

Fast forward feature selection of hyperspectral images for classification with Gaussian mixture models

Mathieu Fauvel, Clément Dechesne, Anthony Zullo and Frédéric Ferraty

Abstract—A fast forward feature selection algorithm is presented in this paper. It is based on a Gaussian mixture model (GMM) classifier. GMM are used for classifying hyperspectral images. The algorithm selects iteratively spectral features that maximizes an estimation of the classification rate. The estimation is done using the k-fold cross validation. In order to perform fast in terms of computing time, an efficient implementation is proposed. First, the GMM can be updated when the estimation of the classification rate is computed, rather than re-estimate the full model. Secondly, using marginalization of the GMM, sub models can be directly obtained from the full model learned with all the spectral features. Experimental results for two real hyperspectral data sets show that the method performs very well in terms of classification accuracy and processing time. Furthermore, the extracted model contains very few spectral channels.

Index Terms—Hyperspectral image classification, nonlinear feature selection, Gaussian mixture model, parsimony.

I. INTRODUCTION

Since the pioneer paper of J. Jimenez and D. Landgrebe [1], it is well known that hyperspectral images need specific processing techniques because conventional ones made for multispectral/panchromatic images do not adapt well to hyperspectral images. Generally speaking, the increasing number of spectral channels poses theoretical and practical problems [2]. In particular, for the purpose of pixel classification, the spectral dimension needs to be handled carefully because of the “Hughes phenomenon” [3]: with a limited training set, beyond a certain number of spectral features, a reliable estimation of the model parameters is not possible.

Many works have been published since the 2000s to address the problem of classifying hyperspectral images. A non-exhaustive list should include techniques from the machine learning theory (Support Vector Machines, Random Forest, neural networks) [4], statistical models [1] and dimension reduction [5]. SVM, and kernel methods in general, have shown remarkable performances on hyperspectral data in terms of classification accuracy [6]. However, these methods may

suffer from a high computational load and the interpretation of the model is usually not trivial.

In parallel to the emergence of kernel methods, the reduction of the spectral dimension has received a lot of attention. According to the absence or presence of training set, the dimension reduction can be unsupervised or supervised. The former try to describe the data with a lower number of features that minimize a reconstruction error measure, while the latter try to extract features that maximize the separability of the classes. One of the most used unsupervised feature extraction method is the principal component analysis (PCA) [1]. But it has been demonstrated that PCA is not optimal for the purpose of classification [7]. Supervised method, such as the Fisher discriminant analysis or the non-weighted feature extraction have shown to perform better for the purpose of classification. Other feature extraction techniques, such as independent component analysis [8], have been applied successfully and demonstrate that even SVM can benefit from feature reduction [9], [10]. However, conventional supervised techniques suffer from similar problems than classification algorithms in high dimensional space.

Rather than supervised and unsupervised techniques, one can also distinguish dimension reduction techniques into *feature extraction* and *feature selection*. Feature extraction returns a linear/nonlinear combination of the original features, while feature selection returns a subset of the original features. While feature extraction and feature selection both reduce the dimensionality of the data, the latter is much more interpretable for the end-user. The extracted subset corresponds to the most important features for the classification, i.e., the most important wavelengths. For some applications, these spectral channels can be used to infer mineralogical and chemical properties [11].

Feature selection techniques generally need a criterion, that evaluates how the model built with a given subset of features performs, and an optimization procedure that tries to find the subset of features that maximizes/minimizes the criterion [12]. Several methods have been proposed according to that setting. For instance, an entropy measure and a genetic algorithm have been proposed in [13, Chapter 9], but the band selection was done independently of the classifier, i.e., the criterion was not directly related to the classification accuracy. Jeffries Matusita (JM) distance and steepest-ascent like algorithms were proposed in [14]. The method starts with a conventional sequential forward selection algorithm, then the obtained set of features is updated using local search. The method has

M. Fauvel, C. Deschene and A. Zullo are with the Université de Toulouse, INP-ENSAT, UMR 1201 DYNAFOR, France and with the INRA, UMR 1201 DYNAFOR, France

A. Zullo and F. Ferraty are with Institut de Mathématiques de Toulouse - UMR 5219 IMT & Université de Toulouse, France

This work was supported by the French National Research Agency (ANR) under Project Grant ANR-13-JS02-0005-01 (Asterix project), by the EFPA department of the French National Institute for Agricultural Research (INRA) through an Innovative Project grant, by the French National Spatial Agency (CNES) and the Midi-Pyrénées region.

been extended in [15] where a multiobjective criterion was used to take into account the class separability and the spatial variability of the features. JM distance and exhaustive search as well as some refinement techniques have been proposed also in [12]. However rather than extracting spectral features, the algorithm returns the average over a certain bandwidth of contiguous channels, which can make the interpretation difficult and often leads to select a large part of the electromagnetic spectrum. Similarly, spectral intervals selection was proposed in [16], where the criterion used was the square representation error (square error between the approximate spectra and the original spectra) and the optimization problem was solved using dynamic programming. These two methods reduce the dimensionality of the data, but cannot be used to extract spectral variables. Recently, forward selection and genetic algorithm driven by the classification error minimization has been proposed in [17].

