Multilayer Perceptron with Functional Inputs: an Inverse Regression Approach

Louis Ferré, Nathalie Villa

Multi-layer perceptron with functional inputs: an inverse regression approach

Louis Ferré Nathalie Villa

Équipe GRIMM, Université Toulouse Le Mirail, France

Abstract

Functional data analysis is a growing research field since more and more pratical applications involve functional data. In this paper, we focus on the problem of regression and classification with functional predictors: the model suggested combines an efficient dimension reduction procedure (functional SIR, first introduced by Ferré & Yao (2003)), for which we give a regularized version, with the accuracy of a neural network. Some consistency results are given and the method is successfully confronted to real life data.

Keywords: classification, dimension reduction, functional data analysis, multi-layer perceptron, prediction.

1 Introduction

Functional regression is now a very important part of statistics as functional variables occur frequently in practical applications. We present two examples that take place in functional data analysis (FDA). First, a regression problem where the regressor are curves is introduced (see Figure 1): the Tecator data problem (available at `http://lib.stat.cmu.edu/datasets/tecator`) consists in predicting the fat content of pieces of meat from a near infrared absorbance spectrum. This data set first appears in Borggaard & Thodberg (1992) and has also already been studied, among others, in Thodberg (1996), Ferré & Yao (2003) (with an inverse regression approach) and Ferraty & Vieu (2003).

[Figure 1 about here.]

Secondly, in the phoneme data set, the data are log-periodograms of a 32 ms duration corresponding to recorded speakers and we expect to determine which one of the five phonemes, [sh] as in "she", [dcl] as in "dark", [iy] as in "she", [aa] as in "dark" and [ao] as in "water", corresponds to this recording (extracted from the TIMIT database and available at `http://www-stat.stanford.edu/~tibs/ElemStatLearn/data.html`). It has already been described by Hastie *et al.* (1995) and by Ferraty & Vieu (2003). Clearly, here, functional data is also involved but we face now a classification

problem. However, we will see that both - regression and classification - can be tackled via a common modelling.

An extensive review of the numerous studies developped for functional data analysis can be found in Ramsay & Silverman (1997) including regression and classification but also many factorial methods. A particularity of functional regression is that it often leads to ill-posed problems because of the infinite dimension of the feature space. Then original solutions have been introduced to overcome this problem: for example, Cardot *et al.* (1999) studied the functional linear regression. At the same time, Dauxois *et al.* (2001) and then Ferré & Yao (2003), Ferré & Yao (2005) have proposed a semi-parametric model for Hilbertian variables which corresponds to the functional version of Li's Sliced Inverse Regression, Li (1991).

On a classification point of view, many solutions have been proposed to overcome ill-posed functional problems including the popular penalization methods. Friedman (1989) presents the RDA model based on regularization and shrinkage while Hastie *et al.* (1994) and Hastie *et al.* (1995) propose a discriminant analysis penalized by smoothing functionals. On the other hand, it has been used for Canonical Correlation Analysis in Leurgans *et al.* (1993) and other examples of the regularization use are given in Ramsay & Silverman (1997).

Nonlinear methods for functional data analysis have also been developped: for instance, neural network models (Rossi & Conan-Guez (2005) for multilayer perceptrons and Rossi *et al.* (2004) for the SOM algorithm), k-nearest neighbour models (Biau *et al.* (2005)) or non parametric discrimination (Ferraty & Vieu (2003)).

In this paper, we propose a new way to achieve functional regression: the idea is to join the efficiency of a dimension reduction method using smoothing penalization, to the strong adaptability of a neural network which can provide highly non linear solutions even if the number of predictors is too large for classical nonparametric methods such as kernels smoothing. The functional SIR dimension reduction method is first presented in Section 2. For this penalized version, consistency results are given in Section 3. Section 4 discusses Neural Network and gives consistency results for the proposed model combining FSIR and Neural Networks (which will be called SIR-NNr). Section 5 is devoted to applications: Section 5.1 deals with the Tecator data set and Section 5.2 with the phoneme data set. In Appendix, we give a sketch of the proofs. All programs have been made using Matlab and are available on request.

2 Sliced Inverse Regression

Let Y be a real random variable and X be a multivariate variable assumed to have a fourth moment. To overcome the curse of dimensionality in the nonparametric regression of Y on X, Li (1991) introduced the Sliced Inverse Regression. He considers the following model

$$Y = f(a_1'X, a_2'X, \ldots, a_q'X, \epsilon),$$

where ϵ is centered and independent of X, f is an unknown function and $(a_j)_{j=1,\dots,q}$ are lineary independent vectors.

The space spanned by $(a_j)_{j=1,\dots,q}$ is called EDR (Effective Dimension Reduction) space. SIR deals with the estimation of this EDR space and the aim of sliced inverse regression is to estimate it by means of the eigenvectors of the matrix $Var(X)^{-1}Var(E(X|Y))$.

In the multivariate context, numerous works deal with SIR. In particular, methods have been proposed to improve SIR: different estimates of the covariance of the conditional mean have been built (in Hsing & Carroll (1992) and Zhu & Fang (1996)) while other methods have been proposed to estimate the EDR space (for example, PHD proposed by Li (1992), SAVE by Cook & Weisberg (1991) or MAVE by Xia *et al.* (2002)). The main interest of this model is that, once the EDR space is estimated, the estimation of f is obtained very easily with traditional techniques provided that q is not too large.

2.1 Functional SIR

Now consider a real random variable Y and X a random variable taking its values in $\mathcal{L}^2_{\mathcal{T}}$, the space of squared intregrable functions from a compact interval $\mathcal{T}$ into $\mathbb{R}$. With the usual inner product defined by, for all f, g in $\mathcal{L}^2_{\mathcal{T}}$, $\langle f, g\rangle = \int_{\mathcal{T}} f(t)g(t)dt$, $\mathcal{L}^2_{\mathcal{T}}$ is a Hilbert space. We will assume that the random variable X is centered, without loss of generality, and has a fourth moment. Then, the covariance operator of X exists and is defined by $\Gamma_X = E(X \otimes X)$ where $X \otimes X$ denotes the operator which associates to any f in $\mathcal{L}^2_{\mathcal{T}}$, $\langle f, X\rangle X$. We also get that $E(X|Y)$ and $\Gamma_{E(X|Y)} = Var(E(X|Y))$ exist. Ferré and Yao (2003) have proposed to investigate the following model for functional inverse regression:

$$Y = f(\langle X, a_1\rangle, \dots, \langle X, a_q\rangle, \epsilon) \tag{1}$$

where f is an unknown function, ϵ a random variable which is centered and independent of X and $(a_j)_{j=1,\dots,q}$ are lineary independent functions of $\mathcal{L}^2_{\mathcal{T}}$.

