

Remote Sensing Image Classification With Large-Scale Gaussian Processes

Pablo Morales-Álvarez^{1b}, Adrián Pérez-Suay, *Member, IEEE*, Rafael Molina, *Senior Member, IEEE*,
and Gustau Camps-Valls^{1b}, *Senior Member, IEEE*

Abstract—Current remote sensing image classification problems have to deal with an unprecedented amount of heterogeneous and complex data sources. Upcoming missions will soon provide large data streams that will make land cover/use classification difficult. Machine-learning classifiers can help at this, and many methods are currently available. A popular kernel classifier is the Gaussian process classifier (GPC), since it approaches the classification problem with a solid probabilistic treatment, thus yielding confidence intervals for the predictions as well as very competitive results to the state-of-the-art neural networks and support vector machines. However, its computational cost is prohibitive for large-scale applications, and constitutes the main obstacle precluding wide adoption. This paper tackles this problem by introducing two novel efficient methodologies for GP classification. We first include the standard random Fourier features approximation into GPC, which largely decreases its computational cost and permits large-scale remote sensing image classification. In addition, we propose a model which avoids randomly sampling a number of Fourier frequencies and alternatively *learns* the optimal ones within a variational Bayes approach. The performance of the proposed methods is illustrated in complex problems of cloud detection from multispectral imagery and infrared sounding data. Excellent empirical results support the proposal in both computational cost and accuracy.

Index Terms—Cloud detection, Gaussian process classification (GPC), IAVISA, Infrared Atmospheric Sounding Interferometer (IASI)/Advanced Very High Resolution Radiometer (AVHRR), random Fourier features (RFFs), Spinning Enhanced Visible and Infrared Imager/Meteosat Second Generation (SEVIRI/MSG), variational inference.

I. INTRODUCTION

“...Nature almost surely operates by combining chance with necessity, randomness with determinism...”

—Eric Chaisson, *Epic of Evolution: Seven Ages of the Cosmos*

Manuscript received June 9, 2017; revised August 24, 2017; accepted September 23, 2017. Date of publication October 24, 2017; date of current version January 26, 2018. This work was supported in part by the European Research Council through the ERC-CoG-2014 SEDAL Project under Grant 647423, in part by the Spanish Ministry of Economy and Competitiveness under Project TIN2013-43880-R, Project TIN2015-64210-R, Project DPI2016-77869-C2-2-R, and Project TEC2016-77741-R, in part by EUMETSAT under Contract EUM/RSP/SOW/14/762293, and in part by the Spanish Excellence Network TEC2016-81900-REDT. The work of P. Morales-Álvarez was supported by the *La Caixa* Foundation. (Corresponding author: Pablo Morales-Álvarez.)

P. Morales-Álvarez and R. Molina are with the Department of Computer Science and Artificial Intelligence, University of Granada, 18010 Granada, Spain (e-mail: pabloramora@decsai.ugr.es; rms@decsai.ugr.es).

A. Pérez-Suay and G. Camps-Valls are with the Image Processing Laboratory, Universitat de València, 46010 València, Spain (e-mail: adrian.perez@uv.es; gustau.camps@uv.es).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TGRS.2017.2758922

EARTH-observation (EO) satellites provide a unique source of information to address some of the challenges of the Earth system science [1]. Current EO applications for image classification have to deal with a huge amount of heterogeneous and complex data sources.

The superspectral Copernicus Sentinels [2], [3], as well as the planned EnMAP [4], HypSIRI [5], PRISMA [6], and FLEX [7], will soon provide unprecedented data streams to be analyzed. Very high resolution (VHR) sensors like Quickbird, Worldview-2, and the recent Worldview-3 [8] 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) [9] sensor on board the MetOp satellite series, impose even larger constraints: the orbital period of Metop satellites (101 min), the large spectral resolution (8461 spectral channels between 645 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. The IASI mission delivers approximately 1.3×10^6 spectra per day, which gives a rate of about 29 GB/day to be processed. EO radar images also increased in resolution, and 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 archives¹ already collected by several past missions. In addition, we should also be prepared for the near future in diversity and complementarity of sensors.² These large-scale data problems require enhanced processing techniques that should be accurate, robust, and fast. Standard classification algorithms cannot cope with this new scenario efficiently.

In the last decade, kernel methods have dominated the field of remote sensing image classification [10], [11]. In particular, a kernel method called support vector machine (SVM) [12]–[16] 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 [17]–[21], the richness of hyperspectral imagery [16], [22], and exploited the power of clusters of

¹The Earth Observing System Data and Information System, for example, is managing around 4 TB daily, and the flow of data to users is about 20 TB daily.

²Follow the links for an up-to-date list of current ESA, EUMETSAT, JAXA, CNSA, and NASA EO missions.

computers [23], [24]. Undoubtedly, kernel methods have been the most widely studied classifiers, and became the preferred choice for users and practitioners. However, they are still not widely adopted in real practice because of the high computational cost when dealing with large-scale problems. Roughly speaking, given n examples available for training, kernel machines need to store kernel matrices of size $n \times n$, and to process them using standard linear algebra tools (matrix inversion, factorization, eigendecomposition, and so on) that typically scale cubically, $\mathcal{O}(n^3)$. This is an important constraint that hampers their applicability to large-scale EO data processing.

An alternative kernel classifier to SVM is the Gaussian process classifier (GPC) [25]. GPC has appealing theoretical properties, as it approaches the classification problem with a solid probabilistic treatment, and a very good performance in practice. The GPC method was originally introduced in the field of remote sensing in [26], where very good capabilities for land cover classification from multi/hyperspectral imagery were illustrated. Since then, GPC has been widely used in practice and extended to many settings: hyperspectral image classification [27], semantic annotation of high-resolution remote sensing images [28], change detection problems with semisupervised GPC [29], or classification of images with the help of user's intervention in active learning schemes [30], [31]. Unfortunately, like any other kernel method, its computational cost is very large. This is probably the reason why GPC has not yet been widely adopted by the geoscience and remote sensing community in large-scale classification scenarios, despite its powerful theoretical background and excellent performance in practice.

In GP for classification, we face two main problems. First, the nonconjugate observation model for classification (usually based on the sigmoid, probit, or Heaviside step functions) renders the calculation of the marginal distribution needed for inference impossible. The involved integrals are not analytically tractable, so one has to resort to numerical methods or approximations [25]. One could rely on Markov chain Monte Carlo methods, but they are computationally far too expensive. By assuming a Gaussian approximation to the posterior of the latent variables, one can use the Laplace approximation and the (more accurate) expectation propagation [32], [33]. The observation model can also be bounded, leading to the variational inference approach [34] that we use in this paper. Once the nonconjugacy of the observation model has been solved, the second problem is the inversion of huge matrices, which yields the unbearable $\mathcal{O}(n^3)$ complexity. Note that this is the only difficulty that appears when GP is used for regression, where the observation model can be analytically integrated out. This efficiency problem could be addressed with recent sparse GP approximations based on inducing points and approximate inference [35], but they come at the price of a huge number of parameters to estimate.

In this paper, we introduce two alternative pathways to perform large-scale remote sensing image classification with GPC. First, following the ideas in [36], we approximate the squared exponential (SE) kernel matrix of GPC by a linear one based on projections over a reduced set of random

Fourier features (RFFs). This novel method is referred to as RFF-GPC. It allows us to work in the primal space of features, which significantly reduces the computational cost of large-scale applications. In fact, a recent similar approach allows for using millions of examples in SVM-based land cover classification and regression problems [37]. The solid theoretical ground and the good empirical RFF-GPC performance make it a very useful method to tackle large-scale problems in Earth observation. However, RFF-GPC can only approximate (theoretically and in practice) a predefined kernel (the SE in this paper), and the approximation does not necessarily lead to a discriminative kernel. These shortcomings motivate our second methodological proposal: we introduce a novel approach to avoid *randomly sampling* a number of Fourier frequencies, and instead we propose *learning* the optimal ones within the variational approach. Therefore, Fourier frequencies are no longer *randomly sampled* and *fixed*, but *latent variables* estimated directly from data via variational inference. We refer to this method as variational Fourier features (VFFs)-GPC. The performance of RFF-GPC and VFF-GPC is illustrated in large- and medium-size real-world remote sensing image classification problems: 1) classification of clouds over landmarks from a long time series of Spinning Enhanced Visible and Infrared Imager/Meteosat Second Generation (SEVIRI/MSG) remote sensing images and 2) cloud detection using IASI and Advanced Very High Resolution Radiometer (AVHRR) infrared sounding data, respectively. Excellent empirical results support the proposed large-scale methods in both accuracy and computational efficiency. In particular, the extraordinary performance of VFF-GPC in the medium-size data set justifies its use not only as a large-scale method, but also as a general-purpose and scalable classification tool capable of learning an appropriate discriminative kernel.

The remainder of this paper is organized as follows. Section II reviews the RFF approximation and introduces it into GPC, deriving RFF-GPC and the more sophisticated VFF-GPC. Section III introduces the two real-world remote sensing data sets used for the experimental validation. Section IV presents the experimental results comparing the two proposed methods and standard GPC in terms of accuracy and efficiency. Section V concludes this paper with some remarks and future outlook.

II. LARGE-SCALE GAUSSIAN PROCESS CLASSIFICATION

GPs [38] are a probabilistic state-of-the-art model for regression and classification tasks. In the geostatistic community, GP for regression is usually referred to as *kriging*. For input–output data pairs $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$, a GP models the underlying dependence from a function-space perspective, i.e., introducing latent variables $\{f_i = f(\mathbf{x}_i) \in \mathbb{R}\}_{i=1}^n$ that jointly follow a normal distribution $\mathcal{N}(\mathbf{0}, \mathbf{K}) = (k(\mathbf{x}_i, \mathbf{x}_j))_{1 \leq i, j \leq n}$. The kernel function k encodes the sort of functions f favored, and \mathbf{K} is the so-called kernel matrix. The observation model of the output y given the latent variable f depends on the problem at hand. In binary classification (i.e., when $y \in \{0, 1\}$), the (nonconjugate) logistic observation

model is widely used. It is given by the sigmoid function as $p(y = 1|f) = \psi(f) = (1 + \exp(-f))^{-1} \in (0, 1)$.