Feature selection has been also proposed for kernel methods. A recursive scheme used to remove features that exhibit few influence on the decision function of a nonlinear SVM was discussed in [18]. Alternatively, a shrinkage method based on ℓ_1 -norm and linear SVM has been investigated by Tuia *et al.* [19]. The authors proposed a method where the features are extracted during the training process. However, to make the method tractable in terms of computational load, a linear model is used for the classification, which can limit the discriminating power of the classifier. In [20], a dependence measure between spectral features and thematic classes is proposed using kernel evaluation. The measure has the advantage to be applicable to multiclass problem making the interpretation of the extracted features easier.

Feature selection usually provides good results in terms of classification accuracy. However, several drawbacks can be identified from the above mentioned literature:

- It can be very time consuming, in particular when nonlinear classification models are used.
- When linear models are used for the selection of features, performances in terms of classification accuracy are not satisfying and therefore another nonlinear classifier should be used after the feature extraction.
- For multiclass problem, it is sometimes difficult to interpret the extracted features when a collection of binary classifiers is used (e.g., SVM).

In this work, it is proposed to use a forward strategy, based on [21], that uses an efficient implementation scheme and allows to process a large amount of data, both in terms of number of samples and variables. The method, called *nonlinear parsimonious feature selection* (NPFS), selects iteratively a spectral feature from the original set of features and adds it to a pool of selected features. This pool is used to learn a Gaussian mixture model (GMM) and each feature is selected according to a classification rate. The iteration stops when the increased in terms of classification rate is lower than a user defined threshold or when the maximum number of features is reached. In comparison to other feature extraction algorithms, the main contributions of NPFS is the ability to select spectral features through a nonlinear classification model and its high

computational efficiency. Furthermore, NPFS usually extracts a very few number of features (lower than 5 % of the original number of spectral features).

The remaining of the paper is organized as follows. Section II presents the algorithm with the Gaussian mixture model and the efficient implementation. Experimental results on three hyperspectral data sets are presented and discussed in Section III. Conclusions and perspectives conclude the paper in Section IV.

II. NON LINEAR PARSIMONIOUS FEATURE SELECTION

The following notations are used in the remaining of the paper. $\mathcal{S} = \{\mathbf{x}_i, y_i\}_{i=1}^n$ denotes the set of training pixels, where $\mathbf{x}_i \in \mathbb{R}^d$ is a d -dimensional pixel vector, $y_i = 1, \dots, C$ is its corresponding class, C the number of classes, n the total number of training pixels and n_c the number of training pixels in class c .

A. Gaussian mixture model

For a Gaussian mixture model, it is supposed that the observed pixel is a realization of a d -dimensional random vector such as

$$p(\mathbf{x}) = \sum_{c=1}^C \pi_c p(\mathbf{x}|c), \quad (1)$$

where π_c is the proportion of class c ($0 \leq \pi_c \leq 1$ and $\sum_{c=1}^C \pi_c = 1$) and $p(\mathbf{x}|c)$ is a d -dimensional Gaussian distribution, i.e.,

$$p(\mathbf{x}|c) = \frac{1}{(2\pi)^{d/2} |\Sigma_c|^{1/2}} \exp \left(-\frac{1}{2} (\mathbf{x} - \boldsymbol{\mu}_c)^\top \Sigma_c^{-1} (\mathbf{x} - \boldsymbol{\mu}_c) \right).$$

with $\boldsymbol{\mu}_c$ being the mean vector of class c , Σ_c being the covariance matrix of class c and $|\Sigma_c|$ its determinant. Following the maximum *a posteriori* rule, a given pixel is classified to the class c if $p(c|\mathbf{x}) \geq p(k|\mathbf{x})$ for all $k = 1, \dots, C$. Using the Bayes formula, the posterior probability can be written as

$$p(c|\mathbf{x}) = \frac{\pi_c p(\mathbf{x}|c)}{\sum_{k=1}^C \pi_k p(\mathbf{x}|k)}. \quad (2)$$

Therefore, the maximum *a posteriori* rule can be written as

$$\mathbf{x} \text{ belongs to } c \Leftrightarrow c = \arg \max_{k=1, \dots, C} \pi_k p(\mathbf{x}|k). \quad (3)$$

By taking the log of eq. (3) the final decision rule is obtained (also known as quadratic discriminant function)

$$Q_c(\mathbf{x}) = -(\mathbf{x} - \boldsymbol{\mu}_c)^\top \Sigma_c^{-1} (\mathbf{x} - \boldsymbol{\mu}_c) - \ln(|\Sigma_c|) + 2 \ln(\pi_c). \quad (4)$$

Using standard maximization of the log-likelihood, the estimator of the model parameters are given by

$$\hat{\pi}_c = \frac{n_c}{n}, \quad (5)$$

$$\hat{\boldsymbol{\mu}}_c = \frac{1}{n_c} \sum_{i=1}^{n_c} \mathbf{x}_i, \quad (6)$$

$$\hat{\Sigma}_c = \frac{1}{n_c} \sum_{i=1}^{n_c} (\mathbf{x}_i - \hat{\boldsymbol{\mu}}_c)(\mathbf{x}_i - \hat{\boldsymbol{\mu}}_c)^\top. \quad (7)$$

with n_c is the number of sample of class c .