The crucial point of functional SIR is that, unlike the multivariate case, Γ_X^{-1} is not defined since we have to assume that Γ_X is a positive definite operator which implies that it is not invertible as defined from $\mathcal{L}^2_{\mathcal{T}}$ to $\mathcal{L}^2_{\mathcal{T}}$. However, if we call $(\delta_i)_{i=1,\dots,\infty}$ its sequence of eigenvalues and $(u_i)_{i=1,\dots,\infty}$ those of orthonormed eigenvectors, R_Γ the image of Γ_X and $R_\Gamma^{-1} = \left\{h \in \mathcal{L}^2_\tau : \exists f \in R_\Gamma, h = \sum_i (1/\delta_i)(u_i \otimes u_i)(f)\right\}$, Γ_X is a one-to-one mapping from R_Γ^{-1} to R_Γ whose inverse, called Γ_X^{-1}, is defined by $\Gamma_X^{-1} = \sum_i (1/\delta_i) u_i \otimes u_i$.

We focus on the estimation of the estimation of the EDR space spanned by the vectors $(a_j)_{j=1,\dots,q}$. Now, the key of the method comes from the following theorem:

Theorem 1 (Ferré & Yao (2003)). *Writing* $A = (\langle X, a_1\rangle, \dots, \langle X, a_q\rangle)^T$, *if*

(A1) *for all* u *in* $\mathcal{L}^2_{\mathcal{T}}$ *there exists* v *in* $\mathbb{R}^q$ *such that:* $E(\langle u, X\rangle|A) = v^T A$

then $E(X|Y)$ *belongs to the subspace spanned by* $\Gamma_X a_1, \dots, \Gamma_X a_q$.

Remark 1. Note that Cook & Weisberg (1991) show that elliptically distributed variables satisfy condition **(A1)** in the multidimensional context but this can be transposed in infinite dimensional Hilbert spaces (see Yao (2001)).

By using the result of Dauxois *et al.* (2001), a consequence of Theorem 1 is that the EDR subspace contains the Γ_X-orthonormed eigenvectors of $\Gamma_X^{-1}\Gamma_{E(X|Y)}$ associated with the q positive eigenvalues. Then, in the following, $(a_j)_{j=1,\dots,q}$ will denote those eigenvectors. This is the generalization of Li (1991) on SIR to infinite dimensional case.

A basis of the EDR space is thus given by the eigenvector of $\Gamma_X^{-1}\Gamma_{E(X|Y)}$ but to ensure that these eigenvectors exist in $\mathcal{L}^2_{\mathcal{I}}$, we have to assume that (see Ferré & Yao (2005) for details) $\sum_i \sum_j 1/(\delta_i\delta_j)E(E(\zeta_i|Y)E(\zeta_j|Y))^2 < +\infty$, where $X = \sum_i \zeta_i u_i$ is the Karhunen-Loève decomposition of X.

Let $\{(X^n, Y^n)\}_{n=1,\dots,N}$ be an i.i.d. sample. In order to estimate the EDR space, we have to choose an estimate for $\Gamma_{E(X|Y)}$. We propose a slicing approach: in Ferré & Yao (2003), the estimate is obtained by partitionning the domain of Y in $(I_h)_{h=1,\dots,H}$ and by setting $\Gamma^N_{E(X|Y)} = \sum_{h=1}^H (N_h/N)\mu_h \otimes \mu_h - \overline{X} \otimes \overline{X}$, where, if $\mathbb{I}$ is the indicator function, $N_h = \sum_{n=1}^N \mathbb{I}_{\{Y^n \in I_h\}}$, $\mu_h = (1/N_h)\sum_{n=1}^N X^n \mathbb{I}_{\{Y^n \in I_h\}}$ and $\overline{X}$ is the empirical mean. Another approach, based on a kernel estimate, has been developped in Ferré & Yao (2005). Although this could be used in our context, we focus on a slicing approach for the sake of simplicity.

A usual estimate of Γ_X is $\Gamma^N_X = (1/N)\sum_{n=1}^N X^n \otimes X^n - \overline{X} \otimes \overline{X}$, but this estimate is ill conditionned (because Γ_X^{-1} is not a bounded operator) so the eigenvectors of $(\Gamma^N_X)^{-1}\Gamma^N_{E(X|Y)}$ do not converge to the eigenvectors of $\Gamma_X^{-1}\Gamma_{E(X|Y)}$. That is the reason why penalization or regularization is needed.

Ferré & Yao (2003) suggest to proceed like Bosq (1991) by considering, instead of Γ_X, a sequence of finite rank operators with bounded inverses and converging to Γ_X. This leads to the estimates $(a_j^N)_{j=1,\dots,q}$ of $(a_j)_{j=1,\dots,q}$ that, under some conditions, satisfy $\| a_j^N - a_j \| \to_p 0$.

The authors also suggest a way of estimating the EDR space for functional data without inverting the covariance operator of the regressor (Ferré & Yao (2005)).

We propose, in Section 3, a regularized approach by penalization.

2.2 SIR for classification

Let $\mathcal{C}_1, \dots, \mathcal{C}_H$ be H groups. When Y is multidimensional, the results of Dauxois *et al.* (2001) are still available and by setting $Y = (\mathbb{I}_{\mathcal{C}_1}, \dots, \mathbb{I}_{\mathcal{C}_H})$, where $\mathbb{I}_{\mathcal{C}_h}$ is the indicator function of the hth group, Model (1) remains valid and we get a natural way to include classification problems into FSIR, see Ferré & Villa (2005). Note that, in the functional case, multivariate methods for discrimination have been extended, mainly inspired from Linear Discriminant Analysis (LDA). In this area, let us mention the works of Hastie *et al.* (1994), Hastie *et al.* (1995) and James & Sugar (2003).

Now, by estimating $\Gamma_{E(X|Y)}$ by

$$\Gamma^N_{E(X|Y)} = \frac{1}{N}\sum_{h=1}^{H} N_h \widehat{E}(X|Y=h) \otimes \widehat{E}(X|Y=h) - \overline{X} \otimes \overline{X}$$

where $N_h = \sum_{n=1}^{N} \mathbb{I}_{\{Y^n=h\}}$ and $\widehat{E}(X|Y=h) = (1/N_h)\sum_{n=1}^{N} X^n \mathbb{I}_{\{Y^n=h\}}$, FSIR leads to a discriminant analysis. The estimation of the EDR space is identical to the discriminant space in linear discriminant analysis. However, the estimation of f leads to a natural classification rule. Indeed, since we have, for all x, $f(x) = E(Y|X=x) = (P(C_1|X=x), ..., P(C_H|X=x))$, the estimation of f coincides with the estimation of the probabilities of the groups conditionally to X.

3 Regularized functional SIR

In Section 2, we saw that the EDR space contains the eigenvalues of the operator $\Gamma_X^{-1}\Gamma_{E(X|Y)}$. Thus, as it is the case for Discriminant Analysis, the estimator of the first direction of the EDR space can be found by maximizing a Rayleigh criterion: $\max_a \langle \Gamma_{E(X|Y)} a, a\rangle / \langle \Gamma_X a, a\rangle$. Unfortunately, as Γ_X^N is ill conditionned, the maximization of the empirical Rayleigh expression does not lead to a good estimate of the EDR space: that is the reason why a regularization is needed.

Provided that we have smooth functions, a relevant method for functional data is to penalize the covariance operator in the Rayleigh expression by introducing smoothing constraints on the estimated functions. This method has already proved its great efficiency (see Hastie *et al.* (1995) for an example of the penalized discriminant analysis).

