CENTER FOR
MACHINE PERCEPTION


CZECH TECHNICAL UNIVERSITY

# Statistical Pattern Recognition Toolbox for Matlab User's guide 

Vojtěch Franc and Václav Hlaváč
\{xfrancv,hlavac\}@cmp.felk.cvut.cz

CTU-CMP-2004-08
June 24, 2004

Available at
ftp://cmp.felk.cvut.cz/pub/cmp/articles/Franc-TR-2004-08.pdf

The authors were supported by the European Union under project IST-2001-35454 ECVision, the project IST-2001-32184 ActIPret and by the Czech Grant Agency under project GACR 102/03/0440 and by the Austrian Ministry of Education under project CONEX GZ 45.535.

Research Reports of CMP, Czech Technical University in Prague, No. 8, 2004
Published by
Center for Machine Perception, Department of Cybernetics
Faculty of Electrical Engineering, Czech Technical University
Technická 2, 16627 Prague 6, Czech Republic
fax +420224357385 , phone +420224357637 , www: http://cmp.felk.cvut.cz

# Statistical Pattern Recognition Toolbox for Matlab 

Vojtěch Franc and Václav Hlaváč

June 24, 2004

## Contents

1 Introduction ..... 3
1.1 What is the Statistical Pattern Recognition Toolbox and how it has been developed? ..... 3
1.2 The STPRtool purpose and its potential users ..... 3
1.3 A digest of the implemented methods ..... 4
1.4 Relevant Matlab toolboxes of others ..... 5
1.5 STPRtool document organization ..... 5
1.6 How to contribute to STPRtool? ..... 5
1.7 Acknowledgements ..... 6
2 Linear Discriminant Function ..... 8
2.1 Linear classifier ..... 9
2.2 Linear separation of finite sets of vectors ..... 10
2.2.1 Perceptron algorithm ..... 11
2.2.2 Kozinec's algorithm ..... 12
2.2.3 Multi-class linear classifier ..... 14
2.3 Fisher Linear Discriminant ..... 15
2.4 Generalized Anderson's task ..... 17
2.4.1 Original Anderson's algorithm ..... 19
2.4.2 General algorithm framework ..... 21
2.4.3 $\varepsilon$-solution using the Kozinec's algorithm ..... 21
2.4.4 Generalized gradient optimization ..... 22
3 Feature extraction ..... 25
3.1 Principal Component Analysis ..... 26
3.2 Linear Discriminant Analysis ..... 29
3.3 Kernel Principal Component Analysis ..... 31
3.4 Greedy kernel PCA algorithm ..... 32
3.5 Generalized Discriminant Analysis ..... 33
3.6 Combining feature extraction and linear classifier ..... 34
4 Density estimation and clustering ..... 41
4.1 Gaussian distribution and Gaussian mixture model ..... 41
4.2 Maximum-Likelihood estimation of GMM ..... 43
4.2.1 Complete data ..... 44
4.2.2 Incomplete data ..... 45
4.3 Minimax estimation ..... 47
4.4 Probabilistic output of classifier ..... 48
4.5 $K$-means clustering ..... 50
5 Support vector and other kernel machines ..... 57
5.1 Support vector classifiers ..... 57
5.2 Binary Support Vector Machines ..... 58
5.3 One-Against-All decomposition for SVM ..... 63
5.4 One-Against-One decomposition for SVM ..... 64
5.5 Multi-class BSVM formulation ..... 66
5.6 Kernel Fisher Discriminant ..... 68
5.7 Pre-image problem ..... 70
5.8 Reduced set method ..... 73
6 Miscellaneous ..... 76
6.1 Data sets ..... 76
6.2 Visualization ..... 77
6.3 Bayesian classifier ..... 78
6.4 Quadratic classifier ..... 80
6.5 $K$-nearest neighbors classifier ..... 82
7 Examples of applications ..... 85
7.1 Optical Character Recognition system ..... 85
7.2 Image denoising ..... 87
8 Conclusions and future planes ..... 94

## Chapter 1

## Introduction

### 1.1 What is the Statistical Pattern Recognition Toolbox and how it has been developed?

The Statistical Pattern Recognition Toolbox (abbreviated STPRtool) is a collection of pattern recognition (PR) methods implemented in Matlab.

The core of the STPRtool is comprised of statistical PR algorithms described in the monograph Schlesinger, M.I., Hlaváć, V: Ten Lectures on Statistical and Structural Pattern Recognition, Kluwer Academic Publishers, 2002 [26]. The STPRtool has been developed since 1999. The first version was a result of a diploma project of Vojtěch Franc [8]. Its author became a PhD student in CMP and has been extending the STPRtool since. Currently, the STPRtool contains much wider collection of algorithms.

The STPRtool is available at
http://cmp.felk.cvut.cz/cmp/cmp_software.html.

We tried to create concise and complete documentation of the STPRtool. This document and associated help files in .html should give all information and parameters the user needs when she/he likes to call appropriate method (typically a Matlab mfile). The intention was not to write another textbook. If the user does not know the method then she/he will likely have to read the description of the method in a referenced original paper or in a textbook.

### 1.2 The STPRtool purpose and its potential users

- The principal purpose of STPRtool is to provide basic statistical pattern recognition methods to (a) researchers, (b) teachers, and (c) students. The STPRtool is likely to help both in research and teaching. It can be found useful by practitioners who design PR applications too.
- The STPRtool offers a collection of the basic and the state-of-the art statistical PR methods.
- The STPRtool contains demonstration programs, working examples and visualization tools which help to understand the PR methods. These teaching aids are intended for students of the PR courses and their teachers who can easily create examples and incorporate them into their lectures.
- The STPRtool is also intended to serve a designer of a new PR method by providing tools for evaluation and comparison of PR methods. Many standard and state-of-the art methods have been implemented in STPRtool.
- The STPRtool is open to contributions by others. If you develop a method which you consider useful for the community then contact us. We are willing to incorporate the method into the toolbox. We intend to moderate this process to keep the toolbox consistent.


### 1.3 A digest of the implemented methods

This section should give the reader a quick overview of the methods implemented in STPRtool.

- Analysis of linear discriminant function: Perceptron algorithm and multiclass modification. Kozinec's algorithm. Fisher Linear Discriminant. A collection of known algorithms solving the Generalized Anderson's Task.
- Feature extraction: Linear Discriminant Analysis. Principal Component Analysis (PCA). Kernel PCA. Greedy Kernel PCA. Generalized Discriminant Analysis.
- Probability distribution estimation and clustering: Gaussian Mixture Models. Expectation-Maximization algorithm. Minimax probability estimation. $K$-means clustering.
- Support Vector and other Kernel Machines: Sequential Minimal Optimizer (SMO). Matlab Optimization toolbox based algorithms. Interface to the $S V M^{\text {light }}$ software. Decomposition approaches to train the Multi-class SVM classifiers. Multi-class BSVM formulation trained by Kozinec's algorithm, Mitchell-Demyanov-Molozenov algorithm and Nearest Point Algorithm. Kernel Fisher Discriminant.


### 1.4 Relevant Matlab toolboxes of others

Some researchers in statistical PR provide code of their methods. The Internet offers many software resources for Matlab. However, there are not many compact and well documented packages comprised of PR methods. We found two useful packages which are briefly introduced below.

NETLAB is focused to the Neural Networks applied to PR problems. The classical PR methods are included as well. This Matlab toolbox is closely related to the popular PR textbook [3]. The documentation published in book [21] contains detailed description of the implemented methods and many working examples. The toolbox includes: Probabilistic Principal Component Analysis, Generative Topographic Mapping, methods based on the Bayesian inference, Gaussian processes, Radial Basis Functions (RBF) networks and many others.
The toolbox can be downloaded from
http://www.ncrg.aston.ac.uk/netlab/.
PRTools is a Matlab Toolbox for Pattern Recognition [7]. Its implementation is based on the object oriented programming principles supported by the Matlab language. The toolbox contains a wide range of PR methods including: analysis of linear and non-linear classifiers, feature extraction, feature selection, methods for combining classifiers and methods for clustering.

The toolbox can be downloaded from
http://www.ph.tn.tudelft.nl/prtools/.

### 1.5 STPRtool document organization

The document is divided into chapters describing the implemented methods. The last chapter is devoted to more complex examples. Most chapters begin with definition of the models to be analyzed (e.g., linear classifier, probability densities, etc.) and the list of the implemented methods. Individual sections have the following fixed structure: (i) definition of the problem to be solved, (ii) brief description and list of the implemented methods, (iii) reference to the literature where the problem is treaded in details and (iv) an example demonstrating how to solve the problem using the STPRtool.

### 1.6 How to contribute to STPRtool?

If you have an implementation which is consistent with the STPRtool spirit and you want to put it into public domain through the STPRtool channel then send an email to xfrancv@cmp.felk.cvut.cz.

Please prepare your code according to the pattern you find in STPRtool implementations. Provide piece of documentation too which will allow us to incorporate it to STPRtool easily.

Of course, your method will be bound to your name in STPRtool.

### 1.7 Acknowledgements

First, we like to thank to Prof. M.I. Schlesinger from the Ukrainian Academy of Sciences, Kiev, Ukraine who showed us the beauty of pattern recognition and gave us many valuable advices.

Second, we like to thank for discussions to our colleagues in the Center for Machine Perception, the Czech Technical University in Prague, our visitors, other STPRtool users and students who used the toolbox in their laboratory exercises.

Third, thanks go to several Czech, EU and industrial grants which indirectly contributed to STPRtool development.

Forth, the STPRtool was rewritten and its documentation radically improved in May and June 2004 with the financial help of ECVision European Research Network on Cognitive Computer Vision Systems (IST-2001-35454) lead by D. Vernon.

## Notation

Upper-case bold letters denote matrices, for instance A. Vectors are implicitly column vectors. Vectors are denoted by lower-case bold italic letters, for instance $\boldsymbol{x}$. The concatenation of column vectors $\boldsymbol{w} \in \mathbb{R}^{n}, \boldsymbol{z} \in \mathbb{R}^{m}$ is denoted as $\boldsymbol{v}=[\boldsymbol{w} ; \boldsymbol{z}]$ where the column vector $\boldsymbol{v}$ is $(n+m)$-dimensional.

| $\left\langle\boldsymbol{x} \cdot \boldsymbol{x}^{\prime}\right\rangle$ | Dot product between vectors $\boldsymbol{x}$ and $\boldsymbol{x}^{\prime}$. |
| :--- | :--- |
| $\boldsymbol{\Sigma}$ | Covariance matrix. |
| $\boldsymbol{\mu}$ | Mean vector (mathematical expectation). |
| S | Scatter matrix. |
| $\mathcal{X} \subseteq \mathbb{R}^{n}$ | $n$-dimensional input space (space of observations). |
| $\mathcal{Y}$ | Set of $c$ hidden states - class labels $\mathcal{Y}=\{1, \ldots, c\}$. |
| $\mathcal{F}$ | Feature space. |
| $\boldsymbol{\phi}: \mathcal{X} \rightarrow \mathcal{F}$ | Mapping from input space $\mathcal{X}$ to the feature space $\mathcal{F}$. |
| $k: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ | Kernel function (Mercer kernel). |
| $\mathcal{T}_{X Y}$ | Labeled training set (more precisely multi-set) $\mathcal{T}_{X Y}=$ |
| $\mathcal{T}_{X}$ | $\left\{\left(\boldsymbol{x}_{1}, y_{1}\right), \ldots,\left(\boldsymbol{x}_{l}, y_{l}\right)\right\}$. |
| $q: \mathcal{X} \rightarrow \mathcal{Y}$ | Unlabeled training set $\mathcal{T}_{X}=\left\{\boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{l}\right\}$. |
| $f_{y}: \mathcal{X} \rightarrow \mathbb{R}$ | Classification rule. |
| $f: \mathcal{X} \rightarrow \mathbb{R}$ | Discriminant function associated with the class $y \in \mathcal{Y}$. |
| $\|\cdot\|$ | Discriminant function of the binary classifier. |
| $\\|\cdot\\|$ | Cardinality of a set. |
| $\operatorname{det}(() \cdot)$ | Euclidean norm. |
| $\delta(i, j)$ | Matrix determinant. |
| Kronecker delta $\delta(i, j)=1$ for $i=j$ and $\delta(i, j)=0$ |  |
| $V$ | otherwise. |