For GMM, the ‘‘Hughes phenomenon’’ is related to the estimation of the covariance matrix. If the number of training samples is not sufficient for a good estimation the computation of the inverse and of the determinant in eq.(4) will be very numerically unstable, leading to poor classification accuracy. For instance for the covariance matrix, the number of parameters to estimate is equal to $d(d+1)/2$: if $d = 100$ then 5050 parameters have to be estimated then the minimum number of training samples for the considered class should be at least 5050. Note in that case the estimation will be possible but not accurate. Feature selection tackles this problem by allowing the construction of GMM with a reduced number p of variables, with $p \ll d$ and $p(p+1)/2 < n_c$.

B. Forward feature selection

The forward feature selection works as follow [22, Chapter 3]. It starts with an empty pool of selected features. At each step, the feature that most improves an estimation of the classification rate is added to the pool. The algorithm stops either if the increase of the estimated classification rate is too low or if the maximum number of features is reached.

The k -fold cross-validation (k -CV) is used in this work to estimate the classification rate. To compute the k -CV, a subset is removed from \mathcal{S} and the GMM is learned with the remaining training samples. A test error is computed with the removed training samples used as validation samples. The process is iterated k times and the estimated classification rate is computed as the mean test error over the k subsets of \mathcal{S} .

The efficient implementation of the NPFS relies on a fast estimation of the parameters of the GMM when the k -CV is computed. In the following, it will be shown that by using update rules of the parameters and the marginalization properties of the Gaussian distribution, it is possible to perform k -CV and forward selection quickly. As a consequence, the GMM model is learned only one time during the whole training step. The algorithm 1 presents a pseudo code of the proposed method.

1) *Fast estimation of the model on $\mathcal{S}^{n-\nu}$* : In this subsection, it is shown that each parameter can be easily updated when a subset is taken off \mathcal{S} .

Proposition 1 (Proportion): The update rule for the proportion is

$$\hat{\pi}_c^{n-\nu} = \frac{n\hat{\pi}_c^n - \nu_c}{n - \nu} \quad (8)$$

where $\hat{\pi}_c^{n-\nu}$ and $\hat{\pi}_c^n$ are the proportions of class c computed over $n - \nu$ and n respectively, ν is the number of removed samples from \mathcal{S} , ν_c is the number of removed samples from class c such as $\sum_{c=1}^C \nu_c = \nu$.

Proposition 2 (Mean vector): The update rule for the mean vector is

$$\hat{\mu}_c^{n-\nu_c} = \frac{n_c \hat{\mu}_c^{n_c} - \nu_c \hat{\mu}_c^{\nu_c}}{n_c - \nu_c} \quad (9)$$

where $\hat{\mu}_c^{n_c}$ and $\hat{\mu}_c^{n_c-\nu_c}$ are the mean vectors of class c computed over the n_c and $n_c - \nu_c$ training samples respectively, $\hat{\mu}_c^{\nu_c}$ is the mean vector of the ν_c removed samples from class c .

Proposition 3 (Covariance matrix): The update rule for the covariance matrix is

$$\hat{\Sigma}_c^{n_c-\nu_c} = \frac{n_c}{(n_c - \nu_c)} \hat{\Sigma}_c^{n_c} - \frac{\nu_c}{(n_c - \nu_c)} \hat{\Sigma}_c^{\nu_c} - \frac{n_c \nu_c}{(n_c - \nu_c)^2} (\hat{\mu}_c^{n_c} - \hat{\mu}_c^{\nu_c})(\hat{\mu}_c^{n_c} - \hat{\mu}_c^{\nu_c})^\top \quad (10)$$

where $\hat{\Sigma}_c^{n_c}$ and $\hat{\Sigma}_c^{n_c-\nu_c}$ are the covariance matrices of class c computed over the n_c and $n_c - \nu_c$ training samples respectively.

2) *Particular case of leave-one-out cross-validation*: When very few training samples are available, it is sometimes necessary to resort to leave-one-out cross-validation ($k = n$). Update rules are still valid, but it is also possible to get a fast update of the decision function. If the removed sample does not belong to class c , only the proportion term in eq. (4) change, therefore the updated decision rule can be written as:

$$Q_c^{n_c-1}(\mathbf{x}_n) = Q_c^{n_c}(\mathbf{x}_n) + 2 \ln \left(\frac{n-1}{n} \right). \quad (11)$$

where $Q_c^{n_c}$ and $Q_c^{n_c-1}$ are the decision rules for class c computed with n_c and $n_c - 1$ samples respectively. If the removed sample \mathbf{x}_n belongs to class c then updates rules become:

Proposition 4 (Proportion-loocv):

$$\hat{\pi}_c^{n-1} = \frac{n\hat{\pi}_c^n - 1}{n - 1} \quad (12)$$

Proposition 5 (Mean vector-loocv):

$$\hat{\mu}_c^{n-1} = \frac{n_c \hat{\mu}_c^{n_c} - \mathbf{x}_n}{n_c - 1} \quad (13)$$

Proposition 6 (Covariance matrix-loocv):

$$\hat{\Sigma}_c^{n-1} = \frac{n_c}{n_c - 1} \hat{\Sigma}_c^{n_c} - \frac{n_c}{(n_c - 1)^2} (\mathbf{x}_n - \hat{\mu}_c^{n_c})(\mathbf{x}_n - \hat{\mu}_c^{n_c})^\top. \quad (14)$$

where $n_c - 1$ denotes that the estimation is done with only $n_c - 1$ samples rather than the n_c samples of the class.

