h}_i) = \sum_{r=1}^n \alpha_r^j k(\mathbf{h}_r, \mathbf{h}_i), \forall j \in \{1, \dots, n\}$  where the coefficient are obtained by  $\alpha = (\mathbf{K}^\top \mathbf{K} + \lambda \mathbb{I})^{-1} \mathbf{K} \mathbf{Y}$ .

**Algorithm 1:** MV-RKM hyperparameter tuning with grid-search

---

**Input:** Time series  $\{\mathbf{x}_i\}_{i=1}^n, \mathbf{x}_i \in \mathbb{R}^d$   
**Output:** Optimal kernel parameters

```

1 while validation error  $\geq 0$  do
2   Init. kernels  $k_x, k_y$ ; lag  $p$  // Sample from grid
3    $\mathbf{X} \leftarrow \{[\mathbf{x}_i^{(1)}, \dots, \mathbf{x}_{i-p}^{(1)}, \mathbf{x}_i^{(d)}, \dots, \mathbf{x}_{i-p}^{(d)}]\}_{i=p+1}^{n-1}$  // AR format
4    $\mathbf{Y} \leftarrow \{\mathbf{x}_{i+1}\}_{i=p+1}^{n-1}$  // One-step ahead
5    $\mathbf{K}(\mathcal{X}) \leftarrow \mathbf{X}, \mathbf{K}(\mathcal{Y}) \leftarrow \mathbf{Y}$  // kernel matrices
6    $\mathbf{H}, \mathbf{\Lambda} \leftarrow (\mathbf{K}(\mathcal{X}), \mathbf{K}(\mathcal{Y}))$  // eigen-decomposition Eq. (7)

```

---

**Algorithm 2:** MV-RKM recursive prediction

---

**Input:** Initial lagged vector:  $\mathbf{x}_0 \in \mathbb{R}^{(p+1)d}$ , time-steps  $n'$   
**Output:** Recursive forecasts  $\{\mathbf{y}_0, \dots, \mathbf{y}_{n'}\}$

```

1 while  $i \leq n'$  do
2   if  $i \neq 0$  then
3      $\mathbf{x}_i \leftarrow \{\mathbf{y}_{i-1}, \mathbf{x}_{i-1}, \dots, \mathbf{x}_{i-p}\}$ 
4     Compute  $\mathbf{k}_x(\mathbf{x}_i)$  // similarity vector  $\mathcal{O}(n)$ 
5     Get  $\mathbf{h}_i$  (see Eq. (12)) // latent vector  $\mathcal{O}(1)$ 
6     if  $\psi(\cdot)$  is known & invertible then
7        $\mathbf{y}_i \leftarrow$  solve Eq. (13) //  $\mathcal{O}(1)$ 
8     else
9       Compute  $\mathbf{k}_y(\mathbf{y}_i)$  // similarity vector  $\mathcal{O}(n)$ 
10       $\mathbf{y}_i \leftarrow$  solve Eq. (15) // pre-image problem
11    i++

```

---

### 3.3 Computational complexity

The eigendecomposition for the training (Eq. (7)) requires slightly less than  $\mathcal{O}(n^3)$  operations, since the matrices are hermitian. The complexity of the predictions in the latent space is  $\mathcal{O}(n)$  (Eq. (12)), since it requires kernels  $\mathbf{k}_x(\mathbf{x}_i)$ . Further, if the  $k_y$  is a linear kernel, then the cost of prediction is constant i.e.  $\mathcal{O}(1)$  (Eq. (13)). Otherwise, one has to solve the pre-image problem and the complexity depends on the choice of pre-image method used. For instance with the kernel-smoother, it is  $\mathcal{O}(n_r)$  and with kernel ridge regression, it is  $\mathcal{O}(n^3)$ , since it requires matrix inversion.