A. Random Fourier Features

The main issue with large-scale applications of GP is its $\mathcal{O}(n^3)$ cost at the training phase, which comes from the $n \times n$ kernel matrix inversion. The work in [36] presents a general methodology (based on Bochner's theorem [39]) to approximate any positive-definite shift-invariant kernel k by a linear one. This is achieved by explicitly projecting the original d -dimensional data \mathbf{x} onto $\mathcal{O}(D)$ RFFs $\mathbf{z}(\mathbf{x})$, whose linear kernel k_L approximates k . This linearity will enable us to work in the primal space of features and substitute $n \times n$ matrix inversions by $\mathcal{O}(D) \times \mathcal{O}(D)$ ones, resulting in a total $\mathcal{O}(nD^2 + D^3)$ computational cost. In large-scale applications, one can set $D \ll n$, and thus the obtained $\mathcal{O}(nD^2)$ complexity represents an important benefit over the original $\mathcal{O}(n^3)$. Moreover, the complexity at test is also reduced from $\mathcal{O}(n^2)$ to $\mathcal{O}(D^2)$, even becoming independent on n .

In this paper, we use the well-known SE (or Gaussian) kernel $k(\mathbf{x}, \mathbf{x}') = \gamma \cdot \exp(-\|\mathbf{x} - \mathbf{x}'\|^2 / (2\sigma^2))$. Following the methodology in [36], this kernel can be linearly approximated as:

$$k(\mathbf{x}, \mathbf{x}') \approx k_L(\mathbf{x}, \mathbf{x}') = \gamma \cdot \mathbf{z}(\mathbf{x})^\top \mathbf{z}(\mathbf{x}') \quad (1)$$

where

$$\mathbf{z}(\mathbf{x})^\top = D^{-1/2} \cdot (\cos(\mathbf{w}_1^\top \mathbf{x}), \sin(\mathbf{w}_1^\top \mathbf{x}), \dots, \cos(\mathbf{w}_D^\top \mathbf{x}), \sin(\mathbf{w}_D^\top \mathbf{x})) \in \mathbb{R}^{2D} \quad (2)$$

and the *Fourier frequencies* \mathbf{w}_i must be sampled from a normal distribution $\mathcal{N}(\mathbf{0}, \sigma^{-2}\mathbf{I})$. As explained in [36, Claim 1], this approximation exponentially improves with the number D of Fourier frequencies used (and also exponentially worsens with d , the original dimension of \mathbf{x}). However, obviously, increasing D in our methods will go at the cost of increasing the $\mathcal{O}(nD^2)$ and $\mathcal{O}(D^2)$ complexities. Other kernels different from the SE one could be used, but that would imply sampling from a different distribution.

Our novel RFF-GPC method considers a standard Bayesian linear model over these new features \mathbf{z} . Such a linear model corresponds to GP classification with the linear kernel k_L [40, Ch. 6]. Since k_L approximates the SE kernel k , our RFF-GPC constitutes an approximation to GP classification with the SE kernel. However, RFF-GPC is well suited for large-scale applications, as it works in the primal space of $\mathcal{O}(D)$ features and thus presents a $\mathcal{O}(nD^2)$ [resp. $\mathcal{O}(D^2)$] train (resp. test) complexity.

Note that RFF-GPC needs to sample the Fourier frequencies \mathbf{w}_i from $\mathcal{N}(\mathbf{0}, \sigma^{-2}\mathbf{I})$ from the beginning, whereas hyperparameters σ and γ must be estimated during the learning process (just as in standard GP classification). In order to uncouple \mathbf{w}_i and σ , in the sequel, we consider the equivalent features

$$\mathbf{z}(\mathbf{x}|\sigma, \mathbf{W})^\top = D^{-1/2} \cdot (\cos(\sigma^{-1}\mathbf{w}_1^\top \mathbf{x}), \sin(\sigma^{-1}\mathbf{w}_1^\top \mathbf{x}), \dots, \cos(\sigma^{-1}\mathbf{w}_D^\top \mathbf{x}), \sin(\sigma^{-1}\mathbf{w}_D^\top \mathbf{x})) \quad (3)$$

with \mathbf{w}_i now sampled from $\mathcal{N}(\mathbf{0}, \mathbf{I})$ and fixed. Note that we have collectively denoted $\mathbf{W} = [\mathbf{w}_1, \dots, \mathbf{w}_D]^\top \in \mathbb{R}^{D \times d}$.

At this point, it is natural to consider other possibilities for the Fourier frequencies \mathbf{W} , rather than just randomly sample and fix them from the beginning. The proposed VFF-GPC model treats them as hyperparameters to be estimated (so as to maximize the likelihood of the observed data), just like σ and γ in the case of RFF-GPC. This makes VFF-GPC more expressive, flexible, and tailored to the data, although it may no longer constitute an approximation to the SE kernel (for which \mathbf{w}_i must be normally distributed). More specifically, VFF-GPC does start with an approximated SE kernel (since \mathbf{w}_i values are initialized with a normal distribution), but the maximum *a posteriori* (MAP) optimization on \mathbf{w}_i makes it *learn* a new kernel which may no longer approximate an SE one. Therefore, for VFF-GPC, we also use $\mathbf{z}(\mathbf{x}|\sigma, \mathbf{W})$ as in (3), with now both σ and \mathbf{W} to be estimated. Interestingly, VFF-GPC extends the sparse spectrum GP model originally introduced for regression in [41], to GP classification.

More specifically, the authors there also sparsify the SE kernel by working on the primal space of cos and sin Fourier features [41, eq. (5)]. However, our classification setting involves a sigmoid-based (logistic) observation model for the output given the latent variable [see (4)], whereas in regression, this is just given by a normal distribution. Therefore, VFF-GPC needs to additionally deal with the nonconjugacy of the sigmoid, which motivates the variational bound of (6) and the consequent variational inference procedure described in Section II-C. It is interesting to realize that VFF-GPC is introduced here as a natural extension of RFF-GPC, whereas there is not a regression analogous for RFF-GPC in [41].

Another possibility for the Fourier frequencies would be to estimate them (just as in VFF-GPC) but considering alternative prior distributions $p(\mathbf{W})$ (which means utilizing alternative kernels). Moreover, instead of MAP inference, we could address the marginalization of the Fourier frequencies \mathbf{W} . Alternatively, to promote sparsity, the use of Gaussian scale models [42] could also be investigated. These possibilities will be explored in the future work, and here we will concentrate on RFF-GPC and VFF-GPC.

B. Models Formulation

As anticipated in Section II-A, RFF-GPC and VFF-GPC are standard Bayesian linear models working on the explicitly mapped features $\mathbf{z}(\mathbf{x}|\sigma, \mathbf{W})$ of (3). In the case of RFF-GPC, \mathbf{W} is sampled from $\mathcal{N}(\mathbf{0}, \mathbf{I})$ at the beginning and fixed, with σ to be estimated. In the case of VFF-GPC, both \mathbf{W} and σ are estimated, with $\mathcal{N}(\mathbf{0}, \mathbf{I})$ prior over \mathbf{W} . In order to derive both methods in a unified way, Φ will denote σ for RFF-GPC and both (\mathbf{W}, σ) for VFF-GPC.

Since we are dealing with binary classification, we consider the standard logistic observation model

$$p(y = 1|\boldsymbol{\beta}, \Phi, \mathbf{x}) = \psi(\boldsymbol{\beta}^\top \mathbf{z}) = (1 + \exp(-\boldsymbol{\beta}^\top \mathbf{z}))^{-1} \quad (4)$$

where $\mathbf{z} = \mathbf{z}(\mathbf{x}|\Phi)$. For the weights $\boldsymbol{\beta} \in \mathbb{R}^{2D}$, we utilize the prior normal distribution $p(\boldsymbol{\beta}|\gamma) = \mathcal{N}(\boldsymbol{\beta}|\mathbf{0}, \gamma\mathbf{I})$, with γ to be estimated [see (1)].

For an observed data set $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n \subset \mathbb{R}^d \times \{0, 1\}$, the joint pdf reads

$$\begin{aligned} p(\mathbf{y}, \boldsymbol{\beta} | \Phi, \gamma, \mathbf{X}) &= p(\mathbf{y} | \boldsymbol{\beta}, \Phi, \mathbf{X}) p(\boldsymbol{\beta} | \gamma) \\ &= \left(\prod_{i=1}^n p(y_i | \boldsymbol{\beta}, \Phi, \mathbf{x}_i) \right) p(\boldsymbol{\beta} | \gamma) \end{aligned} \quad (5)$$

where we collectively denote $\mathbf{y} = (y_1, \dots, y_n)^\top$ and $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]^\top$. For the sake of brevity, from now on, we will systematically omit the conditioning on \mathbf{X} .

C. Variational Inference

Given the observed data set \mathcal{D} , in this section, we seek point estimates of γ and Φ by maximizing the marginal likelihood $p(\mathbf{y} | \Phi, \gamma)$ (in VFF-GPC, the additional prior $p(\mathbf{W}) = \mathcal{N}(\mathbf{W} | \mathbf{0}, \mathbf{I})$ yields MAP inference, instead of maximum likelihood one, for \mathbf{W}). After that, we obtain (an approximation to) the posterior distribution $p(\boldsymbol{\beta} | \mathbf{y}, \Phi, \gamma)$. Due to the non-conjugate observation model, the required integrals will be mathematically intractable, and we will resort to the variational inference approximation [40, Sec. 10.6].

