

Parsimonious Mahalanobis Kernel for the Classification of High Dimensional Data

M. Fauvel¹, A. Villa², J. Chanussot³ and J.A. Benediktsson⁴

¹ INRA, DYNAFOR, BP 32607 - Auzeville-Tolosane 31326 - Castanet Tolosan - FRANCE

² Aresys srl, via Bistolfi 49, 20134 Milano and Dipartimento di Elettronica ed Informazione, Politecnico di Milano, 20133 Milano, Italy

³ GIPSA-lab, Departement Image Signal, BP 46 - 38402 Saint Martin d'Hères - FRANCE

⁴ Dept. of Electrical and Computer Engineering, University of Iceland Hjardarhagi 2-6, 107 Reykjavik - ICELAND

Abstract

The classification of high dimensional data with kernel methods is considered in this article. Exploiting the *emptiness* property of high dimensional spaces, a kernel based on the Mahalanobis distance is proposed. The computation of the Mahalanobis distance requires the inversion of a covariance matrix. In high dimensional spaces, the estimated covariance matrix is ill-conditioned and its inversion is unstable or impossible. Using a parsimonious statistical model, namely the *High Dimensional Discriminant Analysis* model, the specific signal and noise subspaces are estimated for each considered class making the inverse of the class specific covariance matrix explicit and stable, leading to the definition of a *parsimonious Mahalanobis* kernel. A SVM based framework is used for selecting the hyperparameters of the parsimonious Mahalanobis kernel by optimizing the so-called *radius-margin* bound. Experimental results on three high dimensional data sets show that the proposed kernel is suitable for classifying high dimensional data, providing better classification accuracies than the conventional Gaussian kernel.

1 Introduction

High Dimensional (HD) data sets are commonly available for fully or partially automatic processing: For a relatively low number of samples, n , a huge number of variables, d , is simultaneously accessible. For instance, in hyperspectral imagery, hundreds of spectral wavelengths are recorded for a given pixel; in gene expression analysis, the measure of expression level of thousands of genes is typical; in customer recommendation systems for web services, to each potential client a high number of variables is associated (his past choices, his personal information ...) [1, 2, 3]. For each sample, it is possible to have either numerical or alphabetical variables which can be sparse or with a different signal to noise ratio.

In terms of processing, such data may need to be either classified, clustered, filtered or inverted in a supervised or unsupervised way. Although many algorithms exist in the literature for small or moderate dimensions (from Bayesian methods to Machine Learning techniques) most of them are not well suited to HD data. Actually, HD data pose critical theoretical and practical problems that need to be addressed specifically [2].

Indeed, HD spaces exhibit non intuitive geometrical and statistical properties when compared to lower dimensional spaces. Most of them do not behave in a similar way as in three dimensional Euclidean spaces (Table 1 summarizes the main properties of HD spaces) [4]. For instance, samples following a uniform law will have a tendency to have a high concentration in the corners [5]. The same property holds for normally distributed data: samples tend to have a high concentration in the tails [6], making density estimation a difficult task. This problem can be related to the number of parameters t to be estimated to fit a Gaussian

distribution which grows quadratically with the space dimensionality, $t = d(d + 3)/2$ (5150 for $d = 100$). Because of this, conventional generative methods are not suitable for analyzing this type of data.

Unfortunately, discriminative methods also suffer if the dimensionality is high, due to the “*concentration of measure phenomenon*” [2]. In HD spaces, samples tend to be equally distant from each other [7]. Hence, it is clear that Nearest Neighbors methods will definitively fail to process such data. Moreover, the Euclidean distance will not be appropriate to assess the similarity between two samples. In fact, it has been shown that every Minkowski norm ($\|\mathbf{x}\|_m = (\sum_{i=1}^d |\mathbf{x}_i|^m)^{1/m}$, $m = 1, 2, \dots$) is affected by this phenomenon [8]. Therefore, every method based on the distance between samples [9] (SVM with Gaussian kernel, neural network, Nearest Neighbors, Locally Linear Embedding. . .) are potentially affected by this phenomenon [10, 11].

An additional property, for which the consequences are more practical than theoretical, is the “*empty space phenomenon*” [12]: In HD spaces, the available samples usually fill a very small part of the space. Therefore, most of the space is empty. Note that if originally the empty space phenomenon was considered as a problem, it will be seen in the following that it is actually the basis of several useful statistical models.

Today, the phrasing “*curse of dimensionality*”, originally from R. Bellman [12], refers to the aforementioned problems of HD data and reflects how processing HD data is difficult. However, as D. Donoho has noticed [2], there is also a “*Blessing of dimensionality*”: For instance in classification, the class separability is improved when the dimensionality of the data increases. Consider for example a comparison between hyperspectral (hundreds of spectral wavelengths) and multispectral (tens of spectral wavelengths) remote sensing images [13]. The former contains much more information, and enables a more accurate distinction of the land cover classes. However, if conventional methods are used, the additional information contained in hyperspectral images will not lead to an increase of the classification accuracy [5]. Hence, using conventional methods, classification accuracies remains low.

Several methods have been proposed in the literature to deal with HD data for the purpose of classification. A highly used strategy is *Dimension Reduction* (DR). DR aims at reducing the dimensionality of data by mapping them onto another space of a lower dimension, without discarding any, or as less as possible, of the meaningful information. Recent overviews of DR can be found in [14, 15, 16]. Two main approaches can be defined. 1) Unsupervised DR: The algorithms are applied directly on the data without exploiting any prior information, and project the data into a lower dimensional space, according to some criterion (data variance maximization for PCA, independence for ICA . . .). 2) Supervised DR: Training samples are available and are exploited to find a lower dimensional subspace where the class separability is improved. Fisher Discriminant Analysis (FDA) is surely one of the most famous supervised DR method. However, FDA maximizes the ratio of the “between classes” scatter matrix, S_b , and the “within classes” scatter matrix, S_w . The optimal solution is given by the eigenvector corresponding to the first eigenvalues of $S_w^{-1}S_b$. In HD, S_w^{-1} is in general ill-conditioned which limits the effectiveness of the method. Other popular DR methods such as Laplacian eigenmaps, Isomap or Locally Linear Embedding [15, Chapter 4 and 5] may be also limited by the dimensionality since they are based on the Euclidean distance between the samples. One last drawback of DR methods is the risk of losing relevant information. In general, DR methods act globally, which can be a problem for classification purpose: Different classes may be mapped onto the same subspace, even if the global discrimination criteria is maximized.

An alternative strategy to DR has been recently proposed, i.e., the subspace models [17]. These models assume that each class is located in a specific subspace and consider the original space without DR for the processing. For instance, the Probabilistic Principal Component Analysis (PPCA) [18] assumes that the classes are normally distributed in a lower dimensional subspace and are linearly embedded in the original subspace with additive white noise. Such models exploit the *empty space* property of HD data without discarding any dimension of the data [19, 20]. A general subspace model that encompasses several other models is the High Dimensional Discriminant Analysis (HDDA) model, proposed by Bouveyron *et al.* [21, 22].

Conversely, kernel based methods do not reduce the dimensionality but rather work with the full HD data [23]. These discriminative methods are known to be more robust to size of the dimensionality

Table 1: Summary of HD spaces properties.

High Dimensional Spaces	
<i>Curse</i>	<i>Blessing</i>
Poor statistical estimation	Emptiness
Concentration of measure	Class separability

than conventional generative methods. However, local kernel methods are sensitive to the size of the dimensionality [24]. A kernel method is said to be local if the decision function value for a new sample depends on the neighbors of that sample in the training set. Since in HD data the neighborhood of a sample is mostly empty, such local methods are negatively impacted by the dimension. For instance, SVM with Gaussian kernel

$$k_g(\mathbf{x}, \mathbf{z}) = \exp\left(-\frac{\|\mathbf{x} - \mathbf{z}\|^2}{2\sigma^2}\right) \quad (1)$$

is such a local kernel method.

