



Contents lists available at ScienceDirect

Remote Sensing of Environment

journal homepage: www.elsevier.com/locate/rse

Randomized kernels for large scale Earth observation applications

Adrián Pérez-Suay*, Julia Amorós-López, Luis Gómez-Chova, Valero Laparra, Jordi Muñoz-Marí, Gustau Camps-Valls

Image Processing Laboratory (IPL), Parc Científic, Universitat de València, C/ Catedrático José Beltrán, Paterna 2. 46980, València, Spain

ARTICLE INFO

Article history:

Received 1 July 2016

Received in revised form 13 January 2017

Accepted 11 February 2017

Available online xxxxx

Keywords:

Machine learning

Image classification

Biophysical parameter estimation

Kernel machines

Random features

Cloud screening

IASI/MetOp

SEVIRI/MSG

Copernicus Sentinel-2

Emulation

Model inversion

ABSTRACT

Current remote sensing applications of bio-geophysical parameter estimation and image classification have to deal with an unprecedented big amount of heterogeneous and complex data sources. New satellite sensors involving a high number of improved time, space and wavelength resolutions give rise to challenging computational problems. Standard physical inversion techniques cannot cope efficiently with this new scenario. Dealing with land cover classification of the new image sources has also turned to be a complex problem requiring large amount of memory and processing time. In order to cope with these problems, statistical learning has greatly helped in the last years to develop statistical retrieval and classification models that can ingest large amounts of Earth observation data. Kernel methods constitute a family of powerful machine learning algorithms, which have found wide use in remote sensing and geosciences. However, kernel methods are still not widely adopted because of the high computational cost when dealing with large scale problems, such as the inversion of radiative transfer models or the classification of high spatial-spectral-temporal resolution data. This paper introduces to the remote sensing community an efficient kernel method for fast statistical retrieval of atmospheric and biophysical parameters and image classification problems. We rely on a recently presented approximation to shift-invariant kernels using projections on random Fourier features. The method proposes an explicit mapping function defined through a set of projections randomly sampled from the Fourier domain. It is proved to approximate the implicit mapping of a kernel function. This allows to deal with large-scale data but taking advantage of kernel methods. The method is simple, computationally very efficient in both memory and processing costs, and easily parallelizable. We show that kernel regression and classification is now possible for datasets with millions of samples. Examples on atmospheric parameter retrieval from hyperspectral infrared sounders like IASI/Metop; large scale emulation and inversion of the familiar PROSAIL radiative transfer model on Sentinel-2 data; and the identification of clouds over landmarks in time series of MSG/Seviri images show the efficiency and effectiveness of the proposed technique.

© 2017 Elsevier Inc. All rights reserved.

1. Introduction

1.1. Remote sensing and the big data challenge

Earth-observation (EO) satellites provide a unique source of information to address some of the challenges of the Earth system science (Berger et al., 2012). EO deals with the important objective of monitoring and modelling the processes on the Earth surface and their interaction with the atmosphere. To accomplish this ambitious goal, EO deploys both data acquired by remote sensing airborne and

satellite sensors, as well as quantitative *in situ* measurements of biophysical parameters (Camps-Valls et al., 2011). Predictive models of biophysical parameters and classification of remotely sensed images are thus relevant outputs to monitor our Planet.

In this context, current EO applications of image classification and biophysical parameter estimation, have to deal with an unprecedented big amount of heterogeneous and complex data sources. Spatio-temporally explicit quantitative methods are a requirement in a variety of Earth system data processing applications. Optical Earth observing satellites for example, endowed with a high temporal resolution, enable the retrieval and hence monitoring of climate and bio-geophysical variables (Dorigo et al., 2007; Schaepman et al., 2009). The super-spectral Copernicus Sentinel-2 (S2) (Drusch et al., 2012) and the forthcoming Sentinel-3 mission (Donlon et al., 2012), as well as the planned EnMAP (Steffler et al., 2007),

* Corresponding author.

E-mail address: adrian.perez@uv.es (A. Pérez-Suay).

HyspIRI (Roberts et al., 2012), PRISMA (Labate et al., 2009) and FLEX (Kraft et al., 2013), will soon provide unprecedented data streams. Very high resolution (VHR) sensors like Quickbird, Worldview-2 and the recent Worldview-3 (Longbotham et al., 2014) also pose big challenges to data processing. The challenge is not only attached to optical sensors. Infrared sounders, like the Infrared Atmospheric Sounding Interferometer (IASI) (Tournier et al., 2002) sensor on board the Metop satellite series, impose even larger constraints: the orbit time of Metop satellites (101 min), the large spectral resolution (8461 spectral channels between 645 cm^{-1} and 2760 cm^{-1}), and the spatial resolution (60×1530 samples) of the IASI instrument yield several hundreds of gigabytes of data to be processed daily. EO radar images also increased in resolution, and the current platforms, such as ERS-1/2, ENVISAT, RadarSAT-1/2, TerraSAR-X, and Cosmo-SkyMED give raise to extremely fine resolution data that call for advanced scalable processing methods. Besides, we should not forget the availability of the extremely large remote sensing data archives already collected by several past missions, such as ENVISAT, Seviri/MSG, Cosmo-SkyMED, Landsat, or SPOT.

These large-scale data problems require enhanced processing techniques that should be accurate, robust and fast. Standard physical inversion techniques and parametric classification algorithms cannot cope (nor adapt to) this new scenario efficiently. Over the last few decades a wide diversity of methods have been developed to tackle particular EO processing tasks, but only a few of them made it into operational processing chains, and many of them are only in its infancy.

1.2. Machine learning for Earth observation data analysis

In order to cope with these problems, statistical learning (also known as machine learning) has greatly helped in the last years to develop statistical retrieval and classification models that can ingest large amount of Earth observation data. Machine learning has become a standard paradigm for the analysis of remote sensing and geoscience data, at both local and global scales (Camps-Valls et al., 2011). Machine learning actually constitute a relevant alternative to parametric and physically-based models, which rely on established physical relations and implement complex combinations of scientific hypotheses, and give rise to too rigid solutions and eventual model discrepancies (see Berger et al., 2012 and references therein).

Alternatively, the framework of *statistical inference and machine learning* is concerned about developing *data-driven models* and they solely rely on the “unreasonable effectiveness of data” (Halevy et al., 2009). The field has proven successful in many disciplines of Science and Engineering (Hastie et al., 2009) and, in general, nonlinear and nonparametric model instantiations typically lead to more flexible and improved performance over physically-based approximations.

In the last decade, machine learning has attained outstanding results in the estimation of climate variables and related bio-geo-physical parameters at local and global scales, and on the classification of remote sensing images (Camps-Valls et al., 2011). Current operational vegetation products, like leaf area index (LAI), are typically produced with neural networks (Bacour et al., 2006; Baret et al., 2013; Duveiller et al., 2011). Gross Primary Production (GPP) as the largest global CO_2 flux is estimated using ensembles of random forests and neural networks (Beer et al., 2010; Jung et al., 2011). Similarly, the contribution of supervised classifiers has been improving the efficacy of the land cover/use mapping methods since the 1970s: Gaussian models such as Linear Discriminant Analysis (LDA) were replaced in the 1990s by non-parametric models able to fit the distribution observed in data of increasing dimensionality, which were later superseded by decision trees (Friedl and Brodley, 1997; Hansen et al., 1996) and then by neural networks (NN, Bischof and Leona, 1998, Bischof et al., 1992, Bruzzone and Fernández Prieto, 1999).

1.3. Kernel machines and random features for efficient EO data processing

The last decade kernel methods emerged as a family of powerful machine learning algorithms, and found wide use in remote sensing and geosciences (Camps-Valls and Bruzzone, 2009; Camps-Valls et al., 2011). In the last decade, a kernel method called support vector machines (SVM, Camps-Valls and Bruzzone, 2005, Camps-Valls et al., 2004, Foody and Mathur, 2004, Huang et al., 2002, Melgani and Bruzzone, 2004) was gradually introduced in the field, and quickly became a standard for image classification. Further SVM developments considered the simultaneous integration of spatial, spectral and temporal information (Benediktsson et al., 2005; Camps-Valls et al., 2008; Fauvel et al., 2008; Pacifici et al., 2009; Tuia et al., 2009), the richness of hyperspectral imagery (Camps-Valls and Bruzzone, 2005; Plaza et al., 2009), and exploiting the power of clusters of computers (Muñoz-Marí et al., 2009; Plaza et al., 2008). We observed a similar adoption of kernel machines for biophysical parameter retrieval: support vector regression showed high efficiency in modelling LAI, fCOVER and evapotranspiration (Durbha et al., 2007; Yang et al., 2006), and kernel methods like Gaussian Processes (GPs) (Rasmussen and Williams, 2006) recently provided excellent results in retrieving vegetation parameters (Camps-Valls et al., 2016; Lázaro-Gredilla et al., 2014; Pasolli et al., 2010; Verrelst et al., 2013a, 2012, 2013b).

However, kernel methods are still not widely adopted because of the high computational cost when dealing with large scale problems, such as the inversion of radiative transfer models or the classification of high spatial-spectral-temporal resolution data. Roughly speaking, given n examples available to develop the models, kernel methods typically need to store in memory *kernel matrices* \mathbf{K} of size $n \times n$ and to process them using standard linear algebra tools (matrix inversion, factorization, eigendecomposition, etc.). This is an important constraint that hamper its applicability to large scale EO data processing.

In this paper, we introduce a kernel method for efficiently approximate kernels, that make nonlinear classification possibly with millions of examples. We will focus on the two most relevant EO data problems: statistical retrieval of bio-geo-physical parameters and image classification problems. The method proposes an explicit mapping function defined through a set of projections randomly sampled from the Fourier domain. It is proved to approximate the implicit mapping of a kernel function, \mathbf{K} . This allows to deal with large-scale data but taking advantage of kernel methods. The method is simple, computationally very efficient in both memory and processing costs, and easily parallelizable through standard divide-and-conquer strategies as the ones proposed in Zhang et al. (2013).