First, note that integrating out $\boldsymbol{\beta}$ in (5) is not analytically possible due to the sigmoid functions in the observation model $p(\mathbf{y} | \boldsymbol{\beta}, \Phi)$. To overcome this problem, we use the variational bound

$$\log(1 + e^x) \leq \lambda(\xi)(x^2 - \xi^2) + \frac{x - \xi}{2} + \log(1 + e^\xi) \quad (6)$$

which is true for any real numbers x, ξ and where $\lambda(\xi) = (1/2\xi)(\psi(\xi) - 1/2)$ [40, Sec. 10.6]. Applying it to every factor of $p(\mathbf{y} | \boldsymbol{\beta}, \Phi)$, we have the following lower bound:

$$p(\mathbf{y} | \boldsymbol{\beta}, \Phi) \geq \exp(-\boldsymbol{\beta}^\top \mathbf{Z}^\top \boldsymbol{\Lambda} \mathbf{Z} \boldsymbol{\beta} + \mathbf{v}^\top \mathbf{Z} \boldsymbol{\beta}) \cdot C(\boldsymbol{\xi}). \quad (7)$$

Here, we write $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_n]^\top \in \mathbb{R}^{n \times 2D}$ for the projected-data matrix (which depends on Φ), $\boldsymbol{\Lambda}$ is the diagonal matrix $\text{diag}(\lambda(\xi_1), \dots, \lambda(\xi_n))$, $\mathbf{v} = \mathbf{y} - (1/2) \cdot \mathbf{1}_{n \times 1}$, and the term $C(\boldsymbol{\xi}) = \prod_i \exp(\lambda(\xi_i) \xi_i^2 + (1/2) \xi_i - \log(1 + \exp(\xi_i)))$ only depends on $\boldsymbol{\xi}$. The key is that this lower bound for $p(\mathbf{y} | \boldsymbol{\beta}, \Phi)$ is conjugate with the normal prior $p(\boldsymbol{\beta} | \gamma)$ (since it is the exponential of a quadratic function on $\boldsymbol{\beta}$), and thus it allows us integrating out $\boldsymbol{\beta}$ in (5). In exchange, we have introduced n additional hyperparameters $\boldsymbol{\xi} = (\xi_1, \dots, \xi_n)^\top$ that will need to be estimated along with Φ and γ .

Therefore, substituting $p(\mathbf{y} | \boldsymbol{\beta}, \Phi)$ for its bound, (5) can be lower bounded as

$$p(\mathbf{y}, \boldsymbol{\beta} | \Phi, \gamma) \geq F(\mathbf{y}, \Phi, \gamma, \boldsymbol{\xi}) \cdot \mathcal{N}(\boldsymbol{\beta} | \boldsymbol{\mu}, \boldsymbol{\Sigma}) \quad (8)$$

where we have denoted

$$\begin{aligned} F(\mathbf{y}, \Phi, \gamma, \boldsymbol{\xi}) &= C(\boldsymbol{\xi}) |\gamma^{-1} \boldsymbol{\Sigma}|^{1/2} \exp\left(\frac{1}{2} \boldsymbol{\mu}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}\right) \\ \boldsymbol{\Sigma} &= (\mathbf{Z}^\top (2\boldsymbol{\Lambda}) \mathbf{Z} + \gamma^{-1} \mathbf{I})^{-1}, \quad \boldsymbol{\mu} = \boldsymbol{\Sigma} \mathbf{Z}^\top \mathbf{v}. \end{aligned} \quad (9)$$

Now, it is clear that we can marginalize out $\boldsymbol{\beta}$, and iteratively estimate the optimal values of Φ , γ , and $\boldsymbol{\xi}$ as those that maximize $F(\mathbf{y}, \Phi, \gamma, \boldsymbol{\xi})$ [or equivalently $\log F(\mathbf{y}, \Phi, \gamma, \boldsymbol{\xi})$] for the observed \mathbf{y} . Starting at $\boldsymbol{\xi}^{(1)}$, $\Phi^{(1)}$, and $\gamma^{(1)}$, we can calculate $\boldsymbol{\xi}^{(k)}$, $\Phi^{(k)}$, and $\gamma^{(k)}$ for $k \geq 1$. In the case of $\boldsymbol{\xi}$, we make use of the local maximum condition $\partial F / \partial \boldsymbol{\xi} = 0$.

From there, it is not difficult to prove that the optimal value satisfies [40]

$$\boldsymbol{\xi}^{(k+1)} = \sqrt{\text{diag}(\mathbf{Z}^{(k)} \boldsymbol{\Sigma}^{(k)} (\mathbf{Z}^{(k)})^\top) + (\mathbf{Z}^{(k)} \boldsymbol{\mu}^{(k)})^2} \quad (10)$$

where the square and square root of a vector are understood as elementwise. In the case of Φ and γ , we use nonlinear conjugate gradient methods [43] and obtain (note that for VFF-GPC, we can collapse \mathbf{W} and σ , removing the prior on \mathbf{W})

$$\begin{aligned} (\Phi^{(k+1)}, \gamma^{(k+1)}) &= \arg \max_{\Phi, \gamma} \{-\log |2\gamma \mathbf{Z}^\top \boldsymbol{\Lambda}^{(k+1)} \mathbf{Z} + \mathbf{I}| \\ &\quad + \mathbf{v}^\top \mathbf{Z} (2\mathbf{Z}^\top \boldsymbol{\Lambda}^{(k+1)} \mathbf{Z} + \gamma^{-1} \mathbf{I})^{-1} \mathbf{Z}^\top \mathbf{v}\}. \end{aligned} \quad (11)$$

Once the hyperparameters Φ , γ , and $\boldsymbol{\xi}$ have been estimated by $\hat{\Phi}$, $\hat{\gamma}$, and $\hat{\boldsymbol{\xi}}$, respectively, we need to compute the posterior $p(\boldsymbol{\beta} | \mathbf{y}, \hat{\Phi}, \hat{\gamma})$. As before, this is mathematically intractable due to the sigmoids in the observation model. Therefore, we again resort to the variational bound in (8) to get an optimal approximation $\hat{p}(\boldsymbol{\beta})$ to the posterior $p(\boldsymbol{\beta} | \mathbf{y}, \hat{\Phi}, \hat{\gamma})$. Namely, we do it by minimizing (an upper bound of) the Kullback-Leibler (KL) divergence between both distributions

$$\begin{aligned} &\text{KL}(\hat{p}(\boldsymbol{\beta}) || p(\boldsymbol{\beta} | \mathbf{y}, \hat{\Phi}, \hat{\gamma})) \\ &= \int \hat{p}(\boldsymbol{\beta}) \log \frac{\hat{p}(\boldsymbol{\beta})}{p(\boldsymbol{\beta} | \mathbf{y}, \hat{\Phi}, \hat{\gamma})} d\boldsymbol{\beta} \\ &= \log p(\mathbf{y} | \hat{\Phi}, \hat{\gamma}) + \int \hat{p}(\boldsymbol{\beta}) \log \frac{\hat{p}(\boldsymbol{\beta})}{p(\mathbf{y}, \boldsymbol{\beta} | \hat{\Phi}, \hat{\gamma})} d\boldsymbol{\beta} \\ &\leq \log p(\mathbf{y} | \hat{\Phi}, \hat{\gamma}) - \log F(\mathbf{y}, \hat{\Phi}, \hat{\gamma}, \hat{\boldsymbol{\xi}}) + \text{KL}(\hat{p}(\boldsymbol{\beta}) || \mathcal{N}(\boldsymbol{\beta} | \hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\Sigma}})). \end{aligned}$$

Thus, the minimum is reached for $\hat{p}(\boldsymbol{\beta}) = \mathcal{N}(\boldsymbol{\beta} | \hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\Sigma}})$, with $\hat{\boldsymbol{\mu}}$ and $\hat{\boldsymbol{\Sigma}}$ calculated in (9) using $\hat{\Phi}$, $\hat{\gamma}$, and $\hat{\boldsymbol{\xi}}$.

In summary, at training time, our methods RFF-GPC and VFF-GPC run iteratively until convergence of the hyperparameters Φ , γ , and $\boldsymbol{\xi}$ to their optimal values $\hat{\Phi}$, $\hat{\gamma}$, and $\hat{\boldsymbol{\xi}}$ (see Algorithm 1). The computations involved there suppose a computational complexity of $\mathcal{O}(nD^2 + D^3)$ [which equals $\mathcal{O}(nD^2)$ when $n \gg D$], whereas standard GPC scales as $\mathcal{O}(n^3)$. At test time, the probability of class 1 for a previously unseen instance $\mathbf{x}_* \in \mathbb{R}^d$ is

$$\begin{aligned} p(y_* = 1) &\approx \int p(y = 1 | \boldsymbol{\beta}, \hat{\Phi}, \mathbf{x}_*) \hat{p}(\boldsymbol{\beta}) d\boldsymbol{\beta} \\ &\approx \psi(\hat{\mathbf{z}}_*^\top \hat{\boldsymbol{\mu}} \cdot (1 + (\pi/8) \hat{\mathbf{z}}_*^\top \hat{\boldsymbol{\Sigma}} \hat{\mathbf{z}}_*)^{-1/2}) \end{aligned} \quad (12)$$

with ψ being the sigmoid function. Whereas GPC presents a computational cost of $\mathcal{O}(n^2)$ for each test instance, (12) implies a complexity of $\mathcal{O}(D^2)$ in the case of our methods. In particular, note that this is independent on the number of training instances. These significant reductions in computational cost (both at training and test) make our proposal suitable for large-scale and real-time applications in general, and in EO applications in particular.

Finally, regarding the convergence of the proposed methods, we cannot theoretically guarantee the convergence to a global optimum (only a local one), since we are using conjugate gradient methods to solve the nonconvex optimization problem in (11). However, from a practical viewpoint, we have

Algorithm 1 Training RFF-GPC and VFF-GPC

Require: Data set $\mathcal{D} = \{(\mathbf{x}_i, y_i)\}_{i=1}^n \subset \mathbb{R}^d \times \{0, 1\}$ and the number D of Fourier frequencies.