## Chapter 2

## Linear Discriminant Function

The linear classification rule $q: \mathcal{X} \subseteq \mathbb{R}^{n} \rightarrow \mathcal{Y}=\{1,2, \ldots, c\}$ is composed of a set of discriminant functions

$$
f_{y}(\boldsymbol{x})=\left\langle\boldsymbol{w}_{y} \cdot \boldsymbol{x}\right\rangle+b_{y}, \quad \forall y \in \mathcal{Y}
$$

which are linear with respect to both the input vector $\boldsymbol{x} \in \mathbb{R}^{n}$ and their parameter vectors $\boldsymbol{w} \in \mathbb{R}^{n}$. The scalars $b_{y}, \forall y \in \mathcal{Y}$ introduce bias to the discriminant functions. The input vector $\boldsymbol{x} \in \mathbb{R}^{n}$ is assigned to the class $y \in \mathcal{Y}$ its corresponding discriminant function $f_{y}$ attains maximal value

$$
\begin{equation*}
y=\underset{y^{\prime} \in \mathcal{Y}}{\operatorname{argmax}} f_{y^{\prime}}(\boldsymbol{x})=\underset{y^{\prime} \in \mathcal{Y}}{\operatorname{argmax}}\left(\left\langle\boldsymbol{w}_{y^{\prime}} \cdot \boldsymbol{x}\right\rangle+b_{y^{\prime}}\right) . \tag{2.1}
\end{equation*}
$$

In the particular binary case $\mathcal{Y}=\{1,2\}$, the linear classifier is represented by a single discriminant function

$$
f(\boldsymbol{x})=\langle\boldsymbol{w} \cdot \boldsymbol{x}\rangle+b
$$

given by the parameter vector $\boldsymbol{w} \in \mathbb{R}^{n}$ and bias $b \in \mathbb{R}$. The input vector $\boldsymbol{x} \in \mathbb{R}^{n}$ is assigned to class $y \in\{1,2\}$ as follows

$$
q(\boldsymbol{x})=\left\{\begin{array}{lll}
1 & \text { if } & f(\boldsymbol{x})=\langle\boldsymbol{w} \cdot \boldsymbol{x}\rangle+b \geq 0,  \tag{2.2}\\
2 & \text { if } & f(\boldsymbol{x})=\langle\boldsymbol{w} \cdot \boldsymbol{x}\rangle+b<0 .
\end{array}\right.
$$

The data-type used to describe the linear classifier and the implementation of the classifier is described in Section 2.1. Following sections describe supervised learning methods which determine the parameters of the linear classifier based on available knowledge. The implemented learning methods can be distinguished according to the type of the training data:
I. The problem is described by a finite set $\mathcal{T}=\left\{\left(\boldsymbol{x}_{1}, y_{1}\right), \ldots,\left(\boldsymbol{x}_{l}, y_{l}\right)\right\}$ containing pairs of observations $\boldsymbol{x}_{i} \in \mathbb{R}^{n}$ and corresponding class labels $y_{i} \in \mathcal{Y}$. The methods using this type of training data is described in Section 2.2 and Section 2.3.
II. The parameters of class conditional distributions are (completely or partially) known. The Anderson's task and its generalization are representatives of such methods (see Section 2.4).

The summary of implemented methods is given in Table 2.1.
References: The linear discriminant function is treated throughout for instance in books [26, 6].

Table 2.1: Implemented methods: Linear discriminant function

| linclass | Linear classifier (linear machine). |
| :--- | :--- |
| andrerr | Classification error of the Generalized Anderson's <br> task. |
| androrig | Original method to solve the Anderson-Bahadur's <br> task. |
| eanders | Epsilon-solution of the Generalized Anderson's task. <br> ganders <br> ggradander |
| Solves the Generalized Anderson's task. <br> Gradient method to solve the Generalized Anderson's <br> task. |  |
| ekozinec | Kozinec's algorithm for eps-optimal separating hyper- <br> plane. |
| mperceptr | Perceptron algorithm to train multi-class linear ma- <br> chine. |
| perceptron | Perceptron algorithm to train binary linear classifier. <br> fld |
| fldqp | Fisher Linear Discriminant. |
| Fisher Linear Discriminant using Quadratic Program- |  |
| demo_linclass | ming. |
| demo on the algorithms learning linear classifiers. |  |

### 2.1 Linear classifier

The STPRtool uses a specific structure array to describe the binary (see Table 2.2) and the multi-class (see Table 2.3) linear classifier. The implementation of the linear classifier itself provides function linclass.

## Example: Linear classifier

The example shows training of the Fisher Linear Discriminant which is the classical example of the linear classifier. The Riply's data set riply_trn.mat is used for training. The resulting linear classifier is then evaluated on the testing data riply_tst.mat.

Table 2.2: Data-type used to describe binary linear classifier.

| Binary linear classifier (structure array): |  |
| :--- | :--- |
| . $\mathrm{W}[n \times 1]$ | The normal vector $\boldsymbol{w} \in \mathbb{R}^{n}$ of the separating hyper- |
|  | plane $f(\boldsymbol{x})=\langle\boldsymbol{w} \cdot \boldsymbol{x}\rangle+b$. |
| . $\mathrm{b}[1 \times 1]$ | Bias $b \in \mathbb{R}$ of the hyperplane. |
| . fun $=$ 'linclass' | Identifies function associated with this data type. |

Table 2.3: Data-type used to describe multi-class linear classifier.

| Multi-class linear classifier (structure array): |  |
| :--- | :--- |
| . $\mathrm{W}[n \times c]$ | Matrix of parameter vectors $\boldsymbol{w}_{y} \in \mathbb{R}^{n}, y=1, \ldots, c$ of |
|  | the linear discriminant functions $f_{y}(\boldsymbol{x})=\left\langle\boldsymbol{w}_{y} \cdot \boldsymbol{x}\right\rangle+b$. |
| . $\mathrm{b}[c \times 1]$ | Parameters $b_{y} \in \mathbb{R}, y=1, \ldots, c$. |
| $. f u n={ }^{\prime}$ 'linclass' | Identifies function associated with this data type. |

```
trn = load('riply_trn'); % load training data
tst = load('riply_tst'); % load testing data
model = fld( trn ); % train FLD classifier
ypred = linclass( tst.X, model ); % classify testing data
cerror( ypred, tst.y ) % evaluate testing error
```

ans $=$
0.1080

### 2.2 Linear separation of finite sets of vectors

The input training data $\mathcal{T}=\left\{\left(\boldsymbol{x}_{1}, y_{1}\right), \ldots,\left(\boldsymbol{x}_{l}, y_{l}\right)\right\}$ consists of pairs of observations $\boldsymbol{x}_{i} \in \mathbb{R}^{n}$ and corresponding class labels $y_{i} \in \mathcal{Y}=\{1, \ldots, c\}$. The implemented methods solve the task of (i) training separating hyperplane for binary case $c=2$, (ii) training $\varepsilon$-optimal hyperplane and (iii) training the multi-class linear classifier. The definitions of these tasks are given below.

The problem of training the binary $(c=2)$ linear classifier (2.2) with zero training error is equivalent to finding the hyperplane $\mathcal{H}=\{\boldsymbol{x}:\langle\boldsymbol{w} \cdot \boldsymbol{x}\rangle+b=0\}$ which separates the training vectors of the first $y=1$ and the second $y=2$ class. The problem is formally defined as solving the set of linear inequalities

$$
\begin{array}{ll}
\left\langle\boldsymbol{w} \cdot \boldsymbol{x}_{i}\right\rangle+b \geq 0, & y_{i}=1,  \tag{2.3}\\
\left\langle\boldsymbol{w} \cdot \boldsymbol{x}_{i}\right\rangle+b<0, & y_{i}=2 .
\end{array}
$$

with respect to the vector $\boldsymbol{w} \in \mathbb{R}^{n}$ and the scalar $b \in \mathbb{R}$. If the inequalities (2.3) have a solution then the training data $\mathcal{T}$ is linearly separable. The Perceptron (Section 2.2.1)
and Kozinec's (Section 2.2.2) algorithm are useful to train the binary linear classifiers from the linearly separable data.

If the training data is linearly separable then the optimal separating hyperplane can be defined the solution of the following task

$$
\begin{align*}
\left(\boldsymbol{w}^{*}, b^{*}\right) & =\underset{\boldsymbol{w}, b}{\operatorname{argmax}} m(\boldsymbol{w}, b) \\
& =\underset{\boldsymbol{w}, b}{\operatorname{argmax}} \min \left(\min _{i \in \mathcal{I}_{1}} \frac{\left\langle\boldsymbol{w} \cdot \boldsymbol{x}_{i}\right\rangle+b}{\|\boldsymbol{w}\|}, \min _{i \in \mathcal{I}_{2}}-\frac{\left\langle\boldsymbol{w} \cdot \boldsymbol{x}_{i}\right\rangle+b}{\|\boldsymbol{w}\|}\right), \tag{2.4}
\end{align*}
$$

where $\mathcal{I}_{1}=\left\{i: y_{i}=1\right\}$ and $\mathcal{I}_{2}=\left\{i: y_{i}=2\right\}$ are sets of indices. The optimal separating hyperplane $\mathcal{H}^{*}=\left\{\boldsymbol{x} \in \mathbb{R}^{n}:\left\langle\boldsymbol{w}^{*} \cdot \boldsymbol{x}\right\rangle+b^{*}=0\right\}$ separates training data $\mathcal{T}$ with maximal margin $m\left(\boldsymbol{w}^{*}, b^{*}\right)$. This task is equivalent to training the linear Support Vector Machines (see Section 5). The optimal separating hyperplane cannot be found exactly except special cases. Therefore the numerical algorithms seeking the approximate solution are applied instead. The $\varepsilon$-optimal solution is defined as the vector $\boldsymbol{w}$ and scalar $b$ such that the inequality

$$
\begin{equation*}
m\left(\boldsymbol{w}^{*}, b^{*}\right)-m(\boldsymbol{w}, b) \leq \varepsilon, \tag{2.5}
\end{equation*}
$$

holds. The parameter $\varepsilon \geq 0$ defines closeness to the optimal solution in terms of the margin. The $\varepsilon$-optimal solution can be found by Kozinec's algorithm (Section 2.2.2).

The problem of training the multi-class linear classifier $c>2$ with zero training error is formally stated as the problem of solving the set of linear inequalities

$$
\begin{equation*}
\left\langle\boldsymbol{w}_{y_{i}} \cdot \boldsymbol{x}_{i}\right\rangle+b_{y_{i}}>\left\langle\boldsymbol{w}_{y} \cdot \boldsymbol{x}_{i}\right\rangle+b_{y}, \quad i=1, \ldots, l, \quad y_{i} \neq y \tag{2.6}
\end{equation*}
$$

with respect to the vectors $\boldsymbol{w}_{y} \in \mathbb{R}^{n}, y \in \mathcal{Y}$ and scalars $b_{y} \in \mathbb{R}, y \in \mathcal{Y}$. The task (2.6) for linearly separable training data $\mathcal{T}$ can be solved by the modified Perceptron (Section 2.2.3) or Kozinec's algorithm. An interactive demo on the algorithms separating the finite sets of vectors by a hyperplane is implemented in demo_linclass.