Update rules have been proposed in [23] for the leave-one-out case. Authors have proposed a way to compute the inverse of the covariance matrix with a low computational cost when one sample is removed. It is based on the Sherman-Morrison-Woodbury formula. In their approach, the inverse of the covariance matrix is computed explicitly in 4. In this work, we choose to not compute the inverse but rather solve the linear problem $\Sigma^{-1}\mathbf{x}$. This approach is more demanding in terms of processing time (still fast when the number of variables is low ~ 10 -15) but far more robust in terms of numerical stability. An update rule for the case where the sample belongs the class c can be written by using the Cholesky decomposition of the covariance matrix and rank-one downdate, but the downdate step is not numerically stable and not used here.

3) *Marginalization of Gaussian distribution*: To get the GMM model over a subset of the original set of features, it is only necessary to drop the non-selected features from the mean vector and the covariance matrix [24]. For instance, let $\mathbf{x} = [\mathbf{x}_s, \mathbf{x}_{ns}]$ where \mathbf{x}_s and \mathbf{x}_{ns} are the selected variables and

Algorithm 1 NPFS pseudo code

Require: \mathcal{S} , k , δ , maxvariable

```

1: Randomly cut  $\mathcal{S}$  into  $k$  subsets such as  $\mathcal{S}_1 \cup \dots \cup \mathcal{S}_k = \mathcal{S}$  and  $\mathcal{S}_i \cap \mathcal{S}_j = \emptyset$ 
2: Learn the full GMM with  $\mathcal{S}$ 
3: Initialize the set of selected variables  $\varphi_s$  to empty set ( $|\varphi_s| = 0$ ) and
   available variables  $\varphi_a$  to the original set of variables ( $|\varphi_a| = d$ )
4: while  $|\varphi_s| \leq \text{maxvariable}$  do
5:   for all  $\mathcal{S}_u \subset \mathcal{S}$  do
6:     Update the model using eq. (8), (9) and (10) (or their loocv
       counterparts) according to  $\mathcal{S}_u$ 
7:   for all  $s \in \varphi_a$  do
8:     Compute the classification rate on  $\mathcal{S}_u$  for each set of variables
        $\varphi_s \cup s$  using the marginalization properties
9:   end for
10: end for
11: Average the classification rate over the  $k$ -fold
12: if Improvement in terms of classification rate w.r.t. previous iteration
    is lower than  $\delta$  then
13:   break
14: else
15:   Add the variable  $s$  corresponding to the maximum classification
     rate to  $\varphi_s$  and remove it from  $\varphi_a$ 
16: end if
17: end while
  
```

the non-selected variables respectively, the mean vector can be written as

$$\hat{\boldsymbol{\mu}} = [\boldsymbol{\mu}_s, \boldsymbol{\mu}_{ns}]^\top \quad (15)$$

and the covariance matrix as

$$\boldsymbol{\Sigma} = \begin{bmatrix} \boldsymbol{\Sigma}_{s,s} & \boldsymbol{\Sigma}_{s,ns} \\ \boldsymbol{\Sigma}_{ns,s} & \boldsymbol{\Sigma}_{ns,ns} \end{bmatrix}. \quad (16)$$

The marginalization over the non-selected variables shows that \mathbf{x}_s is also a Gaussian distribution with mean vector $\boldsymbol{\mu}_s$ and covariance matrix $\boldsymbol{\Sigma}_{s,s}$. Hence, once the full model is learned, all the sub-models built with a subset of the original variables are available at no computational cost.

III. EXPERIMENTAL RESULTS

A. Data

Two data sets have been used in the experiments. The first data set has been acquired in the region surrounding the volcano Hekla in Iceland by the AVIRIS sensor. 157 spectral channels from 400 to 1,840 nm were recorded. 12 classes have been defined for a total of 10,227 referenced pixels. The second data set has been acquired by the ROSIS sensor during a flight campaign over Pavia, northern Italy. 103 spectral channels were recorded from 430 to 860 nm. 9 classes have been defined for a total of 42,776 referenced pixels.

For each data set, 50, 100 and 200 training pixels per class were randomly selected and the remaining referenced pixels were used for the validation. 50 repetitions were done for which a new training set have been generated randomly.

B. Competitive methods

Several conventional feature selection methods have been used as baseline.

- Recursive Feature Elimination (RFE) for nonlinear SVM [18]. In the experiment, a Gaussian kernel was used.




Fig. 1. Mean number \bar{n}_s of selected features for the different methods for Hekla data set. The red line indicates the original number of spectral features. Projected ℓ_1 SVM is not reported because the mean number of extracted features was too high (e.g., 6531 for $n_c=50$).