The contributions of this paper are as follows: (1) the introduction to the remote sensing community of this new efficient method to perform nonlinear regression and classification; (2) the extension of the method to work with other than Fourier bases, such as wavelets, stumps, and Walsh expansions that can cope with other data characteristics; and (3) to give experimental evidences in several illustrative and challenging real problems in EO data processing: atmospheric and biophysical parameter retrieval, model inversion and emulation of radiative transfer models, and remote sensing image classification using data from multispectral data. In particular, we will show that kernel regression/classification is now possible in applications involving datasets with millions of examples and high dimensionality. The efficiency, accuracy and effectiveness of the technique is illustrated in atmospheric parameter retrieval from hyperspectral infrared sounders like IASI, large scale emulation and inversion of the familiar PROSAIL radiative transfer model on Sentinel-2 data, and the identification of clouds over landmarks in time series of MSG/Seviri images. In addition, we will show that the method is very simple to implement, computationally very efficient

in both memory and processing cost, and easily parallelizable for operational services and product generation.

1.4. Outline

The remainder of the paper is organized as follows. Next Section 2 briefly reviews the field of kernel methods and the proposed RKS approximation, discusses implementation issues, gives intuition about the involved parameters in the model, and introduces the different novel extensions introduced in this work. Section 3 presents and discusses the experimental results in three challenging problems of bio-geo-physical parameter retrieval and data classification. Section 4 concludes the paper with some discussion and remarks, outlines the future work, and notes some research opportunities in related EO fields.

2. Random feature kernels for large scale EO data processing

Kernel methods constitute an appropriate framework to approach many statistical inference problems (Shawe-Taylor and Cristianini, 2004). In the last decade these methods have replaced other techniques in many fields of science and engineering, and have become the new standard in remote sensing data analysis (Camps-Valls and Bruzzone, 2009; Camps-Valls et al., 2011). Kernel methods allow treating in the very same framework different problems, from feature extraction (Arenas-García et al., 2013) to classification (Camps-Valls et al., 2014) and regression (CampsValls et al., 2012). The fundamental building block of the theory of kernel learning is the *kernel function*, which compares multidimensional data objects. In a nutshell, given n data points, all kernel methods have to operate with a squared (eventually huge) matrix of size $n \times n$, which contains all pairwise sample similarities. Designing an appropriate kernel function that captures data dependencies is, nevertheless, not easy in general. Many approaches have been followed so far to tackle this problem: from learning the metric implicit in the kernel (Weinberger and Saul, 2008; Weinberger and Tesauro, 2007) to learning compositions of simpler kernels (Duvenaud et al., 2013; Rakotomamonjy et al., 2008). Selecting and optimizing a kernel function is very challenging even with moderate amounts of data. Many efforts have been done to deliver large-scale versions of kernel machines able to work with several thousands of examples (Bottou et al., 2002). They typically resort to reduce the dimensionality of the problem by decomposing the kernel matrix using a subset of bases: for instance using Nyström eigendecompositions (Kumar et al., 2012), sparse and low rank approximations (Arenas-García et al., 2013; Fine and Scheinberg, 2001), or smart sample selection (Bordes et al., 2005). However, there is no clear evidence that these approaches work in general, given that they are simple heuristic approximations to the kernel.

2.1. From linear to kernel least squares regression and classification

Here we start by reviewing the standard linear regression (LR) (Geladi and Kowalski, 1986) and its kernel version, the kernel ridge regression (KRR) (Schölkopf and Smola, 2002; Shawe-Taylor and Cristianini, 2004). Note that KRR is also known as Least Squares SVM (LS-SVM), and it shares formulation with Gaussian processes (GPs) (Camps-Valls et al., 2016; Rasmussen and Williams, 2006). The main difference between KRR and GPs is the parameter training procedure: in KRR the training is usually done by cross-validation and in GPs the hyperparameters are inferred using gradient descent procedures on the marginal maximum likelihood. Reviewing LR and KRR formulation will help to understand the ideas behind the proposed RKS method.

Let $\mathbf{x}_i \in \mathbb{R}^d$ (inputs) and $\mathbf{y}_i \in \mathbb{R}^o$ (outputs), where $i = 1, \dots, n$ indicates the index of the n training samples. In the LR case we want to

fit a linear function to predict the output sample \mathbf{y}_* given an input sample $\mathbf{x}_* \in \mathbb{R}^d$, i.e. $\mathbf{y}_* = \mathbf{W}^T \mathbf{x}_*$ (we assume here that both input and output data are bias corrected). In order to fit the weights \mathbf{W} we can use the training samples (where we have access to \mathbf{y}_i examples) and perform the least squares solution: $\mathbf{W} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{Y}$, where \mathbf{X} and \mathbf{Y} contain all the training input and output samples in a matrix form respectively, i.e. \mathbf{Y} is $n \times o$ and \mathbf{X} is $n \times d$, thus \mathbf{W} is $d \times o$. Note that in general the inversion of the matrix $\mathbf{X}^T \mathbf{X}$ could be unstable, which is typically solved via Tikhonov's regularization, and then $\mathbf{W} = (\mathbf{X}^T \mathbf{X} + \lambda \mathbf{I})^{-1} \mathbf{X}^T \mathbf{Y}$. Either case, linear regression is computationally very efficient, as it only implies inverting a $d \times d$ matrix. Note that under a Bayesian interpretation, parameter λ is related to the amount of noise in the input data. In our case, we fit the parameter λ by minimizing the prediction error in a out-of-sample test set via a standard cross-validation procedure.

In the KRR case, we want to perform a linear least squares regression in a Hilbert space, \mathcal{H} , of very high (possibly infinite) dimensionality $D_{\mathcal{H}}$, where samples have been mapped through a mapping $\phi(\cdot)$. In matrix notation, the model is given by $\mathbf{Y} = \Phi \mathbf{W}_{\mathcal{H}}$, where now the weights are defined in the representation space of unknown dimensionality, $\mathbf{W}_{\mathcal{H}} \in \mathcal{H}$. Then, as in the regularized linear regression setting, we want to minimize the regularized squared loss function in \mathcal{H} ,

$$\mathcal{L}_p = \|\mathbf{Y} - \Phi \mathbf{W}_{\mathcal{H}}\|^2 + \lambda \|\mathbf{W}_{\mathcal{H}}\|^2,$$

with respect to model weights $\mathbf{W}_{\mathcal{H}}$ and the regularization term λ . To fit λ we are going to follow the same procedure as in LR, a classical cross-validation with the training samples. To find the solution for $\mathbf{W}_{\mathcal{H}}$ first we take derivatives with respect to $\mathbf{W}_{\mathcal{H}}$ and equating them to zero, by doing so we obtain $\mathbf{W}_{\mathcal{H}} = (\Phi^T \Phi + \lambda \mathbf{I})^{-1} \Phi^T \mathbf{Y}$, where Φ is the matrix of mapped samples, $[\phi(\mathbf{x}_1), \phi(\mathbf{x}_2), \dots, \phi(\mathbf{x}_n)]^T$, whose size is $n \times D_{\mathcal{H}}$. Note that this problem is not solvable as the inverse runs on matrix $\Phi^T \Phi$ which is of size $D_{\mathcal{H}} \times D_{\mathcal{H}}$, and Φ is in principle unknown. However, by applying the Representer's theorem we can express the solution as a linear combination of mapped samples, $\mathbf{W}_{\mathcal{H}} = \Phi^T \mathbf{\Lambda}$, and then the solution is expressed as a function of the dual weights $\mathbf{\Lambda} \in \mathbb{R}^{n \times o}$ (one per sample and variable), $\mathbf{\Lambda} = (\Phi \Phi^T + \lambda \mathbf{I})^{-1} \mathbf{Y}$. Now the problem is solvable as we only need to compute the inverse of the (regularized) Gram matrix $\mathbf{K} := \Phi \Phi^T$ of size $n \times n$. Even though the mapping is unknown, one can replace this inner product matrix with a similarity matrix between samples, which is known as the *kernel matrix*, \mathbf{K} .

We finally need to show that we never actually require access to the mapped feature vectors into \mathcal{H} . In practice, for a new test example \mathbf{x}_* , we only want the predicted value (\mathbf{y}_*), which is computed by projecting it onto the solution $\mathbf{W}_{\mathcal{H}}$, and then replacing the dot product with the kernel function:

$$\hat{\mathbf{y}}_* = \phi_* \mathbf{W}_{\mathcal{H}} = \phi_* \Phi^T \mathbf{\Lambda} = \mathbf{K}_* \mathbf{\Lambda}, \quad (1)$$

where the matrix \mathbf{K}_* contains the similarities between the test example \mathbf{x}_* and all training samples, \mathbf{X} . The important message here is of course that we only need access to the kernel function $K(\cdot, \cdot)$ that measures the similarity between two feature vectors, not the nonlinear mapping function $\phi(\cdot)$. Examples of typical kernel functions are the linear $K(\mathbf{x}_i, \mathbf{x}_j) = \mathbf{x}_i^T \mathbf{x}_j$, the polynomial $K(\mathbf{x}_i, \mathbf{x}_j) = (\alpha \mathbf{x}_i^T \mathbf{x}_j + b)^p$, and the one used here, the Gaussian Function (Radial Basis Function, RBF) kernel $K(\mathbf{x}_i, \mathbf{x}_j) = \exp(-\|\mathbf{x}_i - \mathbf{x}_j\|^2 / (2\sigma^2))$. Therefore, we only have two free parameters to tune: the regularization parameter λ and the kernel parameter σ . As for the λ parameter, we used a cross-validation strategy for optimizing σ . For the interested reader, a MATLAB implementation of KRR and other regression algorithms can be found at our web page <http://isp.uv.es/>.