### 2.2.1 Perceptron algorithm

The input is a set $\mathcal{T}=\left\{\left(\boldsymbol{x}_{1}, y_{1}\right), \ldots,\left(\boldsymbol{x}_{l}, y_{l}\right)\right\}$ of binary labeled $y_{i} \in\{1,2\}$ training vectors $\boldsymbol{x}_{i} \in \mathbb{R}^{n}$. The problem of training the separating hyperplane (2.3) can be formally rewritten to a simpler form

$$
\begin{equation*}
\left\langle\boldsymbol{v} \cdot \boldsymbol{z}_{i}\right\rangle>0, \quad i=1, \ldots, l, \tag{2.7}
\end{equation*}
$$

where the vector $\boldsymbol{v} \in \mathbb{R}^{n+1}$ is constructed as

$$
\begin{equation*}
\boldsymbol{v}=[\boldsymbol{w} ; b] \tag{2.8}
\end{equation*}
$$

and transformed training data $\mathcal{Z}=\left\{\boldsymbol{z}_{1}, \ldots, \boldsymbol{z}_{l}\right\}$ are defined as

$$
\boldsymbol{z}_{i}=\left\{\begin{array}{rll}
{\left[\boldsymbol{x}_{i} ; 1\right]} & \text { if } & y_{i}=1  \tag{2.9}\\
-\left[\boldsymbol{x}_{i} ; 1\right] & \text { if } & y_{i}=2 .
\end{array}\right.
$$

The problem of solving (2.7) with respect to the unknown vector $\boldsymbol{v} \in \mathbb{R}^{n+1}$ is equivalent to the original task (2.3). The parameters $(\boldsymbol{w}, b)$ of the linear classifier are obtained from the found vector $\boldsymbol{v}$ by inverting the transformation (2.8).

The Perceptron algorithm is an iterative procedure which builds a series of vectors $\boldsymbol{v}^{(0)}, \boldsymbol{v}^{(1)}, \ldots, \boldsymbol{v}^{(t)}$ until the set of inequalities (2.7) is satisfied. The initial vector $\boldsymbol{v}^{(0)}$ can be set arbitrarily (usually $\boldsymbol{v}=\mathbf{0}$ ). The Novikoff theorem ensures that the Perceptron algorithm stops after finite number of iterations $t$ if the training data are linearly separable. Perceptron algorithm is implemented in function perceptron.

References: Analysis of the Perceptron algorithm including the Novikoff's proof of convergence can be found in Chapter 5 of the book [26]. The Perceptron from the neural network perspective is described in the book [3].

## Example: Training binary linear classifier with Perceptron

The example shows the application of the Perceptron algorithm to find the binary linear classifier for synthetically generated 2-dimensional training data. The generated training data contain 50 labeled vectors which are linearly separable with margin 1. The found classifier is visualized as a separating hyperplane (line in this case) $\mathcal{H}=$ $\left\{\boldsymbol{x} \in \mathbb{R}^{2}:\langle\boldsymbol{w} \cdot \boldsymbol{x}\rangle+b=0\right\}$. See Figure 2.1.

```
data = genlsdata( 2, 50, 1);
% generate training data
model = perceptron( data ); % call perceptron
figure;
ppatterns( data ); % plot training data
pline( model ); % plot separating hyperplane
```


### 2.2.2 Kozinec's algorithm

The input is a set $\mathcal{T}=\left\{\left(\boldsymbol{x}_{1}, y_{1}\right), \ldots,\left(\boldsymbol{x}_{l}, y_{l}\right)\right\}$ of binary labeled $y_{i} \in\{1,2\}$ training vectors $\boldsymbol{x}_{i} \in \mathbb{R}^{n}$. The Kozinec's algorithm builds a series of vectors $\boldsymbol{w}_{1}^{(0)}, \boldsymbol{w}_{1}^{(1)}, \ldots, \boldsymbol{w}_{1}^{(t)}$ and $\boldsymbol{w}_{2}^{(0)}, \boldsymbol{w}_{2}^{(1)}, \ldots, \boldsymbol{w}_{2}^{(t)}$ which converge to the vector $\boldsymbol{w}_{1}^{*}$ and $\boldsymbol{w}_{2}^{*}$ respectively. The vectors $\boldsymbol{w}_{1}^{*}$ and $\boldsymbol{w}_{2}^{*}$ are the solution of the following task

$$
\boldsymbol{w}_{1}^{*}, \boldsymbol{w}_{2}^{*}=\underset{\boldsymbol{w}_{1} \in \overline{\mathcal{X}_{1}}, \boldsymbol{w}_{2} \in \overline{\mathcal{X}_{2}}}{\operatorname{argmin}}\left\|\boldsymbol{w}_{1}-\boldsymbol{w}_{2}\right\|,
$$

where $\overline{\mathcal{X}_{1}}$ stands for the convex hull of the training vectors of the first class $\mathcal{X}_{1}=$ $\left\{\boldsymbol{x}_{i}: y_{i}=1\right\}$ and $\overline{\mathcal{X}_{2}}$ for the convex hull of the second class likewise. The vector $\boldsymbol{w}^{*}=$


Figure 2.1: Linear classifier trained by the Perceptron algorithm
$\boldsymbol{w}_{1}^{*}-\boldsymbol{w}_{2}^{*}$ and the bias $b^{*}=\frac{1}{2}\left(\left\|\boldsymbol{w}_{2}^{*}\right\|^{2}-\left\|\boldsymbol{w}_{1}^{*}\right\|^{2}\right)$ determine the optimal hyperplane (2.4) separating the training data with the maximal margin.

The Kozinec's algorithm is proven to converge to the vectors $\boldsymbol{w}_{1}^{*}$ and $\boldsymbol{w}_{2}^{*}$ in infinite number of iterations $t=\infty$. If the $\varepsilon$-optimal optimality stopping condition is used then the Kozinec's algorithm converges in finite number of iterations. The Kozinec's algorithm can be also used to solve a simpler problem of finding the separating hyperplane (2.3). Therefore the following two stopping conditions are implemented:
I. The separating hyperplane (2.3) is sought for $(\varepsilon<0)$. The Kozinec's algorithm is proven to converge in a finite number of iterations if the separating hyperplane exists.
II. The $\varepsilon$-optimal hyperplane (2.5) is sought for $(\varepsilon \geq 0)$. Notice that setting $\varepsilon=0$ force the algorithm to seek the optimal hyperplane which is generally ensured to be found in an infinite number of iterations $t=\infty$.

The Kozinec's algorithm which trains the binary linear classifier is implemented in the function ekozinec.
References: Analysis of the Kozinec's algorithm including the proof of convergence can be found in Chapter 5 of the book [26].

## Example: Training $\varepsilon$-optimal hyperplane by the Kozinec's algorithm

The example shows application of the Kozinec's algorithm to find the $(\varepsilon=0.01)$ optimal hyperplane for synthetically generated 2-dimensional training data. The generated training data contains 50 labeled vectors which are linearly separable with
margin 1. The found classifier is visualized (Figure 2.2) as a separating hyperplane $\mathcal{H}=\left\{\boldsymbol{x} \in \mathbb{R}^{2}:\langle\boldsymbol{w} \cdot \boldsymbol{x}\rangle+b=0\right\}$ (straight line in this case). It could be verified that the found hyperplane has margin model.margin which satisfies the $\varepsilon$-optimality condition.

```
data = genlsdata(2,50,1); % generate data
% run ekozinec
model = ekozinec(data, struct('eps',0.01));
figure;
ppatterns(data); % plot data
pline(model); % plot found hyperplane
```



Figure 2.2: $\varepsilon$-optimal hyperplane trained with the Kozinec's algorithm.

### 2.2.3 Multi-class linear classifier

The input training data $\mathcal{T}=\left\{\left(\boldsymbol{x}_{1}, y_{1}\right), \ldots,\left(\boldsymbol{x}_{l}, y_{l}\right)\right\}$ consists of pairs of observations $\boldsymbol{x}_{i} \in \mathbb{R}^{n}$ and corresponding class labels $y_{i} \in \mathcal{Y}=\{1, \ldots, c\}$. The task is to design the linear classifier (2.1) which classifies the training data $\mathcal{T}$ without error. The problem of training (2.1) is equivalent to solving the set of linear inequalities (2.6). The inequalities (2.6) can be formally transformed to a simpler set

$$
\begin{equation*}
\left\langle\boldsymbol{v} \cdot \boldsymbol{z}_{i}^{y}\right\rangle>0, \quad i=1, \ldots, l, y=\mathcal{Y} \backslash\left\{y_{i}\right\} \tag{2.10}
\end{equation*}
$$

where the vector $\boldsymbol{v} \in \mathbb{R}^{(n+1) c}$ contains the originally optimized parameters

$$
\boldsymbol{v}=\left[\boldsymbol{w}_{1} ; b_{1} ; \boldsymbol{w}_{2} ; b_{2} ; \ldots ; \boldsymbol{w}_{c} ; b_{c}\right] .
$$

The set of vectors $\boldsymbol{z}_{i}^{y} \in \mathbb{R}^{(n+1) c}, i=1, \ldots, l, y \in \mathcal{Y} \backslash\left\{y_{i}\right\}$ is created from the training set $\mathcal{T}$ such that

$$
\boldsymbol{z}_{i}^{y}(j)=\left\{\begin{align*}
{\left[\boldsymbol{x}_{i} ; 1\right], } & \text { for } j=y_{i}  \tag{2.11}\\
-\left[\boldsymbol{x}_{i} ; 1\right], & \text { for } j=y, \\
\mathbf{0}, & \text { otherwise }
\end{align*}\right.
$$

where $\boldsymbol{z}_{i}^{y}(j)$ stands for the $j$-th slot between coordinates $(n+1)(j-1)+1$ and $(n+1) j$. The described transformation is known as the Kesler's construction.

The transformed task (2.10) can be solved by the Perceptron algorithm. The simple form of the Perceptron updating rule allows to implement the transformation implicitly without mapping the training data $\mathcal{T}$ into $(n+1) c$-dimensional space. The described method is implemented in function mperceptron.
References: The Kesler's construction is described in books [26, 6].

## Example: Training multi-class linear classifier by the Perceptron

The example shows application of the Perceptron rule to train the multi-class linear classifier for the synthetically generated 2 -dimensional training data pentagon.mat. A user's own data can be generated interactively using function createdata('finite', 10) (number 10 stands for number of classes). The found classifier is visualized as its separating surface which is known to be piecewise linear and the class regions should be convex (see Figure 2.3).

```
data = load('pentagon'); % load training data
model = mperceptron(data); % run training algorithm
figure;
ppatterns(data); % plot training data
pboundary(model); % plot decision boundary
```


### 2.3 Fisher Linear Discriminant

The input is a set $\mathcal{T}=\left\{\left(\boldsymbol{x}_{1}, y_{1}\right), \ldots,\left(\boldsymbol{x}_{l}, y_{l}\right)\right\}$ of binary labeled $y_{i} \in\{1,2\}$ training vectors $\boldsymbol{x}_{i} \in \mathbb{R}^{n}$. Let $\mathcal{I}_{y}=\left\{i: y_{i}=y\right\}, y \in\{1,2\}$ be sets of indices of training vectors belonging to the first $y=1$ and the second $y=2$ class, respectively. The class separability in a direction $\boldsymbol{w} \in \mathbb{R}^{n}$ is defined as

$$
\begin{equation*}
F(\boldsymbol{w})=\frac{\left\langle\boldsymbol{w} \cdot \mathbf{S}_{\mathbf{B}} \boldsymbol{w}\right\rangle}{\left\langle\boldsymbol{w} \cdot \mathbf{S}_{\mathbf{W}} \boldsymbol{w}\right\rangle}, \tag{2.12}
\end{equation*}
$$

where $\mathbf{S}_{\mathbf{B}}$ is the between-class scatter matrix

$$
\mathbf{S}_{\mathbf{B}}=\left(\boldsymbol{\mu}_{1}-\boldsymbol{\mu}_{2}\right)\left(\boldsymbol{\mu}_{1}-\boldsymbol{\mu}_{2}\right)^{T}, \quad \boldsymbol{\mu}_{y}=\frac{1}{\left|\mathcal{I}_{y}\right|} \sum_{i \in \mathcal{I}_{y}} \boldsymbol{x}_{i}, y \in\{1,2\}
$$