Only for RFF-GPC, sample the Fourier frequencies \mathbf{w}_i from $\mathcal{N}(\mathbf{0}, \mathbf{I})$ and fix them.

Initialize $\xi^{(1)} = \mathbf{1}_{n \times 1}$, $\gamma^{(1)} = 1$, and $\Phi^{(1)}$. For RFF-GPC (where $\Phi = \sigma$), $\Phi^{(1)}$ is initialized as the mean distance between (a subset of) the training data points \mathbf{x}_i . For VFF-GPC (where $\Phi = (\sigma, \mathbf{W})$), σ is initialized as described for RFF-GPC and \mathbf{W} with a random sample from $\mathcal{N}(\mathbf{0}, \mathbf{I})$.

repeat

 Update $\xi^{(k+1)}$ with eq. (10).

 Update $\Phi^{(k+1)}$ and $\gamma^{(k+1)}$ with eq. (11), using $\Phi^{(k)}$ and $\gamma^{(k)}$ as initial values for the conjugate gradient method.

until convergence

return Optimal hyperparameter $\hat{\Phi}$ and the posterior distribution $\hat{p}(\beta) = \mathcal{N}(\beta | \hat{\mu}, \hat{\Sigma})$.

experimentally checked that both methods have a satisfactory similar convergence pattern. Namely, in the first iterations, the hyperparameters experiment more pronounced changes, widely exploring the hyperparameters space. Then, once they reach a local optimum vicinity, these variations become smaller. Eventually, the hyperparameters' values hardly change and the stop criterion is satisfied.

III. DATA COLLECTION AND PREPROCESSING

This section introduces the data sets used for comparison purposes in the experiments. We considered: 1) a continuous year of MSG data involving several hundred thousands of labeled pixels for cloud classification and 2) a medium-size manually labeled data set used to create the operational IASI cloud mask.

A. Cloud Detection Over Landmarks With SEVIRI/MSG

We focus on the problem of cloud identification over landmarks using SEVIRI 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 models and geometric quality assessment in Level 1 instrument processing chain. Detection of clouds over landmarks is an essential step in the MSG processing chain, as undetected clouds are one of the most significant sources of error in landmark matching (see Fig. 1).

The data set used in the experiments was provided by EUMETSAT, and contains SEVIRI/MSG Level 1.5 acquisitions for 200 landmarks of variable size for a whole year (2010). Landmarks mainly cover coastlines, islands, or inland waters. We selected all multispectral images from a particular landmark location, Dakhla (Western Sahara), which involves 35 040 MSG acquisitions with a fixed resolution of 20×26 pixels. In addition, Level 2 cloud products were provided for each landmark observation, so Level 2 cloud mask [44] is used as the best available “ground truth”




Fig. 1. Landmarks are essential in image registration and geometric quality assessment. Misclassification of cloud contamination in landmarks degrades the correlation matching, which is a cornerstone for the image navigation and registration algorithms.

to validate the results. We framed the problem for this particular landmark as a pixelwise classification one.

A total amount of 16 features were extracted from the images, involving band ratios, spatial, contextual, and morphological features, and discriminative cloud detection scores. In particular, we considered: seven channels converted to top of atmosphere reflectance (R1, R2, R3, and R4) and brightness temperature (BT7, BT9, and BT10), three band ratios, and six spatial features. On the one hand, the three informative band ratios were: 1) a cloud detection ratio, $R_{0.8 \mu\text{m}}/R_{0.6 \mu\text{m}}$; 2) 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 3) the NDVI $(R_{0.8 \mu\text{m}} - R_{0.6 \mu\text{m}})/(R_{0.8 \mu\text{m}} + R_{0.6 \mu\text{m}})$. On the other hand, the six spatial features were obtained by applying average filters of sizes 3×3 and 5×5 , as well as a standard deviation filter of size 3×3 , on both bands R1 and BT9.

Based on previous studies [44], [45], and in order to simplify the classification task, the different illumination conditions (and hence difficulty) over the landmarks are studied by splitting the day into four ranges (subproblems) according to the solar zenith angle (SZA) values: *high* ($SZA < SZA_{\text{median}}$), *mid* ($SZA_{\text{median}} < SZA < 80^\circ$), *low* ($80^\circ < SZA < 90^\circ$), and *night* ($SZA > 90^\circ$). Therefore, different classifiers are developed for each SZA range.

The final amount of pixels available for each illumination condition is $n = 1\,500\,000$ for *high*, *mid*, and *night*, and $n = 1\,365\,083$ for *low*. Moreover, each problem has different dimensionalities: all the $d = 16$ features were used for the three daylight problems, and $d = 6$ was used for the night one (some bands and ratios are meaningless at night).

B. Cloud Detection With the IASI/AVHRR Data

The IAVISA data set is part of the study “IASI/AVHRR Visual Scenes Analysis and Cloud Detection,”³ whose aim is to improve the IASI cloud detection by optimizing the coefficients used for a predefined set of cloud tests performed in the IASI Level 2 processing. The data set was derived by visual analysis of globally distributed data, and served as input for the optimization and validation of the IASI cloud detection. Each collected IASI sample was classified concerning its cloudiness based on the visual inspection of the AVHRR Level 1B inside the IASI footprint. Each sample classifies a single IASI instantaneous field of view as being cloud free (clear sky, 28%), partly cloudy low (26%), partly cloudy high (26%),

³<http://www.brockmann-consult.de/iavisa-info-web/>

TABLE I
SAMPLES DISTRIBUTION PER SURFACE TYPES

ID	Surface Type	Samples
1	Evergreen Needle Forrest	603
2	Evergreen Broad Forrest	876
3	Deciduous Needle Forrest	89
4	Deciduous Broad Forrest	324
5	Mixed Forest	602
6	Closed Shrubs	329
7	Open Shrubs	1484
8	Woody Savannas	768
9	Savannas	656
10	Grassland	886
11	Wetlands	147
12	Crops	1294
13	Urban	44
14	Crop/Mosaic	1436
15	Snow/Ice	1443
16	Barren/Desert	1379
17	Water	12150
18	Tundra	413

or cloudy (20%). For the sake of simplicity, here we focus on discriminating between cloudy and cloud-free pixels.

In order to ensure the representativeness of the data set for the natural variability of clouds, labeling further considered additional conditions depending on: 1) the surface type (see Table I); 2) the climate zone (Köppen classification over land and geographical bands over sea); 3) the season; and 4) day/night discrimination. First, the surface type database used as ancillary information was the International Geosphere-Biosphere Program scene types in the Clouds and the Earth's Radiant Energy System/Surface and Atmospheric Radiation Budget surface map. The 18-class map was used to identify surface properties of a given region. The distribution of the surface types in the map is given in Table I, showing the even distribution of clouds across land cover types which ensures representativeness (natural variability) of the database. Second, the different climate zones were sampled as follows: tropical ($n = 7499$), dry ($n = 3237$), temperate ($n = 8150$), cold ($n = 2205$), and polar zones ($n = 3832$). Third, seasonality was also taken into account, and yielded the following distribution: spring ($n = 5862$), summer ($n = 6930$), autumn ($n = 6662$), and winter ($n = 5469$). The global sampling and some cloudy and cloud-free chips are shown in Fig. 2. The final database consists of $n = 24923$ instances and $d = 8461$ original features, which have been summarized to $d = 100$ more informative directions through a standard principal component analysis.

IV. EXPERIMENTS

In this section, we empirically demonstrate the performance of the proposed methods in the two real problems described earlier. Moreover, we carry out an exhaustive comparison to GPC with the SE kernel (in the sequel, GPC for brevity).

In order to provide a deeper understanding of our methods, different values of D (a number of Fourier frequencies) will be used. Different sizes n of the training data set will also be considered, in order to analyze the scalability of the methods




Fig. 2. (Left) Global sample coverage for all seasons, times of day, and cloud cases. (Right) Examples of cloud-free and cloudy samples.

and to explore the tradeoff between accuracy and computational cost. When increasing n (respectively, D), we will add training instances (respectively, Fourier frequencies) to those already used for lower values of n (respectively, D). We provide numerical (in terms of recognition rate), computational (training and test times), and visual (by inspecting classification maps) assessments.

A. Cloud Detection in the Landmarks Data Set

In this large-scale problem, the number of training examples is selected as $n \in \{10000, 50000, 100000, 300000\}$ for RFF-GPC and VFF-GPC, and $n \in \{5000, 10000, 15000\}$ for GPC. Note that the improvement at training computational cost [$\mathcal{O}(nD^2)$ for the proposed methods against $\mathcal{O}(n^3)$ for GPC] enables us to consider much greater training data sets for our methods. In fact, as we will see in Fig. 3, even with $n = 300000$, RFF-GPC and VFF-GPC are computationally cheaper than GPC with $n = 15000$. Indeed, higher values of n are not considered for GPC to avoid exceeding the (already expensive) 10^6 s of training CPU time needed with just $n = 15000$. Regarding the number D of Fourier frequencies, we use $D \in \{10, 25, 50, 100, 150, 200\}$.

The experimental results, which include predictive performance [test overall accuracy (OA)], training CPU time, and test CPU time, are shown in Fig. 3. Every single value is the average of five independent runs under the same setting. Namely, for each illumination condition, a test data set is fixed and five different balanced training data sets are defined with the remaining data. Note that a general first observation across Fig. 3 suggests that higher accuracies are obtained for higher illumination conditions (SZA).

Fig. 3 reveals an overwhelming superiority of RFF-GPC and VFF-GPC over standard GPC: our proposed methods achieve a higher predictive performance while investing substantially lower training CPU time. This is very clear from the third column plots, where for any blue point we can find orange and yellow points which are placed more north-west (i.e., higher test OA and lower CPU training time). Furthermore, the fourth column shows an equally extraordinary reduction in test CPU time (production time), where the proposed methods are more than 100 times faster than GPC. In particular, this makes RFF-GPC and VFF-GPC better suited than standard GPC for real-time EO applications.