The KRR can be readily used for binary classification as well, in which case is known as the least squares SVM. As noted in [Suykens and Vandewalle \(1999\)](#), the least squares classification problem is essentially the same as the regression problem by (1) considering the signal model $f(\mathbf{x}_i) = \text{sign}[\mathbf{w}^\top \phi(\mathbf{x}_i) + b]$; and (2) introducing equality constraints $\mathbf{y}_i(\mathbf{w}^\top \phi(\mathbf{x}_i) + b) = 1 - \mathbf{e}_i$, where \mathbf{e}_i is the error variable, $i = 1, \dots, n$. The optimization problem is exactly the same as for regression, and one only needs to encode the binary class labels as $\{-1, +1\}$ and then apply the $\text{sign}[\cdot]$ operator on the model predictions.

In conclusion, the KRR/LS-SVM can be seen as a regularized linear regression in a (possibly) infinite feature space. For doing this regression, one only needs to compute the kernel (Gram) matrix $\mathbf{K} \in \mathbb{R}^{n \times n}$, and solve the normal equations. This can be a huge computational challenge depending on the amount of available training data examples n . This is why very often kernel methods in remote sensing did not make it to be operational in problems with more than a few thousand labeled training data. In addition, note that in KRR all the training examples receive a weight, and in the test (prediction) phase, one needs to compare (compute the similarity through the kernel function) of all test data to all training data, \mathbf{K}^* . This is also a problem when millions of instances arrive in the production phase. Both problems can be addressed by explicitly defining the feature mapping, as we will see next in the proposed RKS method.

2.2. Approximating kernels with projections on random features

Instead of adopting a standard kernel function K , in this paper we explore an alternative pathway: rather than *optimization* we will follow *randomization*. While odd at a first glance, the approach has surprisingly yielded competitive results in last years, being able to exploit many samples at a fraction of the computational cost. Besides its practical convenience, the approximation of the kernel with random bases is also theoretically consistent.

The seminal work in [Rahimi and Recht \(2007\)](#) presented the randomization framework. The RKS methodology exploits a property by which a continuous kernel $K(\mathbf{x}, \mathbf{x}') = K(\mathbf{x} - \mathbf{x}', 0)$ on \mathbb{R}^d is positive definite (p.d.) if and only if K is the Fourier transform of a non-negative measure ([Reed and Simon, 1981](#); [Rudin, 1987](#)). If a shift-invariant kernel K is properly scaled, its Fourier transform $p(\boldsymbol{\omega})$ is a proper probability distribution. This property is used to approximate kernel functions and matrices with linear projections on a number of D random features, as follows:

$$K(\mathbf{x}, \mathbf{x}') = \int_{\mathcal{d}} p(\boldsymbol{\omega}) e^{-j\boldsymbol{\omega}^\top (\mathbf{x} - \mathbf{x}')} d\boldsymbol{\omega} \approx \frac{1}{D} \sum_{r=1}^D e^{-j\boldsymbol{\omega}_r^\top \mathbf{x}} e^{j\boldsymbol{\omega}_r^\top \mathbf{x}'}$$

where $p(\boldsymbol{\omega})$ is set to be the inverse Fourier transform of K , $j = \sqrt{-1}$, and $\boldsymbol{\omega}_r \in \mathbb{R}^d$ is randomly sampled from a data-independent distribution $p(\boldsymbol{\omega})$ ([Rahimi and Recht, 2009](#)).

Note that we can define a D -dimensional *randomized* feature map $\mathbf{z}(\mathbf{x}) : \mathbb{R}^d \rightarrow \mathbb{C}^D$, which can be explicitly constructed as

$$\mathbf{z}(\mathbf{x}) := [\exp(j\boldsymbol{\omega}_1^\top \mathbf{x}), \dots, \exp(j\boldsymbol{\omega}_D^\top \mathbf{x})]^\top. \quad (2)$$

Other definitions are possible: one could for instance expand the exponentials in pairs $[\cos(\boldsymbol{\omega}_r^\top \mathbf{x}), \sin(\boldsymbol{\omega}_r^\top \mathbf{x})]$, but this increases the mapped data dimensionality to \mathbb{R}^{2D} , while approximating exponentials by $[\cos(\boldsymbol{\omega}_r^\top \mathbf{x} + b_r)]$, where $b_r \sim \mathcal{U}(0, 2\pi)$, is more efficient (still mapping to \mathbb{R}^D) but has revealed less accurate, as studied in [Sutherland and Schneider \(2015\)](#). In matrix notation, given n data points, the kernel matrix $\mathbf{K} \in \mathbb{R}^{n \times n}$ can be approximated with the explicitly mapped data, $\mathbf{Z} = [\mathbf{z}_1 \dots \mathbf{z}_n]^\top \in \mathbb{R}^{n \times D}$, and will be denoted

as $\hat{\mathbf{K}} \approx \mathbf{Z}\mathbf{Z}^\top$. This property can be used to approximate any shift-invariant kernel. For instance, the familiar squared exponential (SE) Gaussian kernel $K(\mathbf{x}, \mathbf{x}') = \exp(-\|\mathbf{x} - \mathbf{x}'\|^2 / (2\sigma^2))$ can be approximated using $\boldsymbol{\omega}_r \sim \mathcal{N}(\mathbf{0}, \sigma^{-2}\mathbf{I})$, $1 \leq r \leq D$. An illustrative example of how to approximate the kernel K with random bases is given in [Fig. 1](#).

2.3. RKS in practice for classification and regression problems

The RKS algorithm reduces to three simple steps: first draw D i.i.d. samples $\boldsymbol{\omega}_1, \dots, \boldsymbol{\omega}_D \in \mathbb{R}^d$ from p , and $b_1, \dots, b_D \in \mathbb{R}$ from the uniform distribution $[0, 2\pi]$; then construct the low-dimensional feature map \mathbf{z} (e.g. $\mathbf{z} = \sqrt{\frac{2}{D}} [\cos(\boldsymbol{\omega}_1^\top \mathbf{x} + b_1), \dots, \cos(\boldsymbol{\omega}_D^\top \mathbf{x} + b_D)]^\top$), and finally solving a classical least squares regression model using \mathbf{z} as input data:

$$\mathbf{W} = (\mathbf{Z}^\top \mathbf{Z})^{-1} \mathbf{Z}^\top \mathbf{Y}. \quad (3)$$

Instead of doing this last step, one could actually use the explicit mapping \mathbf{z} to approximate the kernel function $K(\mathbf{x}_i, \mathbf{x}_j) \approx \mathbf{z}_i^\top \mathbf{z}_j$ and its associated kernel matrix, $\mathbf{K} = \mathbf{Z}\mathbf{Z}^\top$, where $\mathbf{Z} := [\mathbf{z}_1, \dots, \mathbf{z}_n]^\top \in \mathbb{R}^{n \times D}$ for optimization and prediction, and solving the problem in the dual space as classical kernel methods do. However, this would result in the same computational burden. The RKS proceeds in solving the primal problem directly: one first replaces the implicitly mapped data Φ with the explicit maps \mathbf{Z} , and solves the normal equations as for LR. Therefore the cost of the RKS reduces to invert matrix $\mathbf{Z}^\top \mathbf{Z}$ of size $D \times D$. This is very convenient for large scale problems compared to the standard kernel methods (KRR, SVM...) that involve matrices of size $n \times n$. In many current problems in EO data processing this is important since n (the number of samples available) is huge. On the other hand feature redundancy is typically present (either spatial, spectral, temporal) which justifies the use of relatively low number of random features D , and thus using RKS reduces the memory requirements drastically. The latter is in turn important for test phases involving large data streams. Efficiency in both speed and memory requirements, and for both training and testing phases, are summarized in [Table 1](#).

Finally, we would like to highlight that the RKS algorithm can actually exploit other approximating functions besides Fourier expansions. Note that actually any shift-invariant kernel, i.e. $K(\mathbf{x}, \mathbf{x}') = K(\mathbf{x} - \mathbf{x}', 0)$, can be represented using random cosine features. Randomly sampling distribution functions impacts the definition of the corresponding reproducing kernel Hilbert space: sampling the Fourier bases with $z_{\boldsymbol{\omega}}(\mathbf{x}) = \sqrt{2} \cos(\boldsymbol{\omega}^\top \mathbf{x} + b)$ actually leads to the Gaussian RBF kernel $K(\mathbf{x}, \mathbf{x}') = \exp(-\|\mathbf{x} - \mathbf{x}'\|^2 / (2\sigma^2))$, while a random stump (i.e. sigmoid-shaped functions) sampling defined by $z_{\boldsymbol{\omega}}(\mathbf{x}) = \text{sign}(\mathbf{x} - \boldsymbol{\omega})$ leads to the kernel $K(\mathbf{x}, \mathbf{x}') = 1 - \frac{1}{\pi} \|\mathbf{x} - \mathbf{x}'\|_1$. Another possibility is to resort to binning bases functions, which partition the input space using an axis-aligned grid, and assign a binary indicator to each partition, which is shown to approximate a Laplacian kernel, $K(\mathbf{x}, \mathbf{x}') = \exp(-\|\mathbf{x} - \mathbf{x}'\|_1 / (2\sigma^2))$ ([Rahimi and Recht, 2007](#)). In this paper, we will also explore the possibility of using Walsh and the Gabor basis functions widely used in signal and image processing.

3. Experimental results

This section presents experimental results on the use in RKS in several remote sensing applications: atmospheric parameter retrieval from IASI infrared sounding data; emulation and inversion of the PROSAIL radiative transfer model; and cloud detection over landmarks in MSG/SEVIRI image time series. MATLAB source code, and illustrative examples are provided for the interested reader in <http://isp.uv.es/code/rks2017.html>.