- Linear SVM with ℓ_1 (SVM $_{\ell_1}$) constraint on the feature vector [19] based on the LIBLINEAR implementation [25].
- To overcome the limitation of the linear model used in LIBLINEAR, an explicit computation of order 2 polynomial feature space has been used together with LIBLINEAR (SVM $_{\ell_1}^p$). Formally, a nonlinear transformation ϕ has been apply on the original samples:

$$\begin{aligned} \mathbb{R}^d &\rightarrow \mathbb{R}^p \\ \mathbf{x} = [x_1, \dots, x_d] &\mapsto \phi(\mathbf{x}) = [x_1, \dots, x_d, x_1^2, x_1x_2, \dots, \\ &\quad x_1x_d, x_2^2, x_2x_3, \dots, x_d^2] \end{aligned}$$

with $p = \binom{2+d}{2}$. For Hekla data and University of Pavia data, the dimension p of the projected space is 12561 and 5460, respectively.

For comparison, a SVM with a Gaussian kernel and a order 2 polynomial kernel classifier, based on the LIBSVM [26], with all the variables have been used too.

For the linear/nonlinear SVM, the penalty parameter and the kernel hyperparameters were selected using 5-fold cross-validation. For NPFS, the threshold (δ in Algorithm 1) was set to 0.5% and the maximum number of extracted features was set to 20. The estimation of the error has been computed with a leave-one-out CV (n -NPFS) and a 5-fold CV (5-NPFS). Each variable has been standardized before the processing (i.e., zero mean and unit variance).

C. Results

The mean accuracies and the variance over the 50 runs are reported in Table I and Table II. The mean numbers of extracted features for the different methods are reported in Figure 1 and Figure 2.

From the tables, it can be seen that there is no difference in the results obtained with n -NPFS or 5-NPFS. They perform equally on both data sets in terms of classification accuracy or number of extracted features. However, 5-NPFS is much faster in terms of computation time.

RFE and SVM $_{\text{gauss}}$ provide the best results in terms of classification accuracy, except for the Hekla data set and

TABLE I

CLASSIFICATION ACCURACIES FOR HEKLA DATA SET. THE RESULTS CORRESPOND TO THE MEAN VALUE AND VARIANCE OF THE OVERALL ACCURACY OVER THE 50 REPETITIONS. THE BEST RESULT FOR EACH TRAINING SETUP IS REPORTED IN BOLD FACE. n -NPFS AND 5-NPFS CORRESPOND TO THE NPFS COMPUTED WITH THE LEAVE-ONE-OUT AND 5-FOLD CROSS-VALIDATION, RESPECTIVELY. RFE, SVM_{ℓ_1} AND $SVM_{\ell_1}^p$ CORRESPOND TO THE RECURSIVE FEATURE EXTRACTION SVM, THE LINEAR SVM WITH ℓ_1 CONSTRAINT AND THE LINEAR SVM WITH ℓ_1 WITH THE EXPLICIT ORDER 2 POLYNOMIAL FEATURE SPACE, RESPECTIVELY. SVM_{poly} AND SVM_{gauss} CORRESPOND TO THE CONVENTIONAL NONLINEAR SVM WITH A ORDER 2 POLYNOMIAL KERNEL AND GAUSSIAN KERNEL, RESPECTIVELY.

n_c	n -NPFS	5-NPFS	RFE	SVM_{ℓ_1}	$SVM_{\ell_1}^p$	SVM_{poly}	SVM_{gauss}
50	92.5 \pm 1.2	92.4 \pm 1.2	90.2 \pm 1.8	90.3 \pm 1.0	91.6 \pm 0.6	84.6 \pm 1.6	90.4 \pm 1.6
100	94.8 \pm 0.7	94.6 \pm 0.6	95.6 \pm 0.3	93.9 \pm 0.5	94.8 \pm 0.1	91.4 \pm 0.4	95.6 \pm 0.3
200	95.9 \pm 0.3	95.8 \pm 0.3	96.8 \pm 1.1	95.6 \pm 0.1	96.3 \pm 0.1	95.5 \pm 0.1	96.8 \pm 1.1

TABLE II

CLASSIFICATION ACCURACIES FOR UNIVERSITY OF PAVIA DATA SET. SAME NOTATIONS AS IN TABLE I.

n_c	n -NPFS	5-NPFS	RFE	SVM_{ℓ_1}	$SVM_{\ell_1}^p$	SVM_{poly}	SVM_{gauss}
50	82.2 \pm 4.4	83.4 \pm 7.6	84.7 \pm 4.0	75.1 \pm 2.5	81.0 \pm 2.8	82.9 \pm 3.4	84.8 \pm 3.4
100	86.3 \pm 3.2	85.9 \pm 3.1	88.4 \pm 0.9	77.3 \pm 1.4	83.6 \pm 1.3	86.5 \pm 1.6	88.4 \pm 1.4
200	87.7 \pm 3.1	87.9 \pm 1.9	90.8 \pm 0.3	78.5 \pm 0.7	85.5 \pm 0.4	88.8 \pm 0.6	90.8 \pm 0.3




Fig. 2. Mean number \bar{n}_s of selected features for the different methods for University of Pavia data set. The red line indicates the original number of spectral features. Projected ℓ_1 SVM is not reported because the mean number of extracted features was too high (e.g., 5110 for $n_c=50$).