3.1 Main result

Let $\mathcal{S}$ be the subspace of $\mathcal{L}^2_{\mathcal{I}}$ of functions with a squared integrable second derivative. We introduce a penalty through a bilinear form defined on $\mathcal{S} \times \mathcal{S}$ by, for all f, g in $\in \mathcal{S}$, $[f,g] = \int_{\mathcal{I}} D^2 f(t) D^2 g(t) dt$. We also define the penalized bilinear form associated with empirical operators Γ_X and Γ_X^N:

$$Q_\alpha(f,g) = \langle \Gamma_X f, g\rangle + \alpha[f,g] \qquad \text{and} \qquad Q^N_\alpha(f,g) = \langle \Gamma^N_X f, g\rangle + \alpha[f,g]$$

where α is a regularization parameter. The solutions of the regularized FIR are given by maximizing, under orthogonal constraints, the function

$$\gamma^N(a) = \frac{\langle \Gamma^N_{E(X|Y)} a, a\rangle}{\langle \Gamma^N_X a, a\rangle + \alpha[a,a]}.$$

In order to obtain consistency results for the estimates of $(a_j)_{j=1,...,q}$, we make the following assumptions:

(A2) $E(\| X \|^4) < +\infty$;

(A3) for all $\alpha > 0$, $\inf_{\|a\|=1,\ a\in\mathcal{S}} Q_\alpha(a,a) = \rho_\alpha > 0$;

(A4) $\Gamma^N_{E(X|Y)}$ is a continuous operator which converges in probability to $\Gamma_{E(X|Y)}$ with $\sqrt{N}$ rate;

(A5) $\lim_{N\to+\infty} \alpha = 0$, $\lim_{N\to+\infty} \sqrt{N}\alpha = +\infty$;

(A6) $(a_j)_{j=1,\dots,q}$ belong to $\mathcal{S}$ and verify, for all u such that $\langle \Gamma_X u, a_1\rangle = 0$ and that $\langle \Gamma_X u, u\rangle = 1$, $\langle \Gamma_{E(X|Y)} u, u\rangle \leq \langle \Gamma_{E(X|Y)} a_2, a_2\rangle = \lambda_2 < \lambda_1$.

Since, $\mathcal{S}$ is not a closed subset, γ^N could not reach a maximum on $\mathcal{S}$. However, the following result holds:

Theorem 2. *Under assumptions* **(A1)-(A6)**, *with probability converging to 1, the function γ^N reaches its maximum on $\mathcal{S}$ when N grows to $+\infty$.*
In this case, let then a_1^N be a vector of $\mathcal{S}$ for which γ^N is maximum and which is such that $\langle \Gamma_X a_1^N, a_1\rangle = 1$. Then,

$$\langle \Gamma_X(a_1^N - a_1), a_1^N - a_1\rangle \to_p 0,$$

when N tends to $+\infty$.

Remark 2. For an understandable presentation, we introduce a particular type of penalization but previous results can be found for other regularization functionals satisfying the assumptions. For example, we can replace the bilinear form $[.,.]$ by another one which is similar to the one used in Ridge-PDA (Hastie *et al.* (1995)).

Remark 3. Assumptions **(A2)**, **(A3)** and **(A5)** are technical assumptions that ensure the existence and convergence for $(a_j^N)_{j=1,\dots,q}$: **(A2)** implies that Γ_X^N will converge to Γ_X at the $\sqrt{N}$ rate; we can find in Leurgans *et al.* (1993) conditions that involve **(A3)**. This assumption shows the purpose of regularization: it controls the scaling of Q_α and, thanks to **(A5)**, ensures that the denominator of γ^N doesn't go too fast to 0. Finally **(A5)** gives a way of choosing regularization parameter α (for pratical aspects see section 3.2).

Remark 4. When working with a compact operator T, the ridge regularization $T + \alpha I$ (where I denotes the identity operator) always leads to $\inf_{\|\alpha\|=1}\langle (T + \alpha I)a, a\rangle = \rho_\alpha > 0$ which is exactly assumption **(A3)**. Here, the regularization applied to Γ_X is not the ridge one but is more adapted to the smoothness of the data; an intuitive meaning of this is the ridge regularization of a $D^2\Gamma_X D^{-2}$ type operator (see also section 3.2 for a consequence of this penalization and the link with assumption **(A3)**).

Remark 5. Assumption **(A5)** is fullfilled by the usual estimates introduced above: Li (1991) emphasized the fact that the sliced estimate is consistant, with rate $\sqrt{N}$, for the variable $(Y \in \mathcal{I}_h)_{h=1,\dots,H}$ which satisfies assumption

(A1) as Y does. Ferré & Yao (2005) proved the consistency of the Nadaraya-Watson estimate of $\Gamma_{E(X|Y)}$ and the hilbertian Central Limit theorem ensures the consistency of the estimate given for the classification case.

3.2 Practical aspects

On a practical point of view, X has been observed at some points t_1, t_2, ..., t_D (for an understandable presentation, we suppose that these observations have been centered). The optimization of the penalized Rayleigh expression described in Section 3.1 can be performed by using, for example, B-Splines $(B_i)_i$ to parametrize a_1^N:

$$a_1^N(t) = \sum_i A_{1i} B_i(t) = A_1 B$$

where B is the matrix containing the values of $(B_i(t))_i$ at the points t_1, t_2, ..., t_D. Similarly, the matrix of observations $\mathbf{X} = (X^n(t_d))_{n=1,\ldots,N,\ d=1,\ldots,D}$ can be written in the form of B-Splines: $\mathbf{X} = CB$ with $C = \left[C^1, \ \ldots, \ C^N\right]'$. Let $B^{(2)}$ be the vector containing the values $D^2B(t)$.

If we use the slicing estimate of $\Gamma_{E(X|Y)}$ for regression, we introduce, for all $h = 1, \ldots, H$, $Y_h = \left[\mathbb{I}_{\{Y^1 \in I_h\}}, \ \ldots, \ \mathbb{I}_{\{Y^N \in I_h\}}\right]'$. Then, the problem of maximizing γ^N is equivalent to maximizing $(A'M_eA)/(A'M_{X,\alpha}A)$ where M_e is the estimator of $\Gamma_{E(X|Y)}$ obtained by the slicing approach: $M_e = \sum_{h=1}^{H} (N_h/N) BB'C'Y_hY_h'CBB'$ and where $M_{X,\alpha} = (1/N)BB'C'CBB' + \alpha B^{(2)\prime} B^{(2)}$. This expression underlines the role of the penalization: the matrix $(1/N)BB'C'CBB'$ is usually ill-conditionned (because of the high-dimension of the data) and have tiny eigenvalues (that can even be equal to 0). Provided that $B^{(2)\prime}B^{(2)}$ is invertible, the eigenvalues are rescaled in a basis depending on $B^{(2)}$ and are minored by a strictly positive number depending on α: assumption **(A3)** is then practically fullfilled.

The first solution is the eigenvector, with $M_{X,\alpha}$-norm equal to 1, associated with the largest eigenvalue of the matrix $M_{X,\alpha}^{-1}M_e$. By pursuing the procedure under othogonality constraints, we get that the other solutions are the $M_{X,\alpha}$-orthonormal eigenvectors of $M_{X,\alpha}^{-1}M_e$.