## 4 Experiments

In this section, we perform ablation studies on our proposed model on synthetic and real-world datasets to study the influence of different hyperparameters on the embeddings  $\mathbf{H}$  and the predictions. Additionally, we show how visualizing the latent embeddings can help with explaining the data and model predictions. Finally, the model is quantitatively compared against other standard time series models.

### 4.1 Ablation study

**Sine wave.** A simple sine wave is used to illustrate how the models learns the dynamics present in the input space, as latent representations. After hyperparameter tuning, we obtain a closed-loop attractor in the latent space (ref. Fig. 1). The forecasted embeddings of the unseen test set, indicated by the green curve, follow the same manifold as the embeddings of training set as indicated with blue curve. When the forecasted embeddings are transformed back to input space, it becomes apparent that the model has learned the data correctly, reflected by the high accuracy of the forecasted data points.

**Sum of sines.** A synthetic dataset was created as the sum of two sine waves as follows:  $\sum_{i=1}^n \alpha_i \sin(2\pi\omega_i t + a_i)$ , where  $n = 2$ ,  $\alpha_i \in \{1, 0.2\}$ ,  $\omega_i \in \{1, 20\}$ ,  $a_i \in \{0, 0\}$ . Using this dataset, we illustrate how the predictions of the model can be explained using the orthogonal latent components. On Fig. 2, the latent components are visualised on the right. On the left, the predictions of the model are shown where iteratively more and more of the plotted components



Table 1: Mean squared error on the forecasted data. Standard deviation for 10 iterations between brackets for the stochastic models.

Data	RNN	ARMA	KRLS	LSTM	LS-SVM	RRKM	RKM-MV	
							( $k_y$ rbf)	( $k_y$ lin)
SantaFe	3037.17 ( $\pm 625.36$ )	2224.55	416.17	2272.13 ( $\pm 328.87$ )	113.78	119.06	<b>90.23</b>	127.83
Lorenz	386.01 ( $\pm 51.53$ )	412.28	273.00	331.49 ( $\pm 61.30$ )	<b>131.87</b>	247.19	182.65	163.659
Chickenpox	34329.95 ( $\pm 9513.07$ )	23571.35	19505.99	22296.24 ( $\pm 3097.18$ )	15541.28	20716.91	<b>15479.93</b>	16807.56
Energy	16002.11 ( $\pm 1809.89$ )	24797.40	11994.98	14781.65 ( $\pm 3126.80$ )	10635.00	12764.097	<b>10385.07</b>	10474.06
Turbine	2401.12 ( $\pm 644.53$ )	1317.67	1193.82	1817.38 ( $\pm 552.21$ )	1226.87	1299.915	<b>1126.55</b>	1291.72

are used. Visually, it can be seen that more components correspond with a more accurate forecast. This is further shown by plotting the error, in terms of the mean squared error, as done on Fig. 3a. When visualised, the orthogonal components offer a measure of explainability and show how the model decomposes the signal into uncorrelated components. Generally, we observe that higher frequency components are extracted for lower lag values and lower frequency components for high lag values (see first and second row in Fig. 2). Empirically, it was found that a lower lag means that a higher number of components are required to fit the model sufficiently well.

**SantaFe.** The same experiment was performed on the SantaFe data, which is a real world chaotic uni-variate time series dataset. We can observe the same effects as with the sine experiments in Fig. 2 (last row). The predictions improve given increasingly more of the visualised components. Further, the components can be observed to be able to capture the jumps in the chaotic time series, given a certain lag. It can be observed that, given a sufficiently large lag, more components correlate with a lower prediction error (see Fig. 3b). This effect is similar as in principal component analysis where only the top components are needed for a good reconstruction of the data.