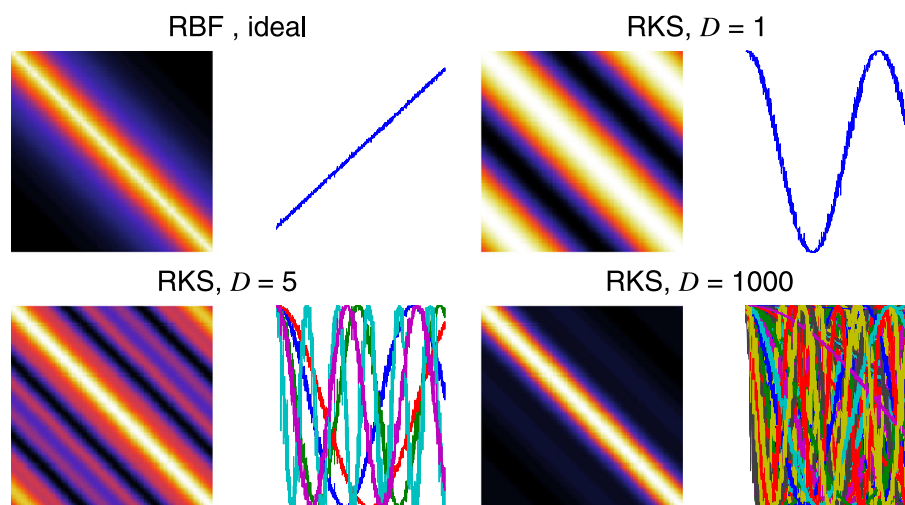


Fig. 1. Illustration of the effect of randomly sampling D bases from the Fourier domain on the kernel matrix. With sufficiently large D , the kernel matrix generated by RKS approximates that of the RBF kernel, at a fraction of the time.

3.1. Experiment 1: atmospheric parameter retrieval from MetOp/IASI infrared sounding data

In this first experiment, we exploit random feature kernels in a challenging regression problem in remote sensing: the estimation of atmospheric profiles from large scale hyperspectral infrared sounders. Temperature and water vapor are atmospheric parameters of high importance for weather forecast and atmospheric chemistry studies (Hilton et al., 2009; Liou, 2002). Observations from space-borne high spectral resolution infrared sounding instruments can be used to calculate the profiles of such atmospheric parameters with unprecedented accuracy and vertical resolution (Huang et al., 1992). In this work, we focus on the Infrared Atmospheric Sounding Interferometer (IASI) onboard Metop. The use of Metop data in Numerical Weather prediction (NWP) accounts for 40% of the impact of all space based observations in NWP forecasts.

Products obtained from IASI data are a significant improvement in the quality of the measurements used for meteorological models. In particular, IASI collects rich spectral information to derive temperature and moisture profiles, which are essential to the understanding of weather and to derive atmospheric forecasts. The sensor provides infrared spectra with high resolution between 645 cm^{-1} and 2760 cm^{-1} , from which temperature and humidity (or related dew point temperature) profiles with high vertical resolution and accuracy are derived. Additionally, it is used for the determination of trace gases such as ozone, nitrous oxide, carbon dioxide and methane, as well as land and sea surface temperature and emissivity and cloud properties. In summary, IASI provides radiances in 8461 spectral channels, between 3.62 and $15.5\text{ }\mu\text{m}$ with a spectral resolution of 0.5 cm^{-1} after apodization (Chalon et al., 2001; Siméoni et al., 1997). Its spatial resolution is 25 km at nadir with an Instantaneous Field of View (IFOV) size of 12 km at an altitude of 819 km . This huge data

dimensionality typically requires simple and computationally efficient processing techniques that can exploit the wealth of available observations provided by ECMWF re-analysis.

Actually, EUMETSAT, NOAA, NASA and other operational agencies are continuously developing product processing facilities to obtain L2 atmospheric profile products from infrared hyperspectral radiance instruments, such as IASI. One of the retrieval techniques commonly used in L2 processing is based on the canonical linear regression (LR), which is a valuable and very computationally efficient method. It consists of performing a canonical least squares linear regression on top of the data projected onto the first principal components or Empirical Orthogonal Functions (EOF) of the measured brightness temperature spectra (or radiances) and the atmospheric state parameters. To further improve the results of this scheme for retrieval, *nonlinear statistical retrieval methods* can be applied as an efficient alternative to more costly optimal estimation (OE) schemes. These methods have proven to be valid in retrieval of temperature, dew point temperature, and ozone atmospheric profiles when the original data are used (CampsValls et al., 2012).

In this experiment, we followed the same procedure as in CampsValls et al. (2012) where LR and KRR were applied using hyperpixels (i.e. all the spectral components at a particular spatial position) of IASI data to predict temperature and dew point temperature at different pressure levels. First, the dimensionality of the IASI FOVs ($d = 8461$) is reduced to 100 principal components by using the classical principal component analysis (EOF/PCA) transformation in the spectral domain. Therefore, this leads to a lower dimensionality $d = 100$ of the feature vectors. Then, LR and KRR models are trained using a fraction of the data and the rest of the data is employed to assess the models' performance. In addition to the standard LR and KRR, we also incorporate the proposed RKS method for retrieval, which allows us to train nonlinear regression efficiently.

Table 1

Computational time and memory costs for different linear, (approximate) kernel methods and random feature kernels in problems with d dimensions, D features, and n samples.

Method	Train time	Test time	Train memory	Test memory
LR (Geladi and Kowalski, 1986)	$\mathcal{O}(d^2n)$	$\mathcal{O}(d^2)$	$\mathcal{O}(d^2)$	$\mathcal{O}(d)$
Naive (Shawe-Taylor and Cristianini, 2004)	$\mathcal{O}(n^2d)$	$\mathcal{O}(nd)$	$\mathcal{O}(nd)$	$\mathcal{O}(nd)$
Low rank (Fine and Scheinberg, 2001)	$\mathcal{O}(nDd)$	$\mathcal{O}(Dd)$	$\mathcal{O}(Dd)$	$\mathcal{O}(Dd)$
RKS (Rahimi and Recht, 2007)	$\mathcal{O}(Ddn)$	$\mathcal{O}(Dd)$	$\mathcal{O}(Dd)$	$\mathcal{O}(Dd)$

For the RKS we use the recipe given in Section 2.3: first we draw D i.i.d. samples $\omega_1, \dots, \omega_D \in \mathbb{R}^d$, then we generate the projected data (Eq. (2)) and finally we solve the least squares problem (Eq. (3)). In addition to using Fourier basis, in this experiment we explore the use of different basis to generate the projections for the RKS (i.e. when applying Eq. (2)). In particular we used: Walsh, Wavelet, Stump and Sinc.

Fig. 2 shows results for the prediction of the temperature atmospheric profile. We trained LR and KRR models using $n = 5000$ examples from an IASI orbit (2008–07–17), while the RKS approximations were trained with an ensemble of 100,000 examples. All models were then tested on the same independent test set of 20,000 examples. Experiments were performed in a standard laptop using MATLAB on an Intel 3.3 GHz processor with 8 GB RAM memory under Ubuntu 14.4. Fig. 2 (a) shows that LR cannot cope with the nonlinearity of the problem, which can be addressed by using the kernel least squares regression method, KRR. However, training the KRR with more than 5000 samples turns out to be hard in regular machines. Using RKS instead is beneficial. It is actually observed that a sufficiently large number of randomly sampled bases for kernel approximation can improve the results in terms of accuracy and computational efficiency: in this case $D > 600$ random features were enough to beat the full 5000-samples KRR (thus using $n = 5000 \gg D$). The big leap in computational cost is observed in Fig. 2 (b) (note the log-scale). A trade-off comparison in Fig. 2 (c) reveals that the best accuracy-cost compromise in this particular example is to sample from the traditional squared-shaped Haar wavelet.

3.2. Experiment 2: emulation and inversion of the PROSAIL radiative transfer model for Sentinel-2

The second experiment deals with both the *statistical emulation and inversion* of the familiar PROSAIL radiative transfer model. PROSAIL is the combination of the PROSPECT leaf optical properties model and the SAIL canopy bidirectional reflectance model. PROSAIL has been used to develop new methods for retrieval of vegetation biophysical properties. Essentially, PROSAIL links the spectral variation of canopy reflectance, which is mainly related to leaf biochemical contents, with its directional variation, which is primarily related to canopy architecture and soil/vegetation contrast. This link is the key to the simultaneous estimation of canopy biophysical/structural variables for applications in agriculture, plant physiology, and ecology at different scales. PROSAIL has become one of the most popular radiative transfer tools due to its ease of use, robustness, and consistent validation by lab/field/space experiments over the years. Our

aim in this set of experiments is to both mimic and invert PROSAIL with kernel methods: these two processes imply solving the forward and inverse problems involving multidimensional large-scale datasets.

A standard approach in vegetation parameter retrieval consists of inverting PROSAIL. This *hybrid* approach essentially implies simulating radiances using PROSAIL for a set of state vectors and observation conditions. The inversion can be done in several ways, either using numerical optimization, look-up-tables, or statistical approaches (as we are interested herein). Either way, the critical question about the representativity of the created dataset still remains: typically no more than a few thousand points are generated using PROSAIL with the hope that these variable-radiances pairs explain the problem variability well. Even for fast RTMs such as PROSAIL, generating millions of pairs is challenging computationally. In recent years, machine learning techniques have been used not only for model inversion but for *emulation* of RTMs, i.e. statistical models act as fast approximations to complex physical models. This approach has a long and successful story in statistics (Kennedy and O'Hagan, 2001; O'Hagan and Kingman, 1978; Sacks et al., 1989), but efficient implementations were not accessible to the large audience because of the high computational burden involved in training the algorithms. Emulators are essentially surrogate models or metamodels: they are generally orders of magnitude faster than the original RTM, and can then be used *in lieu* of it, opening the door to more advanced biophysical parameter estimation methods, using e.g. data assimilation (DA) concepts (Lewis et al., 2012; Quaife et al., 2008).