If we deal with classification, the same procedure is achieved by letting $Y_h = \left[\mathbb{I}_{\{Y^1=h\}}, \ \ldots, \ \mathbb{I}_{\{Y^N=h\}}\right]'$.

Finally we have to find the optimal value for α. This can be done, if the sample is large enough (which is the case in the presented applications), by dividing it into two parts: we apply the previous procedure on the first part to find $(a_j^N)_j$ and evaluate the error committed by Model (1) on the second part; the best parameter is then chosen to minimize this error.

4 Multilayer perceptrons

4.1 Approximation by multilayer perceptrons

After the EDR space is estimated, the goal is to get an estimation of the function f in (1): we propose to use a feedforward neural network with one hidden layer. This method (see, e.g., Bishop (1995) for a review on Neural Networks) is an alternative to other nonparametric regressions if the dimension of the EDR space is too large. It has the advantage of working in any cases while some nonparametric methods, such as kernel smoothing or splines smoothing, face the curse of dimensionality.

The main interest of neural networks is their ability to approximate any function with the desired precision (universal approximation); see, for instance, Hornik (1993) for the multivariate context and Stinchcombe (1999) and Rossi & Conan-Guez (2005) in the infinite dimensional one.

4.2 A consistency result

Multi-layer perceptrons approximations of functionals in infinite dimensional spaces have been studied in Chen & Chen (1995), Sandberg & Xu (1996) and Rossi & Conan-Guez (2005). Several strategies are available either by directly using the curves as inputs of the feedforward neural networks or by first projecting the data onto a classical functional basis (such as a spline basis, a Fourier basis, wavelets) or a basis derived from the PCA of X. This latter approach is used by Thodberg (1996).

Our approach is similar but, instead of projecting the data onto a fixed basis or a principal component basis, we project them onto the EDR space. The EDR space behaves as an efficient subspace for the regression of Y on X and it is a way to get a basis which takes into account the relationship between Y and X. In fact, the data are projected onto an estimation of the EDR space, so the accuracy of the projection and then the estimation of the optimal weights for the neural network also depend on how good the EDR space is estimated.

We construct a perceptron (see Figure 2) with one hidden layer having

- as inputs, the coordinates of the projection of X onto $\mathrm{Span}\{(a_j)_{j=1,\ldots,q}\}$: $\langle X, a_1\rangle, \ldots, \langle X, a_q\rangle$;
- q_2 neurons on the hidden layer (where q_2 is a parameter to be estimated);
- as outputs, one neuron for regression and H neurons for classification, representing target Y.

[Figure 2 about here.]

The output of such a neural network is then $\sum_{i=1}^{q_2} w_i^{(2)} g\left(\sum_{j=1}^{q} w_{i,j}^{(1)} \langle X, a_j\rangle + w_i^{(0)}\right)$ where g is the activation function (for example a sigmoid). The purpose of the training step is then to find w^*

which minimizes a loss function L between the output of the neural network with weights $w = \left((w_i^{(2)})_{i=1,\dots,q_2}, (w_{i,j}^{(1)})_{i=1,\dots,q_2}^{j=1,\dots,q}, (w_i^{(0)})_{i=1,\dots,q_2}\right)$, and the target Y:

$$w^* = \arg\min \left\{ E\left[L\left(\sum_{i=1}^{q_2} w_i^{(2)} g\left(\sum_{j=1}^{q} w_{i,j}^{(1)} \langle X, a_j \rangle + w_i^{(0)} \right), Y \right) \right] \right\}. \quad (2)$$

Actually, we obtain an estimation w_N^* of w^* by

$$w_N^* = \arg\min \left\{ \sum_{n=1}^{N} L\left(\sum_{i=1}^{q_2} w_i^{(2)} g\left(\sum_{j=1}^{q} w_{i,j}^{(1)} \langle X^n, a_j^N \rangle + w_i^{(0)} \right), Y^n \right) \right\}.$$

White (1989) gives a consistency theorem for the weights of a neural networks estimated by a set of iid observations. Since $(a_j^N)_j$ is an estimation of the EDR space deduced from the whole data set $\{(X^n, Y^n)\}_n$, the inputs of our functional perceptron used to determine w_N^* do not satisfy the iid assumption and a proper consistency result is then needed.

Let us introduce some notations: ζ is the function from $\mathcal{O} \times \mathcal{W}$ ($\mathcal{O}$ is an open set of $\mathbb{R}^{q+1}$ and $\mathcal{W}$ is a compact set of $\mathbb{R}^{(q+2)q_2}$) such as for all $z = (u, y)$ in $\mathcal{O}$, $\zeta(z, w) = L\left(\sum_{i=1}^{q_2} w_i^{(2)} g\left(\sum_{j=1}^{q} w_{i,j}^{(1)} u_j + w_i^{(0)}\right), y\right)$; Z is the couple of random variables $(\{\langle X, a_j \rangle\}_j, Y)$ and $(Z_n)_{n=1,\dots,N}$ are observations of Z; finally, $(\tilde{Z}_N^n)_{n=1,\dots,N}$ are the couples of $(\{\langle X^n, a_j^N \rangle\}_j, Y^n)$. In our context, the consistency of the Multi-layer Perceptron is given by the following theorem:

Theorem 3. *Under assumptions* **(A1)-(A6)** *and the following assumptions*

(A7) *for all z in $\mathcal{O}$, $\zeta(z, .)$ is continuous;*

(A8) *there is a measurable function $\tilde{\zeta}$ from $\mathcal{O}$ into $\mathbb{R}$ such that, for all z in $\mathcal{O}$, for all w in $\mathcal{W}$, $|\zeta(z, w)| < \tilde{\zeta}(z)$ and $E(\tilde{\zeta}(Z)) < +\infty$;*

(A9) *for all w in $\mathcal{W}$, there exists $C(w) > 0$ such that, for all (x, y) and (x', y') in $\mathcal{O}$, $|\zeta((x, y), w) - \zeta((x', y), w)| \leq C(w) \parallel x - x' \parallel$*

(A10) *for all w in $\mathcal{W}$, $\zeta(., w)$ is measurable.*

If $\mathcal{W}^*$ is the set of minimizers of the problem (2) then

$$d(w_N^*, \mathcal{W}^*) \to_p 0$$

as N tends to $+\infty$ with d defined by: $d(w, \mathcal{W}) = \inf_{\tilde{w} \in \mathcal{W}} \parallel w - \tilde{w} \parallel$ where $\parallel . \parallel$ is the usual euclidean distance.

Remark 6. This list of assumptions is, for example, verified by a perceptron with one hidden layer and a sigmoid function $g(x) = e^x/(1+e^x)$ on the hidden layer associated with the square error $L(\psi, y) = \| \psi - y \|^2$ provided that Y is bounded.

Remark 7. Assumptions **(A1)-(A6)** ensure the convergence of $(a_j^N)_{j=1,\dots,q}$ to $(a_j)_{j=1,\dots,q}$ but they can be replaced by a list of assumptions implying the same result. For example, we would have the same consistency result by projecting the data on the estimated EDR space found by the functional SIR presented in Ferré & Yao (2003) and Ferré & Yao (2005).