In this paper, it is proposed to use subspace models to construct a kernel adapted to high dimensional data. The chosen approach for including subspace models in a kernel function is to consider the Mahalanobis distance, d_{Σ_c} , between two samples for a given class, c , with covariance matrix, Σ_c :

$$d_{\Sigma_c}(\mathbf{x}, \mathbf{z}) = \sqrt{(\mathbf{x} - \mathbf{z})^t \Sigma_c^{-1} (\mathbf{x} - \mathbf{z})}.$$

Previous works on the Mahalanobis kernel [25, 26, 27, 28] were limited by the effect of dimensionality on the matrix inversion. In [25], the covariance matrix was computed on the whole training set. The associated implicit model is that the classes share the same covariance matrix, which is not true in practice. Diagonal and full covariance matrices were investigated in [26] for the purpose of classification and in [27] for the purpose of regression. However, in a similar way, the covariance matrix was computed for all the training samples. Computing the covariance matrix for the Mahalanobis distance with all the training samples is equivalent to project the data on all the principal components, scale the variance to one, and then applying the Euclidean distance. By doing so, classes could overlap more than in the original input space and the discrimination between them would be decreased.

In this work, the HDDA model is used for the definition of a class specific covariance matrix adapted for HD data. The specific signal and noise subspaces are estimated for each considered class, ensuring a parsimonious characterization of the classes. Following the HDDA model it is then possible to derive an explicit formulation of the inverse of the covariance matrix, without any regularization or dimension reduction. The parsimonious Mahalanobis kernel is constructed by substituting the Euclidean distance with the Mahalanobis distance computed using the HDDA model. It is proposed in this work to define several hyperparameters in the kernel to control the influence of the signal and noise subspaces in the classification process. These hyperparameters are optimized during the training process by the minimization of the so-called *radius margin* bound of the SVM classifier. Compared to the previous works on the Mahalanobis kernel for HD data, the proposed method allows the use of a more complex model, a separate covariance matrix per class, with higher efficiency in terms of accuracy. The remainder of the paper is organized as follows. The subspace model and the proposed kernel are discussed in Section 2. The problem of selecting the hyperparameters for classification with SVM is addressed in Section 3. The Section 4 details the estimation of the size of the signal subspace. Results on simulated and real high dimensional data are reported in Section 5. Conclusions and perspectives conclude the paper.

2 Regularized Mahalanobis Kernel

2.1 Review of HDDA model

The most general HDDA sub-model, refers to $[a_{ij}b_iQ_id_i]$ in [21, 22], is used in this work. Here, we will review the HDDA model but restricted to the problem of the covariance matrix inversion. However HDDA was originally proposed for classification or clustering with Gaussian mixture model. Interested readers can find a detailed presentation of HDDA in [21, 22].

In subspace models, it is assumed that the data from each class are clustered in the vector space. This cluster does not need to have an elliptic shape but it is generally assumed that the data follow a Gaussian distribution. The covariance matrix of the class c can be written through its eigenvalue decomposition:

$$\Sigma_c = \mathbf{Q}_c \mathbf{\Lambda}_c \mathbf{Q}_c^t$$

where $\mathbf{\Lambda}_c$ is the diagonal matrix of eigenvalues λ_{ci} , $i \in \{1, \dots, d\}$, of Σ_c and \mathbf{Q}_c is the matrix that contains the corresponding eigenvectors \mathbf{q}_{ci} . The HDDA model assumes the p_c first eigenvalues are different and the remaining $d - p_c$ eigenvalues are identical. The model is similar to PPCA, but more general in the sense that additional sub-models can be defined. In particular, the intrinsic dimension p_c are not constrained in HDDA whereas there are assumed to be equal for each class in PPCA.

Under the HDDA framework, the covariance matrix has the following expression:

$$\Sigma_c = \sum_{i=1}^{p_c} \lambda_{ci} \mathbf{q}_{ci} \mathbf{q}_{ci}^t + b_c \sum_{i=p_c+1}^d \mathbf{q}_{ci} \mathbf{q}_{ci}^t$$

where the last $d - p_c$ eigenvalue are equal to b_c . The inverse can be computed explicitly by

$$\Sigma_c^{-1} = \underbrace{\sum_{i=1}^{p_c} \frac{1}{\lambda_{ci}} \mathbf{q}_{ci} \mathbf{q}_{ci}^t}_{\mathcal{A}_c} + \frac{1}{b_c} \underbrace{\sum_{i=p_c+1}^d \mathbf{q}_{ci} \mathbf{q}_{ci}^t}_{\bar{\mathcal{A}}_c}.$$

This statistical model can be understood equivalently by a geometrical assumption: For each class, the data belong to a cluster that lives in a lower dimensional space \mathcal{A}_c , namely the signal subspace. The original input space can be decomposed as $\mathbb{R}^d = \mathcal{A}_c \oplus \bar{\mathcal{A}}_c$ (by construction $\bar{\mathcal{A}}$ is the noise subspace which contains only white noise). Figure 1 gives an illustration of that in \mathbb{R}^3 .

Using $\mathbf{I} = \sum_{i=1}^d \mathbf{q}_{ci} \mathbf{q}_{ci}^t$, \mathbf{I} being the identity matrix, the inverse can be finally written as

$$\Sigma_c^{-1} = \sum_{i=1}^{p_c} \left(\frac{1}{\lambda_{ci}} - \frac{1}{b_c} \right) \mathbf{q}_{ci} \mathbf{q}_{ci}^t + \frac{1}{b_c} \mathbf{I}. \quad (2)$$

Standard likelihood maximization shows that the parameters $(\lambda_{ci}, \mathbf{q}_{ci})_{i=1, \dots, p_c}$ and b_c can be computed from the sample covariance matrix [21]:

$$\hat{\Sigma}_c = \frac{1}{n_c} \sum_{i=1}^{n_c} (\mathbf{x}_i - \bar{\mathbf{x}}_c) (\mathbf{x}_i - \bar{\mathbf{x}}_c)^t$$

where $\bar{\mathbf{x}}_c$ is the sample mean for class c and n_c the number of samples of the class. λ_{ci} is estimated by the i first eigenvalue of $\hat{\Sigma}_c$, \mathbf{q}_{ci} by the corresponding eigenvector and $\hat{b}_c = (\text{trace}(\hat{\Sigma}_c) - \sum_{i=1}^{\hat{p}_c} \hat{\lambda}_{ci}) / (d - \hat{p}_c)$ (the estimation of the dimension p_c of the subspace is discussed later). The last $d - p_c$ eigenvalues and their corresponding eigenvectors are not needed for the computation of the inverse in (2).

The major advantage of such a model is that it reduces drastically the number of parameters to estimate for computing the inverse matrix. Indeed, with the full covariance matrix, $d(d+3)/2$ parameters are to

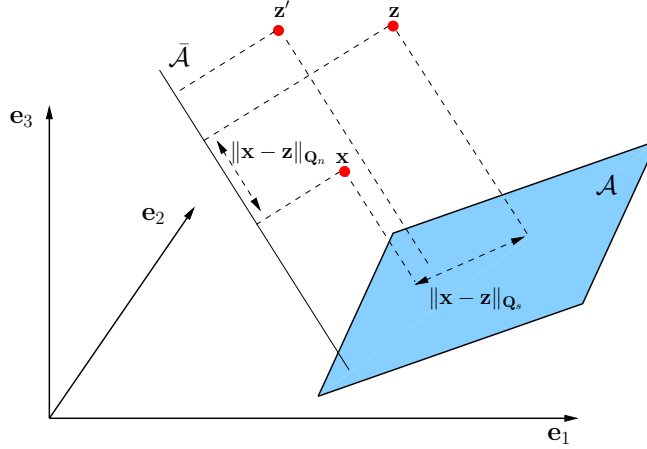


Figure 1: Cluster-based model. The distance between \mathbf{x} and \mathbf{z} is computed both in the signal subspace and in the noise subspace. Note that in this example $\dim(\bar{\mathcal{A}}_c) < \dim(\mathcal{A}_c)$, but for real data it is usually the opposite. $\|\cdot\|_{\mathbf{Q}_s}$ is the dot product in \mathcal{A}_c and $\|\cdot\|_{\mathbf{Q}_n}$ is the dot product in $\bar{\mathcal{A}}_c$.