Importantly for our interests, we should note that, once trained, the emulator can be used to generate new reflectances from new state vectors extremely fast: note that the cost is linear with the new points. We developed a KRR-based emulator of PROSAIL to generate $n = 1,000,000$ pairs of Sentinel-2 spectra ($d = 13$ spectral channels) and 7 associated parameters: Total Leaf Area Index (LAI), Leaf angle distribution (LAD), Solar Zenit Angle (SZA), Azimuth Angle (PSI), Chlorophyll a+b content C_{ab} [$\mu\text{g}/\text{cm}^2$], equivalent water thickness C_w [g/cm^2] and dry matter content, C_m [g/cm^2]. See Table 2 for some configuration details of the emulation runs. Some numbers should be given here to truly appreciate the power of emulators: even for the fast PROSAIL, generating the original $n = 5000$ training samples took 20 min, training the KRR model/emulator took around 5 min, and generating the one million dataset from the trained model (emulator) took few seconds in a standard laptop.

The developed emulator mimics the complex RTM by learning the input-output nonlinear relations directly from data. By doing so, emulators encode in a set of weights the physical rules governing the

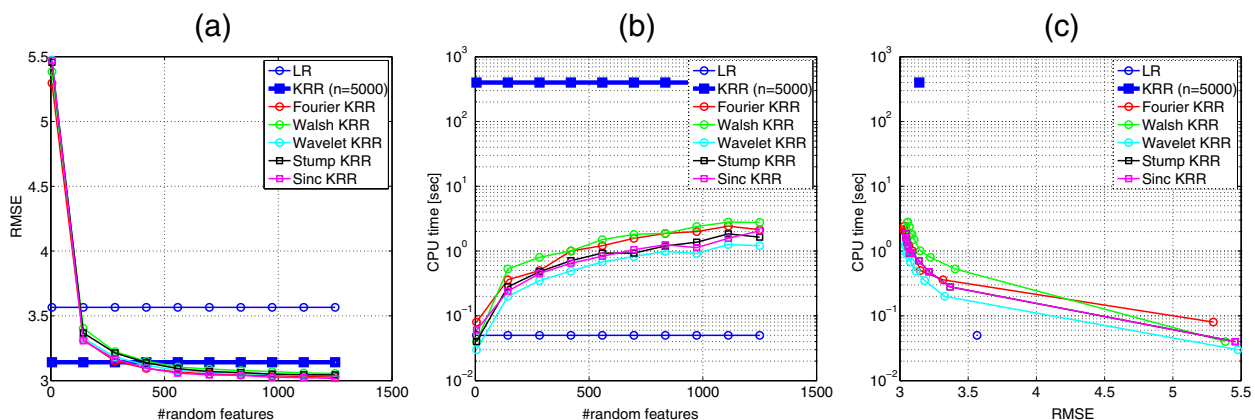


Fig. 2. Results of the RKS approach for different random sinks: (a) RMSE [K] and (b) training time [s] versus the number of random features drawn; and (c) RMSE [K] versus training time [s].

Table 2
Configuration parameters of the simulated data.

Parameter	Sampling	Min	Max
RTM model: Prospect 4			
Leaf structural parameter	Fixed	1.50	1.50
C_{ab} , chlorophyll a+b [$\mu\text{g}/\text{cm}^2$]	$\mathcal{U}(14, 49)$	0.067	79.97
C_w , equivalent water thickness [mg/cm^2]	$\mathcal{U}(10, 31)$	2	50
C_m , dry matter [mg/cm^2]	$\mathcal{U}(5.9, 19)$	1.0	3.0
RTM model: 4SAIL			
Diffuse/direct light	Fixed	10	10
Soil coefficient	Fixed	0	0
Hot spot	Fixed	0.01	0.01
Observer zenith angle	Fixed	0	0
LAI, Leaf Area Index	$\mathcal{U}(1.2, 4.3)$	0.01	6.99
LAD, Leaf Angle Distribution	$\mathcal{U}(28, 51)$	20.04	69.93
SZA, Solar Zenit Angle	$\mathcal{U}(8.5, 31)$	0.082	49.96
PSI, Azimut Angle	$\mathcal{U}(30, 100)$	0.099	179.83

vegetation-canopy interactions in PROSAIL. The one-million spectra dataset generated was now used for model inversion. Fig. 3 shows the obtained results for the inversion of the PROSAIL emulator. We show both the normalized RMSE and the computational cost of a regularized linear regression, KRR and RKS. In all cases we predict the seven parameters with a single multiple-output regression model. In this experiment, we trained KRR with $n = 2000$ samples, and consequently trained RKS for a maximum of $D = 2000$ random features for the sake of a fair comparison. We execute RKS with a total of $n = 400,000$ samples for training and a set of projections generated using Eq. (2). Several conclusions can be derived: 1) RKS yields in general competitive performance versus KRR; and 2) RKS largely improves predictions for LAD, SZA, and PSI estimation, while similar in accuracy to KRR for the rest of parameters.

3.3. Experiment 3: cloud detection over landmarks in MSG/SEVIRI

In the third experiment, we cast the problem of cloud identification over landmarks on Meteosat Second Generation (MSG) data. This satellite mission constitutes a fundamental tool for weather forecasting, providing images of the full Earth disk every 15 min. Matching the landmarks accurately is of paramount importance in image navigation and registration (INR) models and geometric

quality assessment (GQA) in the Level 1 instrument processing chain. Cloud contamination detection over landmarks is a essential step in the MSG processing chain, as undetected clouds are one of the most significant sources of error in landmark matching (see Fig. 4).

The landmark matching application requires only a binary detection. Furthermore, the cloud detection has to be carried out in real-time to be included in the landmark matching MSG processing chain, which implies efficient and robust detection schemes. Therefore, the proposed classification scheme is designed evaluating the complexity, the scalability, and parallelization of computations.

The dataset provided by EUMETSAT contains MSG/SEVIRI Level 1.5 acquisitions for 200 landmarks of variable size for a whole year, which are mainly located over the coastline, islands, or inland waters. A full Earth disk image is acquired every 15 min, which produces 96 images per day and results in 35,040 images (or chips) per landmark in 2010. Additionally, Level 2 cloud products were provided for each landmark observation so the Level 2 cloud mask (Derrien and Le Gléau, 2005) is used as the best available 'ground truth' to validate the results. Summarizing, in this problem, we have to deal with near 7 million MSG/SEVIRI multispectral images acquired during 2010.

The proposed cloud detection methodology is based on an ensemble of dedicated classifiers. We follow a divide-and-conquer strategy where specific classifiers per landmark and illumination conditions are developed. This strategy allows us to train pixel-based classifiers with millions of samples at computational affordable times. Standard steps in a pattern recognition problem are as follows: (1) pre-processing, (2) feature extraction, (3) sample selection, (4) classification, and (5) eventual combination of the individual decisions of a set of trained classifiers in order to obtain the optimal classification ensemble.

First, we perform a conversion from observed calibrated radiance to top of atmosphere (TOA) reflectance and brightness temperature units. Although in statistical retrieval these types of transformations do not dramatically help the learning models, we correct all images to compensate for illumination differences due to diurnal or seasonal cycles. It was necessary because we develop specific classification models for different ranges of sun zenith angle (SZA).

The next step is feature extraction, where different features are selected as inputs to the classifiers: 7 channels converted to TOA reflectance (R1, R2, R3, R4) and brightness temperature (BT7, BT9, BT10), 3 band ratios, and 6 spatial features (mean and standard

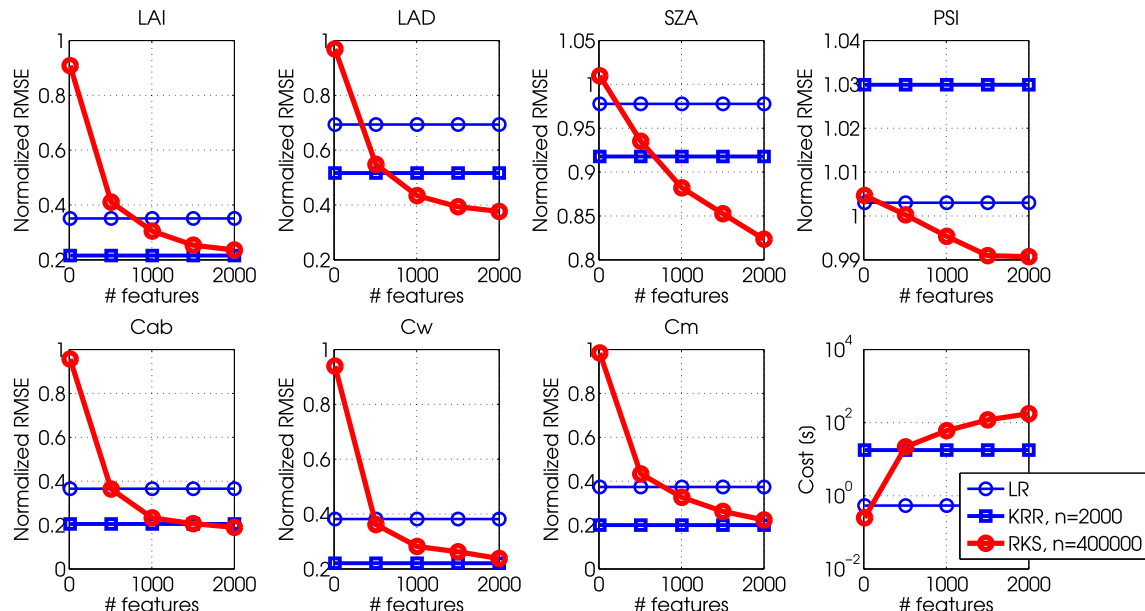


Fig. 3. RMSE results in the PROSAIL inversion experiment for the seven parameters and the computational cost (bottom right).