5 Applications

5.1 Tecator data

As already said, the Tecator data problem consists in predicting the fat content of pieces of meat from a near infrared absorbance spectrum. We have $N = 215$ observations of (X, Y) where X is the spectrum of absorbance discretized at one hundred points and Y is the fat content.

In order to compute the procedure described in section 3.2, we project the data onto a cubic Spline basis. Because of their smoothness, these data are very well projected onto a basis with 40 equally spaced knots (actually, when using 40 equally spaced knots, or more, the interpolation of the observations by the Spline basis is exact); then, for simplicity reasons, we used this projection for the computation when needed and used the original data in the other cases. We tried several classical methods in order to test the efficiency of SIR-NNr. The competitors are:

- **SIR-NNr**: the functional SIR regularized by penalization, presented in Section 3, precedes a neural network. The neural network training step is made by early stopping procedure: the learning sample is divided into 3 samples (training / validation / test); the training sample is used to train the neural network, the validation sample for an early stopping procedure (when the validation error increases, training is stopped) and this training step is performed 10 times. The best performance of the test sample gives the optimal weights;

- **SIR-NNk**: here we use the smoothed functional inverse regression method presented in Ferré & Yao (2003) as pre-processing to a neural network; the purpose is to show the benefit of the regularization. The neural network is also trained by early stopping;

- **PCA-NN**: in order to show the advantage of SIR, we compute a principal component analysis (as Thodberg (1996)) before a neural network procedure is used (a classical neural network while Thodberg uses a sophisticated bayesian neural network);

- **NNf**: this method is the functional neural network (the Spline projections are used to represent the functional weights and inputs) described by Rossi & Conan-Guez (2005). In this paper, B-Spline basis projection is selected by cross-validation which leads to a huge computational time: we do not follow this approach and use the cubic basis with 40 knots;
- **SIR-L**: after projecting the data onto the EDR space determined by regularized SIR, we compute a linear regression in order to show the efficiency of a neural network compared to a classical parametric method.

We also have to notice that some classical nonparametric methods, such as kernel estimates which depend on the euclidean norm, can not be used for this data set as the dimensionality of the EDR space is too large compared with the number of data (the value of q is given in Table 1).

Before we compare the different methods and in order to limit computational time, we determined the best parameters for each one. Our sample is divided into two parts: on the first one, we determine the values of $(a_j^N)_j$ and of the weights of the neural network for various values of α, q and q_2. On the second part, we determine the standard error of prediction (SEP): the "best" parameters are those which minimize this SEP (see Table 1).

[Table 1 about here.]

Then, in order to see, not only the error made by each method, but also its variability, we randomly build 50 samples divided as follows: the learning sample contains 172 observations and the test sample contains 43. All five methods are first trained on the learning sample (with their optimal parameters pre-determined as described above) and the standard error of prediction (SEP) is then performed on the test sample.

Figure 3 gives the boxplot of the test errors for the 50 samples.

[Figure 3 about here.]

These results show the excellent performances obtained by SIR-NNr: its SEP average over the 50 samples is twice lower than any of the other competitors. Moreover, this method garantees a good stability unlike the others. SIR seems to be a very good pre-processing stage, as SIR-NNk also obtains good performances. Then we have NNf but its rather good results suffer from a very slow computational time. To show this, we give the computational time of each method: when SIR-NNr takes 100 seconds per sample, NNf takes 350 and SIR-L only 1. Clearly NNf is very expensive while SIR-L is very fast but works poorly. Actually, it is closely related to the number of inputs: 42 for NNf and 20 for SIR-NNr.

5.2 Phoneme data

In this section, we compare our methodology with other approaches on a classification problem, namely the phoneme data. The data are log-periodograms

of a 32 ms duration corresponding to recorded speakers; it deals with the discrimination of five speech frames corresponding to five phonemes transcribed as follow: [sh] as in "she", [dcl] as in "dark", [iy] as in "she", [aa] as in "dark" and [ao] as in "water". Finally, the data consist in 4 509 log-periodograms of a 256 length (see Figure 4).

[Figure 4 about here.]

We tried several classical methods in order to test the efficiency of SIR-NNr which is compared with:

- **SIR-NNp**: a classical SIR as presented in Ferré & Yao (2003) as preprocessing of a neural network;
- **SIR-K**: a regularized functional SIR where the function f is estimated by a nonparametric kernel method;
- **Ridge-PDA**: the penalized discriminant analysis introduced in Hastie *et al.* (1995) which uses ridge penalty;
- **NPCD-PCA**: a nonparametric method using kernels and semi-metrics based on Principal Component Analysis and introduced by Ferraty & Vieu (2003).

The optimal parameters for these methods, choosen as in the previous example, are shown in Table 2.

[Table 2 about here.]

For the SIR stage, the optimal dimension of the EDR space is set to 4: it is the maximum dimension possible as the operator $\Gamma^N_{E(X|Y)}$ is of rank $H-1$. We can also see that this dimension is relevant by looking at the projection of the data onto the EDR space (for SIR-NNr, for example, see Figure 5): only the fourth axis is able to separate the phonems [aa] and [ao].

[Figure 5 about here.]

Then we randomly build 50 samples divided as follows: the learning sample contains 1 735 log-periodograms (347 for each class) and the test sample contains also 1 735 (347 for each class). All five methods are first trained on the learning sample and the test error rate is then computed on the test sample. Figure 6 proposes the boxplot of the test error rates.

[Figure 6 about here.]

The results of SIR-NNr, SIR-NNp and SIR-K are very close. The benefit of SIR is highlighted since those three methods work better than others based on different projections of data. The advantage of regularization is also revealed since it leads again to the best results. Then comes RPDA and finally NPCD-PCA which provides the poorest performances. On the contrary, due to a low

dimensionality, neural networks seem to be less performant than kernels and to have a bigger variability (standard deviation is 0.56 for SIR-NNr and only 0.40 for SIR-K): this problem can be removed by increasing the number of training steps, by using more sophisticated architecture or a regularization technique (such as weight decay) but at the price of a larger computational cost. Finally, if SIR-K obtains the best mean (8.09 % versus 8.21 % for SIR-NNr), SIR-NNr is the method which reaches the best minimum which shows its great potential.

In conclusion, both on regression and classification problems, regularized SIR-NN is a competitive solution for functional problems: we can explain these good results by noting that the procedure combines an efficient dimension reduction model and the great accuracy of a neural network, which is able to approximate almost every function. Thus this model can be efficient both for ill-posed problems thanks to the penalized functional and for problems with a large dimensionality thanks to the neural network step. Finally it has another great advantage: computational time is rather short and does not increase too much with the number of observation points for the curves.

6 Acknowledgments

The authors are grateful to the two referees, the Associate Editor and the Editor for their detailed and constructive comments and suggestions.

References

Biau, G., Bunea, F., & Wegkamp, M. (2005). Functional classification in Hilbert spaces. *IEEE Trans. Inform. Theory*, **51**, 2163–2172.

Bishop, C. (1995). *Neural Networks for Pattern Recognition*. Oxford University Press, New York.

Borggaard, C. & Thodberg, H. (1992). Optimal minimal neural interpretation of spectra. *Analytical Chemistry*, **64**, 545–551.