Figure 2.3: Multi-class linear classifier trained by the Perceptron algorithm.
and $\mathbf{S}_{W}$ is the within class scatter matrix defined as

$$
\mathbf{S}_{\mathbf{W}}=\mathbf{S}_{1}+\mathbf{S}_{2}, \quad \mathbf{S}_{y}=\sum_{i \in \mathcal{Y}_{y}}\left(\boldsymbol{x}_{i}-\boldsymbol{\mu}_{y}\right)\left(\boldsymbol{x}_{i}-\boldsymbol{\mu}_{y}\right)^{T}, y \in\{1,2\}
$$

In the case of the Fisher Linear Discriminant (FLD), the parameter vector $\boldsymbol{w}$ of the linear discriminant function $f(\boldsymbol{x})=\langle\boldsymbol{w} \cdot \boldsymbol{x}\rangle+b$ is determined to maximize the class separability criterion (2.12)

$$
\begin{equation*}
\boldsymbol{w}=\underset{\boldsymbol{w}^{\prime}}{\operatorname{argmax}} F\left(\boldsymbol{w}^{\prime}\right)=\underset{\boldsymbol{w}^{\prime}}{\operatorname{argmax}} \frac{\left\langle\boldsymbol{w}^{\prime} \cdot \mathbf{S}_{\mathbf{B}} \boldsymbol{w}^{\prime}\right\rangle}{\left\langle\boldsymbol{w}^{\prime} \cdot \mathbf{S}_{\mathbf{W}} \boldsymbol{w}^{\prime}\right\rangle} . \tag{2.13}
\end{equation*}
$$

The bias $b$ of the linear rule must be determined based on another principle. The STPRtool implementation, for instance, computes such $b$ that the equality

$$
\left\langle\boldsymbol{w} \cdot \mu_{1}\right\rangle+b=-\left(\left\langle\boldsymbol{w} \cdot \mu_{2}\right\rangle+b\right),
$$

holds. The STPRtool contains two implementations of FLD:
I. The classical solution of the problem (2.13) using the matrix inversion

$$
\boldsymbol{w}=\mathbf{S}_{\mathbf{W}}^{-1}\left(\boldsymbol{\mu}_{1}-\boldsymbol{\mu}_{2}\right) .
$$

This case is implemented in the function fld.
II. The problem (2.13) can be reformulated as the quadratic programming (QP) task

$$
\boldsymbol{w}=\underset{\boldsymbol{w}^{\prime}}{\operatorname{argmin}}\left\langle\boldsymbol{w}^{\prime} \cdot \mathbf{S}_{\mathbf{W}} \boldsymbol{w}^{\prime}\right\rangle,
$$

subject to

$$
\left\langle\boldsymbol{w}^{\prime} \cdot\left(\boldsymbol{\mu}_{1}-\boldsymbol{\mu}_{2}\right)\right\rangle=2 .
$$

The QP solver quadprog of the Matlab Optimization Toolbox is used in the toolbox implementation. This method can be useful when the matrix inversion is hard to compute. The method is implemented in the function fldqp.

References: The Fisher Linear Discriminant is described in Chapter 3 of the book [6] or in the book [3].

## Example: Fisher Linear Discriminant

The binary linear classifier based on the Fisher Linear Discriminant is trained on the Riply's data set riply_trn.mat. The QP formulation fldqp of FLD is used. However, the function fldqp can be replaced by the standard approach implemented in the function fld which would yield the same solution. The found linear classifier is visualized (see Figure 2.4) and evaluated on the testing data riply_tst.mat.

```
trn = load('riply_trn'); % load training data
model = fldqp(trn); % compute FLD
figure;
ppatterns(trn); pline(model); % plot data and solution
tst = load('riply_tst'); % load testing data
ypred = linclass(tst.X,model); % classify testing data
cerror(ypred,tst.y) % compute testing error
```

ans =
0.1080

### 2.4 Generalized Anderson's task

The classified object is described by the vector of observations $\boldsymbol{x} \in \mathcal{X} \subseteq \mathbb{R}^{n}$ and a binary hidden state $y \in\{1,2\}$. The class conditional distributions $p_{X \mid Y}(\boldsymbol{x} \mid y), y \in\{1,2\}$ are known to be multi-variate Gaussian distributions. The parameters $\left(\boldsymbol{\mu}_{1}, \boldsymbol{\Sigma}_{1}\right)$ and $\left(\boldsymbol{\mu}_{2}, \boldsymbol{\Sigma}_{2}\right)$ of these class distributions are unknown. However, it is known that the parameters $\left(\boldsymbol{\mu}_{1}, \boldsymbol{\Sigma}_{1}\right)$ belong to a certain finite set of parameters $\left\{\left(\boldsymbol{\mu}^{i}, \boldsymbol{\Sigma}^{i}\right): i \in \mathcal{I}_{1}\right\}$. Similarly $\left(\boldsymbol{\mu}_{2}, \boldsymbol{\Sigma}_{2}\right)$ belong to a finite set $\left\{\left(\boldsymbol{\mu}^{i}, \boldsymbol{\Sigma}^{i}\right): i \in \mathcal{I}_{2}\right\}$. Let $q: \mathcal{X} \subseteq \mathbb{R}^{n} \rightarrow\{1,2\}$


Figure 2.4: Linear classifier based on the Fisher Linear Discriminant.
be a binary linear classifier (2.2) with discriminant function $f(\boldsymbol{x})=\langle\boldsymbol{w} \cdot \boldsymbol{x}\rangle+b$. The probability of misclassification is defined as

$$
\operatorname{Err}(\boldsymbol{w}, b)=\max _{i \in \mathcal{I}_{1} \cup \mathcal{I}_{2}} \varepsilon\left(\boldsymbol{w}, b, \boldsymbol{\mu}^{i}, \boldsymbol{\Sigma}^{i}\right),
$$

where $\varepsilon\left(\boldsymbol{w}, b, \boldsymbol{\mu}^{i}, \boldsymbol{\Sigma}^{\mathbf{i}}\right)$ is probability that the Gaussian random vector $\boldsymbol{x}$ with mean vector $\boldsymbol{\mu}^{i}$ and the covariance matrix $\boldsymbol{\Sigma}^{i}$ satisfies $q(\boldsymbol{x})=1$ for $i \in \mathcal{I}_{2}$ or $q(\boldsymbol{x})=2$ for $i \in \mathcal{I}_{1}$. In other words, it is the probability that the vector $\boldsymbol{x}$ will be misclassified by the linear rule $q$.

The Generalized Anderson's task (GAT) is to find the parameters $\left(\boldsymbol{w}^{*}, b^{*}\right)$ of the linear classifier

$$
q(\boldsymbol{x})= \begin{cases}1, & \text { for } \quad f(\boldsymbol{x})=\left\langle\boldsymbol{w}^{*} \cdot \boldsymbol{x}\right\rangle+b^{*} \geq 0 \\ 2, & \text { for } \quad f(\boldsymbol{x})=\left\langle\boldsymbol{w}^{*} \cdot \boldsymbol{x}\right\rangle+b^{*}<0,\end{cases}
$$

such that the error $\operatorname{Err}\left(\boldsymbol{w}^{*}, b^{*}\right)$ is minimal

$$
\begin{equation*}
\left(\boldsymbol{w}^{*}, b^{*}\right)=\underset{\boldsymbol{w}, b}{\operatorname{argmin}} \operatorname{Err}(\boldsymbol{w}, b)=\underset{\boldsymbol{w}, b}{\operatorname{argmin}} \max _{i \in \mathcal{I}_{1} \cup \mathcal{I}_{2}} \varepsilon\left(\boldsymbol{w}, b, \boldsymbol{\mu}^{i}, \boldsymbol{\Sigma}^{i}\right) . \tag{2.14}
\end{equation*}
$$

The original Anderson's task is a special case of (2.14) when $\left|\mathcal{I}_{1}\right|=1$ and $\left|\mathcal{I}_{2}\right|=1$.
The probability $\varepsilon\left(\boldsymbol{w}, b, \boldsymbol{\mu}^{i}, \boldsymbol{\Sigma}^{i}\right)$ is proportional to the reciprocal of the Mahalanobis distance $r^{i}$ between the ( $\boldsymbol{\mu}^{i}, \boldsymbol{\Sigma}^{i}$ ) and the nearest vector of the separating hyperplane $\mathcal{H}=\left\{\boldsymbol{x} \in \mathbb{R}^{n}:\langle\boldsymbol{w} \cdot \boldsymbol{x}\rangle+b=0\right\}$, i.e.,

$$
r^{i}=\min _{\boldsymbol{x} \in \mathcal{H}}\left\langle\left(\boldsymbol{\mu}^{i}-\boldsymbol{x}\right) \cdot\left(\boldsymbol{\Sigma}^{i}\right)^{-1}\left(\boldsymbol{\mu}^{i}-\boldsymbol{x}\right)\right\rangle=\frac{\left\langle\boldsymbol{w} \cdot \boldsymbol{\mu}^{i}\right\rangle+b}{\sqrt{\left\langle\boldsymbol{w} \cdot \boldsymbol{\Sigma}^{i} \boldsymbol{w}\right\rangle}} .
$$

The exact relation between the probability $\varepsilon\left(\boldsymbol{w}, b, \boldsymbol{\mu}^{i}, \boldsymbol{\Sigma}^{\mathbf{i}}\right)$ and the corresponding Mahalanobis distance $r^{i}$ is given by the integral

$$
\begin{equation*}
\varepsilon\left(\boldsymbol{w}, b, \boldsymbol{\mu}^{i}, \boldsymbol{\Sigma}^{\mathbf{i}}\right)=\int_{r^{i}}^{\infty} \frac{1}{\sqrt{2 \pi}} e^{-\frac{1}{2} t^{2}} d t \tag{2.15}
\end{equation*}
$$

The optimization problem (2.14) can be equivalently rewritten as

$$
\left(\boldsymbol{w}^{*}, b^{*}\right)=\underset{\boldsymbol{w}, b}{\operatorname{argmax}} F(\boldsymbol{w}, b)=\underset{\boldsymbol{w}, b}{\operatorname{argmax}} \min _{i \in \mathcal{I}_{1} \cup \mathcal{I}_{2}} \frac{\left\langle\boldsymbol{w} \cdot \boldsymbol{\mu}^{i}\right\rangle+b}{\sqrt{\left\langle\boldsymbol{w} \cdot \boldsymbol{\Sigma}^{i} \boldsymbol{w}\right\rangle}} .
$$

which is more suitable for optimization. The objective function $F(\boldsymbol{w}, b)$ is proven to be convex in the region where the probability of misclassification $\operatorname{Err}(\boldsymbol{w}, b)$ is less than 0.5 . However, the objective function $F(\boldsymbol{w}, b)$ is not differentiable.

The STPRtool contains implementations of the algorithm solving the original Anderson's task as well as implementations of three different approaches to solve the Generalized Anderson's task which are described bellow. An interactive demo on the algorithms solving the Generalized Anderson's task is implemented in demo_anderson.
References: The original Anderson's task was published in [1]. A detailed description of the Generalized Anderson's task and all the methods implemented in the STPRtool is given in book [26].

### 2.4.1 Original Anderson's algorithm

The algorithm is implemented in the function androrig. This algorithm solves the original Anderson's task defined for two Gaussians only $\left|\mathcal{I}_{1}\right|=\left|\mathcal{I}_{2}\right|=1$. In this case, the optimal solution vector $\boldsymbol{w} \in \mathbb{R}^{n}$ is obtained by solving the following problem

$$
\begin{aligned}
\boldsymbol{w} & =\left((1-\lambda) \boldsymbol{\Sigma}_{1}+\lambda \boldsymbol{\Sigma}_{2}\right)^{-1}\left(\boldsymbol{\mu}_{1}-\boldsymbol{\mu}_{2}\right) \\
\frac{1-\lambda}{\lambda} & =\sqrt{\frac{\left\langle\boldsymbol{w} \cdot \boldsymbol{\Sigma}_{2} \boldsymbol{w}\right\rangle}{\left\langle\boldsymbol{w} \cdot \boldsymbol{\Sigma}_{1} \boldsymbol{w}\right\rangle}}
\end{aligned}
$$

where the scalar $0<\lambda<1$ is unknown. The problem is solved by an iterative algorithm which stops while the condition

$$
\left|\gamma-\frac{1-\lambda}{\lambda}\right| \leq \varepsilon, \quad \gamma=\sqrt{\frac{\left\langle\boldsymbol{w} \cdot \boldsymbol{\Sigma}_{2} \boldsymbol{w}\right\rangle}{\left\langle\boldsymbol{w} \cdot \boldsymbol{\Sigma}_{1} \boldsymbol{w}\right\rangle}}
$$

is satisfied. The parameter $\varepsilon$ defines the closeness to the optimal solution.