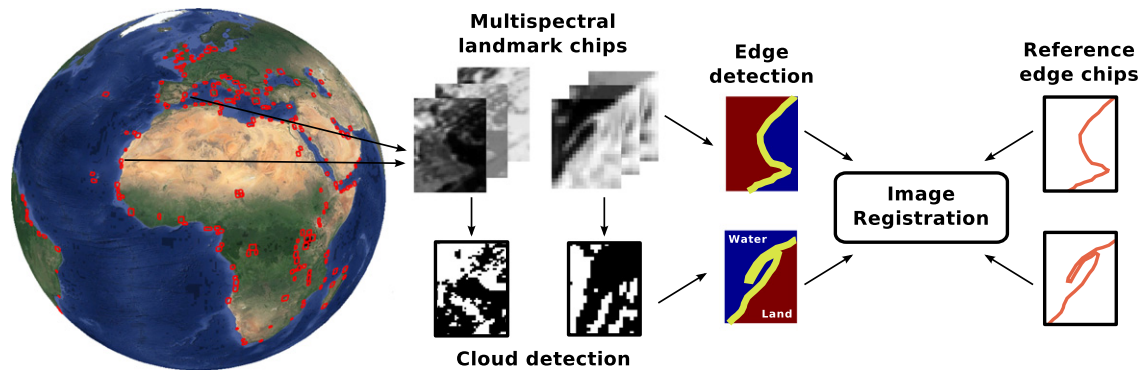


Fig. 4. Landmarks are essential in image registration and geometric quality assessment. Any misclassification of a landmark due to cloud contamination degrades the correlation matching which is a cornerstone for the image navigation and registration (INR) algorithms.

deviation of bands R1 and BT9). The informative band ratios are as follows (Derrien and Le Gléau, 2005; Hocking et al., 2011): a cloud detection ratio, $R_{0.8 \mu\text{m}}/R_{0.6 \mu\text{m}}$; a snow index, $(R_{0.6 \mu\text{m}} - R_{1.7 \mu\text{m}})/(R_{0.6 \mu\text{m}} + R_{1.7 \mu\text{m}})$; and the NDVI, $(R_{0.8 \mu\text{m}} - R_{0.6 \mu\text{m}})/(R_{0.8 \mu\text{m}} + R_{0.6 \mu\text{m}})$.

Some of the selected classifiers for benchmarking can only ingest a reduced number of training samples due to computational constraints. Hence, sample selection is a critical issue that directly affects the performance of the trained classifiers. We adopt different strategies to alleviate this issue, and also accounting for the land-cover types in each landmark. For each landmark, we select samples that cover all months/dates with a balanced number of cloud-free and cloudy over land and water. For all the analyzed sub-problems, we split the labeled dataset into two disjoint sets: the so-called training and testing sets with different sizes for training (between $n = 5000$ to $n = 10^6$) and 10^6 pixels for testing.

Finally, in order to simplify the classification task, the different illumination conditions have been independently analyzed splitting the day in four ranges (sub-problems) according to the solar zenith angle (SZA) values: high light conditions (midday), medium light conditions, low light conditions (sunrise/twilight), and night. Therefore, the final implemented classification scheme considers a pool of 4 classifiers per landmark, each one of them dedicated to different SZA ranges. We selected the SVM and RKS classifiers trained through the standard v -fold cross-validation. For all the cloud detection experiments we choose the values $D = \{5000, 7000\}$ to map from the original data (both $d = 16$ and $d = 6$ in the night scenario) to the D -dimensional space through the complex exponential mapping (i.e. Fourier basis functions) which appears in Eq. (2). To obtain the solution, a standard linear regression is performed in the D -dimensional transformed space. This solution approximates of the RBF kernel and allows the method to deal with a million of samples. Note that the SVM classifier is very computationally demanding and we restricted the training set to a limited number of samples ($n \leq 10^4$), while the proposed RKS classifier allow us to train the models with higher number of samples ($n = 10^6$), becoming an excellent alternative to SVM in large-scale classification problems.

Classification results for the different time-of-day specific classifiers for the landmark site of Ad Dakhla, Morocco, is shown on Fig. 5. We train the SVM classifiers with $[10^3-10^4]$, and RKS with $[10^3-10^6]$ data points. We can observe that the best results are reached by RKS classifiers with higher number of training points, despite the SVM outperforms RKS for less than 10^4 training samples in daytime sub-problems and 20,000 in the night case. Hence, the benefits of using RKS are more visible when more information is included ($n > 20,000$ training samples). It is worth noting that the performance of SVM is good enough and it reach a relatively high overall classification accuracy with a little amount of training data. But the efficiency of RKS dealing with large-scale datasets allows obtaining better results as increases the training set size. Moreover,

the use of more random features, D , also provides more flexible solutions increasing the detection accuracy. However, this increased dimensionality directly affects the CPU time and memory storage, leading to a better results in classification but more computationally demanding. The low light conditions (twilight) case and night case are more difficult to solve since one can not rely on visible and near infrared channels. In particular, twilight case has more variability on results due to its complexity. Night is a special case since less features are available to solve the problem (only thermal channels are feed to the classifier), which explain the lower classification accuracy compared to the high and medium light conditions.

4. Discussion and conclusions

This paper explored the use of randomly-generated bases for large scale kernel regression in several remote sensing data processing problems. We focused on the relevant issues of biophysical parameter estimation, model inversion and emulation, and multispectral time series image classification. We exploited the approximation of the kernel function via random sampling from Fourier, wavelets, Walsh and stump functions, and showed results in three relevant problems in Earth observation. First, we tackled a high-dimensional large-scale problem very common in remote sensing: the estimation of atmospheric profiles from large scale hyperspectral infrared sounding IASI radiances. Second, we explored the proposed method for the inversion of the widely used PROSAIL radiative transfer model for which we used 1 million pairs of Sentinel-2 simulations generated via a kernel emulator. Both settings induce multi-input and multi-output problems. The third application dealt with the classification of clouds over landmarks in time series of MSG/Seviri images: we exploited the methodology to train dedicated classifiers for 200 landmarks sites and with different illumination conditions. Results in all problems showed that we can train kernel regression and classification models with several hundreds of thousands of data points, which is not possible in standard kernel optimization strategies, such as support vector machines or kernel ridge regression.

We noted however that the method has two main shortcomings. First, the memory bottleneck is still present as one has to store the \mathbf{Z} matrix, which is $n \times D$. This will be addressed in the future through low-rank and block-wise approximations of \mathbf{Z} . And second, other (sparser) bases can be more appropriate. In this work, we used the Walsh basis but results did not improve those of standard Fourier bases. Alternatives to Hadamard expansions, much in line of Fast-food (Le et al., 2013), could eventually improve further the results and efficiency.

In conclusion, the proposed method produced noticeable gains in accuracy and computational efficiency in all examples. Now it is possible to train sophisticated nonlinear regression methods and classifiers using great many points in a standard laptop.

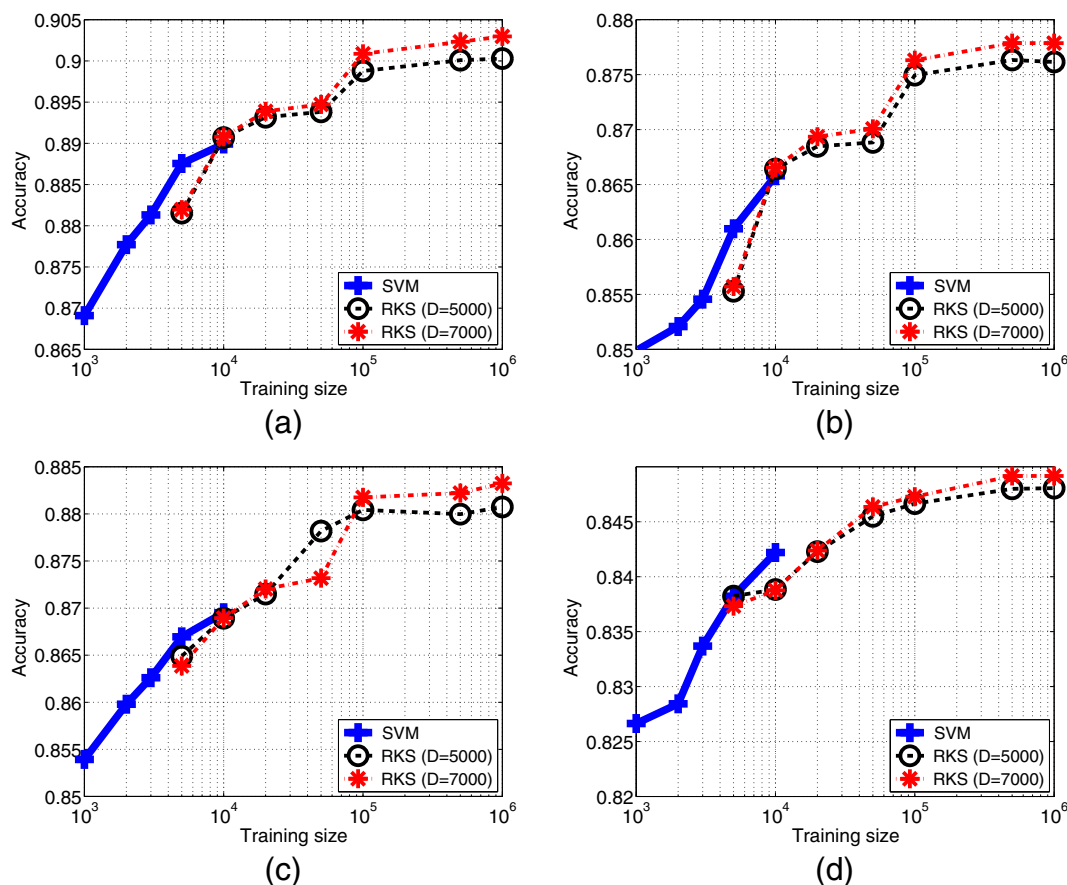


Fig. 5. SVM and RKS classification results (overall accuracy) for the four cloud detection subproblems depending on the light conditions (time-of-day): (a) high light (midday), (b) medium light, (c) low (sunrise/twilight), and (d) night.