be estimated. With the HDDA model, only $d(p_c + 1) + 1 - p_c(p_c - 1)/2$ parameters are to be estimated. For instance, if $d = 100$ and $p_c = 10$, 5150 parameters are needed for the full covariance and only 1056 for the HDDA model. Furthermore, the stability is improved since the smallest eigenvalues of the covariance matrix and their corresponding eigenvectors, which are difficult to compute accurately, are not used in (2). Finally, using the HDDA model, the square Mahalanobis distance for class c is approximated by

$$d_{\hat{\Sigma}_c}^2(\mathbf{x}, \mathbf{z}) = \sum_{i=1}^{\hat{p}_c} \left(\frac{1}{\hat{\lambda}_{ci}} - \frac{1}{\hat{b}_c} \right) \|\hat{\mathbf{q}}_{ci}^t(\mathbf{x} - \mathbf{z})\|^2 + \frac{\|\mathbf{x} - \mathbf{z}\|^2}{\hat{b}_c}. \quad (3)$$

This approach relies on the analysis of the empirical covariance matrix, as with PCA. But instead of keeping only significant eigenvalues, (3) considers all the original space, without discarding any dimension. This has two main theoretical advantages over the conventional PCA:

1. Two samples may be close in the signal subspace but far apart in the original space, which is a problem for classification tasks. It can be handled by considering the noise subspace together with the signal subspace. Consider for instance \mathbf{z} , \mathbf{z}' and \mathbf{x} in Figure 1. In \mathcal{A} , \mathbf{z}' seems closer to \mathbf{x} than \mathbf{z} , while it is not as it can be seen by adding $\bar{\mathcal{A}}$ in the distance computation.
2. An accurate estimation of the signal subspace size \hat{p}_c is necessary if PCA is applied: The worst scenario being $\hat{p}_c \ll p_c$, *i.e.*, relevant eigenvectors are discarded. By considering both the signal and the noise subspaces, the method becomes less sensitive to \hat{p}_c . Even in the worst case scenario, the eigenvectors are still considered.

2.2 Mahalanobis Kernel

The regularized Mahalanobis kernel for class c is constructed by substituting (3) to the Euclidean distance in the Gaussian kernel (1) and switching eigenvalues $(\hat{\lambda}_{ci}, \hat{b}_c)$ to hyperparameters $(\sigma_{ci}^2, \sigma_{c\hat{p}_c+1}^2)$ that are optimized during the training step:

$$k_m(\mathbf{x}, \mathbf{z}|c) = \exp \left(-\frac{1}{2} \left(\sum_{i=1}^{\hat{p}_c} \frac{\|\hat{\mathbf{q}}_{ci}^t(\mathbf{x} - \mathbf{z})\|^2}{\sigma_{ci}^2} + \frac{\|\mathbf{x} - \mathbf{z}\|^2}{\sigma_{c\hat{p}_c+1}^2} \right) \right) \quad (4)$$

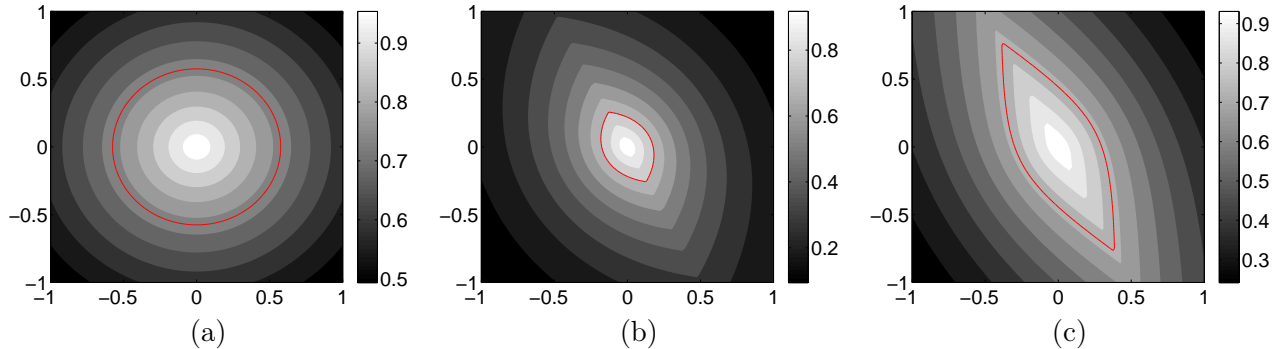


Figure 2: Values of the kernel function $k(\mathbf{0}, \mathbf{x})$ with $\mathbf{0} = [0, 0]$ and $\mathbf{x} \in [-1, 1] \times [-1, 1]$. The red line represents the contour line for the value 0.75. (a) is the Gaussian kernel, (b) is the kernel (4) with $\sigma_1^2 = \sigma_2^2 = 0.5$, (c) is the kernel (4) with $\sigma_1^2 = 1.5$ and $\sigma_2^2 = 0.5$. The covariance matrix used was $\begin{bmatrix} 0.6 & -0.2 \\ -0.2 & 0.6 \end{bmatrix}$ and the signal subspace was of dimension 1, spanned by the first eigenvector of the covariance matrix.

where σ_{ci} , $i \in \{1, \dots, \hat{p}_c + 1\}$ are the hyperparameters of the kernel. As described in Section 3, these parameters are tuned during the training step. The hyperparameters have been introduced for the following reason. It is known that the principal directions are not optimal for classification since they do not maximize any discrimination criterion. However, they still span a subspace where there are variations in the data of the considered class. The hyperparameters σ_{ci} allow to control which directions are more relevant (or discriminative) for the classification process: The feature space is modified during the training process to ensure a better discrimination between samples.

It is interesting to note that the regularized Mahalanobis kernel can be expressed as the product of Gaussian kernels:

$$k_m(\mathbf{x}, \mathbf{z}|c) = k_g(\mathbf{x}, \mathbf{z}) \times \prod_{i=1}^{\hat{p}_c} k_g(\hat{\mathbf{q}}_{ci}^t \mathbf{x}, \hat{\mathbf{q}}_{ci}^t \mathbf{z}). \quad (5)$$

The feature space induced by the kernel and the influence of the hyperparameters is analyzed in the next section.

2.3 Geometry of the induced feature space

Working with a kernel function is equivalent to work with samples mapped onto a feature space \mathcal{H} , where the dot product is equivalent to the kernel evaluation in the input space [23, 29]:

$$k(\mathbf{x}, \mathbf{z}) = \langle \phi(\mathbf{x}), \phi(\mathbf{z}) \rangle_{\mathcal{H}},$$

ϕ being the feature map. Under some weak conditions, the projected samples in the feature space live on a Riemannian manifold [30, 31]. The metric tensor is

$$g_{ij}(\mathbf{x}) = \left. \frac{\partial^2 k(\mathbf{x}, \mathbf{z})}{\partial x_i \partial z_j} \right|_{\mathbf{z}=\mathbf{x}} \quad (6)$$

which is, for the Gaussian kernel, $g_{ij}(\mathbf{x}) = \sigma^{-2} \delta_{ij}$ with $\delta_{ij} = 1$ if $i = j$ and 0 otherwise. This metric stretches or compresses the Euclidean distance between \mathbf{x} and \mathbf{z} by a factor σ^{-2} . Each variable is assumed equally relevant for the given task, *e.g.*, classification or regression.

For the kernel (4) the metric tensor is:

$$g_{ij}(\mathbf{x}|c) = \sum_{l=1}^{\hat{p}_c} \frac{q_{cli} q_{clj}}{\sigma_{cl}^2} + \frac{\delta_{ij}}{\sigma_{c\hat{p}_c+1}^2} \quad (7)$$

with q_{cli} the i^{th} element of \mathbf{q}_{cl} . The distance between two samples is stretched (if $\sigma_{cl}^2 \geq 1$) or compressed (if $\sigma_{cl}^2 \leq 1$) along the \hat{p}_c first principal components of class c (first term of the right part of the equation) and along the original components (last term of the equation). In other words, each principal component is weighted according to its relevance for the processing.