Bosq, D. (1991). Modelization, non-parametric estimation and prediction for continuous time processes, In *Nonparametric functional estimation and related topics, Nato ASI Series C* (ed. Roussas, G.), **335**, 509–529. Kluwer Academic Publishers, Dortrecht.

Cardot, H., Ferraty, F., & Sarda, P. (1999). Functional Linear Model. *Statist. Probab. Lett.*, **45**, 11–22.

Chen, T. & Chen, H. (1995). Universal approximation to nonlinear operators by neural networks with arbitrary activation functions and its application to dynamical systems. *IEEE Transactions on Neural Networks*, **6**(4), 911–917.

Cook, R. & Weisberg, S. (1991). Comment on sliced inverse regression for dimension reduction by K.C. Li. *J. Amer. Statist. Assoc.*, **86**, 328–332.

Dauxois, J., Ferré, L., & Yao, A. (2001). Un modèle semi-paramétrique pour variable aléatoire hilbertienne. *C. R. Math. Acad. Sci. Paris*, **327**(I), 947–952.

Ferraty, F. & Vieu, P. (2003). Curves discrimination: a non parametric approach. *Comput. Statist. Data Anal.*, **44**, 161–173.

Ferré, L. & Villa, N. (2005). Discrimination de courbes par régression inverse fonctionnelle. *Revue de Statistique Appliquée*, **LIII**(1), 39–57.

Ferré, L. & Yao, A. (2003). Functional sliced inverse regression analysis. *Statistics*, **37**, 475–488.

Ferré, L. & Yao, A. (2005). Smoothed functional inverse regression. *Statist. Sinica*, **15**(3), 665–683.

Friedman, J. (1989). Regularized discriminant analysis. *J. Amer. Statist. Assoc.*, **84**, 165–175.

Hastie, T., Buja, A., & Tibshirani, R. (1995). Penalized discriminant analysis. *Ann. Statist.*, **23**, 73–102.

Hastie, T., Tibshirani, R., & Buja, A. (1994). Flexible discriminant analysis by optimal scoring. *J. Amer. Statist. Assoc.*, **89**, 1255–1270.

Hornik, K. (1993). Some new results on neural network approximation. *Neural Networks*, **6**(8), 1069–1072.

Hsing, T. & Carroll, R. (1992). An asymptotic theory for sliced inverse regression. *Ann. Statist.*, **20**, 1040–1061.

James, G. & Sugar, C. (2003). Clustering for sparsely sampled functional data. *J. Amer. Statist. Assoc.*, **98**, 397–408.

Leurgans, S., Moyeed, R., & Silverman, B. (1993). Canonical correlation analysis when the data are curves. *J. R. Statist. Soc. Ser. B*, **55**, 725–740.

Li, K. (1991). Sliced inverse regression for dimension reduction. *J. Amer. Statist. Assoc.*, **86**, 316–342.

Li, K. (1992). On principal hessian directions for data visualisation and dimension reduction: another application of Stein's lemma. *Ann. Statist.*, **87**, 1025–1039.

Ramsay, J. & Silverman, B. (1997). *Functional Data Analysis*. Springer Verlag, New York.

Rossi, F. & Conan-Guez, B. (2005). Functional multi-layer perceptron: a non-linear tool for functional data anlysis. *Neural Networks*, **18**(1), 45–60.

Rossi, F., Conan-Guez, B., & El Golli, A. (2004). Clustering functional data with the som algorithm. In *ESANN'2004 proceedings* 305–312, Bruges, Belgique.

Sandberg, I. & Xu, L. (1996). Network approximation of input-output maps and functionals. *Circuits Systems Signal Process*, **15**(6), 711–725.

Stinchcombe, M. (1999). Neural network approximation of continuous functionals and continuous functions on compactifications. *Neural Networks*, **12**(3), 467–477.

Thodberg, H. (1996). A review of bayesian neural network with an application to near infrared spectroscopy. *IEEE Transaction on Neural Networks*, **7**(1), 56–72.

White, H. (1989). Learning in Artificial Neural Network: A Statistical Perspective. *Neural Computation*, **1**, 425–464.

Xia, Y., Tong, H., Li, W., & Zhu, L. (2002). An adaptative estimation of dimension reduction space. *J. R. Statist. Soc. Ser. B*, **64**, 363–410.

Yao, A. (2001). Un modèle semi-paramérique pour variables fonctionnelles : la régression inverse fonctionnelle. PhD thesis, Université Toulouse III, France.

Zhu, L. & Fang, K. (1996). Asymptotics for kernel estimate of sliced inverse regression. *Ann. Statist.*, **24**, 1053–1068.

Nathalie Villa, Équipe GRIMM, Université Toulouse Le Mirail, 5 allées A. Machado, F-31058 Toulouse cedex 1, France.
E-mail: villa@univ-tlse2.fr.

A Appendix

Here we give the main lines of the proofs of Theorems 2 and 3.

A.1 Theorem 2

The proof of this theorem is related to the one of Theorem 1 in Leurgans *et al.* (1993) and only sketches are given.

Lemma 1: Using Central Limit Theorem, it is easy to show that if $\delta^N = \max\{|||\Gamma_X^N - \Gamma_X|||; |||\Gamma_{E(X|Y)}^N - \Gamma_{E(X|Y)}|||\}$ and if the sequence $(k_N)_N$ satisfies $\sqrt{N}k_N \to +\infty$ then $k_N^{-1}\delta^N \to_p 0$.

Existence: We have for α in $[0,1]$, $Q_\alpha = (1-\alpha)\langle\Gamma_X\ .,.\rangle + \alpha Q_1$ and then, for all u such that $\parallel u \parallel = 1$, $(1/\alpha)Q_\alpha(u,u) \geq (1/\alpha - 1)\langle\Gamma_X u, u\rangle + Q_1 > \rho_1$ by the positiveness of Γ_X. Then, $\sqrt{N}\rho_\alpha > \alpha\sqrt{N}\rho_1$ and we have

$$\sqrt{N}\rho_\alpha \to +\infty\ . \tag{3}$$

Then, by Lemma 1, noting $\Delta_1^N = \Gamma_X^N - \Gamma_X$, $\lim_{N\to+\infty} P\left(\{\omega \in \Omega : |||\Delta_1^N||| \leq (1/2)\rho_\alpha\}\right) = 1$ (where Ω denotes the probability space on which X and Y are defined). But, we have

$$\{\omega \in \Omega : |||\Delta_1^N||| \leq \frac{1}{2}\rho_\alpha\} \subset \left\{\omega : \forall\ a \in \mathcal{S},\ \parallel a \parallel = 1,\ Q_\alpha^N(a,a) \geq \frac{1}{2}\rho_\alpha > 0\right\}$$

and finally the right hand part of the previous equation has a probability converging to 1 when N converges to $+\infty$.