Regarding the practical differences between RFF-GPC and VFF-GPC, we observe that RFF-GPC is faster (at training) whereas VFF-GPC is more accurate. This is a natural consequence of their theoretical formulations: the estimation of the Fourier frequencies \mathbf{W} in VFF-GPC makes it more flexible and

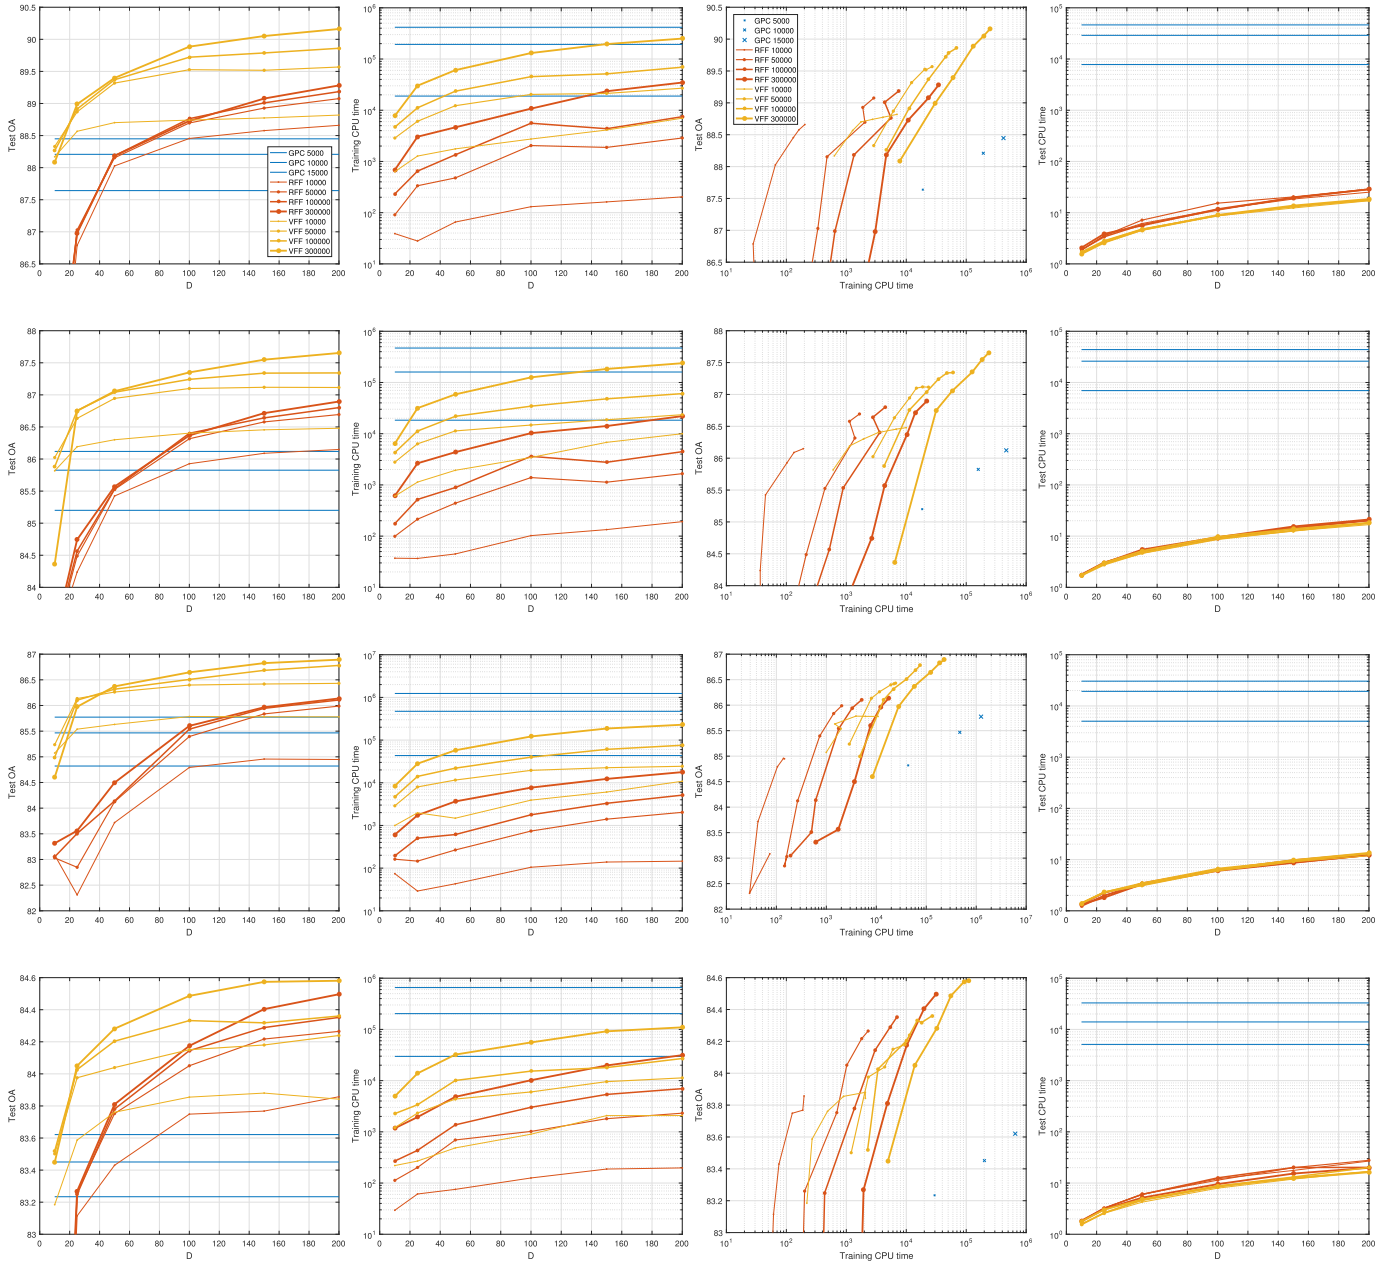


Fig. 3. Experimental results for Landmarks data set. (From top to bottom) Rows correspond with the described *high*, *mid*, *low*, and *night* illumination conditions. The first column shows the test OA of RFF-GPC, VFF-GPC, and GPC for the different values of n (the number of training examples) and D (the number of Fourier frequencies) considered. The second column is analogous, but displays the CPU time (in seconds) needed to train each method (instead of the test OA). The third column summarizes the two previous ones, providing a tradeoff between test OA and training CPU time. The last column is analogous to the first and second ones, but showing the CPU time used at the test step (production time). The legend for the second and fourth columns is the same as in the first one. However, note that in the third column plots, the GPC lines degenerate into single points (since GPC does not depend on D). In both legends, the numbers indicate the amount n of training examples used, which determines the width/size of the lines/points too. As further explained in the main text, shown results are the mean over five independent runs.

expressive, but involves a heavier training. Therefore, in this particular problem, the final practical choice between the two proposed methods would depend on the relative importance that the user assigns to test accuracy (where VFF-GPC stands out) and training cost (where RFF-GPC does so). In terms of test cost, both methods are very similar, as expected from the identical $\mathcal{O}(D^2)$ theoretical test complexity. The independence of this quantity on n is also intuitively reflected in the experiments, with all the RFF-GPC and VFF-GPC

lines collapsing onto a single one in the fourth column plots of Fig. 3.

At this point, it is worth to analyze a bit further the role of D in the performance of our methods. Recall (Section II-A) that RFF-GPC is an approximation to GPC, with an error that exponentially decreases with the ratio D/d between the dimensions of the projected Fourier features space and the original one. Therefore, it is theoretically expected that the performance of RFF-GPC increases with D , becoming equivalent

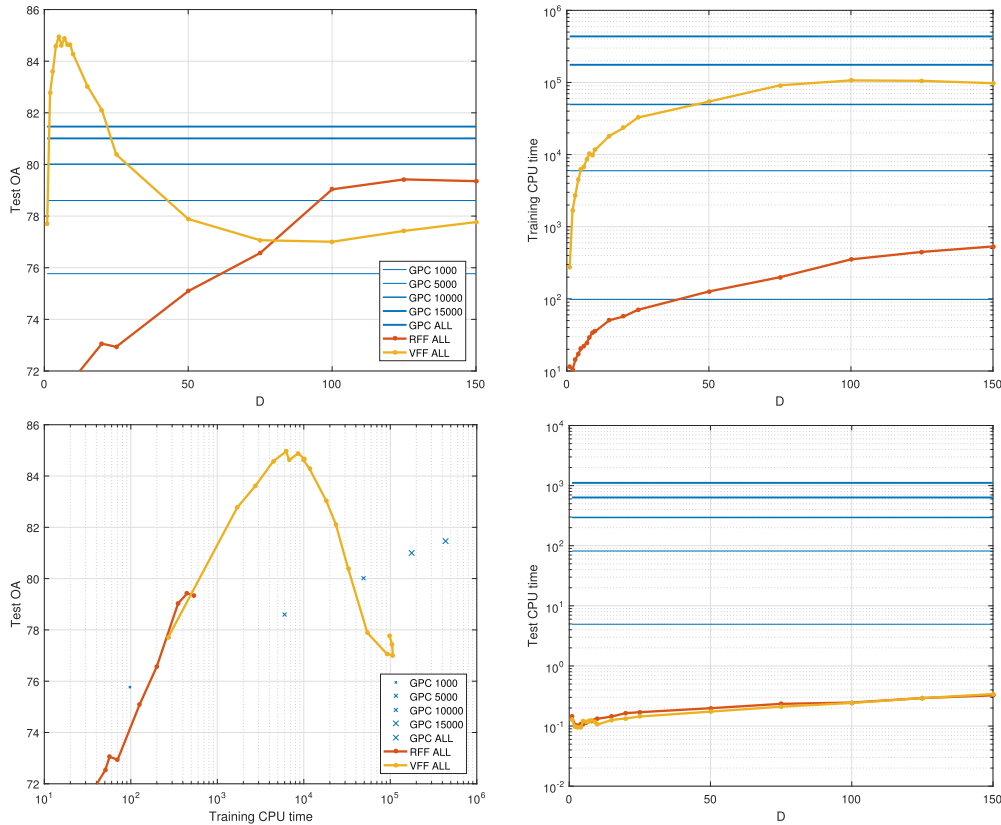


Fig. 4. Experimental results for the IAVISA data set. From left to right and top to bottom, The first plot shows the test OA of RFF-GPC, VFF-GPC, and GPC for the different values of n (number of training examples) and D (number of Fourier frequencies) considered. The second plot is analogous, but displays the CPU time needed to train each method (instead of the test OA). The third plot summarizes the two previous ones, providing a tradeoff between test OA and training CPU time. The last plot is analogous to the first and second ones, but showing the CPU time used at the test step. The legend for second and fourth plots is the same as the one in the first plot. However, in the third plot the GPC lines degenerate into single points (since GPC does not depend on D). In both legends, the numbers indicate the amount n of training examples used, which determines the width/size of the lines/points too (ALL means the whole training data set, i.e., $n \approx 20\,000$). As explained in the main text, the results are the mean over five independent runs.