$n_c = 50$. From the figure, it can be seen that the number of extracted features is almost equal to original number of spectral features, meaning that in these experiments RFE is equivalent to SVM_{gauss} . Hence, RFE was not able to extract few relevant spectral features.

ℓ_1 SVM applied on the original features or the projected features is not able to extract relevant features. In terms of classification accuracy, the linear SVM does not perform well for the University of Pavia data set. Nonlinear ℓ_1 SVM provides much better results for both data sets. In comparison to the non sparse nonlinear SVM computed with an order 2 polynomial kernel, ℓ_1 nonlinear SVM performs better in terms of classification accuracy for the Hekla data while it performs worst for the University of Pavia data.

In terms of number of extracted features, NPFS provides the best results, by far, with an average number of extracted features equal to 5% of the original number. All the others methods were not able to extract few features without de-

creasing drastically the overall accuracy. For instance, for the Hekla data set and $n_c = 50$, only 7 spectral features are used to build the GMM and leads to the best classification accuracy. A discussion on the extracted features is given in the next subsection.

The figure 3 presents the average classification rate of 5-NPFS, SVM with a Gaussian kernel and a linear SVM applied on the features selected with 5-NPFS. 20 repetitions have been done on the University data set with $n_c=50$. The optimal parameters for SVM and linear SVM have been selected using 5-fold cross-validation. From the figure, it can be seen that the three algorithms have similar trends. When the number of features is relatively low (here lower than 15) GMM performs the best, but when the number of features increases too much, SVM (non linear and linear) performs better in terms of classification accuracy. It is worth noting that such observations are coherent with the literature: SVM are known to perform well in high dimensional space, while GMM is more affected by the dimension.

The mean processing time for the University of Pavia data set for several training set sizes is reported in Table III. It includes parameter optimization for SVM_{gauss} and SVM_{ℓ_1} . Note that the RFE consists in several SVM_{gauss} optimizations, one for each feature removed (hence, if 3 features are removed, the mean processing time is approximately multiplied by 3). It can be seen that the 5-NPFS method is a little influenced by the size of the training set: what is important is the number of (extracted) variables. For $n_s = 50$, the processing time is slightly higher because of overload due to parallelization procedure. n -NPFS is the more demanding in terms of processing time and thus should be used only when the number of training samples is very limited. Finally, it is important to underline that the NPFS is implemented in Python while SVM used a state of the art implementation in C++ [26].

From these experiments, and from a practical viewpoint, NPFS is a good compromise between high classification




Fig. 3. Classification rate in function of the number of extracted features. Continuous line corresponds to 5-NPFS, dashed line to SVM with a Gaussian kernel and dash-dotted line to a linear SVM.

TABLE III
MEAN PROCESSING TIME IN SECONDS IN FUNCTION OF THE NUMBER OF SAMPLES PER CLASS FOR THE UNIVERSITY OF PAVIA DATA SET. 20 REPETITIONS HAVE BEEN DONE ON LAPTOP WITH 8GB OF RAM AND INTEL(R) CORE(TM) I7-3667U CPU @ 2.00GHZ PROCESSOR.

n_s	50	100	200	400
SVM _{gauss}	11	40	140	505
SVM _{ℓ_1}	52	115	234	498
n -NPFS	242	310	472	883
5-NPFS	35	31	29	43

accuracy and sparse modeling.

D. Discussion

The extracted channels by 5-NPFS and n -NPFS were compared for one training set of the University of Pavia data set: two channels were the same for both methods, 780nm and 776nm; two channels were very close, 555nm and 847nm for 5-NPFS and 551nm and 855nm for n -NPFS; one channel was close, 521nm for 5-NPFS and 501nm for n -NPFS. The other channel selected with n -NPFS is 772nm. If the process is repeated, the result is terms of features selected with n -NPFS and 5-NPFS is similar: on average 35% of the selected features are identical (not necessarily the first ones) and the others selected features are close in terms of wavelength.

The influence of the parameter `delta` has been investigated on the University of Pavia data set. 20 repetitions have been done with $n_c = 50$ for several values of `delta`. Results are reported on figure 4. From the figure, it can be seen that when `delta` is set to a value larger than approximately 1%, the algorithm stops too early and the number of selected features is too low to perform well. Conversely, setting `delta` to a small value does not change the classification rate, a plateau being reached for `delta` lower than 0.5%. In fact, because of the ‘‘Hughes phenomenon’’, adding spectral features to the GMM will first lead to an increase of the classification rate but then (after a possible plateau) the classification rate will decrease,




Fig. 4. The dotted line and the crossed line represent the mean error rate and the average number of selected features, respectively, as a function of `delta`. The simulation was done on the University of Pavia data set, with $n_c = 50$ and for the 5-NPFS algorithm.

i.e., the improvement after two iterations of the algorithm will be negative.