Let $\overline{\mathcal{B}(0,1)}$ be the weak closure of $\{a \in \mathcal{S} \;\; Q_\alpha^N(a,a) = 1\}$ and ζ be the functional defined on $\{a \in \mathcal{S} \;\; Q_\alpha^N(a,a) = 1\}$ by $\zeta(a) = \langle \Gamma_{E(X|Y)}^N a, a\rangle$, then ζ can be extended to a uniformly continuous functional $\tilde{\zeta}$ defined on $\overline{\mathcal{B}(0,1)}$ for the weak topology. Finally, provided that $Q_\alpha^N(a,a) \geq (1/2)\rho_\alpha$, $\tilde{\zeta}$ reaches its maximum on weak compact $\overline{\mathcal{B}(0,1)}$ which concludes the proof of the existence of $(a_j^N)_{j=1,\ldots,q}$.

Consistency: For the following, we suppose that we consider a $\tilde{\omega} \in \Omega$ such that $\tilde{\omega} \in \{\omega \in \Omega : \; \gamma^N$ has a maximum on $\mathcal{S}$ and reaches it$\}$. Let $\lambda_1^N = \lambda_1^N(\tilde{w})$ be this maximum and λ_1^α be the maximum of $\gamma_\alpha(a) = \langle \Gamma_{E(X|Y)} a, a\rangle / (\langle \Gamma_X a, a\rangle + \alpha[a,a])$ on $\mathcal{S}$; λ_1^α is well defined thanks to assumption **(A3)**.

Considering $\gamma_\alpha(a)/\gamma_0(a)$, we easily show that

$$\lambda_1^\alpha \to \lambda_1. \tag{4}$$

Then, by proving that $\sup_{a\in\mathcal{S}} |\gamma^N(a) - \gamma_\alpha(a)| \to_p 0$, we can show that

$$\left|\lambda_1^N - \lambda_1^\alpha\right| \to_p 0. \tag{5}$$

Finally, by combining (4) and (5), we conclude that

$$\lambda_1^N \to_p \lambda_1 \tag{6}$$

Then, by using (6), we demonstrate that

$$\gamma(a_1^N) \to_p \lambda_1 = \gamma(a_1). \tag{7}$$

Thanks to the conclusion of Theorem 1 we show that $\lim_{N\to+\infty} \mathbb{P}(\langle \Gamma_{E(X|Y)} a_1, a_1^N - a_1\rangle = \langle \Gamma_X a_1, a_1^N - a_1\rangle = 0) = 1$. Let μ_N be $\langle \Gamma_X(a_1^N - a_1), a_1^N - a_1\rangle$; if $\langle \Gamma_{E(X|Y)} a_1, a_1^N - a_1\rangle = 0$, we have $\lambda_1^{-1}\gamma(a_1^N) \leq (1+\lambda_1^{-1}\lambda_2\mu_N)/(1+\mu_N)$. As $\lambda_1^{-1}\lambda_2 < 1$, the right hand side of the previous inequality is less than 1; but $\lambda_1^{-1}\gamma(a_1^N)$ converges in probability to 1 by (7) so $(1+\lambda_1^{-1}\lambda_2\mu_N)/(1+\mu_N) \to_p 1$ and then we conclude with $\mu_N \to_p 0$.

A.2 Theorem 3

The proof of this theorem is close to the one found in Rossi & Conan-Guez (2005); the main difference is that the projection for the data is a random variable. The proof will be divided into two parts:

We first prove that

$$\sup_{w\in\mathcal{W}} \left| \frac{1}{N} \sum_{n=1}^{N} \zeta(\tilde{Z}_N^n, w) - E(\zeta(Z,w)) \right| \to_p 0. \tag{8}$$

Forall w in $\mathcal{W}$, we have

$$\left| \frac{1}{N} \sum_{n=1}^{N} \zeta(\tilde{Z}_N^n, w) - E(\zeta(Z,w)) \right|$$

$$\leq \left| \frac{1}{N} \sum_{n=1}^{N} \zeta(\tilde{Z}_N^n, w) - \frac{1}{N} \sum_{n=1}^{N} \zeta(Z_n, w) \right| + \left| \frac{1}{N} \sum_{n=1}^{N} \zeta(Z_n, w) - E(\zeta(Z, w)) \right|.$$

For proving that $\left|(1/N)\sum_{n=1}^{N} \zeta(Z_n, w) - E(\zeta(Z, w))\right| \rightarrow_{a.s.} 0$, we need a general Uniform Strong Law of Large Numbers. Such a result is given in Rossi & Conan-Guez (2005) and, by assumptions **(A7)**, **(A8)** and **(A10)**, Corollary 3 of Rossi & Conan-Guez (2005) directly implies that $\sup_{w \in \mathcal{W}} \left|(1/N)\sum_{n=1}^{N} \zeta(Z_n, w) - E(\zeta(Z, w))\right| \rightarrow_{a.s.} 0$.

Using assumption **(A9)** we see that

$$\left|\tfrac{1}{N}\sum_{n=1}^{N} \left(\zeta(\tilde{Z}_N^n, w) - \zeta(Z_n, w)\right)\right| \leq C(w) \left[\sum_{j=1}^{q} \langle \Gamma_X^N(a_j^N - a_j), a_j^N - a_j \rangle\right]^{1/2}$$

As $|||\Gamma_X^N - \Gamma_X||| \rightarrow_p 0$ and as, for all $j = 1, \ldots, q$, $\langle \Gamma_X(a_j^N - a_j), a_j^N - a_j \rangle \rightarrow_p 0$, we then conclude that $\sup_{w \in \mathcal{W}} \left|(1/N)\sum_{n=1}^{N} \left(\zeta(\tilde{Z}_N^n, w) - \zeta(Z_n, w)\right)\right| \rightarrow_p 0$ (by the same reference as above), which finally implies (8).

Secondly, let ϵ be a positive real. According to the Dominated Convergence Theorem, $E(\zeta(Z, .))$ is a continuous function which reaches its minimum m on compact set $\mathcal{W}$. Then we can show that there is a $\eta(\epsilon) > 0$ such that, for all w in $\mathcal{W}$,

$$|E(\zeta(Z, w)) - m| \leq \eta \;\Rightarrow\; d(w, \mathcal{W}^*) \leq \epsilon. \tag{9}$$

Then let $\Omega_{\eta,N}$ be the following subset of Ω

$$\left\{ \omega \in \Omega : \sup_{w \in \mathcal{W}} \left| \frac{1}{N} \sum_{n=1}^{N} \zeta(\tilde{Z}_N^n, w) - E(\zeta(Z, w)) \right| \leq \frac{\eta}{3} \right\}.$$

If $\omega \in \Omega_{\eta,N}$ then, as $\mathcal{W}$ is a compact set, we can find, for all $N \in \mathbb{N}$, $w_N^*(\omega) \in \mathcal{W}$ which minimizes $(1/N)\sum_{n=1}^{N} \zeta(\tilde{Z}_N^n(\omega), w)$. Let w^* be in the closure of $(w_N^*)_N$; then by arguments similar to the ones used in the first part of the proof we show that, for all $\omega \in \Omega_{\eta,N}$ and for all $w \in \mathcal{W}$, $E(\zeta(Z, w^*)) \leq E(\zeta(z, w)) + \eta$, which implies by the use of (9) that $\Omega_{\eta,N} \subset \{\omega \;\; d(w^*(\omega), \mathcal{W}^*) \leq \epsilon\}$ and this concludes the proof as $\lim_{N \rightarrow +\infty} P(\Omega_{\eta,N}) = 1$.