## Example: Solving the original Anderson's task

The original Anderson's task is solved for the Riply's data set riply_trn.mat. The parameters $\left(\boldsymbol{\mu}_{1}, \boldsymbol{\Sigma}_{1}\right)$ and $\left(\boldsymbol{\mu}_{2}, \boldsymbol{\Sigma}_{2}\right)$ of the Gaussian distribution of the first and the second class are obtained by the Maximum-Likelihood estimation. The found linear classifier is visualized and the Gaussian distributions (see Figure 2.5). The classifier is validated on the testing set riply_tst.mat.

```
trn = load('riply_trn'); % load training data
distrib = mlcgmm(data); % estimate Gaussians
model = androrig(distrib); % solve Anderson's task
figure;
pandr( model, distrib ); % plot solution
tst = load('riply_tst'); % load testing data
ypred = linclass( tst.X, model ); % classify testing data
cerror( ypred, tst.y ) % evaluate error
```

ans =
0.1150


Figure 2.5: Solution of the original Anderson's task.

### 2.4.2 General algorithm framework

The algorithm is implemented in the function ganders. This algorithm repeats the following steps:
I. An improving direction $(\Delta \boldsymbol{w}, \Delta b)$ in the parameter space is computed such that moving along this direction increases the objective function $F$. The quadprog function of the Matlab optimization toolbox is used to solve this subtask.
II. The optimal movement $k$ along the direction $(\boldsymbol{\Delta} \boldsymbol{w}, \Delta b)$. The optimal movement $k$ is determined by $1 D$-search based on the cutting interval algorithm according to the Fibonacci series. The number of interval cuttings of is an optional parameter of the algorithm.
III. The new parameters are set

$$
\boldsymbol{w}^{(t+1)}:=\boldsymbol{w}^{(t)}+k \boldsymbol{\Delta} \boldsymbol{w}, \quad b^{(t+1)}:=b^{(t)}+k \Delta b .
$$

The algorithm iterates until the minimal change in the objective function is less than the prescribed $\varepsilon$, i.e., the condition

$$
F\left(\boldsymbol{w}^{(t+1)}, b^{(t+1)}\right)-F\left(\boldsymbol{w}^{(t)}, b^{(t)}\right)>\varepsilon,
$$

is violated.

## Example: Solving the Generalized Anderson's task

The Generalized Anderson's task is solved for the data mars.mat. Both the first and second class are described by three $2 D$-dimensional Gaussian distributions. The found linear classifier as well as the input Gaussian distributions are visualized (see Figure 2.6).

```
distrib = load('mars'); % load input Gaussians
model = ganders(distrib); % solve the GAT
figure;
pandr(model,distrib); % plot the solution
```


### 2.4.3 $\varepsilon$-solution using the Kozinec's algorithm

This algorithm is implemented in the function eanders. The $\varepsilon$-solution is defined as a linear classifier with parameters $(\boldsymbol{w}, b)$ such that

$$
F(\boldsymbol{w}, b)<\varepsilon
$$



Figure 2.6: Solution of the Generalized Anderson's task.
is satisfied. The prescribed $\varepsilon \in(0,0.5)$ is the desired upper bound on the probability of misclassification. The problem is transformed to the task of separation of two sets of ellipsoids which is solved the Kozinec's algorithm. The algorithm converges in finite number of iterations if the given $\varepsilon$-solution exists. Otherwise it can get stuck in an endless loop.

## Example: $\varepsilon$-solution of the Generalized Anderson's task

The $\varepsilon$-solution of the generalized Anderson's task is solved for the data mars.mat. The desired maximal possible probability of misclassification is set to 0.06 (it is $6 \%$ ). Both the first and second class are described by three $2 D$-dimensional Gaussian distributions. The found linear classifier and the input Gaussian distributions are visualized (see Figure 2.7).

```
distrib = load('mars'); % load Gaussians
options = struct('err',0.06'); % maximal desired error
model = eanders(distrib,options); % training
figure; pandr(model,distrib); % visualization
```


### 2.4.4 Generalized gradient optimization

This algorithm is implemented in the function ggradandr. The objective function $F(\boldsymbol{w}, b)$ is convex but not differentiable thus the standard gradient optimization cannot


Figure 2.7: $\varepsilon$-solution of the Generalized Anderson's task with maximal 0.06 probability of misclassification.
be used. However, the generalized gradient optimization method can be applied. The algorithm repeats the following two steps:
I. The generalized gradient $(\Delta \boldsymbol{w}, \Delta b)$ is computed.
II. The optimized parameters are updated

$$
\boldsymbol{w}^{(t+1)}:=\boldsymbol{w}^{(t)}+\Delta \boldsymbol{w}, \quad b^{(t+1)}=b^{(t)}+\Delta b .
$$

The algorithm iterates until the minimal change in the objective function is less than the prescribed $\varepsilon$, i.e., the condition

$$
F\left(\boldsymbol{w}^{(t+1)}, b^{(t+1)}\right)-F\left(\boldsymbol{w}^{(t)}, b^{(t)}\right)>\varepsilon
$$

is violated. The maximal number of iterations can be used as the stopping condition as well.

## Example: Solving the Generalized Anderson's task using the generalized gradient optimization

The generalized Anderson's task is solved for the data mars.mat. Both the first and second class are described by three $2 D$-dimensional Gaussian distributions. The maximal number of iterations is limited to $\operatorname{tmax}=10000$. The found linear classifier as well as the input Gaussian distributions are visualized (see Figure 2.8).

```
distrib = load('mars'); % load Gaussians
options = struct('tmax',10000); % max # of iterations
model = ggradandr(distrib,options); % training
figure;
pandr( model, distrib ); % visualization
```



Figure 2.8: Solution of the Generalized Anderson's task after 10000 iterations of the generalized gradient optimization algorithm.

## Chapter 3

## Feature extraction

The feature extraction step consists of mapping the input vector of observations $\boldsymbol{x} \in \mathbb{R}^{n}$ onto a new feature description $\boldsymbol{z} \in \mathbb{R}^{m}$ which is more suitable for given task. The linear projection

$$
\begin{equation*}
\boldsymbol{z}=\mathbf{W}^{T} \boldsymbol{x}+\boldsymbol{b} \tag{3.1}
\end{equation*}
$$

is commonly used. The projection is given by the matrix $\mathbf{W}[n \times m]$ and bias vector $\boldsymbol{b} \in \mathbb{R}^{m}$. The Principal Component Analysis (PCA) (Section 3.1) is a representative of the unsupervised learning method which yields the linear projection (3.1). The Linear Discriminant Analysis (LDA) (Section 3.2) is a representative of the supervised learning method yielding the linear projection. The linear data projection is implemented in function linproj (See examples in Section 3.1 and Section 3.2). The data type used to describe the linear data projection is defined in Table 3.

The quadratic data mapping projects the input vector $x \in \mathbb{R}^{n}$ onto the feature space $\mathbb{R}^{m}$ where $m=n(n+3) / 2$. The quadratic mapping $\phi: \mathbb{R}^{n} \rightarrow \mathbb{R}^{m}$ is defined as

$$
\begin{gather*}
\boldsymbol{\phi}(\boldsymbol{x})=\left[\begin{array}{cccc}
x_{1}, & x_{2}, & \ldots, & x_{n} \\
x_{1} x_{1}, & x_{1} x_{2}, & \ldots, & x_{1} x_{n} \\
& x_{2} x_{2}, & \ldots, & x_{2} x_{n} \\
& & \vdots & \\
& & & x_{n} x_{n}
\end{array}\right]^{T},
\end{gather*}
$$

where $x_{i}, i=1, \ldots, n$ are entries of the vector $\boldsymbol{x} \in \mathbb{R}^{n}$. The quadratic data mapping is implemented in function qmap.

The kernel functions allow for non-linear extensions of the linear feature extraction methods. In this case, the input vectors $\boldsymbol{x} \in \mathbb{R}^{n}$ are mapped into a new higher dimensional feature space $\mathcal{F}$ in which the linear methods are applied. This yields a non-linear (kernel) projection of data which has generally defined as

$$
\begin{equation*}
\boldsymbol{z}=\mathbf{A}^{T} \boldsymbol{k}(\boldsymbol{x})+\boldsymbol{b} \tag{3.3}
\end{equation*}
$$

where $\mathbf{A}[l \times m]$ is a parameter matrix and $\boldsymbol{b} \in \mathbb{R}^{m}$ a parameter vector. The vector $\boldsymbol{k}(\boldsymbol{x})=\left[k\left(\boldsymbol{x}, \boldsymbol{x}_{1}\right), \ldots, k\left(\boldsymbol{x}, \boldsymbol{x}_{l}\right)\right]^{T}$ is a vector of kernel functions centered in the training data or their subset. The $\boldsymbol{x} \in \mathbb{R}^{n}$ stands for the input vector and $\boldsymbol{z} \in \mathbb{R}^{m}$ is the vector of extracted features. The extracted features are projections of $\boldsymbol{\phi}(\boldsymbol{x})$ onto $m$ vectors (or functions) from the feature space $\mathcal{F}$. The Kernel PCA (Section 3.3) and the Generalized Discriminant Analysis (GDA) (Section 3.5) are representatives of the feature extraction methods yielding the kernel data projection (3.3). The kernel data projection is implemented in the function kernelproj (See examples in Section 3.3 and Section 3.5). The data type used to describe the kernel data projection is defined in Table 3. The summary of implemented methods for feature extraction is given in Table 3.1

Table 3.1: Implemented methods for feature extraction

| linproj | Linear data projection. |
| :--- | :--- |
| kerproj | Kernel data projection. |
| qmap | Quadratic data mapping. |
| lin2quad | Merges linear rule and quadratic mapping. |
| lin2svm | Merges linear rule and kernel projection. |
| lda | Linear Discriminant Analysis. |
| pca | Principal Component Analysis. |
| pcarec | Computes reconstructed vector after PCA projection. |
| gda | Generalized Discriminant Analysis. |
| greedykpca | Greedy Kernel Principal Component Analysis. |
| kpca | Kernel Principal Component Analysis. |
| kpcarec | Reconstructs image after kernel PCA. |
| demo_pcacomp | Demo on image compression using PCA. |

Table 3.2: Data-type used to describe linear data projection.

| Linear data projection (structure array): |  |
| :--- | :--- |
| . $\mathrm{W}[n \times m]$ | Projection matrix $\mathrm{W}=\left[\boldsymbol{w}_{1}, \ldots, \boldsymbol{w}_{m}\right]$. |
| . $\mathrm{b}[m \times 1]$ | Bias vector $b$. |
| . fun $=$ 'linproj' | Identifies function associated with this data type. |

### 3.1 Principal Component Analysis

Let $\mathcal{T}_{X}=\left\{\boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{l}\right\}$ be a set of training vectors from the $n$-dimensional input space $\mathbb{R}^{n}$. The set of vectors $\mathcal{T}_{Z}=\left\{\boldsymbol{z}_{1}, \ldots, \boldsymbol{z}_{l}\right\}$ is a lower dimensional representation of the

Table 3.3: Data-type used to describe kernel data projection.