The analysis of the metric tensor exhibits the nature of the proposed kernel for a given class: It is a mixture of a Gaussian kernel on the original variables and a Gaussian kernel on the \hat{p}_c first principal components of the considered class. The hyperparameters σ_{cl} are tuned during the training process. This allows the optimization of the weight of each kernel. If $\sigma_{cl}^2 = +\infty, \forall l \in \{1, \dots, \hat{p}_c\}$, (4) reduces to the conventional Gaussian kernel. On the contrary, if $\sigma_{c\hat{p}_c+1}^2 = +\infty$, (4) reduces to the Gaussian kernel on the \hat{p}_c first principal components. Figure 2 shows the kernel values for different values of the hyperparameters. Note that opposite to the Gaussian kernel, the kernel in (4) is not isotropic.

The following section reviews the basics of SVM classifier and presents how the hyperparameters are computed.

3 L2-SVM and Radius margin bound optimization

Support vector machines (SVM) is a standard classification kernel methods [32]. It has shown to performs very well on several data sets from moderate dimension to high dimensional data [33, 34]. In the following section, the main results are presented but interested readers could see references [32, 35, 36] for further mathematical details about the SVM framework.

3.1 L2-Support Vector Machines

The L2-SVM is considered in this work rather than the conventional L1-SVM [35]: With L2-SVM it is possible to tune the hyperparameters automatically by optimizing the so called radius-margin bound [37]. The L2-SVM solves the conventional L1-SVM optimization problem with a quadratic penalization of errors [35]. Given a training set $\mathcal{S} = \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)\}$, $(\mathbf{x}_i, y_i) \in \mathbb{R}^d \times \{-1; 1\}$, the parameters $(\alpha_i)_{i=1}^n$ and b of the decision function f ,

$$f(\mathbf{z}) = \sum_{i=1}^n \alpha_i k(\mathbf{x}_i, \mathbf{z}) + b,$$

are found by solving the convex optimization problem:

$$\begin{aligned} \max_{\alpha} g(\alpha) &= \sum_{i=1}^n \alpha_i - \frac{1}{2} \sum_{i,j=1}^n \alpha_i \alpha_j y_i y_j \tilde{k}(\mathbf{x}_i, \mathbf{x}_j) \\ \text{subject to} & \quad 0 \leq \alpha_i \text{ and } \sum_{i=1}^n \alpha_i y_i = 0 \end{aligned} \tag{8}$$

where $\tilde{k}(\mathbf{x}_i, \mathbf{x}_j) = k(\mathbf{x}_i, \mathbf{x}_j) + C^{-1} \delta_{ij}$ with k the kernel function and C a positive hyperparameter that is used to penalize the training errors.

An estimate of the generalization errors is given by an upper bound on the number of errors of the leave-one-out procedure, the radius-margin bound \mathcal{T} [36]:

$$\mathcal{T}(\mathbf{p}) := \mathcal{R}^2 \mathcal{M}^2. \tag{9}$$

\mathcal{R}^2 is the radius of the smallest hypersphere that contains all $\phi(\mathbf{x}_i)$, \mathcal{M}^2 is the margin of the classifier, it is given by the optimal objective function of (8), and \mathbf{p} are the hyperparameters. In our setting

$\mathbf{p} = [\sigma_1^2, \dots, \sigma_{\hat{p}_c+1}^2, C]$. \mathcal{R}^2 is obtained by the optimal objective function of the following constraint optimization problem [36]:

$$\begin{aligned} \max_{\beta} g'(\beta) &= \sum_{i=1}^n \beta_i \tilde{k}(\mathbf{x}_i, \mathbf{x}_i) - \sum_{i,j=1}^n \beta_i \beta_j \tilde{k}(\mathbf{x}_i, \mathbf{x}_j) \\ \text{subject to} & \quad 0 \leq \beta_i \text{ and } \sum_{i=1}^n \beta_i = 1. \end{aligned} \quad (10)$$

Since both \mathcal{R}^2 and \mathcal{M}^2 depend on \mathbf{p} , it is possible to optimize \mathcal{T} to set the hyperparameters. Chapelle *et al.* [37], followed later by S. S. Keerthi [38], have proposed an algorithm based on gradient optimization method. It is discussed in the following section.

With the proposed kernel, if two classes i and j are considered, the classifier for “ i vs j ” is not the same as the classifier for “ j vs i ” since the kernel function is specific to the classes i and j , respectively. Indeed, for a multiclass problem, the “one vs one” approach must not be used and the “one vs all” approach should be preferred [39].

3.2 Radius-margin bound Optimization

Computing the gradient of \mathcal{T} requires the computation of the gradient of the following expressions¹

$$\mathcal{M}^2 = 2 \sum_{i=1}^n \tilde{\alpha}_i - \sum_{i,j=1}^n \tilde{\alpha}_i \tilde{\alpha}_j y_i y_j \tilde{k}(\mathbf{x}_i, \mathbf{x}_j) \quad (11)$$

and of

$$\mathcal{R}^2 = \sum_{i=1}^n \tilde{\beta}_i \tilde{k}(\mathbf{x}_i, \mathbf{x}_i) - \sum_{i,j=1}^n \tilde{\beta}_i \tilde{\beta}_j \tilde{k}(\mathbf{x}_i, \mathbf{x}_j) \quad (12)$$

where $(\tilde{\alpha}_i)_{i=1}^n$ and $(\tilde{\beta}_i)_{i=1}^n$ are the optimal parameters of (8) and (10). The gradient of (11) depends on $\tilde{\alpha}_i$, which depends on \mathbf{p} (similar comments hold for (12)). Chapelle *et al.* have proven that since \mathcal{M}^2 and \mathcal{R}^2 are computed via an optimization problem, the gradients of $\tilde{\alpha}_i$ and $\tilde{\beta}_i$ do not enter into account in the computation of their gradients [37]. Hence, the gradient of (10) can be written as:

$$\nabla \mathcal{T} = \left[\frac{\partial \mathcal{T}}{\partial C}, \frac{\partial \mathcal{T}}{\partial \sigma_1^2}, \dots, \frac{\partial \mathcal{T}}{\partial \sigma_{\hat{p}_c+1}^2} \right]^t \quad (13)$$

with

$$\frac{\partial \mathcal{T}}{\partial C} = \frac{\partial \mathcal{R}^2}{\partial C} \mathcal{M}^2 + \mathcal{R}^2 \frac{\partial \mathcal{M}^2}{\partial C} \quad (14)$$

and

$$\frac{\partial \mathcal{T}}{\partial \sigma_\ell^2} = \frac{\partial \mathcal{R}^2}{\partial \sigma_\ell^2} \mathcal{M}^2 + \mathcal{R}^2 \frac{\partial \mathcal{M}^2}{\partial \sigma_\ell^2} \quad (15)$$

for $\ell \in \{1, \dots, \hat{p}_c + 1\}$. The derivatives of \mathcal{R}^2 are

$$\frac{\partial \mathcal{R}^2}{\partial C} = \frac{1}{C^2} \sum_{i=1}^n \tilde{\beta}_i (\tilde{\beta}_i - 1) \quad (16)$$

$$\frac{\partial \mathcal{R}^2}{\partial \sigma_\ell^2} = - \sum_{i,j=1}^n \tilde{\beta}_i \tilde{\beta}_j \frac{\partial \tilde{k}(\mathbf{x}_i, \mathbf{x}_j)}{\partial \sigma_\ell^2} \quad (17)$$

¹For simplicity, the parameter c of the kernel function is omitted, i.e., $\tilde{k}(\mathbf{x}, \mathbf{z}|c)$ is written as $\tilde{k}(\mathbf{x}, \mathbf{z})$.