## 4.2 Real-world experiments

**Baselines:** To provide a benchmark study, we consider six baseline models: Recurrent Neural Networks (RNN), Long-Short Term Memory networks (LSTM), Autoregressive Moving Average (ARMA), Kernel-Recursive Least-Square [25] (KRLS), non-linear autoregressive formulation of the Least-Squares Support Vector Machines [3] (LS-SVM), Recurrent Restricted Kernel Machine [43] (RRKM). Further we consider two variants of the proposed model i.e. when  $k_y(\mathbf{y}_i, \mathbf{y}_j) = \langle \mathbf{y}_i, \mathbf{y}_j \rangle$  and  $k_y(\mathbf{y}_i, \mathbf{y}_j) = \exp(-\|\mathbf{y}_i - \mathbf{y}_j\|^2 / 2\sigma_y^2)$ ; with  $k_x(\mathbf{x}_i, \mathbf{x}_j) = \exp(-\|\mathbf{x}_i - \mathbf{x}_j\|^2 / 2\sigma_x^2)$  in both cases.




Figure 1: MV-RKM on sin wave. The embedded data points are *temporally* de-correlated in the embedding space.




Figure 2: Predictions (left) and components (right) for the sum of sine and santafe experiments. More components correspond to a better prediction. A lower lag for the same parameters generally generate higher frequency components and more components are required for the same performance.

**Datasets:** We consider both uni-variate and multi-variate time series: two chaotic time series (SantaFe laser<sup>3</sup> [50], Lorenz attractor<sup>4</sup>) and three real world datasets (Chickenpox<sup>5</sup> [51], Belgian appliances energy consumption<sup>6</sup> [52] and Gas turbine<sup>7</sup> [53]). See B for further details on these datasets.

On each dataset and method, hyperparameter tuning has been performed with grid-search and the result of the best set of parameters, quantified as the mean squared error, is shown in Table 1. For all methods, the entire validation set is forecasted, in *recursive* fashion, with the initial vector as the last point of the training set. The scores for RNN, ARMA and RRKM are taken from [43] since the experiment setting remains same. From the table, we see that the proposed model RKM-MV outperforms the baselines on most datasets except Lorenz where the LS-SVM has lower error than RKM-MV. Also note that these two models are closest in performance, compared to the rest, since both

<sup>3</sup><https://rdrr.io/cran/TSPred/man/SantaFe.A.html>

<sup>4</sup><https://www.openml.org/search?type=data&sort=runs&id=42182&status=active>

<sup>5</sup><https://archive.ics.uci.edu/ml/datasets/Hungarian+Chickenpox+Cases>

<sup>6</sup><https://archive.ics.uci.edu/ml/datasets/Appliances+energy+prediction>

<sup>7</sup><https://archive.ics.uci.edu/ml/datasets/Gas+Turbine+CO+and+NOx+Emission+Data+Set>






Figure 3: Prediction error with different number of components on the sum of sines (left) and SantaFe (right) datasets. The blue curve indicates the prediction of the model for the best lag. The gray curve indicates the mean prediction error over multiple lag values with corresponding standard deviation. It can be seen that more components indicate a better reconstruction until a plateau is reached.

are in NAR setting. However the solution to LS-SVM requires solving a linear system (ref. C) in contrast with the eigen-decomposition of the proposed model. Another important difference is in the prediction schemes.

Moreover, LS-SVM outperforms RKM-MV when  $k_y$  is linear on most datasets except the Belgian appliances energy dataset. However, the ability to include non-linearity on  $k_y$  is an added advantage of the RKM-MV model which allows it to outperform other methods.

## 5 Conclusion

In this work, we introduced a time series forecasting model based on multi-view kernel PCA. This framework provides new insights into time series modeling such as latent space dynamics and novel relations between kernel PCA and time series forecasting. For future work, we believe a further extension of this model can be made to a multi-source setting. For example, when the data is coming from multiple weather stations and a forecast has to be made for all stations simultaneously. Custom kernels can be used to better capture the relationship between data-points. Moreover, existing techniques in kernel methods literature can be leveraged to improve the computational complexity such as the Nyström approximation or random Fourier features. Besides the topics mentioned in this work, the framework can be further explored towards other tasks involving time series such as denoising, handling missing values and clustering.