| Kernel data projection | (structure array): |
| :--- | :--- |
| . $\mathrm{Alpha}[l \times m]$ | Parameter matrix $\mathbf{A}[l \times m]$. |
| . $\mathrm{b}[m \times 1]$ | Bias vector $b$. |
| . sv.X $[n \times l]$ | Training vectors $\left\{\boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{l}\right\}$. |
| .options.ker $[$ string $]$ | Kernel identifier. |
| .options.arg $[1 \times p]$ | Kernel argument(s). |
| . fun $=$ 'kernelproj' | Identifies function associated with this data type. |

input training vectors $\mathcal{T}_{X}$ in the $m$-dimensional space $\mathbb{R}^{m}$. The vectors $\mathcal{T}_{Z}$ are obtained by the linear orthonormal projection

$$
\begin{equation*}
\boldsymbol{z}=\mathbf{W}^{T} \boldsymbol{x}+\boldsymbol{b} \tag{3.4}
\end{equation*}
$$

where the matrix $\mathbf{W}[n \times m]$ and the vector $\boldsymbol{b}[m \times 1]$ are parameters of the projection. The reconstructed vectors $\mathcal{T}_{\tilde{X}}=\left\{\tilde{\boldsymbol{x}}_{1}, \ldots, \tilde{\boldsymbol{x}}_{l}\right\}$ are computed by the linear back projection

$$
\begin{equation*}
\tilde{\boldsymbol{x}}=\mathbf{W}(\boldsymbol{z}-\boldsymbol{b}), \tag{3.5}
\end{equation*}
$$

obtained by inverting (3.4). The mean square reconstruction error

$$
\begin{equation*}
\varepsilon_{M S}(\mathbf{W}, \boldsymbol{b})=\frac{1}{l} \sum_{i=1}^{l}\left\|\boldsymbol{x}_{i}-\tilde{\boldsymbol{x}}_{i}\right\|^{2}, \tag{3.6}
\end{equation*}
$$

is a function of the parameters of the linear projections (3.4) and (3.5). The Principal Component Analysis (PCA) is the linear orthonormal projection (3.4) which allows for the minimal mean square reconstruction error (3.6) of the training data $\mathcal{T}_{X}$. The parameters $(\mathbf{W}, \boldsymbol{b})$ of the linear projection are the solution of the optimization task

$$
\begin{equation*}
(\mathbf{W}, \boldsymbol{b})=\underset{\mathbf{W}^{\prime}, b^{\prime}}{\operatorname{argmin}} \varepsilon_{M S}\left(\mathbf{W}^{\prime}, \boldsymbol{b}^{\prime}\right) \tag{3.7}
\end{equation*}
$$

subject to

$$
\left\langle\boldsymbol{w}_{i} \cdot \boldsymbol{w}_{j}\right\rangle=\delta(i, j), \quad \forall i, j,
$$

where $\boldsymbol{w}_{i}, i=1, \ldots, m$ are column vectors of the matrix $\mathbf{W}=\left[\boldsymbol{w}_{1}, \ldots, \boldsymbol{w}_{m}\right]$ and $\delta(i, j)$ is the Kronecker delta function. The solution of the task (3.7) is the matrix $\mathbf{W}=\left[\boldsymbol{w}_{1}, \ldots, \boldsymbol{w}_{m}\right]$ containing the $m$ eigenvectors of the sample covariance matrix which have the largest eigen values. The vector $\boldsymbol{b}$ equals to $\mathbf{W}^{T} \boldsymbol{\mu}$, where $\mu$ is the sample mean of the training data. The PCA is implemented in the function pca. A demo showing the use of PCA for image compression is implemented in the script demo_pcacomp.

References: The PCA is treated throughout for instance in books [13, 3, 6].

## Example: Principal Component Analysis

The input data are synthetically generated from the 2-dimensional Gaussian distribution. The PCA is applied to train the 1-dimensional subspace to approximate the input data with the minimal reconstruction error. The training data, the reconstructed data and the 1-dimensional subspace given by the first principal component are visualized (see Figure 3.1).

```
X = gsamp([5;5],[1 0.8;0.8 1],100); % generate data
model = pca(X,1); % train PCA
Z = linproj(X,model); % lower dim. proj.
XR = pcarec(X,model); % reconstr. data
figure; hold on; axis equal; % visualization
h1 = ppatterns(X,'kx');
h2 = ppatterns(XR,'bo');
[dummy,mn] = min(Z);
[dummy,mx] = max(Z);
h3 = plot([XR(1,mn) XR(1,mx)],[XR(2,mn) XR(2,mx)],'r');
legend([h1 h2 h3], ...
    'Input vectors','Reconstructed', 'PCA subspace');
```



Figure 3.1: Example on the Principal Component Analysis.

### 3.2 Linear Discriminant Analysis

Let $\mathcal{T}_{X Y}=\left\{\left(\boldsymbol{x}_{1}, y_{1}\right), \ldots,\left(\boldsymbol{x}_{l}, y_{l}\right)\right\}, \boldsymbol{x}_{i} \in \mathbb{R}^{n}, y \in \mathcal{Y}=\{1,2, \ldots, c\}$ be labeled set of training vectors. The within-class $\mathbf{S}_{\mathbf{W}}$ and between-class $\mathbf{S}_{\mathbf{B}}$ scatter matrices are described as

$$
\begin{align*}
\mathbf{S}_{\mathbf{W}} & =\sum_{y \in \mathcal{Y}} \mathbf{S}_{y}, \quad \mathbf{S}_{y}=\sum_{i \in \mathcal{I}_{y}}\left(\boldsymbol{x}_{i}-\boldsymbol{\mu}_{i}\right)\left(\boldsymbol{x}_{i}-\boldsymbol{\mu}_{i}\right)^{T}  \tag{3.8}\\
\mathbf{S}_{\mathbf{B}} & =\sum_{y \in \mathcal{Y}}\left|\mathcal{I}_{y}\right|\left(\boldsymbol{\mu}_{y}-\boldsymbol{\mu}\right)\left(\boldsymbol{\mu}_{y}-\boldsymbol{\mu}\right)^{T}
\end{align*}
$$

where the total mean vector $\boldsymbol{\mu}$ and the class mean vectors $\boldsymbol{\mu}_{y}, y \in \mathcal{Y}$ are defined as

$$
\boldsymbol{\mu}=\frac{1}{l} \sum_{i=1}^{l} \boldsymbol{x}_{i}, \quad \boldsymbol{\mu}_{y}=\frac{1}{\left|\mathcal{I}_{y}\right|} \sum_{i \in \mathcal{I}_{y}} \boldsymbol{x}_{i}, y \in \mathcal{Y} .
$$

The goal of the Linear Discriminant Analysis (LDA) is to train the linear data projection

$$
\boldsymbol{z}=\mathbf{W}^{T} \boldsymbol{x}
$$

such that the class separability criterion

$$
F(\mathbf{W})=\frac{\operatorname{det}\left(\tilde{\mathbf{S}}_{\mathbf{B}}\right)}{\operatorname{det}\left(\tilde{\mathbf{S}}_{\mathbf{W}}\right)}=\frac{\operatorname{det}\left(\mathbf{S}_{B}\right)}{\operatorname{det}\left(\mathbf{S}_{W}\right)},
$$

is maximized. The $\tilde{\mathbf{S}}_{\mathrm{B}}$ is the between-class and $\tilde{\mathbf{S}}_{\mathrm{W}}$ is the within-class scatter matrix of projected data which is defined similarly to (3.8).
References: The LDA is treated for instance in books [3, 6].

## Example 1: Linear Discriminant Analysis

The LDA is applied to extract 2 features from the Iris data set iris.mat which consists of the labeled 4-dimensional data. The data after the feature extraction step are visualized in Figure 3.2.

```
orig_data = load('iris'); % load input data
model = lda(orig_data,2); % train LDA
ext_data = linproj(orig_data,model); % feature extraction
figure; ppatterns(ext_data); % plot ext. data
```


## Example 2: PCA versus LDA

This example shows why the LDA is superior to the PCA when features good for classification are to be extracted. The synthetical data are generated from the Gaussian mixture model. The LDA and PCA are trained on the generated data. The LDA and


Figure 3.2: Example shows 2-dimensional data extracted from the originally 4dimensional Iris data set using Linear Discriminant Analysis.

PCA directions are displayed as connection of the vectors projected to the trained LDA and PCA subspaces and re-projected back to the input space (see Figure 3.3a). The extracted data using the LDA and PCA model are displayed with the Gaussians fitted by the Maximum-Likelihood (see Figure 3.3b).

```
% Generate data
distrib.Mean = [[5;4] [4;5]]; % mean vectors
distrib.Cov(:,:,1) = [1 0.9; 0.9 1]; % 1st covariance
distrib.Cov(:,:,2) = [1 0.9; 0.9 1]; % 2nd covariance
distrib.Prior = [0.5 0.5]; % Gaussian weights
data = gmmsamp(distrib,250); % sample data
lda_model = lda(data,1); % train LDA
lda_rec = pcarec(data.X,lda_model);
lda_data = linproj(data,lda_model);
pca_model = pca(data.X,1); % train PCA
pca_rec = pcarec(data.X,pca_model);
pca_data = linproj( data,pca_model);
figure; hold on; axis equal; % visualization
ppatterns(data);
```

```
h1 = plot(lda_rec(1,:),lda_rec(2,:),'r');
h2 = plot(pca_rec(1,:),pca_rec(2,:),'b');
legend([h1 h2],'LDA direction','PCA direction');
figure; hold on;
subplot(2,1,1); title('LDA'); ppatterns(lda_data);
pgauss(mlcgmm(lda_data));
subplot(2,1,2); title('PCA'); ppatterns(pca_data);
pgauss(mlcgmm(pca_data));
```


### 3.3 Kernel Principal Component Analysis

The Kernel Principal Component Analysis (Kernel PCA) is the non-linear extension of the ordinary linear PCA. The input training vectors $\mathcal{T}_{X}=\left\{\boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{l}\right\}, \boldsymbol{x}_{i} \in \mathcal{X} \subseteq \mathbb{R}^{n}$ are mapped by $\phi: \mathcal{X} \rightarrow \mathcal{F}$ to a high dimensional feature space $\mathcal{F}$. The linear PCA is applied on the mapped data $\mathcal{T}_{\Phi}=\left\{\boldsymbol{\phi}\left(\boldsymbol{x}_{1}\right), \ldots, \boldsymbol{\phi}\left(\boldsymbol{x}_{l}\right)\right\}$. The computation of the principal components and the projection on these components can be expressed in terms of dot products thus the kernel functions $k: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ can be employed. The kernel PCA trains the kernel data projection

$$
\begin{equation*}
\boldsymbol{z}=\mathbf{A}^{T} \boldsymbol{k}(\boldsymbol{x})+\boldsymbol{b} \tag{3.9}
\end{equation*}
$$

such that the reconstruction error

$$
\begin{equation*}
\varepsilon_{K M S}(\mathbf{A}, \boldsymbol{b})=\frac{1}{l} \sum_{i=1}^{l}\left\|\boldsymbol{\phi}\left(\boldsymbol{x}_{i}\right)-\tilde{\boldsymbol{\phi}}\left(\boldsymbol{x}_{i}\right)\right\|^{2}, \tag{3.10}
\end{equation*}
$$

is minimized. The reconstructed vector $\tilde{\boldsymbol{\phi}}(\boldsymbol{x})$ is given as a linear combination of the mapped data $\mathcal{T}_{\Phi}$

$$
\begin{equation*}
\tilde{\boldsymbol{\phi}}(\boldsymbol{x})=\sum_{i=1}^{l} \beta_{i} \boldsymbol{\phi}\left(\boldsymbol{x}_{i}\right), \quad \boldsymbol{\beta}=\mathbf{A}(\boldsymbol{z}-\boldsymbol{b}) . \tag{3.11}
\end{equation*}
$$

In contrast to the linear PCA, the explicit projection from the feature space $\mathcal{F}$ to the input space $\mathcal{X}$ usually do not exist. The problem is to find the vector $\boldsymbol{x} \in \mathcal{X}$ its image $\phi(\boldsymbol{x}) \in \mathcal{F}$ well approximates the reconstructed vector $\tilde{\boldsymbol{\phi}}(\boldsymbol{x}) \in \mathcal{F}$. This procedure is implemented in the function kpcarec for the case of the Radial Basis Function (RBF) kernel. The procedure consists of the following step:
I. Project input vector $\boldsymbol{x}_{i n} \in \mathbb{R}^{n}$ onto its lower dimensional representation $\boldsymbol{z} \in \mathbb{R}^{m}$ using (3.9).
II. Computes vector $\boldsymbol{x}_{\text {out }} \in \mathbb{R}^{n}$ which is good pre-image of the reconstructed vector $\tilde{\phi}\left(\boldsymbol{x}_{i n}\right)$, such that

$$
\boldsymbol{x}_{\text {out }}=\underset{\boldsymbol{x}}{\operatorname{argmin}}\left\|\phi(\boldsymbol{x})-\tilde{\boldsymbol{\phi}}\left(\boldsymbol{x}_{i n}\right)\right\|^{2} .
$$

References: The kernel PCA is described at length in book [29, 30].