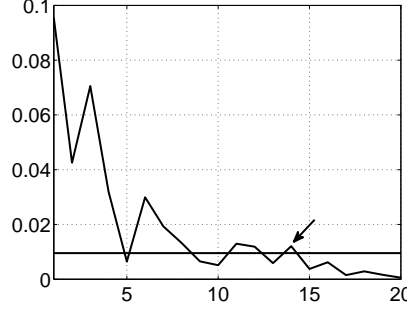


Figure 3: Scree test of Cattell. The threshold s is set to 10%. The estimated \hat{p}_c is 14. See Section 5.1 for a description of the data set. The horizontal axis represents the index i . The vertical axis represents the numerical difference between two consecutive eigenvalues, Δ_i . The curve represents the difference between two consecutive eigenvalues and the horizontal line represents 10 % of the highest difference. The arrow shows the estimated \hat{p}_c .

with

$$\frac{\partial \tilde{k}(\mathbf{x}_i, \mathbf{x}_j)}{\partial \sigma_\ell^2} = \begin{cases} \frac{\|\hat{\mathbf{Q}}_{cl}^t(\mathbf{x}_i - \mathbf{x}_j)\|^2}{\sigma_\ell^4} \tilde{k}(\mathbf{x}_i, \mathbf{x}_j) & \text{if } \ell \in \{1, \dots, \hat{p}_c\} \\ \frac{\|\mathbf{x}_i - \mathbf{x}_j\|^2}{\sigma_\ell^4} \tilde{k}(\mathbf{x}_i, \mathbf{x}_j) & \text{if } \ell = \hat{p}_c + 1. \end{cases} \quad (18)$$

The derivatives of \mathcal{M}^2 are

$$\frac{\partial \mathcal{M}^2}{\partial C} = \frac{1}{C^2} \sum_{i=1}^n \tilde{\alpha}_i \quad (19)$$

$$\frac{\partial \mathcal{M}^2}{\partial \sigma_\ell^2} = - \sum_{i,j=1}^n \tilde{\alpha}_i \tilde{\alpha}_j y_i y_j \frac{\partial \tilde{k}(\mathbf{x}_i, \mathbf{x}_j)}{\partial \sigma_\ell^2}. \quad (20)$$

Once the derivatives have been computed, the optimization of \mathcal{T} is done through a conventional gradient descent, following the framework in [37]. At each iteration t , the set of hyperparameters are updated as with a step proportional to the negative of the gradient of \mathcal{T} :

$$\mathbf{p}^{t+1} = \mathbf{p}^t - \gamma \nabla \mathcal{T}$$

where $\gamma \geq 0$ is a step size parameter. For implementation details, see [40].

4 Estimation of \hat{p}_c

The size of the signal subspace was estimated by the scree test of Cattell [41] using the same methodology as in [21]. The test consists in comparing the difference, Δ_i , between two consecutive eigenvalues λ_i and λ_{i+1} , $\Delta_i = \lambda_i - \lambda_{i+1}$. When the differences Δ_i are below a user-defined threshold s for all i , i.e., $\Delta_j < s, \forall j \in \{i, \dots, d-1\}$, \hat{p}_c is estimated as $\hat{p}_c = i$. In general the threshold is a percentage of the highest difference. Figure 3 shows an example on a simulated data set (see Section 5.1 for a description of the data). The correct value in that case is $p = 10$ but its estimate is $\hat{p} = 14$.

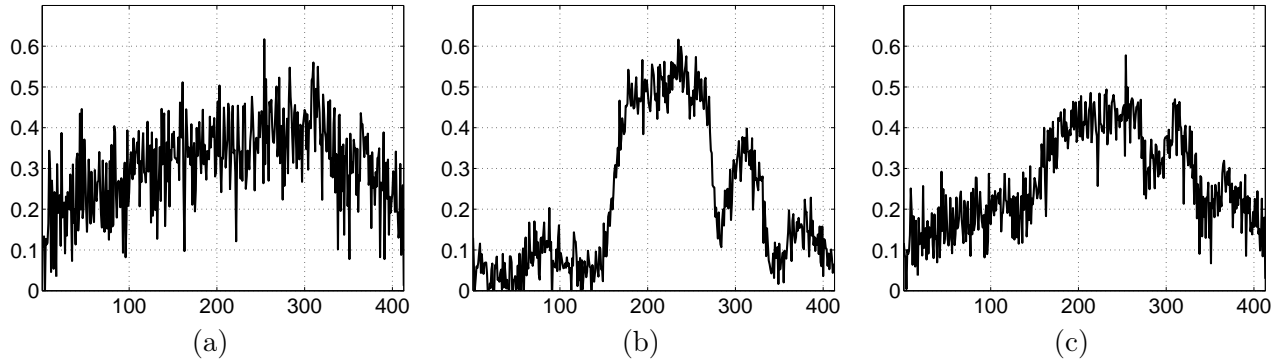


Figure 4: Simulated spectra: (a) \mathbf{s}_1 , (b) \mathbf{s}_2 and (c) \mathbf{x} . The horizontal axis is the variable and the vertical axis is the simulated reflectance. The parameters of the simulation are: $\alpha = [0.6, 0.4]$, $d = 413$, $p = 10$, $N_c = 2$ and $\text{SNR}=1$.

5 Experimental results

Classification results are presented in this section. Regarding the multiclass strategy, the results must be considered as individual binary classification problems: No fusion rules were applied. For instance, in Table 4, the results for the class ‘‘Asphalt’’ should be read as ‘‘Asphalt vs all’’. The reason for the use of this approach is the better interpretation of the results which is obtained because the results are not biased by the multiclass fusion strategy.

5.1 Classification of simulated data following HDDA model

In this section, the proposed kernel, namely the HDDA-Mahalanobis Kernel (HDDA-MK), is used with the SVM in classification and evaluated on simulated data. The performances in terms of classification accuracy have been compared to a SVM with a conventional Gaussian kernel on the original data and on the data projected on the first principal axis of the considered classes, called the PCA-Mahalanobis kernel. The main difference between the HDDA-MK and the PCA-Mahalanobis kernel is that the PCA-Mahalanobis kernel discards the noise subspace while the HDDA-MK also exploits the noise subspace in order to improve the class discrimination. As previously stated, Gaussian kernel and PCA-Mahalanobis kernel correspond to extreme cases of HDDA-MK: $\sigma_{\hat{p}}^2 = +\infty, \forall p \in \{1, \dots, \hat{p}_{N_c}\}$ or $\sigma_{\hat{p}+1}^2 = +\infty$.

Simulated data were constructed using a linear mixture model [42]:

$$\mathbf{x} = \sum_{i=1}^{N_c} \alpha_i \mathbf{s}_i + \mathbf{b} \quad (21)$$

where N_c is the number of classes, $y = j$ such as $\alpha_j = \max_i \alpha_i$, $\mathbf{b} \sim \mathcal{N}(\mathbf{0}, \varepsilon^2 \mathbf{I})$ and \mathbf{s}_i follows the HDDA model. The mean values of \mathbf{s}_i were extracted from a spectral library provided with the ENVI software used in hyperspectral imagery [43]. The number of spectral variables d was set to 413 and p_c was set to 10 for each class. The noise variance ε^2 was adjusted to get a $\text{SNR} = 1$. Three experiments were run for a different number of classes, i.e., 2, 3 and 4 classes, respectively. Figure 4 presents two simulated spectra and their linear mixture. The number of training samples was 1000 and the number of testing samples was 1500. The experiment was repeated 50 times for each configuration. The hyperparameters were estimated using the radius margin bound, for each classifier. Since no difference in terms of classification accuracies were observed, we only report the results of the ‘‘1 vs all’’ classifier.