## Acknowledgements

European Research Council under the European Union’s Horizon 2020 research and innovation programme: ERC Advanced Grants agreements E-DUALITY(No 787960) and Back to the Roots (No 885682). This paper reflects only the authors’ views and the Union is not liable for any use that may be made of the contained information. Research Council KUL: Optimization frameworks for deep kernel machines C14/18/068; Research Fund (projects C16/15/059, C3/19/053, C24/18/022, C3/20/117, C3I-21-00316); Industrial Research Fund (Fellowships 13-0260, IOFm/16/004, IOFm/20/002) and several Leuven Research and Development bilateral industrial projects. Flemish Government Agencies: FWO: projects: GOA4917N (Deep Restricted Kernel Machines: Methods and Foundations), PhD/Postdoc grant, EOS Project no G0F6718N (SeLMA), SBO project S005319N, Infrastructure project I013218N, TBM Project T001919N, PhD Grant (SB/1SA1319N); EWI: the Flanders AI Research Program; VLAIO: CSBO (HBC.2021.0076) Baekeland PhD (HBC.20192204). This research received funding from the Flemish Government (AI Research Program). Other funding: Foundation ‘Kom op tegen Kanker’, CM (Christelijke Mutualiteit).

## A Note on kernel centering

**Centering Kernel Matrix** Let  $\Phi = [\phi(\mathbf{x}_1), \dots, \phi(\mathbf{x}_n)]$  and the sample mean be  $\boldsymbol{\mu}_\phi = \frac{1}{n} \sum_{i=1}^n \phi(\mathbf{x}_i) = \frac{1}{n} \Phi \mathbb{1}_n$ , where  $\mathbb{1}_n = [1, \dots, 1]^\top$ . Let  $\tilde{\phi}(\mathbf{x}_i) = \phi(\mathbf{x}_i) - \boldsymbol{\mu}_\phi$  denote the centered feature vector. Then

$$\begin{aligned} \tilde{k}(\mathbf{x}_i, \mathbf{x}_i) &= k(\mathbf{x}_i, \mathbf{x}_i) - \langle \phi(\mathbf{x}_i), \boldsymbol{\mu}_\phi \rangle_{\mathcal{H}} - \langle \boldsymbol{\mu}_\phi, \phi(\mathbf{x}_i) \rangle_{\mathcal{H}} - \langle \boldsymbol{\mu}_\phi, \boldsymbol{\mu}_\phi \rangle_{\mathcal{H}} \\ \tilde{k}(\mathbf{x}_i, \mathbf{x}_i) &= k(\mathbf{x}_i, \mathbf{x}_i) - \langle \phi(\mathbf{x}_i), \Phi \mathbb{1}_n \rangle_{\mathcal{H}} \frac{1}{n} - \frac{1}{n} \langle \Phi \mathbb{1}_n, \phi(\mathbf{x}_i) \rangle_{\mathcal{H}} + \frac{1}{n} \langle \Phi \mathbb{1}_n, \Phi \mathbb{1}_n \rangle_{\mathcal{H}} \frac{1}{n} \end{aligned}$$

Composing the above in the matrix form  $\forall i$ , gives the following

$$\begin{aligned} \tilde{\mathbf{K}} &= \mathbf{K} - \mathbf{K} \mathbb{1} \mathbb{1}^\top \frac{1}{n} - \frac{1}{n} \mathbb{1} \mathbb{1}^\top \mathbf{K} + \frac{1}{n^2} \mathbb{1} \mathbb{1}^\top \mathbf{K} \mathbb{1} \mathbb{1}^\top \\ \tilde{\mathbf{K}} &= (\mathbb{I} - \mathbb{1} \mathbb{1}^\top \frac{1}{n}) \mathbf{K} (\mathbb{I} - \mathbb{1} \mathbb{1}^\top \frac{1}{n}) \end{aligned} \quad (16)$$