## Example: Kernel Principal Component Analysis

The example shows using the kernel PCA for data denoising. The input data are synthetically generated 2-dimensional vectors which lie on a circle and are corrupted by the Gaussian noise. The kernel PCA model with RBF kernel is trained. The function kpcarec is used to find pre-images of the reconstructed data (3.11). The result is visualized in Figure 3.4.

```
X = gencircledata([1;1],5,250,1); % generate circle data
options.ker = 'rbf'; % use RBF kernel
options.arg = 4; % kernel argument
options.new_dim = 2; % output dimension
model = kpca(X,options); % compute kernel PCA
XR = kpcarec(X,model); % compute reconstruced data
figure; % Visualization
h1 = ppatterns(X);
h2 = ppatterns(XR, '+r');
legend([h1 h2],'Input vectors','Reconstructed');
```


### 3.4 Greedy kernel PCA algorithm

The Greedy Kernel Principal Analysis (Greedy kernel PCA) is an efficient algorithm to compute the ordinary kernel PCA. Let $\mathcal{T}_{X}=\left\{\boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{l}\right\}, \boldsymbol{x}_{i} \in \mathbb{R}^{n}$ be the set of input training vectors. The goal is to train the kernel data projection

$$
\begin{equation*}
\boldsymbol{z}=\mathbf{A}^{T} \boldsymbol{k}_{S}(\boldsymbol{x})+\boldsymbol{b} \tag{3.12}
\end{equation*}
$$

where $\mathbf{A}[d \times m]$ is the parameter matrix, $\boldsymbol{b}[m \times 1]$ is the parameter vector and $\boldsymbol{k}_{S}=\left[k\left(\boldsymbol{x}, \boldsymbol{s}_{1}\right), \ldots, k\left(\boldsymbol{x}, \boldsymbol{s}_{l}\right)\right]^{T}$ are the kernel functions centered in the vectors $\mathcal{T}_{S}=$ $\left\{s_{1}, \ldots, s_{d}\right\}$. The vector set $\mathcal{T}_{S}$ is a subset of training data $\mathcal{T}_{X}$. In contrast to the ordinary kernel PCA, the subset $\mathcal{T}_{S}$ does not contain all the training vectors $\mathcal{T}_{X}$ thus the complexity of the projection (3.12) is reduced compared to (3.9). The objective of the Greedy kernel PCA is to minimize the reconstruction error (3.11) while the size $d$ of the subset $\mathcal{T}_{S}$ is kept small. The Greedy kernel PCA algorithm is implemented in the function greedykpca.

References: The implemented Greedy kernel PCA algorithm is described in [11].

## Example: Greedy Kernel Principal Component Analysis

The example shows using kernel PCA for data denoising. The input data are synthetically generated 2-dimensional vectors which lie on a circle and are corrupted with the Gaussian noise. The Greedy kernel PCA model with RBF kernel is trained. The number of vectors defining the kernel projection (3.12) is limited to 25 . In contrast the ordinary kernel PCA (see experiment in Section 3.3) requires all the 250 training vectors. The function kpcarec is used to find pre-images of the reconstructed data obtained by the kernel PCA projection. The training and reconstructed vectors are visualized in Figure 3.5. The vector of the selected subset $\mathcal{I}_{X}$ are visualized as well.

```
X = gencircledata([1;1],5,250,1); % generate training data
options.ker = 'rbf'; % use RBF kernel
options.arg = 4; % kernel argument
options.new_dim = 2; % output dimension
options.m = 25; % size of sel. subset
model = greedykpca(X,options); % run greedy algorithm
XR = kpcarec(X,model); % reconstructed vectors
figure; % visualization
h1 = ppatterns(X);
h2 = ppatterns(XR,'+r');
h3 = ppatterns(model.sv.X,'ob',12);
legend([h1 h2 h3], ...
    'Training set','Reconstructed set','Selected subset');
```


### 3.5 Generalized Discriminant Analysis

The Generalized Discriminant Analysis (GDA) is the non-linear extension of the ordinary Linear Discriminant Analysis (LDA). The input training data $\mathcal{T}_{X Y}=\left\{\left(\boldsymbol{x}_{1}, y_{1}\right), \ldots,\left(\boldsymbol{x}_{l}, y_{l}\right)\right\}$, $\boldsymbol{x}_{i} \in \mathcal{X} \subseteq \mathbb{R}^{n}, y \in \mathcal{Y}=\{1, \ldots, c\}$ are mapped by $\phi: \mathcal{X} \rightarrow \mathcal{F}$ to a high dimensional feature space $\mathcal{F}$. The ordinary LDA is applied on the mapped data $\mathcal{T}_{\Phi Y}=$ $\left\{\left(\boldsymbol{\phi}\left(\boldsymbol{x}_{1}\right), y_{1}\right), \ldots,\left(\boldsymbol{\phi}\left(\boldsymbol{x}_{l}\right), y_{l}\right)\right\}$. The computation of the projection vectors and the projection on these vectors can be expressed in terms of dot products thus the kernel functions $k: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ can be employed. The resulting kernel data projection is defined as

$$
\boldsymbol{z}=\mathbf{A}^{T} \boldsymbol{k}(\boldsymbol{x})+\boldsymbol{b},
$$

where $\mathbf{A}[l \times m]$ is the matrix and $\boldsymbol{b}[m \times 1]$ the vector trained by the GDA. The parameters $(\mathbf{A}, \boldsymbol{b})$ are trained to increase the between-class scatter and decrease the within-class scatter of the extracted data $\mathcal{T}_{Z Y}=\left\{\left(\boldsymbol{z}_{1}, y_{1}\right), \ldots,\left(\boldsymbol{z}_{l}, y_{l}\right)\right\}, \boldsymbol{z}_{i} \in \mathbb{R}^{m}$. The GDA is implemented in the function gda.
References: The GDA was published in the paper [2].

## Example: Generalized Discriminant Analysis

The GDA is trained on the Iris data set iris.mat which contains 3 classes of 4 -dimensional data. The RBF kernel with kernel width $\sigma=1$ is used. Notice, that setting $\sigma$ too low leads to the perfect separability of the training data $\mathcal{T}_{X Y}$ but the kernel projection is heavily over-fitted. The extracted data $\mathcal{Y}_{Z Y}$ are visualized in Figure 3.6.

```
orig_data = load('iris'); % load data
options.ker = 'rbf'; % use RBF kernel
options.arg = 1; % kernel arg.
options.new_dim = 2; % output dimension
model = gda(orig_data,options); % train GDA
ext_data = kernelproj(orig_data,model); % feature ext.
figure; ppatterns(ext_data); % visualization
```


### 3.6 Combining feature extraction and linear classifier

Let $\mathcal{T}_{X Y}=\left\{\left(\boldsymbol{x}_{1}, y_{1}\right), \ldots,\left(\boldsymbol{x}_{l}, y_{l}\right)\right\}$ be a training set of observations $\boldsymbol{x}_{i} \in \mathcal{X} \subseteq \mathbb{R}^{n}$ and corresponding hidden states $y_{i} \in \mathcal{Y}=\{1, \ldots, c\}$. Let $\boldsymbol{\phi}: \mathcal{X} \rightarrow \mathbb{R}^{m}$ be a mapping

$$
\boldsymbol{\phi}(\boldsymbol{x})=\left[\phi_{1}(\boldsymbol{x}), \ldots, \phi_{m}(\boldsymbol{x})\right]^{T}
$$

defining a non-linear feature extraction step (e.g., quadratic data mapping or kernel projection). The feature extraction step applied to the input data $\mathcal{T}_{X Y}$ yields extracted data $\mathcal{T}_{\Phi Y}=\left\{\left(\boldsymbol{\phi}\left(\boldsymbol{x}_{1}\right), y_{1}\right), \ldots,\left(\boldsymbol{\phi}\left(\boldsymbol{x}_{l}\right), y_{l}\right)\right\}$. Let $f(\boldsymbol{x})=\langle\boldsymbol{w} \cdot \boldsymbol{\phi}(\boldsymbol{x})\rangle+b$ be a linear discriminant function found by an arbitrary training method on the extracted data $\mathcal{T}_{\Phi Y}$. The function $f$ can be seen as a non-linear function of the input vectors $\boldsymbol{x} \in \mathcal{X}$, as

$$
f(\boldsymbol{x})=\langle\boldsymbol{w} \cdot \boldsymbol{\phi}(\boldsymbol{x})\rangle+b=\sum_{i=1}^{m} w_{i} \phi_{i}(\boldsymbol{x})+b .
$$

In the case of the quadratic data mapping (3.2) used as the feature extraction step, the discriminant function can be written in the following form

$$
\begin{equation*}
f(\boldsymbol{x})=\langle\boldsymbol{x} \cdot \mathbf{A} \boldsymbol{x}\rangle+\langle\boldsymbol{x} \cdot \boldsymbol{b}\rangle+c . \tag{3.13}
\end{equation*}
$$

The discriminant function (3.13) defines the quadratic classifier described in Section 6.4. Therefore the quadratic classifier can be trained by (i) applying the quadratic data mapping qmap, (ii) training a linear rule on the mapped data (e.g., perceptron, fld) and (iii) combining the quadratic mapping and the linear rule. The combining is implement in function lin2quad.

In the case of the kernel projection (3.3) applied as the feature extraction step, the resulting discriminant function has the following form

$$
\begin{equation*}
f(\boldsymbol{x})=\langle\boldsymbol{\alpha} \cdot \boldsymbol{k}(\boldsymbol{x})\rangle+b \tag{3.14}
\end{equation*}
$$

The discriminant function (3.14) defines the kernel (SVM type) classifier described in Section 5. This classifier can be trained by (i) applying the kernel projection kernelproj, (ii) training a linear rule on the extracted data and (iii) combing the kernel projection and the linear rule. The combining is implemented in function lin2svm.

## Example: Quadratic classifier trained the Perceptron

This example shows how to train the quadratic classifier using the Perceptron algorithm. The Perceptron algorithm produces the linear rule only. The input training data are linearly non-separable but become separable after the quadratic data mapping is applied. Function lin2quad is used to compute the parameters of the quadratic classifier explicitly based on the found linear rule in the feature space and the mapping. Figure 3.7 shows the decision boundary of the quadratic classifier.

```
% load training data living in the input space
input_data = load('vltava');
% map data to the feature space
map_data = qmap(input_data);
% train linear rule in the feature sapce
lin_model = perceptron(map_data);
% compute parameters of the quadratic classifier
quad_model = lin2quad(lin_model);
% visualize the quadratic classifier
figure; ppatterns(input_data);
pboundary(quad_model);
```


## Example: Kernel classifier from the linear rule

This example shows how to train the kernel classifier using arbitrary training algorithm for linear rule. The Fisher Linear Discriminant was used in this example. The
greedy kernel PCA algorithm was applied as the feature extraction step. The Riply's data set was used for training and evaluation. Figure 3.8 shows the found kernel classifier, the training data and vectors used in the kernel expansion of the discriminant function.

```
    % load input data
    trn = load('riply_trn');
    % train kernel PCA by greedy algorithm
    options = struct('ker','rbf','arg',1,'new_dim',10);
    kpca_model = greedykpca(trn.X,options);
    % project data
map_trn = kernelproj(trn,kpca_model);
% train linear classifier
lin_model = fld(map_trn);
% combine linear rule and kernel PCA
kfd_model = lin2svm(kpca_model,lin_model);
% visualization
figure;
ppatterns(trn); pboundary(kfd_model);
ppatterns(kfd_model.sv.X,'ok',13);
% evaluation on testing data
tst = load('riply_tst');
ypred = svmclass(tst.X,kfd_model);
cerror(ypred,tst.y)
ans =
0.0970
```


(b)

Figure 3.3: Comparison between PCA and LDA. Figure (a) shows input data and found LDA (red) and PCA (blue) directions onto which the data are projected. Figure (b) shows the class conditional Gaussians estimated from data projected data onto LDA and PCA directions.