The presented framework opens a wide venue to develop more efficient kernel machines in the new Era of big EO data. In this work we focused on the two most relevant problems: retrieval and classification. Nevertheless, it does not escape our notice that the proposed methodology can be applied to other fields of EO data processing (anomaly detection, visualization, clustering, unmixing, feature extraction, etc.) and other data modalities and sensory data (SAR, VHR, etc.).

Acknowledgments

The research leading to these results has received funding from EUMETSAT under grant agreement EUM/RSP/SOW/14/762293, the European Research Council (ERC) under the ERC-CoG-2014, SEDAL under grant agreement 647423, and the Spanish Ministry of Economy and Competitiveness (MINECO) and FEDER co-funding through the projects TEC2016-77741-R and TIN2015-64210-R.

References

- Arenas-García, J., Petersen, K.B., Camps-Valls, G., Hansen, L.K., 2013. Kernel multivariate analysis framework for supervised subspace learning: a tutorial on linear and kernel multivariate methods. *IEEE Signal Process. Mag.* 30 (4), 16–29.
- Bacour, C., Baret, F., Béal, D., Weiss, M., Pavageau, K., 2006. Neural network estimation of LAI, fAPAR, fCover and LAI×Cab, from top of canopy MERIS reflectance data: principles and validation. *Remote Sens. Environ.* 105 (4), 313–325.
- Baret, F., Weiss, M., Lacaze, R., Camacho, F., Makhmara, H., Pacholczyk, P., Smets, B., 2013. Geov1: LAI and FAPAR essential climate variables and FCOVER global time series capitalizing over existing products. Part 1: principles of development and production. *Remote Sens. Environ.* 137 (0), 299–309.
- Beer, C., Reichstein, M., Tomelleri, E., Ciais, P., Jung, M., Carvalhais, N., Rödenbeck, C., Arain, M.A., Baldocchi, D., Bonan, G.B., Bondeau, A., Cescatti, A., Lasslop, G., Lindroth, A., Lomas, M., Luyssaert, S., Margolis, H., Oleson, K.W., Rouspard, O.,

- Veenendaal, E., Viovy, N., Williams, C., Woodward, F.I., Papale, D., 2010. Terrestrial gross carbon dioxide uptake: global distribution and covariation with climate. *Science* 329 (834).
- Benediktsson, J., Palmason, J., Sveinsson, J., 2005. Classification of hyperspectral data from urban areas based on extended morphological profiles. *IEEE Trans. Geosci. Rem. Sens.* 43, 480–490.
- Berger, M., Moreno, J., Johannessen, J.A., Levelt, P., Hanssen, R., 2012. ESA's sentinel missions in support of earth system science. *Remote Sens. Environ.* 120, 84–90.
- Bischof, H., Leona, A., 1998. Finding optimal neural networks for land use classification. *IEEE Trans. Geosci. Rem. Sens.* 36 (1), 337–341.
- Bischof, H., Schneider, W., Pinz, A., 1992. May. Multispectral classification of landsat-images using neural networks. *IEEE Trans. Geosci. Remote Sens.* 30 (3), 482–490.
- Bordes, A., Ertekin, S., Weston, J., Bottou, L., 2005. Fast kernel classifiers with online and active learning. *J. Mach. Learn. Res.* 6, 1579–1619.
- Bottou, L., Chapelle, O., DeCoste, D., Weston, J., 2002. *Large-scale Kernel Machines*. MIT Press, USA.
- Bruzzzone, L., Fernández Prieto, F., 1999. A technique for the selection of kernel-function parameters in RBF neural networks for classification of remote-sensing images. *IEEE Trans. Geosci. Rem. Sens.* 37 (2), 1179–1185.
- Camps-Valls, G., Bruzzzone, L., 2005. Jun. Kernel-based methods for hyperspectral image classification. *IEEE Trans. Geosci. Rem. Sens.* 43 (6), 1351–1362.
- Camps-Valls, G., Bruzzzone, L., 2009. *Kernel Methods for Remote Sensing Data Analysis*. Wiley & Sons, UK.
- Camps-Valls, G., Gómez-Chova, L., Calpe, J., Soria, E., Martín, J.D., Alonso, L., Moreno, J., 2004. Jul. Robust support vector method for hyperspectral data classification and knowledge discovery. *IEEE Trans. Geosci. Rem. Sens.* 42 (7), 1530–1542.
- Camps-Valls, G., Gómez-Chova, L., Muñoz-Marí, J., Rojo-Álvarez, J.L., Martínez-Ramón, M., 2008. Kernel-based framework for multi-temporal and multi-source remote sensing data classification and change detection. *IEEE Trans. Geosci. Rem. Sens.* 46 (6), 1822–1835.
- Camps-Valls, G., Muñoz, J., Gómez-Chova, L., Guanter, L., Calbet, X., 2012. Nonlinear statistical retrieval of atmospheric profiles from metop-IASI and MTG-IRS infrared sounding data. *IEEE Trans. Geosci. Rem. Sens.* 50 (5), 1759–1769.
- Camps-Valls, G., Tuia, D., Bruzzzone, L., Atlí Benediktsson, J., 2014. Jan. Advances in hyperspectral image classification: Earth monitoring with statistical learning methods. *IEEE Signal Process. Mag.* 31 (1), 45–54.
- Camps-Valls, G., Tuia, D., Gómez-Chova, L., Jiménez, S., Malo, J. (Eds.), 2011. Sept. *Remote sensing image processing*. Morgan & Claypool.