Figure 5 presents the most selected features for the University of Pavia data set. 1000 random repetitions have been done with $n_c=200$ and the features shaded in the figure have been selected at least 10% times (i.e., 100 times over 1000) using 5-NPFS. Five spectral domains can be identified, two from the visible range and three from the near infrared range. In particular, it can be seen that spectral channels from the red-edge part are selected. The width of the spectral domain indicates the variability of the selection. The high correlation between adjacent spectral bands makes the variable selection ‘‘unstable’’, e.g., for a given training set, the channel t would be selected but for another randomly selected training set it might be the channel $t + 1$ or $t - 1$. It is clearly a limitation of the proposed approach.

To conclude this discussion, similar spectral channels are extracted with n -NPFS and 5-NPFS, while the latter is much times faster. Hence, n -NPFS should be only used when very limited number of samples is available. A certain variability is observed in the selection of the spectral channels due to the high correlation of adjacent spectral channels and the step-wise nature of the method.

IV. CONCLUSION

A nonlinear parsimonious feature selection algorithm for the classification of hyperspectral images and the selection of spectral variables has been presented. Using a Gaussian mixture model classifier, spectral variables are extracted iteratively based on the cross-validation estimate of the classification rate. An efficient implementation is proposed that takes into account some properties of Gaussian mixture model: a fast update of the model parameters and a fast access to the sub-models. Experimental results show that the proposed method is able to select few relevant features, and outperform standard SVM-based sparse algorithms while reaching similar classification rates to those obtained with SVM. Furthermore, in comparison to SVM based feature selection algorithm, multiclass problem




Fig. 5. Most selected spectral domain for the University of Pavia data set. Gray bars correspond to the most selected parts of the spectral domain. Horizontal axis corresponds to the wavelength (in nanometers). The mean value of each class is represented in continuous colored lines.

is handled by the GMM making the interpretation of the extracted channels easier.

More investigation are needed to fully understand which features are extracted, since the method is purely statistical. If the red-edge has been identified, the others extracted features are not clearly interpretable. Moreover, small variability has been observed due to the high correlation between adjacent bands and the step-wise procedure. To overcome this limitation, a continuous interval selection strategy, as in [12], will be investigated. Also, a steepest-ascent search strategy could be used to make the final solution more stable.


The python code of the algorithm is available freely for download: <https://github.com/mfauvel/FFFS>.

V. ACKNOWLEDGMENT

The authors would like to thank Professor P. Gamba, University of Pavia, for providing the University of Pavia data set and Professor J.A. Benediktsson, University of Iceland, for providing the Hekla data set. They would like also to thank the reviewers for their many helpful comments.