**Centering Kernel Vector** Using the same sample statistic  $\boldsymbol{\mu}_\phi$ , a kernel vector representing the similarities  $\mathbf{k}_{\mathbf{x}_i}(\mathbf{x}^*) = [k(\mathbf{x}_1, \mathbf{x}^*), \dots, k(\mathbf{x}_n, \mathbf{x}^*)]^\top$  can be centered as follows :

$$\tilde{\mathbf{k}}_{\mathbf{x}_i}(\mathbf{x}^*) = \mathbf{k}_{\mathbf{x}_i}(\mathbf{x}^*) - \mathbf{K} \mathbb{1} \frac{1}{n} - \frac{1}{n} \mathbb{1} \mathbb{1}^\top \mathbf{k}_{\mathbf{x}_i}(\mathbf{x}^*) + \frac{1}{n^2} \mathbb{1} \mathbb{1}^\top \mathbf{K} \mathbb{1} \quad (17)$$

## B Datasets and Hyperparameters

We refer to Table 2 for more details on model architectures, datasets and hyperparameters used in this paper. The PyTorch library (double precision) in Python was used. See Algorithm 1 for training the MV-RKM model.

Table 2: Datasets used for the experiments.  $N$  is the number of training samples,  $d$  the input dimension,  $s$  the subspace dimension and  $m$  the minibatch size.

Dataset	$N_{\text{train}}$	$N_{\text{test}}$	$d$
SantaFe	1000	100	1
Lorenz	2801	1200	3
Chickenpox	418	104	20
Energy	4000	1000	28
Gas Turbine	5929	1482	11

Lorenz attractor is defined by the following dynamical system:

$$\begin{cases} \dot{x}(t) = -ax + ay \\ \dot{y}(t) = -xz + rx - y \\ \dot{z}(t) = xy - bz. \end{cases}$$

For our simulations, we set  $a = 10$ ,  $r = 28$  and  $b = 2.667$  with initial condition  $x(0) = 1.0$ ,  $y(0) = -1.0$  and  $z(0) = 1.05$  and the above equations are solved using first-order finite difference scheme.

## C Least-Squares Support Vector Machine

We consider the non-linear auto-regressive setting for LS-SVM regression model:  $\mathbf{y} = \sum_{i=1}^n \alpha_i k_x(\mathbf{x}_i, \mathbf{x}) + \mathbf{b}$ , where the parameters  $\alpha, \beta$  are found by solving the following linear system:  $\begin{bmatrix} 0 & \mathbf{1}_n^\top \\ \mathbf{1}_n & \mathbf{K}(\mathcal{X}) + \gamma^{-1} \mathbb{I} \end{bmatrix} \begin{bmatrix} \mathbf{b} \\ \boldsymbol{\alpha} \end{bmatrix} = \begin{bmatrix} 0 \\ \mathbf{y} \end{bmatrix}$ . For details on the derivation from a constrained optimization problem, we refer to [3, p. 98].




Figure 4: MV-RKM (rbf  $k_x(\sigma = 2.1856)$ , linear  $k_y$ ,  $s = 144$ ) on the SantaFe chaotic laser dataset showing short-term and long-term predictions in the first and second row respectively.