to GPC when $D \rightarrow \infty$. Actually, this is supported by the first column of Fig. 3. Moreover, since our problem here presents a low d (16 for *high*, *mid*, and *low*, and 6 for *night*), it is natural that RFF-GPC with just $D = 200$ already gets very similar (even better in some cases) results to standard GPC with the same n (yet far much faster, compare GPC and RFF/VFF for $n = 10\,000$). In the case of VFF-GPC, where the Fourier frequencies are model parameters to be estimated, the number D is directly related to the complexity of the model. Therefore, its increase should not always mean a higher performance in test OA, since large values may provoke overfitting to the training data set (this will be clear in the next data set, whereas it does not occur in Landmarks). Furthermore, unlike RFF-GPC, the performance of VFF-GPC is not directly affected by d .

It is also reasonable to expect that both test OA and training CPU time increase with the training data set size n . More specifically, and from a practical perspective in which the computational resources are finite, the first column in Fig. 3 shows that test OA becomes stalled when only one of n or D increases. However, greater improvements in test OA are achieved when n and D are jointly increased. Note that this is also justifiable from a theoretical viewpoint: the higher the dimensionality of the projected Fourier features

space (which is $2\text{-}D$), the larger number n of examples are required to identify the separation between classes.

B. Cloud Detection With the IAVISA Data Set

As explained in Section III-B, this problem involves a total amount of $n = 24\,923$ instances. We performed fivefold cross validation, which produces five pairs of training/test data sets with (approximately) $20\,000/5\,000$ instances each. Results are then averaged. Since RFF-GPC and VFF-GPC are conceived for large-scale applications [they scale linearly with n , recall their $\mathcal{O}(nD^2)$ training cost], they will not be utilized with the values of n lower than this training data set size of $\approx 20\,000$ (even GPC is able to cope with this size). Indeed, in the case of GPC, we use the values $n \in \{1000, 5000, 10000, 15000, \text{ALL} \approx 20000\}$. Regarding the number of Fourier frequencies D , we consider the grid $D \in \{1 : 10, 15 : 5 : 25, 50 : 25 : 150\}$. The experimental results, which include the same metrics as those used for the previous problem on landmarks, are shown in Fig. 4.

In this case, we again observe a clear outperformance of VFF-GPC against GPC: it achieves higher test OA while requiring less training and test CPU times. Moreover, the improvement in test OA is greater than 3%, and train/test CPU times are around 100 and 1000 times lower, respectively.

However, unlike in the previous problem of cloud detection over landmarks, RFF-GPC does not exhibit such a clear superiority over GPC in this application. Whereas it does drastically decrease the train/test CPU times, it is not able to reach the test OA of GPC with $n = 10\,000$. Therefore, in practice, the optimal choice for this application is VFF-GPC. RFF-GPC would only be recommended if the training CPU time is a very strong limitation.

The main reason why RFF-GPC is not completely competitive in this problem is its theoretical scope: as an efficient approximation to GPC, it is conceived for large-scale applications which are out of the reach of standard GPC. If the size of the problem allows for using GPC (as in this case), then RFF-GPC will only provide a more efficient alternative (less training and test CPU times), but its predictive performance will be always below that of GPC. Moreover, the difference in this performance is directly influenced by the original dimension d of data [recall that the kernel approximation behind RFF-GPC exponentially degrades with d (see Section II-A)]. This is precisely a second hurdle that RFF-GPC finds in IAVISA: the high $d = 100$ makes RFF-GPC with the full data set be quite far from the corresponding GPC at predictive performance (test OA). In conclusion, the ideal setting for RFF-GPC is a large-scale problem (high n) with a few features (low d), precisely the opposite to the IAVISA data set.

Interestingly, VFF-GPC bypasses these limitations of RFF-GPC by *learning a new kernel and not just approximating* the SE one. First, VFF-GPC is not just a GP adaptation well suited for large-scale applications, but a general-purpose, expressive, and very competitive kernel-based classifier that scales well with the number of training instances. Second, as it does not rely on the kernel approximation, VFF-GPC is not affected by the original dimension d of data. Both ideas are empirically supported by the results obtained in IAVISA.

The first plot of Fig. 4 shows that the predictive performance of VFF-GPC does not necessarily improves by increasing D . This is the expected behavior from the theoretical formulation of VFF-GPC, where the Fourier frequencies are D parameters to be estimated. Thus, a higher amount of them confers VFF-GPC a greater flexibility to learn hidden patterns in the training data set, but also the possibility to overfit very particular structures of it which do not generalize to the test set. This is the classical problem of the *model complexity* in machine learning, and it is further illustrated in Fig. 5. Together with the first plot in Fig. 4, it shows the paradigmatic behavior of train and test performance in the presence of overfitting: train OA grows with the model complexity (great flexibility allows for learning very particular structures of the training set, even reaching a 100% of train OA), whereas test OA initially grows (the first patterns are part of the ground truth and thus general to the test set) but then goes down (when the learned information is too specific to the training set). Note that these overfitting phenomena did not occur at Landmarks, where test OA monotonically increased with D . In addition to the different nature of the problems, the training data set size n plays a crucial role at this: smaller data sets (like IAVISA) are




Fig. 5. Train OA in the IAVISA data set for RFF-GPC, VFF-GPC, and GPC with different values of n (the number of training examples) and D (the number of Fourier frequencies). These results complement the first plot in Fig. 4, showing that high values of D make VFF-GPC overfit to the training data set. The legend and its interpretation are the same as there.

more prone to overfitting than larger ones (Landmarks) under the same model complexity.

Finally, it is worth noting that VFF-GPC achieves its maximum test OA when using just $D = 5$ Fourier frequencies. This reflects: 1) a not very sophisticated internal structure of the IAVISA data set (since just five directions are enough to correctly classify 85% of the data) and 2) the VFF-GPC capability to learn those discriminative directions from data. In particular, this shows that VFF-GPC can be used not only as a classifier, but also as a method that learns the most relevant discriminative directions in a data set. Unfortunately, RFF-GPC is not able to benefit from these privileged directions that may exist in some data sets, since it randomly samples and fixes the Fourier frequencies from the beginning.

C. Explicit Classification Maps for Cloud Detection

The Sections IV-A and IV-B were dedicated to thoroughly analyze the performance of the proposed methods, empirically understand their behavior, weaknesses, and strengths, and compare them against GPC. In order to illustrate the explicit cloud detection behind the experiments, here we provide several explanatory classification maps obtained by the best model (in terms of predictive performance) for the Landmarks data set: VFF-GPC with $n = 300\,000$ and $D = 200$.

The classification maps are obtained for the whole year 2010 at the Dakhla landmark, with a total of 34940 satellite acquisitions. The acquired window size is 26×20 pixels. Relying on the proposed feature extraction procedure, we trained the four necessary models (high, mid, low, and night), and then proceed to predict over the whole available amount of chips acquired in the 2010 year.⁴

In Fig. 6, several chips are provided with the aim of illustrating different behaviors. In the first situation (first row), we can see a characteristic error of the L2 cloud mask, which sometimes tends to wrongly label the coastline pixels

⁴A full video with all the classification maps is available at <http://decsai.ugr.es/vip/software.html> and <http://isp.uv.es/code/vff.html>. RFF-GPC and VFF-GPC codes are also provided.