- Camps-Valls, G., Verrelst, J., Muñoz-Marí, J., Laparra, V., Mateo-Jiménez, F., Gomez-Dans, J., 2016, June. A survey on Gaussian processes for earth observation data analysis. *IEEE Geosci. Rem. Sens.* (6).
- Chalon, G., Cayla, F., Diebel, D., 2001. IASI: an advanced sounder for operational meteorology. *Proceedings of the 52nd Congress of IAF.* (Toulouse, France, 1–5 October 2001).
- Derrien, M., Le Gléau, H., 2005. MSG/SEVIRI cloud mask and type from SAFNWC. *Int. J. Remote Sens.* 26 (21), 4707–4732.
- Donlon, C., Berruti, B., Buongiorno, A., Ferreira, M.-H., Féménias, P., Frerick, J., Goryl, P., Klein, U., Laur, H., Mavrocordatos, C., Nieve, J., Rebhan, H., Seitz, B., Stroede, J., Sciarra, R., 2012. The Global Monitoring for Environment and Security (GMES) Sentinel-3 mission. *Remote Sens. Environ.* 120, 37–57.
- Dorigo, W.A., Zurita-Milla, R., de Wit, A.J.W., Brazile, J., Singh, R., Schaepman, M.E., 2007. A review on reflective remote sensing and data assimilation techniques for enhanced agroecosystem modeling. *Int. J. Appl. Earth Obs. Geoinf.* 9 (2), 165–193.
- Drusch, M., Del Bello, U., Carlier, S., Colin, O., Fernandez, V., Gascon, F., Hoersch, B., Isola, C., Laberinti, P., Martimort, P., Meygret, A., Spoto, F., Sy, O., Marchese, F., Bargellini, P., 2012. Sentinel-2: ESA's optical high-resolution mission for GMES operational services. *Remote Sens. Environ.* 120, 25–36.
- Durbha, S., King, R., Younan, N., 2007. Support vector machines regression for retrieval of leaf area index from multiangle imaging spectroradiometer. *Remote Sens. Environ.* 107 (1–2), 348–361.
- Duveiller, G., Weiss, M., Baret, F., Defourny, P., 2011. Retrieving wheat green area index during the growing season from optical time series measurements based on neural network radiative transfer inversion. *Remote Sens. Environ.* 115 (3), 887–896.
- Duvenaud, D., Lloyd, J., Grosse, R., Tenenbaum, J.B., Ghahramani, Z., 2013, June. Structure discovery in nonparametric regression through compositional kernel search. *Proceedings of the 30th International Conference on Machine Learning.*
- Fauvel, M., Benediktsson, J., Chanussot, J., Sveinsson, J., 2008. Spectral and spatial classification of hyperspectral data using SVMs and morphological profiles. *IEEE Trans. Geosci. Rem. Sens.* 46 (11), 3804–3814.
- Fine, S., Scheinberg, K., 2001. Efficient SVM training using low-rank kernel representations. *J. Mach. Learn. Res.* 2, 243–264.
- Food, G.M., Mathur, J., 2004, Jul. A relative evaluation of multiclass image classification by support vector machines. *IEEE Trans. Geosci. Rem. Sens.* 1–9.
- Friedl, M.A., Brodley, C.E., 1997. Decision tree classification of land cover from remotely sensed data. *Remote Sens. Environ.* 61, 399–409.
- Geladi, P., Kowalski, B., 1986. Partial least-squares regression: a tutorial. *Anal. Chim. Acta* 185 (C), 1–17.
- Halevy, A., Norvig, P., Pereira, F., 2009, Mar. The unreasonable effectiveness of data. *IEEE Intell. Syst.* 24 (2), 8–12.
- Hansen, M., Dubayah, R., Defries, R., 1996. Classification trees: an alternative to traditional land cover classifiers. *Int. J. Rem. Sens.* 17 (5), 1075–1081.
- Hastie, T., Tibshirani, R., Friedman, J.H., 2009. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction.* Second, Springer-Verlag, New York, USA.
- Hilton, F., Atkinson, N.C., English, S.J., Eyre, J.R., 2009. Assimilation of IASI at the Met Office and assessment of its impact through observing system experiments. *Q. J. R. Meteorol. Soc.* 135, 495–505.
- Hocking, J., Francis, P.N., Saunders, R., 2011. Cloud detection in meteosat second generation imagery at the met office. *Meteorol. Appl.* 18 (3), 307–323.
- Huang, C., Davis, L.S., Townshend, J.R.G., 2002. An assessment of support vector machines for land cover classification. *Int. J. Rem. Sens.* 23 (4), 725–749.
- Huang, H.L., Smith, W.L., Woolf, H.M., 1992. Vertical resolution and accuracy of atmospheric infrared sounding spectrometers. *J. Appl. Meteor.* 31, 265–274.
- Jung, M., Reichstein, M., Margolis, H.A., Cescatti, A., Richardson, A.D., Arain, M.A., Arneth, A., Bernhofer, C., Bonal, D., Chen, J., Gianelle, D., Gobron, N., Kiely, G., Kutsch, W., Lasslop, G., Law, B.E., Lindroth, A., Merbold, L., Montagnani, L., Moors, E.J., Papale, D., Sottocornola, M., Vaccari, F., Williams, C., 2011. Global patterns of land-atmosphere fluxes of carbon dioxide, latent heat, and sensible heat derived from eddy covariance, satellite, and meteorological observations. *J. Geophys. Res.* Biogeosci. 116 (G3), 1–16.
- Kennedy, M., O'Hagan, A., 2001. Bayesian calibration of computer models. *J. R. Stat. Soc. Ser. B Stat Methodol.* 63 (3), 425–450.
- Kraft, S., Del Bello, U., Drusch, M., Gabriele, A., Harnisch, B., Moreno, J., 2013. On the demands on imaging spectrometry for the monitoring of global vegetation fluorescence from space. *Proceedings of SPIE - The International Society For Optical Engineering.* 8870.
- Kumar, S., Mohri, M., Talwalkar, A., 2012. Sampling methods for the Nyström method. *J. Mach. Learn. Res.* 13, 981–1006.
- Labate, D., Ceccherini, M., Cisbani, A., De Cosmo, V., Galeazzi, C., Giunti, L., Melozzi, M., Pieraccini, S., Stagi, M., 2009. The PRISMA payload optomechanical design, a high performance instrument for a new hyperspectral mission. *Acta Astronaut.* 65 (9–10), 1429–1436.
- Lázaro-Gredilla, M., Titsias, M.K., Verrelst, J., Camps-Valls, G., 2014, April. Retrieval of biophysical parameters with heteroscedastic Gaussian processes. *IEEE Geosci. Remote Sens. Lett.* 11 (4), 838–842.
- Le, Q., Sarló, T., Smola, A., 2013. Fastfood – approximating kernel expansions in loglinear time. *International Conference on Machine Learning.*
- Lewis, P., Gómez-Dans, J., Kaminski, T., Settle, J., Quaife, T., Gobron, N., Styles, J., Berger, M., 2012. An earth observation land data assimilation system (EO-LDAS). *Remote Sens. Environ.* 120, 219–235.
- Liou, K.N., 2002. *An Introduction to Atmospheric Radiation.* Second, Academic Press, Hampton, USA.
- Longbotham, N., Pacifici, F., Baugh, B., CampsValls, G., 2014, June. Pre-launch assessment of worldview-3 information content. pp. 24–27 (Lausanne, Switzerland).
- Melgani, F., Bruzzone, L., 2004. Classification of hyperspectral remote sensing images with support vector machines. *IEEE Trans. Geosci. Rem. Sens.* 42 (8), 1778–1790.
- Muñoz-Marí, J., Plaza, A., Gualtieri, J., Camps-Valls, G., 2009. Parallel programming and applications in grid, P2P and networking systems. In: Xhafa, F. (Ed.), *Parallel Implementation of SVM in Earth Observation Applications.* IOS Press, UK.
- O'Hagan, A., Kingman, J.F.C., 1978. Curve fitting and optimal design for prediction. *J. R. Stat. Soc. Ser. B Methodol.* 40 (1), 1–42.
- Pacifici, F., Chini, M., Emery, W., 2009. A neural network approach using multi-scale textural metrics from very high-resolution panchromatic imagery for urban land-use classification. *Remote Sens. Environ.* 113 (6), 1276–1292.
- Pasolli, L., Melgani, F., Blanzieri, E., 2010. Gaussian process regression for estimating chlorophyll concentration in subsurface waters from remote sensing data. *IEEE Geosci. Rem. Sens. Lett.* 464–468.
- Plaza, A., Benediktsson, J.A., Boardman, J., Brazile, J., Bruzzone, L., Camps-Valls, G., Chanussot, J., Fauvel, M., Gamba, P., Gualtieri, A., Tilton, J., 2009. Recent advances in techniques for hyperspectral image processing. *Remote Sens. Environ.* 113, 110–122.
- Plaza, J., Pérez, R., Plaza, A., Martínez, P., Valencia, D., 2008. Parallel morphological/neural processing of hyperspectral images using heterogeneous and homogeneous platforms. *Cluster Comput.* 11, 17–32.
- Quaife, T., Lewis, P., De Kauwe, M., Williams, M., Law, B.E., Disney, M., Bowyer, P., 2008. Assimilating canopy reflectance data into an ecosystem model with an ensemble kalman filter. *Remote Sens. Environ.* 112 (4), 1347–1364.
- Rahimi, A., Recht, B., 2007. Random features for large-scale kernel machines. *Neural Information Processing Systems.*
- Rahimi, A., Recht, B., 2009. Weighted sums of random kitchen sinks: replacing minimization with randomization in learning. In: Koller, D., Schuurmans, D., Bengio, Y., Bottou, L. (Eds.), *Advances in Neural Information Processing Systems 21.* Curran Associates, Inc., pp. 1313–1320.
- Rakotomamonjy, A., Bach, F., Canu, S., Grandvalet, Y., 2008, nov. SimpleMKL. *J. Mach. Learn. Res.* 9, 2491–2521.
- Rasmussen, C.E., Williams, C.K.I., 2006. *Gaussian Processes for Machine Learning.* The MIT Press, New York.
- Reed, M., Simon, B., 1981, Jan. I: Functional Analysis, Volume 1 (Methods of Modern Mathematical Physics) (Vol 1). First, Academic Press.
- Roberts, D., Quattrochi, D., Hulley, G., Hook, S., Green, R., 2012. Synergies between VSWIR and TIR data for the urban environment: an evaluation of the potential for the hyperspectral infrared imager (hyspIRI) decadal survey mission. *Remote Sens. Environ.* 117, 83–101.
- Rudin, W., 1987. *Real and Complex Analysis.* Third, McGraw-Hill Book Co., New York.
- Sacks, J., Welch, W.J., Mitchell, T.J., Wynn, H.P., 1989. Design and analysis of computer experiments. *Stat. Sci.* 409–423.
- Schaepman, M., Ustin, S., Plaza, A., Painter, T., Verrelst, J., Liang, S., 2009. Earth system science related imaging spectroscopy—an assessment. *Rem. Sens. Environ.* 113 (1), S123–S137.
- Schölkopf, B., Smola, A., 2002. *Learning With Kernels - Support Vector Machines, Regularization, Optimization and Beyond.* MIT Press Series.
- Shawe-Taylor, J., Cristianini, N., 2004. *Kernel Methods for Pattern Analysis.* Cambridge University Press.
- Siméoni, D., Singer, C., Chalon, G., 1997. Infrared atmospheric sounding interferometer. *Acta Astronaut.* 40, 113–118.
- Stuffer, T., Kaufmann, C., Hofer, S., Farster, K., Schreier, G., Mueller, A., Eckardt, A., Bach, H., Penné, B., Benz, U., Haydn, R., 2007. The enMAP hyperspectral imager—an advanced optical payload for future applications in earth observation programmes. *Acta Astronaut.* 61 (1–6), 115–120.
- Sutherland, J., Schneider, J., 2015. On the error of random Fourier features, pp. 862871.
- Suykens, J., Vandewalle, J., 1999. Least squares support vector machine classifiers. *Neural. Process. Lett.* 9 (3), 293–300.
- Tournier, B., Blumstein, D., Cayla, F., Chalon, G., 2002. IASI level 0 and 1 processing algorithms description. *Proc. of ISTCXXII conference.*
- Tuia, D., Pacifici, F., Kanevski, M., Emery, W., 2009, Nov. Classification of very high spatial resolution imagery using mathematical morphology and support vector machines. *IEEE Trans. Geosci. Rem. Sens.* 47 (11), 3866–3879.
- Verrelst, J., Alonso, L., Rivera Caicedo, J., Moreno, J., Camps-Valls, G., 2013. Gaussian process retrieval of chlorophyll content from imaging spectroscopy data. *IEEE J. Sel. Topics Appl. Earth Observations Rem. Sens.* 6 (2), 867–874.
- Verrelst, J., Muñoz, J., Alonso, L., Delegido, J., Rivera, J., Camps-Valls, G., Moreno, J., 2012. Machine learning regression algorithms for biophysical parameter retrieval: opportunities for Sentinel-2 and -3. *Remote Sens. Environ.* 118, 127–139.
- Verrelst, J., Rivera, J., Moreno, J., Camps-Valls, G., 2013. Gaussian processes uncertainty estimates in experimental Sentinel-2 LAI and leaf chlorophyll content retrieval. *ISPRS J. Photogramm. Remote. Sens.* 86, 157–167.
- Weinberger, K., Saul, L., 2008. Fast solvers and efficient implementations for distance metric learning. *Proceedings of the 25th International Conference on Machine Learning.* pp. 1160–1167.
- Weinberger, K., Tesau, G., 2007. Metric learning for kernel regression. *International Conference on Artificial Intelligence and Statistics.* pp. 608–615.
- Yang, F., White, M., Michaelis, A., Ichii, K., Hashimoto, H., Votava, P., Zhu, A.-X., Nemani, R., 2006, Nov. Prediction of continental-scale evapotranspiration by combining MODIS and AmeriFlux data through support vector machine. *IEEE Trans. Geosci. Rem. Sens.* 44 (11), 3452–3461.
- Zhang, Y., Duchi, J.C., Wainwright, M.J., 2013. Divide and conquer kernel ridge regression. *COLT.* pp. 592–617.