## References

- [1] Vladimir N. Vapnik. *The nature of statistical learning theory*. Springer-Verlag New York, Inc., 1995.
- [2] Bernhard Scholkopf and Alexander J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, Cambridge, MA, USA, 2001.
- [3] Johan A. K. Suykens, Tony Van Gestel, Jos De Brabanter, Bart De Moor, and Joos Vandewalle. *Least Squares Support Vector Machines*. World Scientific, River Edge, NJ, January 2002.
- [4] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian processes for machine learning*. Adaptive computation and machine learning. MIT Press, 2006.
- [5] Bernhard Schölkopf, Alexander Smola, and Klaus-Robert Müller. Nonlinear component analysis as a kernel eigenvalue problem. *Neural Computation*, 10(5):1299–1319, 1998.
- [6] S. Mika, G. Ratsch, J. Weston, B. Scholkopf, and K.R. Mullers. Fisher discriminant analysis with kernels. In *Neural Networks for Signal Processing IX: Proceedings of the 1999 IEEE Signal Processing Society Workshop (Cat. No.98TH8468)*, pages 41–48, 1999.
- [7] Boris N. Oreshkin, Dmitri Carпов, Nicolas Chapados, and Yoshua Bengio. N-beats: Neural basis expansion analysis for interpretable time series forecasting. In *International Conference on Learning Representations*, 2020.

- [8] Bryan Lim and Sercan Ö. Arik and Nicolas Loeff and Tomas Pfister. Temporal fusion transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting*, 37(4):1748–1764, 2021.
- [9] Tian Zhou, Ziqing Ma, Qingsong Wen, Xue Wang, Liang Sun, and Rong Jin. FEDformer: Frequency enhanced decomposed transformer for long-term series forecasting. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 27268–27286. PMLR, 17–23 Jul 2022.
- [10] Cristian Challu, Kin G. Olivares, Boris N. Oreshkin, Federico Garza, Max Mergenthaler-Canseco, and Artur Dubrawski. N-hits: Neural hierarchical interpolation for time series forecasting, 2022.
- [11] Youngmin Cho and Lawrence Saul. Kernel methods for deep learning. In Y. Bengio, D. Schuurmans, J. Lafferty, C. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems*, volume 22. Curran Associates, Inc., 2009.
- [12] Andrew Gordon Wilson, Zhiting Hu, Ruslan Salakhutdinov, and Eric P. Xing. Deep kernel learning. In Arthur Gretton and Christian C. Robert, editors, *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *Proceedings of Machine Learning Research*, pages 370–378, Cadiz, Spain, 09–11 May 2016. PMLR.
- [13] Andreas Damianou and Neil D. Lawrence. Deep Gaussian processes. In Carlos M. Carvalho and Pradeep Ravikumar, editors, *Proceedings of the Sixteenth International Conference on Artificial Intelligence and Statistics*, volume 31 of *Proceedings of Machine Learning Research*, pages 207–215, Scottsdale, Arizona, USA, 29 Apr–01 May 2013. PMLR.
- [14] Marco A. Wiering and Lambert R.B. Schomaker. Multi-Layer support vector machines. In J. A. K. Suykens, M. Signoretto, and A. Argyriou, editors, *Regularization, optimization, kernels, and support vector machines*, pages 457–476. Chapman and Hall/CRC, 2014.
- [15] Johan A. K. Suykens. Deep restricted kernel machines using conjugate feature duality. *Neural Computation*, 29(8):2123–2163, August 2017.
- [16] Ruslan Salakhutdinov, Andriy Mnih, and Geoffrey Hinton. Restricted Boltzmann Machines for collaborative filtering. In *ICML '07*, pages 791–798, Corvallis, Oregon, 2007. ACM Press.
- [17] Ilya Sutskever and Geoffrey Hinton. Learning multilevel distributed representations for high-dimensional sequences. In Marina Meila and Xiaotong Shen, editors, *Proceedings of the Eleventh International Conference on Artificial Intelligence and Statistics*, volume 2 of *Proceedings of Machine Learning Research*, pages 548–555, San Juan, Puerto Rico, 21–24 Mar 2007. PMLR.
- [18] Ilya Sutskever, Geoffrey E Hinton, and Graham W Taylor. The recurrent temporal restricted boltzmann machine. In *Advances in Neural Information Processing Systems*, volume 21, 2008.
- [19] Takayuki Osogami. Boltzmann machines and energy-based models. *arXiv:1708.06008 [cs]*, January 2019.
- [20] Takayuki Osogami. Boltzmann machines for time-series. *arXiv:1708.06004 [cs]*, January 2019.
- [21] James Mercer. Functions of positive and negative type, and their connection the theory of integral equations. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 209(441-458):415–446, January 1909.
- [22] Y. Engel, S. Mannor, and R. Meir. The kernel recursive least-squares algorithm. *IEEE Transactions on Signal Processing*, 52(8):2275–2285, 2004.
- [23] S. Van Vaerenbergh, J. Via, and I. Santamaria. A sliding-window kernel rls algorithm and its application to nonlinear channel identification. In *2006 IEEE International Conference on Acoustics Speech and Signal Processing Proceedings*, volume 5, pages V–V, 2006.
- [24] Steven Van Vaerenbergh, Ignacio Santamaría, Weifeng Liu, and José C. Príncipe. Fixed-budget kernel recursive least-squares. In *2010 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 1882–1885, 2010.
- [25] Steven Van Vaerenbergh, Miguel Lázaro-Gredilla, and Ignacio Santamaria. Kernel recursive least-squares tracker for time-varying regression. *Neural Networks and Learning Systems, IEEE Transactions on*, 23:1313–1326, 08 2012.
- [26] K. R. Müller, A. J. Smola, G. Rätsch, B. Schölkopf, J. Kohlmorgen, and V. Vapnik. Predicting time series with support vector machines. In Wulfram Gerstner, Alain Germond, Martin Hasler, and Jean-Daniel Nicoud, editors, *Artificial Neural Networks — ICANN'97*, pages 999–1004, Berlin, Heidelberg, 1997. Springer Berlin Heidelberg.