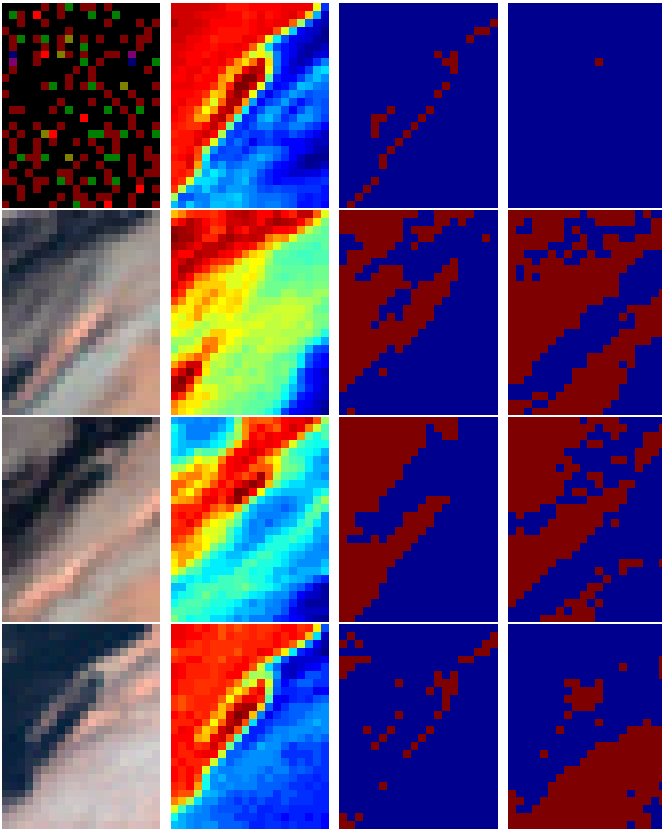


Fig. 6. Explicit classification maps for the Dakhla landmark. The rows correspond with four different acquisitions. The first column shows the visible RGB channels (which are not informative for *night* acquisitions such as the first one), the second column is the infrared $10.8\text{-}\mu\text{m}$ spectral band (very illustrative in night scenarios), the third column represents the ground truth obtained by EUMETSAT (the L2 cloud mask), and the last one is the VFF-GPC classification map. In the last two columns, the red color is used for cloudy pixels and blue for cloud-free ones.

as cloudy. However, VFF-GPC leads to a better classification, identifying just one cloudy pixel and thus avoiding this negative coastline effect.⁵ In the second row, the visual channels show large clouds crossing the landmark. In Fig. 6 (bottom right), a long cloud is unlabeled in the L2 cloud mask but correctly detected by VFF-GPC. While being formally accounted as an error, such discrepancy is actually positive for our method. Moreover, VFF-GPC shows an interesting cloud-sensitive behavior at the top-left cloudy mass, identifying a larger cloudy area than that provided by EUMETSAT. This is a desirable propensity in cloud detection applications, where we prefer to identify larger clouds (and then thoroughly analyze them) rather than missing some of them. In the third row, the RGB channel allows for visually identifying three main cloudy masses at the landmark. The L2 mask poorly labels the central cloudy band, and does not detect the lower cloud. Both deficiencies are overcome by VFF-GPC. Finally, the fourth chip shows a huge cloudy mass that is undetected by the L2 mask but is correctly identified by VFF-GPC.

Therefore, although VFF-GPC is trained with an imperfect ground truth, we observe that it is able to bypass some of these

⁵As a clarification note, the coastline pixels were removed from the training data set by applying a carefully designed morphological filter around coastlines.

deficiencies, and exhibits a desirable cloud-sensitive behavior. This improvement can also be related to the particular design of the training data sets, splitting the problem into four different cases depending on the illumination conditions.

V. CONCLUSION AND FUTURE WORK

We presented two efficient approximations to GP classification to cope with big data classification problems in EO. The first one, RFF-GPC, performs standard GP classification by means of a fast approximation to the kernel (covariance) via D RFFs. The advantage of the method is mainly computational, as the training cost is $\mathcal{O}(nD^2)$ instead of the $\mathcal{O}(n^3)$ induced by the direct inversion of the $n \times n$ kernel matrix [the test cost is also reduced to be independent on n , from $\mathcal{O}(n^2)$ to $\mathcal{O}(D^2)$]. The RFF method approximates the SE covariance with Fourier features randomly sampled in the whole spectral domain. The solid theoretical grounds and good empirical performance make it a very useful method to tackle large-scale problems. Actually, the use of RFF has been exploited before in other settings, from classification with SVMs to regression with the kernel ridge regression. However, we emphasize two main shortcomings. First, the RFF approach can only approximate (theoretically and in practice) a predefined kernel (the SE one in this paper). Second, by sampling the Fourier domain from a Gaussian, one has no control about the expressive power of the representation, since some frequency components of the signal can be better represented than others. As a consequence, the approximated kernel may not have good discrimination capabilities. Noting these two problems, we proposed here our second methodology: a variational GP classifier (VFF-GPC) which goes one step beyond by optimizing over the Fourier frequencies. It is shown to be not just a GP adaptation well suited for large-scale applications, but a whole novel, general-purpose, and very competitive kernel-based classifier that scales well (linearly, as RFF-GPC) with the number of training instances.


We illustrated the performance of the algorithms in two real remote sensing problems of large and medium sizes. In the first case study, a challenging problem dealt with the identification of clouds over landmarks using SEVIRI/MSG imagery. The problem involved several hundred thousand data points for training the classifiers. In the second case study, we used the IAVISA data set, which exploits IASI/AVHRR data to identify clouds with the IASI infrared sounding data. Compared with the original GPC, the experimental results show a high competitiveness in accuracy, a remarkable decrease in computational cost, and an excellent tradeoff between both.

These results encourage us to expand the experimentation to additional problems, trying to exploit the demonstrated potential of VFF-GPC when dealing with any value of n (training data set size) and d (original dimension of the data). Other prior distributions and inference methods, as explained at the end of Section II-A, will also be explored in the future.

REFERENCES

- [1] M. Berger, J. Moreno, J. A. Johannessen, P. F. Levelt, and R. F. Hanssen, "ESA's sentinel missions in support of earth system science," *Remote Sens. Environ.*, vol. 120, pp. 84–90, May 2012.