List of Figures

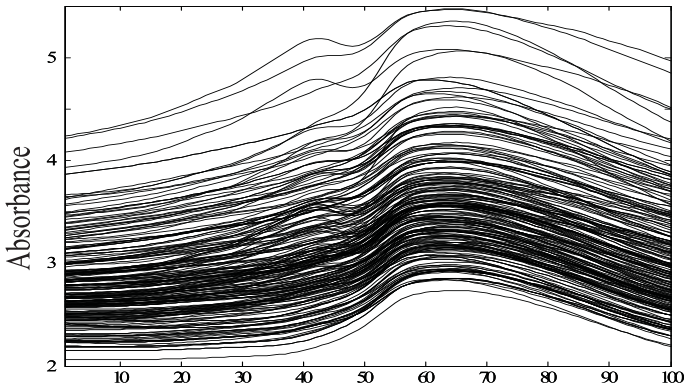

Figure 1: The regressor curves

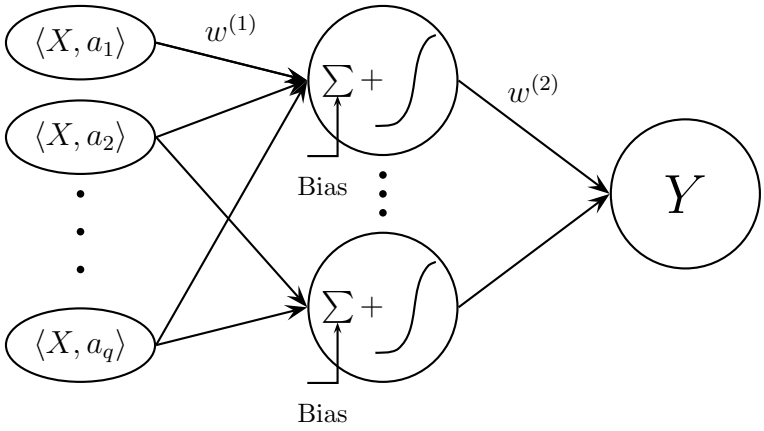

Figure 2: Neural network estimating f

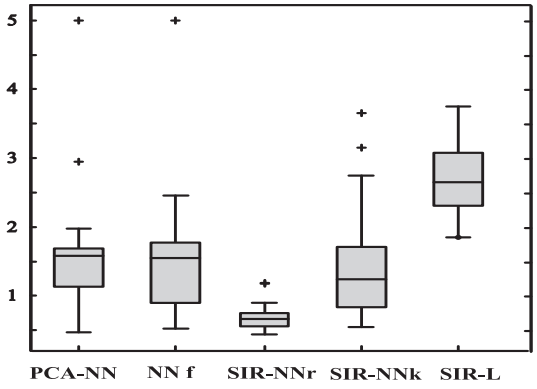

Figure 3: Tecator data set: SEP for 50 samples

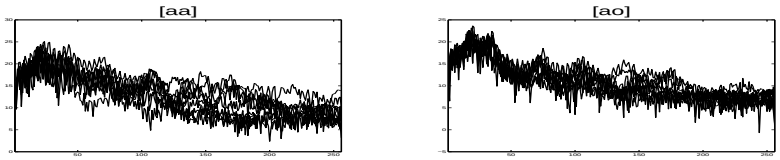

Figure 4: A sample of 10 log-periodograms per class

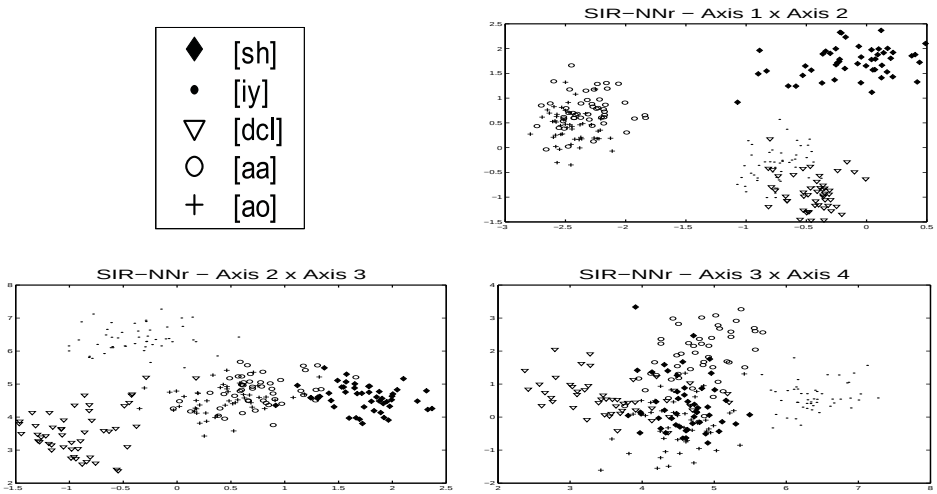

Figure 5: Projection onto the EDR space of 50 log-periodograms by class

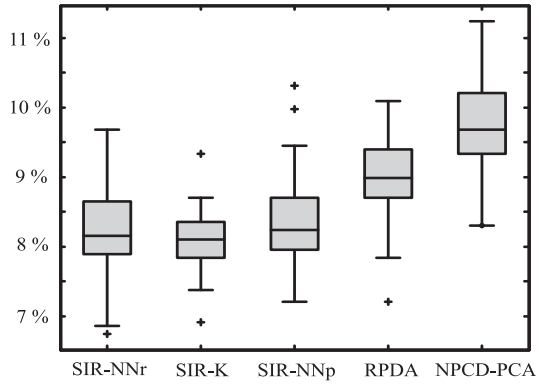

Figure 6: Phoneme Data: Test error rates for 50 samples

List of Tables

Table 1: Best parameters for the five compared methods

	Parameter 1	*Parameter 2*	*Parameter 3*
PCA-NN	$k_n = 25$ (PCA dimension)	$q_2 = 12$ (number of neurons)	
NNf	$q_2 = 18$ (number of neurons)		
SIR-NNr	$\alpha = 5$ (regularization of Γ_X)	$q = 20$ (SIR dimension)	$q_2 = 10$ (number of neurons)
SIR-NNk	$h = 0{,}5$ (kernel window)	$q = 10$ (SIR dimension)	$q_2 = 15$ (number of neurons)
SIR-L	$\alpha = 0{,}5$ (regularization of Γ_X)	$q = 20$ (SIR dimension)	

Table 2: Best parameters for the five compared methods

	Parameter 1	*Parameter 2*	*Parameter 3*
SIR-NNr	$\alpha = 10$ (regularization of Γ_X)	$q = 4$ (SIR dimension)	$q_2 = 15$ (number of neurons)
SIR-NNp	$k_n = 17$ (PCA dimension)	$q = 4$ (SIR dimension)	$q_2 = 12$ (number of neurons)
SIR-K	$\alpha = 10^{-3}$ (regularization of Γ_X)	$q = 4$ (SIR dimension)	$h = 1$ (kernel bandwidth)
RPDA	$\alpha = 5$ (regularization of Γ_X)	q= 4 (PDA dimension)	
NPCD-PCA	$k_n = 7$ (PCA dimension)	$h = 25$ (kernel window)	