5.1.1 Estimation of \hat{p}_c

With simulated data, it is possible to assess how the size of the intrinsic signal subspace is estimated. Figure 5 presents the boxplots of the estimations. After several tries, the threshold for the scree test was

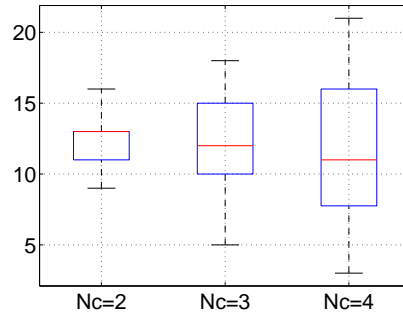


Figure 5: Boxplot of the estimation of \hat{p}_c for the three configurations. The vertical axis represents values of \hat{p}_c .

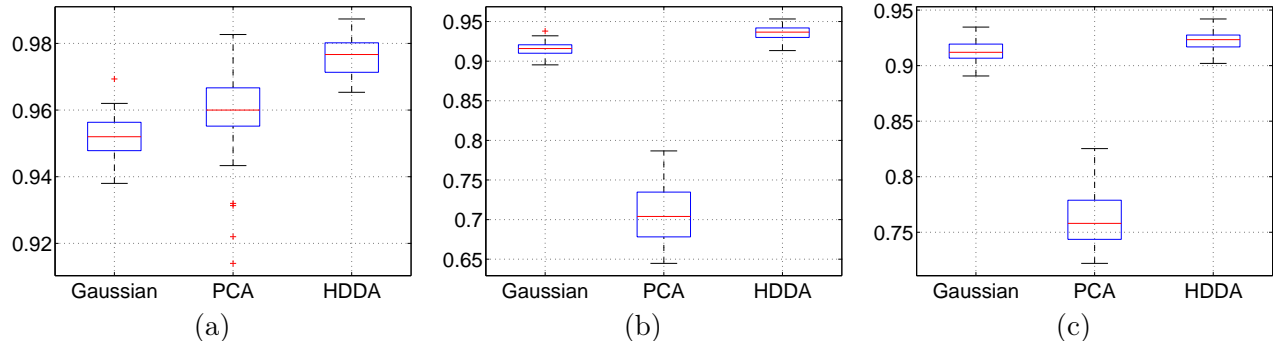


Figure 6: Boxplots of the classification accuracies for the three experiments. (a) $N_c = 2$, (b) $N_c = 3$ and (c) $N_c = 4$. The vertical axis represent the overall classification accuracies.

fixed to 10%. Fixing the threshold to a too high value would lead to underestimate p while a too low value would lead to drastically overestimate it.

From the figure, the scree test overestimates the parameter p , for each configuration. The variance of the estimation is larger when the number of classes is increased while the bias of the estimation decreases. However, the error in estimating p is not too important with regards to the original size of the data ($d = 413$). Furthermore, in previous work [44], another criterion (the BIC [45]) was used to estimate the correct dimension of the subspace where the data live. The BIC criterion showed poor results when the number of training samples for single class n_c was close to the dimension of the data ($d \approx n_c$). From the experiments, the scree test is more robust in such a situation.

5.1.2 Classification accuracies

The percentages of correct classification are reported in Figure 6. For the three experiments, the proposed kernel leads to the best results in terms of accuracies. Although \hat{p}_c was overestimated, it did not penalize the performances of the algorithm in terms of classification accuracies. For $N_c = 2$, the second best result is provided by the PCA-Mahalanobis kernel, while for $N_c = 3$ or 4 it is provided by the Gaussian kernel applied on the original data. For instance, for $N_c = 4$, the mean value of correct classification is 92.2% for the HDDA-Mahalanobis kernel, 91.3% for the conventional Gaussian kernel and only 76.3% for the PCA-Mahalanobis. The results confirm the poor generalization capability of the Mahalanobis kernel when dealing with high dimensional spaces. Although the conventional Gaussian kernel is less sensitive to the problem, the proposed kernel gives a significant improvement of the classification accuracy.

Table 2: Madelon data: Percentage of samples correctly classified.

	Class 1	Class 2	Mean
Gaussian	69.7	69.7	69.7
PCA-Mahalanobis	83.3	81.8	82.5
HDDA-Mahalanobis	84.1	83.8	83.9

Table 3: Arcene data: s is the threshold value in the scree test, (\hat{p}_1, \hat{p}_2) correspond to the estimated size of the signal subspace for each class. The number is the percentage of samples correctly classified.

	s	0.2	0.1	0.05	0.01	0.005	0.001
	(\hat{p}_1, \hat{p}_2)	(2,2)	(2,3)	(3,4)	(8,7)	(11,10)	(22,23)
PCA-Mahalanobis	Class 1	69.0	69.0	70.0	75.0	78.0	51.0
	Class 2	72.0	72.0	62.0	67.0	72.0	75.0
HDDA-Mahalanobis	Class 1	80.0	80.0	80.0	80.0	83.0	51.0
	Class 2	80.0	80.0	81.0	80.0	81.0	74.0

5.2 Classification of Madelon data

Madelon data set is a simulated data set used for the *NIPS Feature Selection Challenge*². It has 5 useful features, 15 redundant features and 480 random probes for a total of 500 features ($d = 500$). It is composed of two classes and the number of training samples is 2000 and the number of testing sample is 600. The threshold for the scree test has been set to 20%. The proposed kernel has been compared to the same kernels as in the previous section.

The classification results in terms of accuracies are reported in Table 2. The conventional Gaussian kernel performs badly on that data set with a global precision of 69,7% (random classifier would achieve 50%). The best accuracy is obtained for the proposed kernel with an average accuracy of 83.9%. The size \hat{p}_c of the signal subspace for the two classes was $\hat{p}_1 = 4$ and $\hat{p}_2 = 4$, respectively. The results obtained with the PCA-Mahalanobis kernel are worse than those obtained with the HDDA-Mahalanobis kernel. Thus, it confirms the pertinence to use both the signal and the noise subspace from the HDDA model in the classification.

5.3 Classification of Arcene data set

Arcene data set is data set used for the *NIPS Feature Selection Challenge*. It has 7000 real variables, 3000 random probes for a total of 10000 features ($d = 10000$). It is composed of two classes. The number of available training samples is 100 and the number of test samples is 100. Therefore, the number of training samples is very small in comparison with the number of variables. About 50% of the data are non zero.

The classification accuracies are reported in Table 3. The selected threshold for the scree test has been set to 0.5%. The results obtained for other values of the threshold are also reported for comparison. The Gaussian kernel achieves a global accuracy of 80%. It performs quite well in terms of classification accuracies related to the dimension of the data. The PCA-Mahalanobis kernel performs worst whatever the threshold value. For the HDDA-Mahalanobis kernel, for the highest value of s the results are equal to those obtain with the Gaussian kernel. Then a slight increased of the accuracy is observed for $s = 0.005$. When the size of the signal subspace is too large ($s = 0.0001$ and $(\hat{p}_1, \hat{p}_2) = (22, 23)$), too many hyperparameters have to be estimated and thus the classification accuracy becomes low.

²<http://www.nipsfsc.ecs.soton.ac.uk/>

Table 4: Classification accuracies for the different kernels in percentage of correctly classified samples. In the first column, the numbers in brackets represent the total number of training and testing samples for each class, respectively.

	\hat{p}_c	Gaussian	PCA-Mahalanobis	HDDA-Mahalanobis
Asphalt (548, 6631)	12	94.8	95.8	95.8
Meadow (540, 18649)	10	79.4	83.6	82.1
Gravel (392, 2099)	9	97.2	97.5	97.2
Tree (524, 3064)	14	94.3	98.2	98.2
Metal Sheet (265, 1345)	7	99.8	99.9	99.9
Bare Soil (532, 5029)	9	87.8	85.9	88.4
Bitumen (375, 1330)	21	98.8	98.7	99.0
Brick (514, 3682)	12	96.7	97.2	97.2
Shadow (231, 947)	14	99.9	99.9	99.9
Average class accuracy		94.3	95.2	95.3

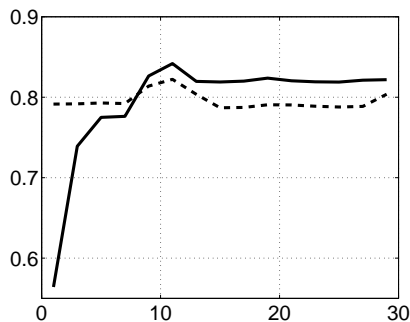


Figure 7: Classification accuracies as a function of the dimension of the signal subspace for the class *meadow*. The horizontal axis represents the value \hat{p}_c and the vertical axis represents the classification accuracies for the class *meadow*. Solid line corresponds to the PCA-Mahalanobis kernel and the dashed line corresponds to the HDDA-Mahalanobis kernel. The selected value with the scree test is 10.