- [2] M. Drusch *et al.*, "Sentinel-2: ESA's optical high-resolution mission for GMES operational services," *Remote Sens. Environ.*, vol. 120, pp. 25–36, May 2012.
- [3] C. Donlon *et al.*, "The Global Monitoring for Environment and Security (GMES) Sentinel-3 mission," *Remote Sens. Environ.*, vol. 120, pp. 37–57, May 2012.
- [4] T. Stuffer *et al.*, "The EnMAP hyperspectral imager—An advanced optical payload for future applications in Earth observation programmes," *Acta Astron.*, vol. 61, nos. 1–6, pp. 115–120, 2007.
- [5] D. A. Roberts, D. A. Quattrochi, G. C. Hulley, S. J. Hook, and R. O. Green, "Synergies between VSWIR and TIR data for the urban environment: An evaluation of the potential for the Hyperspectral Infrared Imager (HypIRI) Decadal Survey mission," *Remote Sens. Environ.*, vol. 117, pp. 83–101, Feb. 2012.
- [6] D. Labate *et al.*, "The PRISMA payload optomechanical design, a high performance instrument for a new hyperspectral mission," *Acta Astron.*, vol. 65, nos. 9–10, pp. 1429–1436, 2009.
- [7] S. Kraft, U. Del Bello, M. Drusch, A. Gabriele, B. Harnisch, and J. Moreno, "On the demands on imaging spectrometry for the monitoring of global vegetation fluorescence from space," *Proc. SPIE*, vol. 8870, p. 88700N, Sep. 2013.
- [8] N. Longbotham, F. Pacifici, B. Baugh, and G. Camps-Valls, "Prelaunch assessment of WorldView-3 information content," in *Proc. Workshop Hyperspectral Image Signal Process., Evol. Remote Sens.*, Lausanne, Switzerland, Jun. 2014, pp. 24–27.
- [9] F. R. Cayla, "IASI level 0 and 1 processing algorithms description," in *Proc. 12th ISTC Conf.*, 2002.
- [10] G. Camps-Valls and L. Bruzzone, Eds., *Kernel Methods for Remote Sensing Data Analysis*. West Sussex, U.K.: Wiley, Dec. 2009.
- [11] G. Camps-Valls, D. Tuia, L. Gómez-Chova, S. Jiménez, and J. Malo, Eds., *Remote Sensing Image Processing*. San Rafael, CA, USA: Morgan & Claypool, Sep. 2011.
- [12] C. Huang, L. S. Davis, and J. R. G. Townshend, "An assessment of support vector machines for land cover classification," *Int. J. Remote Sens.*, vol. 23, no. 4, pp. 725–749, 2002.
- [13] G. Camps-Valls *et al.*, "Robust support vector method for hyperspectral data classification and knowledge discovery," *IEEE Trans. Geosci. Remote Sens.*, vol. 42, no. 7, pp. 1530–1542, Jul. 2004.
- [14] F. Melgani and L. Bruzzone, "Classification of hyperspectral remote sensing images with support vector machines," *IEEE Trans. Geosci. Remote Sens.*, vol. 42, no. 8, pp. 1778–1790, Aug. 2004.
- [15] G. M. Foody and J. Mathur, "A relative evaluation of multiclass image classification by support vector machines," *IEEE Trans. Geosci. Remote Sens.*, vol. 42, no. 6, pp. 1335–1343, Jul. 2004.
- [16] G. Camps-Valls and L. Bruzzone, "Kernel-based methods for hyperspectral image classification," *IEEE Trans. Geosci. Remote Sens.*, vol. 43, no. 6, pp. 1351–1362, Jun. 2005.
- [17] J. A. Benediktsson, J. A. Palmason, and J. R. Sveinsson, "Classification of hyperspectral data from urban areas based on extended morphological profiles," *IEEE Trans. Geosci. Remote Sens.*, vol. 43, no. 3, pp. 480–491, Mar. 2005.
- [18] M. Fauvel, J. A. Benediktsson, J. Chanussot, and J. R. Sveinsson, "Spectral and spatial classification of hyperspectral data using SVMs and morphological profiles," *IEEE Trans. Geosci. Remote Sens.*, vol. 46, no. 11, pp. 3804–3814, Nov. 2008.
- [19] F. Pacifici, M. Chini, and W. J. Emery, "A neural network approach using multi-scale textural metrics from very high-resolution panchromatic imagery for urban land-use classification," *Remote Sens. Environ.*, vol. 113, no. 6, pp. 1276–1292, 2009.
- [20] D. Tuia, F. Pacifici, M. Kanevski, and W. J. Emery, "Classification of very high spatial resolution imagery using mathematical morphology and support vector machines," *IEEE Trans. Geosci. Remote Sens.*, vol. 47, no. 11, pp. 3866–3879, Nov. 2009.
- [21] G. Camps-Valls, L. Gómez-Chova, J. Muñoz-Marí, J. L. Rojo-Álvarez, and M. Martínez-Ramón, "Kernel-based framework for multitemporal and multisource remote sensing data classification and change detection," *IEEE Trans. Geosci. Remote Sens.*, vol. 46, no. 6, pp. 1822–1835, Jun. 2008.
- [22] A. Plaza *et al.*, "Recent advances in techniques for hyperspectral image processing," *Remote Sens. Environ.*, vol. 113, pp. S110–S122, Sep. 2009.
- [23] J. Plaza, R. Pérez, A. Plaza, P. Martínez, and D. Valencia, "Parallel morphological/neural processing of hyperspectral images using heterogeneous and homogeneous platforms," *Cluster Comput.*, vol. 11, no. 1, pp. 17–32, 2008.
- [24] J. Muñoz-Marí, A. Plaza, J. Gualtieri, and G. Camps-Valls, "Parallel programming and applications in grid, P2P and networking systems," in *Parallel Implementation of SVM in Earth Observation Applications*, F. Xhafa, Ed. U.K.: IOS Press, 2009.
- [25] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. New York, NY, USA: MIT Press, 2006.
- [26] Y. Bazi and F. Melgani, "Gaussian process approach to remote sensing image classification," *IEEE Trans. Geosci. Remote Sens.*, vol. 48, no. 1, pp. 186–197, Jan. 2010.
- [27] M. Y. Yang, W. Liao, B. Rosenhahn, and Z. Zhang, "Hyperspectral image classification using Gaussian process models," in *Proc. IEEE Int. Geosci. Remote Sens. Symp. (IGARSS)*, Jul. 2015, pp. 1717–1720.
- [28] K. Chen, P. Jian, Z. Zhou, J. Guo, and D. Zhang, "Semantic annotation of high-resolution remote sensing images via Gaussian process multi-instance multilabel learning," *IEEE Geosci. Remote Sens. Lett.*, vol. 10, no. 6, pp. 1285–1289, Nov. 2013.
- [29] K. Chen, Z. Zhou, C. Huo, X. Sun, and K. Fu, "A semisupervised context-sensitive change detection technique via Gaussian process," *IEEE Geosci. Remote Sens. Lett.*, vol. 10, no. 2, pp. 236–240, Mar. 2013.
- [30] P. Ruiz, J. Mateos, G. Camps-Valls, R. Molina, and A. K. Katsaggelos, "Bayesian active remote sensing image classification," *IEEE Trans. Geosci. Remote Sens.*, vol. 52, no. 4, pp. 2186–2196, Apr. 2014.
- [31] L. Kalantari, P. Gader, S. Graves, and S. A. Bohlman, "One-class Gaussian process for possibilistic classification using imaging spectroscopy," *IEEE Geosci. Remote Sens. Lett.*, vol. 13, no. 7, pp. 967–971, Jul. 2016.
- [32] T. P. Minka, "Expectation propagation for approximate Bayesian inference," in *Proc. 7th Conf. Uncertainty Artif. Intell. (UAI)*, 2001, pp. 362–369.
- [33] M. Kuss and C. E. Rasmussen, "Assessing approximate inference for binary Gaussian process classification," *Mach. Learn. Res.*, vol. 6, pp. 1679–1704, Oct. 2005.
- [34] Z. Chen, S. D. Babacan, R. Molina, and A. K. Katsaggelos, "Variational Bayesian methods for multimedia problems," *IEEE Trans. Multimedia*, vol. 16, no. 4, pp. 1000–1017, Jun. 2014.
- [35] A. Damianou, "Deep Gaussian processes and variational propagation of uncertainty," Ph.D. dissertation, Dept. Neurosci., Univ. Sheffield, Sheffield, U.K., 2015.
- [36] A. Rahimi and B. Recht, "Random features for large-scale kernel machines," in *Proc. Adv. Neural Inf. Process. Syst.*, 2007, pp. 1177–1184.
- [37] A. Pérez-Suay, J. Amorós-López, L. Gómez-Chova, V. Laparra, J. Muñoz-Marí, and G. Camps-Valls, "Randomized kernels for large scale Earth observation applications," *Remote Sens. Environ.*, vol. 202, pp. 54–63, Dec. 2017.
- [38] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. New York, NY, USA: MIT Press, 2006.
- [39] W. Rudin, *Fourier Analysis on Groups*. Hoboken, NJ, USA: Wiley, 2011.
- [40] C. Bishop, *Pattern Recognition and Machine Learning*. New York, NY, USA: Springer, 2006.
- [41] M. Lázaro-Gredilla, J. Quiñero-Candela, C. E. Rasmussen, and A. R. Figueiras-Vidal, "Sparse spectrum Gaussian process regression," *J. Mach. Learn. Res.*, vol. 11, pp. 1865–1881, Jun. 2010.
- [42] S. D. Babacan, R. Molina, M. N. Do, and A. K. Katsaggelos, "Bayesian blind deconvolution with general sparse image priors," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2012, pp. 341–355.
- [43] R. Fletcher and C. M. Reeves, "Function minimization by conjugate gradients," *Comput. J.*, vol. 7, no. 2, pp. 149–154, 1964.
- [44] M. Derrien and H. Le Gléau, "MSG/SEVIRI cloud mask and type from SAFNWC," *Int. J. Remote Sens.*, vol. 26, no. 21, pp. 4707–4732, 2005.
- [45] J. Hocking, P. N. Francis, and R. Saunders, "Cloud detection in Meteosat Second Generation imagery at the Met Office," *Meteorol. Appl.*, vol. 18, no. 3, pp. 307–323, 2011. [Online]. Available: <http://dx.doi.org/10.1002/met.239>, doi: [10.1002/met.239](https://doi.org/10.1002/met.239).



Pablo Morales-Álvarez received the B.Sc. degree in mathematics and the M.Sc. degrees in mathematical physics and data science from the University of Granada, Granada, Spain, in 2014, 2015, and 2016, respectively, where he is currently pursuing the Ph.D. degree with the Department of Computer Science and Artificial Intelligence under the supervision of Prof. R. Molina, funded by a Ph.D. Fellowship from the La Caixa Foundation.

His research interests include probabilistic machine learning methods, specially Gaussian processes, and their application to image processing and classification problems.



Adrián Pérez-Suay (M'17) received the B.Sc. degree in mathematics, the master's degree in advanced computing and intelligent systems, and the Ph.D. degree in computational mathematics and computer science about distance metric learning from the Universitat de València, València, Spain, in 2007, 2010, and 2015, respectively.

He was an Assistant Professor with the Department of Mathematics, Universitat de València, where is currently a Post-Doctoral Researcher with the Image Processing Laboratory, involved in feature

extraction and classification problems in remote sensing data analysis.



Rafael Molina (SM'13) received the B.Sc. degree in mathematics (statistics) and the Ph.D. degree in optimal design in linear models from the University of Granada, Granada, Spain, in 1979 and 1983, respectively.

He was the former Dean of the Computer Engineering School, University of Granada, from 1992 to 2002, where he was the Head of the Computer Science and Artificial Intelligence Department from 2005 to 2007. He has been a Professor of computer science and artificial intelligence with the University

of Granada, since 2000. His research interests include Bayesian modeling and inference in problems, such as image restoration (applications to astronomy and medicine), superresolution of images and video, blind deconvolution, computational photography, source recovery in medicine, compressive sensing, low-rank matrix decomposition, active learning, fusion, and classification.

Dr. Molina was a recipient of the IEEE International Conference on Image Processing Paper Award in 2007 and the ISPA Best Paper Award in 2009. He was a co-author of a paper awarded the Runner-Up Prize at Reception for early stage researchers at the House of Commons. He served as an Associate Editor of the *Applied Signal Processing* from 2005 to 2007, the *IEEE TRANSACTIONS ON IMAGE PROCESSING* from 2010 to 2013, and *Progress in Artificial Intelligence* from 2011 to 2015. He has been an Area Editor of the *Digital Signal Processing* since 2011.



Gustau Camps-Valls (M'04–SM'07) received the B.Sc. degree in physics, the B.Sc. degree in electronics engineering, and the Ph.D. degree in physics from the Universitat de València, València, Spain, in 1996, 1998, and 2002, respectively.

He was a Visiting Researcher with the Remote Sensing Laboratory, University of Trento, Trento, Italy, in 2002, and with the Max Planck Institute for Biological Cybernetics, Tübingen, Germany, in 2009. He was an Invited Professor with the Laboratory of Geographic Information Systems, Ecole

Polytechnique Fédérale de Lausanne, Lausanne, Switzerland, in 2013. He is currently a Professor with the Department of Electrical Engineering and a Research Coordinator with the Image and Signal Processing Group, Universitat de València. He is interested in the development of machine learning algorithms for geoscience and remote sensing data analysis. He has authored 150 journal papers, over 200 conference papers, and 30 international book chapters.

Dr. Camps-Valls has been a member of the Data Fusion Technical Committee of the IEEE Geoscience and Remote Sensing Society since 2007 and the Machine Learning for Signal Processing Technical Committee of the IEEE SPS since 2009. He is currently a member of the MTG-IRS Science Team of EUMETSAT, and a referee and a program committee member of many international journals and conferences. In 2015, he obtained the prestigious European Research Council consolidator grant on statistical learning for earth observation data analysis. He is the Editor of the books *Kernel Methods in Bioengineering, Signal, and Image Processing* (IGI, 2007), *Kernel Methods for Remote Sensing Data Analysis* (Wiley, 2009), and *Remote Sensing Image Processing* (MC, 2011). He is the Co-Editor of the book *Digital Signal Processing with Kernel Methods* (Wiley, 2017). He holds an H-index of 53, entered the ISI list of Highly Cited Researchers in 2011, and Thomson Reuters Science Watch identified one of his papers on kernel-based analysis of hyperspectral images as a Fast Moving Front Research. He is currently an Associate Editor of the *IEEE TRANSACTIONS ON SIGNAL PROCESSING*, the *IEEE SIGNAL PROCESSING LETTERS*, and the *IEEE GEOSCIENCE AND REMOTE SENSING LETTERS*, and an Invited Guest Editor of the *IEEE JOURNAL OF SELECTED TOPICS IN SIGNAL PROCESSING* in 2012 and the *IEEE GEOSCIENCE AND REMOTE SENSING MAGAZINE* in 2015.