REFERENCES

- [1] L. O. Jimenez and D. A. Landgrebe, "Supervised classification in high-dimensional space: geometrical, statistical, and asymptotical properties of multivariate data," *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, vol. 28, no. 1, pp. 39–54, feb 1998.
- [2] D. L. Donoho, "High-dimensional data analysis: the curses and blessing of dimensionality," in *AMS Mathematical challenges of the 21st century*, 2000.
- [3] G. F. Hughes, "On the mean accuracy of statistical pattern recognizers," *IEEE Trans. Inf. Theory*, vol. IT-14, pp. 55–63, January 1968.
- [4] M. Fauvel, Y. Tarabalka, J. A. Benediktsson, J. Chanussot, and J. Tilton, "Advances in Spectral-Spatial Classification of Hyperspectral Images," *Proceedings of the IEEE*, vol. 101, no. 3, pp. 652–675, Mar. 2013.
- [5] C. J. C. Burges, "Dimension reduction: A guided tour," *Foundations and Trends in Machine Learning*, vol. 2, no. 4, pp. 275–365, 2010.
- [6] G. Camps-Valls and L. Bruzzone, Eds., *Kernel Methods for Remote Sensing Data Analysis*, Wiley, 2009.
- [7] A. Cheryadat and L. M. Bruce, "Why principal component analysis is not an appropriate feature extraction method for hyperspectral data," in *Geoscience and Remote Sensing Symposium, 2003. Proceedings*, July 2003, vol. 6, pp. 3420–3422 vol.6.
- [8] A. Villa, J. A. Benediktsson, J. Chanussot, and C. Jutten, "Hyperspectral image classification with independent component discriminant analysis," *IEEE Trans. Geosci. Remote Sens.*, vol. 49, no. 12, pp. 4865–4876, Dec 2011.
- [9] 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 - part 2, pp. 3804–3814, Oct. 2008.
- [10] F.E. Fassnacht, C. Neumann, M. Forster, H. Buddenbaum, A. Ghosh, A. Clasen, P.K. Joshi, and B. Koch, "Comparison of feature reduction algorithms for classifying tree species with hyperspectral data on three central european test sites," *Selected Topics in Applied Earth Observations and Remote Sensing, IEEE Journal of*, vol. 7, no. 6, pp. 2547–2561, June 2014.
- [11] M. Lothode, V. Carrere, and R. Marion, "Identifying industrial processes through vnir-swir reflectance spectroscopy of their waste materials," in *Geoscience and Remote Sensing Symposium (IGARSS), 2014 IEEE International*, July 2014, pp. 3288–3291.
- [12] S.B. Serpico and G. Moser, "Extraction of spectral channels from hyperspectral images for classification purposes," *IEEE Trans. Geosci. Remote Sens.*, vol. 45, no. 2, pp. 484–495, Feb 2007.
- [13] C.-I. Chang, *Hyperspectral Data Exploitation: Theory and Applications*, Wiley, 2007.
- [14] S.B. Serpico and L. Bruzzone, "A new search algorithm for feature selection in hyperspectral remote sensing images," *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 39, no. 7, pp. 1360–1367, Jul 2001.
- [15] L. Bruzzone and C. Persello, "A novel approach to the selection of spatially invariant features for the classification of hyperspectral images with improved generalization capability," *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 47, no. 9, pp. 3180–3191, Sept 2009.
- [16] A.C. Jensen and R. Solberg, "Fast hyperspectral feature reduction using piecewise constant function approximations," *Geoscience and Remote Sensing Letters, IEEE*, vol. 4, no. 4, pp. 547–551, Oct 2007.
- [17] A. Le Bris, N. Chehata, X. Briottet, and N. Paparoditis, "Use intermediate results of wrapper band selection methods : a first step toward the optimisation of spectral configuration for land cover classifications," in *Proc. of the IEEE WHISPERS'14*, 2014.
- [18] D. Tuia, F. Pacifici, M. Kanevski, and W. J. Emery, "Classification of very high spatial resolution imagery using mathematical morphology and support vector machines," *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 47, no. 11, pp. 3866–3879, 2009.
- [19] D. Tuia, M. Volpi, M. Dalla Mura, A. Rakotomamonjy, and R. Flamary, "Automatic feature learning for spatio-spectral image classification with sparse SVM," *IEEE Trans. Geosci. Remote Sens.*, vol. PP, no. 99, pp. 1–13, 2014.
- [20] G. Camps-Valls, J. Mooij, and Bernhard Scholkopf, "Remote sensing feature selection by kernel dependence measures," *Geoscience and Remote Sensing Letters, IEEE*, vol. 7, no. 3, pp. 587–591, July 2010.
- [21] F. Ferraty, P. Hall, and P. Vieu, "Most-predictive design points for functional data predictors," *Biometrika*, vol. 97, no. 4, pp. 807–824, 2010.
- [22] T. Hastie, R. Tibshirani, and J. H. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Springer series in statistics. Springer, 2001.
- [23] J.P. Hoffbeck and D.A. Landgrebe, "Covariance matrix estimation and classification with limited training data," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 18, no. 7, pp. 763–767, Jul 1996.
- [24] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*, The MIT Press, 2005.
- [25] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin, "LIBLINEAR: A library for large linear classification," *Journal of Machine Learning Research*, vol. 9, pp. 1871–1874, 2008.
- [26] Chih-C. Chang and C.-J. Lin, "LIBSVM: A library for support vector machines," *ACM Transactions on Intelligent Systems and Technology*, vol. 2, pp. 27:1–27:27, 2011, Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.



Mathieu Fauvel graduated in electrical engineering from the Grenoble Institut of Technology (Grenoble INP), Grenoble, France, in 2004. He received the M.S and Ph.D. degrees in image and signal processing from the Grenoble INP in 2004 and 2007, respectively. In 2007, he was a teaching assistant in Grenoble INP. From 2008 to 2010, he was a postdoctoral research associate with the MISTIS Team of the National Institute for Research in Computer Science and Control (INRIA). Since 2011, Dr. Fauvel has been an Associate Professor with the

National Polytechnic Institute of Toulouse (ENSAT - University of Toulouse) within the DYNAFOR lab (University of Toulouse - INRA). His research interests are remote sensing, data fusion, pattern recognition, multicomponent signal and image processing. Since 2013, he is the president of the French Chapter of the GRSS.



Clément Dechesne graduated in electrical engineering and signal processing from the Toulouse Institut of Technology (Toulouse INP), Toulouse, France, in 2014. He is with the Matis Laboratory, Institut National de l'Information Gographique et Forestire, where he is currently working toward the Ph.D. degree. His research interests include the field of remote sensing, in particular on the analysis of hyperspectral and LIDAR data for forest monitoring.



Anthony Zullo was a mathematics teacher in the Jean Monnet highschool, Cognac, France, from 2010 to 2011. He graduated in mathematical engineering from the Paul Sabatier University (University of Toulouse), Toulouse, France, in 2013. Since 2013, he has been preparing a Ph.D. degree in statistics applied for remote sensing within the DYNAFOR lab (University of Toulouse - INRA). His research interests are statistical modeling, nonparametric functional data analysis and hyperspectral imaging.



Frédéric Ferraty received the Ph.D. degree in Statistics from the Toulouse III University, Toulouse, France, in 1996 and got Assistant Professor position at Toulouse II University. At the same time, he joined as permanent membership the Toulouse Mathematics Institute. Frederic Ferraty is Full Professor since 2012 at Toulouse II University. His main domain of interest concerns high-dimensional data statistics with a special attention on functional data analysis. Dr. Ferraty is author of many works published in top statistical journals, involved in

numerous international scientific tasks and regularly invited to World Statistics Conferences.