- [27] Johan A. K. Suykens and J. Vandewalle. Recurrent least squares support vector machines. *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, 47(7):1109–1114, 2000.
- [28] Christopher Williams and Carl Rasmussen. Gaussian processes for regression. In D. Touretzky, M.C. Mozer, and M. Hasselmo, editors, *Advances in Neural Information Processing Systems*, volume 8. MIT Press, 1995.
- [29] Michalis Titsias. Variational learning of inducing variables in sparse gaussian processes. In David van Dyk and Max Welling, editors, *Proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics*, volume 5 of *Proceedings of Machine Learning Research*, pages 567–574, Hilton Clearwater Beach Resort, Clearwater Beach, Florida USA, 16–18 Apr 2009. PMLR.
- [30] Marc Deisenroth and Jun Wei Ng. Distributed gaussian processes. In *International Conference on Machine Learning*, pages 1481–1490. PMLR, 2015.
- [31] Neil Lawrence. Probabilistic non-linear principal component analysis with gaussian process latent variable models. *Journal of Machine Learning Research*, 6(60):1783–1816, 2005.
- [32] Geoffrey E. Hinton. Training products of experts by minimizing contrastive divergence. *Neural Comput.*, 14(8):1771–1800, aug 2002.
- [33] Ruslan Salakhutdinov and Geoffrey Hinton. Deep Boltzmann Machines. In *proceedings of the Twelfth International Conference on Artificial Intelligence and Statistics*, volume 5 of JMLR, pages 448–455. PMLR, April 2009.
- [34] Johan A. K. Suykens, Tony Van Gestel, Joos Vandewalle, and Bart De Moor. A support vector machine formulation to PCA analysis and its kernel version. *IEEE Transactions on Neural Networks*, 14(2):447–450, 2003.
- [35] Lynn Houthuys and Johan A.K. Suykens. Tensor-based restricted kernel machines for multi-view classification. *Information Fusion*, 68:54–66, 2021.
- [36] Francesco Tonin, Panagiotis Patrinos, and Johan A.K. Suykens. Unsupervised learning of disentangled representations in deep restricted kernel machines with orthogonality constraints. *Neural Networks*, 142:661–679, 2021.
- [37] David Winant, Joachim Schreurs, and Johan A. K. Suykens. Latent space exploration using generative kernel PCA. In Bart Bogaerts, Gianluca Bontempi, Pierre Geurts, Nick Harley, Bertrand Leblot, Tom Lenaerts, and Gilles Louppe, editors, *Artificial Intelligence and Machine Learning*, pages 70–82, Cham, 2020. Springer International Publishing.
- [38] Francesco Tonin, Arun Pandey, Panagiotis Patrinos, and Johan A. K. Suykens. Unsupervised energy-based out-of-distribution detection using stiefel-restricted kernel machine. In *2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8, 2021.
- [39] Joachim Schreurs and Johan A. K. Suykens. Generative Kernel PCA. In *European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning*, pages 129–134, 2018.