5.4 Classification of real hyperspectral data

The data set considered in this experiment is the *University Area* of Pavia, Italy, acquired with the ROSIS-03 sensor. The image has 103 spectral variables, i.e., each pixel is represented by a vector with 103 features ($d=103$) [46]. Nine classes have been defined by photo-interpretation as seen in first column of Table 4. Here, the threshold was set to 0.01%, because of a very high value of the first principal component (mainly due to the albedo).

Classification results are reported in Table 4. The proposed kernel leads to an increase of the accuracy, compared to the conventional Gaussian kernel. However, for this data set, the PCA-Mahalanobis and HDDA-Mahalanobis kernel perform equally well in terms of accuracies, except for the classes *meadow* and *bare soil*.

To assess the influence of \hat{p}_c on the classification accuracies, the class *meadow* has been classified for a range of values of \hat{p}_c . The results are reported in the Figure 7 for the PCA-Mahalanobis and HDDA-Mahalanobis kernels. From the figure, the optimal \hat{p}_c is 11 which is close to the value selected with the scree test ($\hat{p}_c = 10$). The cumulative variance is 99.72% for $\hat{p}_c = 10$, 99.75% for $\hat{p}_c = 11$ and 99.77% for $\hat{p}_c = 12$. The proposed kernel is slightly influenced by the choice of \hat{p}_c for that class, compared to PCA-Mahalanobis kernel. In particular, the proposed kernel is robust if \hat{p}_c is underestimated. However, if \hat{p}_c is heavily overestimated, too many parameters would need to be estimated and it could degrade the training process.

Table 5: Processing time in seconds of the competitive methods for four different data sets.

	Arcene	Madelon	Asphalt	Meadow
Gaussian	2.4	31.9	113 .0	87.4
PCA-Mahalanobis	14.2	24.0	193.0	185.7
HDDA-Mahalanobis	14.5	80.8	94.2	195.3

5.5 Analysis of the processing time

To assess the computational load of the proposed method, the processing time was computed for four data sets. The two data sets from the *NIPS Feature Selection Challenge* were used and well as the two first classes (Asphalt and Method) of the hyperspectral data set. The results are reported in Table 5. The program runs under Matlab on a two cores 2.67GHz laptop.

For the Arcene data set, the Gaussian kernel shows the lowest computational time. The computation of the first eigenvalues/eigenvectors is demanding since the data have 10000 features. It requires about 14 seconds. The optimization of the hyperparameters is fast since a few number of training samples are available. For the Madelon data sets, the computation of the first eigenvalues/eigenvectors is fast, about 1.3 seconds while the optimization of the kernel hyperparameters is more demanding. For that data set, the HDDA-Mahalanobis is the slowest. Regarding the University data set, for the first two classes, the estimation of the first eigenvalues/eigenvectors is very fast, about 0.3 second. However, the estimation of the kernel hyperparameters is more demanding than with the two others data sets. For the Asphalt problem, the HDDA-Mahalanobis performs the fastest, while it is the slowest for the Meadow problem.

From the above results, it is difficult to point out a clear winner in terms of processing time. From a computational viewpoint, the optimization of the Gaussian kernel is less demanding than PCA- or HDDA-Mahalanobis kernels. However, since the optimization problem is a gradient descent on a non-convex function, a local minimum might be found sooner with one method and make it faster than the others. Nevertheless, the HDDA-Mahalanobis is usually more demanding in terms of processing time. Computational complexity of PCA- and HDDA-Mahalanobis can be assumed to be comparable in terms of processing time.

6 Conclusions

In this paper, a novel kernel adapted to high dimensional data has been proposed. The parsimonious Mahalanobis kernel is based on the *emptiness* property of HD spaces. For each class, the original input space is split into a signal subspace and a noise subspace. Using this assumption, the inversion of the covariance matrix in the Mahalanobis kernel can be accurately computed. The proposed kernel was tested in a SVM framework for the purpose of classification. Experimental results on four data sets demonstrate the potential of the proposed kernel. In each case, the classification accuracy increased compared to the conventional Gaussian kernel and for three cases the proposed kernel showed superior results to simply map the data on the first PCA axes. Consequently, for HD data the HDDA-Mahalanobis kernel should be preferred.

Regarding the computational load, the HDDA-Mahalanobis kernel is more demanding during the training process than the Gaussian kernel since more hyperparameters have to be estimated. Besides that, the HDDA-Mahalanobis kernel is efficient when the dimension of the input space is high. Hence, for moderate or small dimensions, conventional kernels should be preferred.

In the article, only classification is investigated. But other processing could also have been considered, e.g., regression [27]. The critical point is to be able to tune the hyperparameter efficiently, which it is feasible for regression. However, the actual optimization of the hyperparameters is demanding in terms of computations and it is sensitive to local minima. Therefore, a different strategy must be studied.

Perspectives of this work concern the development of new kernels using the HDDA-model. For instance,

it is possible to define a new dot product for the conventional polynomial kernel:

$$k(\mathbf{x}, \mathbf{z}|c) = (\mathbf{x}^t \Sigma_c^{-1} \mathbf{z} + 1)^t. \quad (22)$$

Furthermore, a mixture of kernels using the HDDA could be used. From (5), the HDDA-Mahalanobis kernel can be considered as a product of several kernels. In the future, we will investigate the summation of kernels.

Finally, a free-parameter alternative to the scree test for estimation of the intrinsic dimension must be addressed. For instance, an maximum likelihood estimator for HDDA exists and must be investigated [47].

Acknowledgment

The authors would like to thank the IAPR - TC7 for providing the data and Prof. Paolo Gamba and Prof. Fabio Dell'Acqua of the University of Pavia, Italy, for providing reference data.

References

- [1] H.-P. Kriegel, P. Kröger, and A. Zimek, “Clustering high-dimensional data: A survey on subspace clustering, pattern-based clustering, and correlation clustering,” *ACM Trans. Knowl. Discov. Data*, vol. 3, Mar. 2009.
- [2] D. L. Donoho, “High-dimensional data analysis: the curses and blessing of dimensionality,” in *AMS Mathematical challenges of the 21st century*, 2000.
- [3] A. Plaza, P. Martinez, R. Perez, and J. Plaza, “A new approach to mixed pixel classification of hyperspectral imagery based on extended morphological profiles,” *Pattern Recognition*, vol. 37, no. 6, pp. 1097 – 1116, 2004.
- [4] M. G. Kendall, *A course in the geometry of n-dimensions*. New York: Dover Publication, 1961.
- [5] L. Jimenez and D. A. Landgrebe, “Supervised classification in high dimensional space: geometrical, statistical and asymptotical properties of multivariate data,” *IEEE Trans. Syst., Man, Cybern. B*, vol. 28, pp. 39–54, Feb. 1998.
- [6] J. Herault, A. Guerin-Dugue, and P. Villemain, “Searching for the embedded manifolds in high dimensional data, problems and unsolved questions,” in *European Symposium on Artificial Neural Networks*, pp. 1–12, 2002.
- [7] C. Aggarwal, A. Hinneburg, and D. Keim, “On the surprising behavior of distance metrics in high dimensional space,” in *Database Theory — ICDT 2001*, vol. 1973 of *Lecture Notes in Computer Science*, pp. 420–434, Springer Berlin / Heidelberg, 2001.
- [8] D. Francois, V. Wertz, and M. Verleysen, “The concentration of fractional distances,” *IEEE Trans. Knowl. Data Eng.*, vol. 19, pp. 873–886, July 2007.
- [9] L. Yang, “Distance-preserving projection of high-dimensional data for nonlinear dimensionality reduction,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 26, pp. 1243 –1246, Sept. 2004.
- [10] M. Verleysen, D. Francois, G. Simon, and V. Wertz, “On the effects of dimensionality on data analysis with neural networks,” in *Artificial Neural Nets Problem Solving Methods*, vol. 2687 of *Lecture Notes in Computer Science*, pp. 1044–1044, Springer Berlin / Heidelberg, 2003.
- [11] M. Verleysen, *Learning high dimensional data*. V. Piuri eds, IOS Press, Amsterdam (The Netherlands), 2003.

- [12] R. E. Bellman, *Adaptive control processes - A guided tour*. Princeton, New Jersey, U.S.A., 1961.
- [13] D. A. Landgrebe, *Signal Theory Methods in Multispectral Remote Sensing*. New Jersey: John Wiley and Sons, 2003.
- [14] C. J. C. Burges, “Dimension reduction: A guided tour,” *Foundations and Trends in Machine Learning*, vol. 2, no. 4, pp. 275–365, 2010.
- [15] J. A. Lee and M. Verleysen, *Nonlinear Dimensionality Reduction*. Springer, 2007.
- [16] I. Guyon and A. Elisseeff, “An introduction to variable and feature selection,” *J. Mach. Learn. Res.*, vol. 3, pp. 1157–1182, Mar. 2003.
- [17] L. Parsons, E. Haque, and H. Liu, “Subspace clustering for high dimensional data: a review,” *SIGKDD Explor. Newsl.*, vol. 6, pp. 90–105, June 2004.
- [18] M. E. Tipping and C. M. Bishop, “Probabilistic principal component analysis,” *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, vol. 61, no. 3, pp. 611–622, 1999.
- [19] M. E. Tipping and C. M. Bishop, “Mixtures of probabilistic principal component analyzers,” *Neural Computation*, vol. 11, no. 2, pp. 443–482, 1999.
- [20] J. Baek, G. J. McLachlan, and L. K. Flack, “Mixtures of factor analyzers with common factor loadings: Applications to the clustering and visualization of high-dimensional data,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, pp. 1298–1309, July 2010.
- [21] C. Bouveyron, S. Girard, and C. Schmid, “High-Dimensional Discriminant Analysis,” *Communication in Statistics- Theory and Methods / Communications in Statistics Theory and Methods*, vol. 36, p. 2607–2623, Jan. 2007.
- [22] C. Bouveyron, S. Girard, and C. Schmid, “High-Dimensional Data Clustering,” *Computational Statistics and Data Analysis*, vol. 52, no. 1, pp. 502–519, 2007.
- [23] T. Hofmann, B. Schölkopf, and A. J. Smola, “Kernel methods in machine learning,” *Annals of Statistics*, vol. 36, no. 3, pp. 1171–1220, 2008.
- [24] Y. Bengio, O. Delalleau, and N. Le Roux, “The curse of highly variable functions for local kernel machines,” in *Advances in Neural Information Processing Systems 18 (NIPS’05)* (Y. Weiss, B. Schölkopf, and J. Platt, eds.), pp. 107–114, MIT Press, 2006.
- [25] G. Camps-Valls, A. Rodrigo-Gonzalez, J. Muñoz-Mari, L. Gomez-Chova, and J. Calpe-Maravilla, “Hyperspectral image classification with Mahalanobis relevance vector machines,” in *Geoscience and Remote Sensing Symposium, 2007. IGARSS 2007. IEEE International*, pp. 3802–3805, July 2007.
- [26] S. Abe, “Training of support vector machines with Mahalanobis kernels,” in *Artificial Neural Networks: Formal Models and Their Applications - ICANN 2005*, Lecture Notes in Computer Science, pp. 571–576, Springer Berlin / Heidelberg, 2005.
- [27] Y. Kamada and S. Abe, “Support vector regression using mahalanobis kernels,” in *Artificial Neural Networks in Pattern Recognition* (F. Schwenker and S. Marinai, eds.), vol. 4087 of *Lecture Notes in Computer Science*, pp. 144–152, Springer Berlin / Heidelberg, 2006.
- [28] D. Wang, D. Yeung, and E. Tsang, “Weighted mahalanobis distance kernels for support vector machines,” *Neural Networks, IEEE Transactions on*, vol. 18, pp. 1453–1462, sept. 2007.

- [29] M. Filippone, F. Camastra, F. Masulli, and S. Rovetta, "A survey of kernel and spectral methods for clustering," *Pattern Recognition*, vol. 41, no. 1, pp. 176 – 190, 2008.
- [30] B. Scholkopf, C. Burges, and A. Smola, *Geometry and Invariance in Kernel Based Methods In Advances in Kernel Methods - Support Vector Learning*. MIT Press, 1998.
- [31] P. Williams, S. Li, J. Feng, and S. Wu, "A geometrical method to improve performance of the vector machine," *IEEE Trans. Neural Netw.*, vol. 18, pp. 942–947, May 2007.
- [32] C. J. Burges, "A tutorial on support vector machines for pattern recognition," *Data Mining and Knowledge Discovery*, vol. 2, pp. 121–167, 1998.
- [33] S.-K. Kim, Y. J. Park, K.-A. Toh, and S. Lee, "Svm-based feature extraction for face recognition," *Pattern Recognition*, vol. 43, no. 8, pp. 2871 – 2881, 2010.
- [34] M. Fauvel, J. Chanussot, and J. A. Benediktsson, "A spatial-spectral kernel-based approach for the classification of remote-sensing images," *Pattern Recogn.*, vol. 45, pp. 381–392, Jan. 2012.
- [35] N. Cristianini and J. Shawe-Taylor, *An Introduction to Support Vector Machines and Other Kernel-based Learning Methods*. Cambridge University Press, 2000.
- [36] V. Vapnik, *The Nature of Statistical Learning Theory, Second Edition*. Springer, New York, 1999.
- [37] O. Chapelle, V. Vapnik, O. Bousquet, and S. Mukherjee, "Choosing multiple parameters for support vector machines," *Machine Learning*, vol. 46, pp. 131–159, 2002.
- [38] S. S. Keerthi, "Efficient tuning of svm hyperparameters using radius/margin bound and iterative algorithms," *IEEE Trans. Neural Netw.*, vol. 13, pp. 1225 – 1229, Sept. 2002.
- [39] C.-W. H. and C.-J. L., "A comparison of methods for multiclass support vector machines," *IEEE Trans. Neural Netw.*, vol. 13, pp. 415 –425, Mar. 2002.
- [40] O. Chapelle, "Selection of kernel parameters." <http://olivier.chapelle.cc/ams/>, 2002.
- [41] R. B. Cattell, "The Scree Test For The Number Of Factors," *Multivariate Behavioral Research*, vol. 1, no. 2, pp. 245–276, 1966.
- [42] N. Keshava and J. Mustard, "Spectral unmixing," *Signal Processing Magazine, IEEE*, vol. 19, pp. 44 –57, Jan. 2002.
- [43] Research Systems Inc. (RSI), Boulder, CO, *ENVI User's Guide*, -.
- [44] M. Fauvel, A. Villa, J. Chanussot, and J. Benediktsson, "Mahalanobis kernel for the classification of hyperspectral images," in *Geoscience and Remote Sensing Symposium (IGARSS), 2010 IEEE International*, pp. 3724 –3727, July 2010.
- [45] G. Schwarz, "Estimating the dimension of a model," *The annals of Statistics*, vol. 6, no. 2, pp. 461–464, 1978.
- [46] C. H. Chen and P.-G. P. Ho, "Statistical pattern recognition in remote sensing," *Pattern Recognition*, vol. 41, no. 9, pp. 2731 – 2741, 2008.
- [47] C. Bouveyron, G. Celeux, and S. Girard, "Intrinsic dimension estimation by maximum likelihood in isotropic probabilistic pca," *Pattern Recognition Letters*, vol. 32, no. 14, pp. 1706 – 1713, 2011.