- [40] Arun Pandey, Joachim Schreurs, and Johan A. K. Suykens. Robust generative restricted kernel machines using weighted conjugate feature duality. In *proceedings of the Sixth International Conference on Machine Learning, Optimization, and Data Science (LOD)*, 2020.
- [41] Arun Pandey, Joachim Schreurs, and Johan A.K. Suykens. Generative restricted kernel machines: A framework for multi-view generation and disentangled feature learning. *Neural Networks*, 135:177–191, 2021.
- [42] Arun Pandey, Michaël Fanuel, Joachim Schreurs, and Johan A. K. Suykens. Disentangled representation learning and generation with manifold optimization. *Neural Computation*, 34(10):2009–2036, 09 2022.
- [43] Arun Pandey, Hannes De Meulemeester, Henri De Plaen, Bart De Moor, and Johan Suykens. Recurrent Restricted Kernel Machines for Time-series Forecasting. In *ESANN 2022 proceedings*, pages 399–404, Bruges (Belgium) and online event, 2022. Ciaco - i6doc.com.
- [44] Bernhard Schölkopf, Alexander Smola, and Klaus-Robert Müller. Kernel principal component analysis. In *International conference on artificial neural networks*, pages 583–588. Springer, 1997.
- [45] Sebastian Mika, Bernhard Schölkopf, Alex Smola, Klaus-Robert Müller, Matthias Scholz, and Gunnar Rätsch. Kernel pca and de-noising in feature spaces. In *Advances in Neural Information Processing Systems*, pages 536–542. MIT Press, 1999.
- [46] Anh Tuan Bui, Joon-Ku Im, Daniel W. Apley, and George C. Runger. Projection-free kernel principal component analysis for denoising. *Neurocomputing*, 357:163–176, 2019.
- [47] James T. Kwok and Ivor Wai-Hung Tsang. The pre-image problem in kernel methods. *IEEE Transactions on Neural Networks*, 15:1517–1525, 2003.

- [48] Paul Honeine and Cedric Richard. Preimage problem in kernel-based machine learning. *IEEE Signal Processing Magazine*, 28(2):77–88, March 2011.
- [49] Jason Weston, Bernhard Schölkopf, and Gökhan Bakir. Learning to find pre-images. In *Advances in Neural Information Processing Systems*, volume 16, 2003.
- [50] A.S. Weigend and N.A. Gershenfeld. *Time Series Prediction: Forecasting the Future and Understanding the Past*. Addison-Wesley, 1993.
- [51] Benedek Rozemberczki, Paul Scherer, Oliver Kiss, Rik Sarkar, and Tamas Ferenci. Chickenpox cases in hungary: a benchmark dataset for spatiotemporal signal processing with graph neural networks, 2021.
- [52] Luis M. Candanedo, Véronique Feldheim, and Dominique Deramaix. Data driven prediction models of energy use of appliances in a low-energy house. *Energy and Buildings*, 140:81–97, 2017.
- [53] Heysem Kaya, Pınar Tüfekci, and Erdinç Uzun. Predicting co and nox emissions from gas turbines: Novel data and a benchmark pems. *Turkish Journal of Electrical Engineering and Computer Sciences*, 27(6):4783–4796, 11 2019.