Figure 3.4: Example on the Kernel Principal Component Analysis.


Figure 3.5: Example on the Greedy Kernel Principal Component Analysis.


Figure 3.6: Example shows 2-dimensional data extracted from the originally 4dimensional Iris data set using the Generalized Discriminant Analysis. In contrast to LDA (see Figure 3.2 for comparison), the class clusters are more compact but at the expense of a higher risk of over-fitting.


Figure 3.7: Example shows quadratic classifier found by the Perceptron algorithm on the data mapped to the feature by the quadratic mapping.


Figure 3.8: Example shows the kernel classifier found by the Fisher Linear Discriminant algorithm on the data mapped to the feature by the kernel PCA. The greedy kernel PCA was used therefore only a subset of training data was used to determine the kernel projection.

## Chapter 4

## Density estimation and clustering

The STPRtool represents a probability distribution function $P_{X}(\boldsymbol{x}), \boldsymbol{x} \in \mathbb{R}^{n}$ as a structure array which must contain field .fun (see Table 4.2). The field .fun specifies function which is used to evaluate $P_{X}(\boldsymbol{x})$ for given set of realization $\left\{\boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{l}\right\}$ of a random vector. Let the structure model represent a given distribution function. The Matlab command

$$
y=\text { feval(X,model.fun) }
$$

is used to evaluate $y_{i}=P_{X}\left(\boldsymbol{x}_{i}\right)$. The matrix $\mathrm{X}[n \times l]$ contains column vectors $\left\{\boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{l}\right\}$. The vector y $[1 \times l]$ contains values of the distribution function. In particular, the Gaussian distribution and the Gaussian mixture models are described in Section 4.1. The summary of implemented methods for dealing with probability density estimation and clustering is given in Table 4.1

### 4.1 Gaussian distribution and Gaussian mixture model

The value of the multivariate Gaussian probability distribution function in the vector $\boldsymbol{x} \in \mathbb{R}^{n}$ is given by

$$
\begin{equation*}
P_{X}(\boldsymbol{x})=\frac{1}{(2 \pi)^{\frac{n}{2}} \sqrt{\operatorname{det}(\Sigma)}} \exp \left(-\frac{1}{2}(\boldsymbol{x}-\boldsymbol{\mu})^{T} \Sigma^{-1}(\boldsymbol{x}-\boldsymbol{\mu})\right) \tag{4.1}
\end{equation*}
$$

where $\boldsymbol{\mu}[n \times 1]$ is the mean vector and $\boldsymbol{\Sigma}[n \times n]$ is the covariance matrix. The STPRtool uses specific data-type to describe a set of Gaussians which can be labeled, i.e., each pair $\left(\boldsymbol{\mu}_{i}, \boldsymbol{\Sigma}_{i}\right)$ is assigned to the class $y_{i} \in \mathcal{Y}$. The covariance matrix $\boldsymbol{\Sigma}$ is always represented as the square matrix $[n \times n]$, however, the shape of the matrix can be indicated in the optimal field .cov_type (see Table 4.3). The data-type is described in Table 4.4. The evaluation of the Gaussian distribution is implemented in the function

Table 4.1: Implemented methods: Density estimation and clustering

| gmmsamp | Generates sample from Gaussian mixture model. |
| :--- | :--- |
| gsamp | Generates sample from Gaussian distribution. |
| kmeans | K-means clustering algorithm. |
| pdfgauss | Evaluates for multivariate Gaussian distribution. |
| pdfgauss | Evaluates Gaussian mixture model. |
| sigmoid | Evaluates sigmoid function. |
| emgmm | Expectation-Maximization Algorithm for Gaussian <br>  <br> mixture model. |
| melgmm | Maximizes Expectation of Log-Likelihood for Gaus- <br>  <br> sian mixture. |
| mlcgmm | Maximal Likelihood estimation of Gaussian mixture |
| mlsigmoid | model. |
| Fmgauss | Minting a sigmoid function using ML estimation of Gaussian distribution. |
| rsde | Reduced Set Density Estimator. |
| demo_emgmm | Demo on Expectation-Maximization (EM) algorithm. |
| demo_mmgauss | Demo on minimax estimation for Gaussian. |
| demo_svmpout | Fitting a posteriori probability to SVM output. |

Table 4.2: Data-type used to represent probability distribution function. Probability Distribution Function (structure array):
.fun [string] The name of a function which evaluates probability distribution $y=f e v a l(X$, model.fun ).
pdfgauss. The random sampling from the Gaussian distribution is implemented in the function gsamp.

The Gaussian Mixture Model (GMM) is weighted sum of the Gaussian distributions

$$
\begin{equation*}
P_{X}(\boldsymbol{x})=\sum_{y \in \mathcal{Y}} P_{Y}(y) P_{X \mid Y}(\boldsymbol{x} \mid y) \tag{4.2}
\end{equation*}
$$

where $P_{Y}(y), y \in \mathcal{Y}=\{1, \ldots, c\}$ are weights and $P_{X \mid Y}(\boldsymbol{x} \mid y)$ is the Gaussian distribution (4.1) given by parameters $\left(\boldsymbol{\mu}_{y}, \boldsymbol{\Sigma}_{y}\right)$. The data-type used to represent the GMM is described in Table 4.5. The evaluation of the GMM is implemented in the function pdfgmm. The random sampling from the GMM is implemented in the function gmmsamp.

## Example: Sampling from the Gaussian distribution

The example shows how to represent the Gaussian distribution and how to sample data it. The sampling from univariate and bivariate Gaussians is shown. The distribution function of the univariate Gaussian is visualized and compared to the histogram

Table 4.3: Shape of covariance matrix.

| Parameter cov_type: |  |
| :--- | :--- |
| 'full' | Full covariance matrix. |
| 'diag' | Diagonal covariance matrix. |
| 'spherical' | Spherical covariance matrix. |

Table 4.4: Data-type used to represent the labeled set of Gaussians.

| Labeled set of Gaussians (structure array): |  |
| :--- | :--- |
| . Mean $[n \times k]$ | Mean vectors $\left\{\boldsymbol{\mu}_{1}, \ldots, \boldsymbol{\mu}_{k}\right\}$. |
| . $\operatorname{Cov}[n \times n \times k]$ | Covariance matrices $\left\{\boldsymbol{\Sigma}_{1}, \ldots, \boldsymbol{\Sigma}_{k}\right\}$. |
| . y $[1 \times k]$ | Labels $\left\{y_{1}, \ldots, y_{k}\right\}$ assigned to Gaussians. |
| .Cov_type $[$ string $]$ | Optional field which indicates shape of the covariance |
|  | matrix (see Table 4.3). |
| . fun $=$ 'pdfgauss' | Identifies function associated with this data type. |

of the sample data (see Figure 4.1a). The isoline of the distribution function of the bivariate Gaussian is visualized together with the sample data (see Figure 4.1b).

```
% univariate case
model = struct('Mean',1,'Cov',2); % Gaussian parameters
figure; hold on;
% plot pdf. of the Gaussisan
plot([-4:0.1:5],pdfgauss([-4:0.1:5],model),'r');
[Y,X] = hist(gsamp(model,500),10); % compute histogram
bar(X,Y/500); % plot histogram
% bi-variate case
model.Mean = [1;1]; % Mean vector
model.Cov = [1 0.6; 0.6 1]; % Covariance matrix
figure; hold on;
ppatterns(gsamp(model,250)); % plot sampled data
pgauss(model); % plot shape
```


### 4.2 Maximum-Likelihood estimation of GMM

Let the probability distribution

$$
P_{X Y \mid \Theta}(\boldsymbol{x}, y \mid \boldsymbol{\theta})=P_{X \mid Y \Theta}(\boldsymbol{x} \mid y, \boldsymbol{\theta}) P_{Y \mid \Theta}(y \mid \boldsymbol{\theta}),
$$

Table 4.5: Data-type used to represent the Gaussian Mixture Model.

| Gaussian Mixture | Model (structure array): |
| :--- | :--- |
| . Mean $[n \times c]$ | Mean vectors $\left\{\boldsymbol{\mu}_{1}, \ldots, \boldsymbol{\mu}_{c}\right\}$. |
| .Cov $[n \times n \times c]$ | Covariance matrices $\left\{\boldsymbol{\Sigma}_{1}, \ldots, \boldsymbol{\Sigma}_{c}\right\}$. |
| .Prior $[c \times 1]$ | Weights of Gaussians $\left\{P_{K}(1), \ldots, P_{K}(c)\right\}$. |
| .fun $=$ 'pdfgmm' | Identifies function associated with this data type. |



Figure 4.1: Sampling from the univariate (a) and bivariate (b) Gaussian distribution.
be known up to parameters $\boldsymbol{\theta} \in \Theta$. The vector of observations $\boldsymbol{x} \in \mathbb{R}^{n}$ and the hidden state $y \in \mathcal{Y}=\{1, \ldots, c\}$ are assumed to be realizations of random variables which are independent and identically distributed according to the $P_{X Y \mid \Theta}(\boldsymbol{x}, y \mid \boldsymbol{\theta})$. The conditional probability distribution $P_{X \mid Y \Theta}(\boldsymbol{x} \mid y, \boldsymbol{\theta})$ is assumed to be Gaussian distribution. The $\boldsymbol{\theta}$ involves parameters $\left(\boldsymbol{\mu}_{y}, \boldsymbol{\Sigma}_{y}\right), y \in \mathcal{Y}$ of Gaussian components and the values of the discrete distribution $P_{Y \mid \Theta}(y \mid \boldsymbol{\theta}), y \in \mathcal{Y}$. The marginal probability $P_{X \mid \Theta}(\boldsymbol{x} \mid \boldsymbol{\theta})$ is the Gaussian mixture model (4.2). The Maximum-Likelihood estimation of the parameters $\boldsymbol{\theta}$ of the GMM for the case of complete and incomplete data is described in Section 4.2.1 and Section 4.2.2, respectively.

### 4.2.1 Complete data

The input is a set $\mathcal{T}_{X Y}=\left\{\left(\boldsymbol{x}_{1}, y_{1}\right), \ldots,\left(\boldsymbol{x}_{l}, y_{l}\right)\right\}$ which contains both vectors of observations $\boldsymbol{x}_{i} \in \mathbb{R}^{n}$ and corresponding hidden states $y_{i} \in \mathcal{Y}$. The pairs $\left(\boldsymbol{x}_{i}, y_{i}\right)$ are assumed to samples from random variables which are independent and identically distributed (i.i.d.) according to $P_{X Y \mid \Theta}(\boldsymbol{x}, y \mid \boldsymbol{\theta})$. The logarithm of probability $P\left(\mathcal{T}_{X Y} \mid \boldsymbol{\theta}\right)$ is referred to as the $\log$-likelihood $L\left(\boldsymbol{\theta} \mid \mathcal{T}_{X Y}\right)$ of $\boldsymbol{\theta}$ with respect to $\mathcal{T}_{X Y}$. Thanks to the
i.i.d. assumption the log-likelihood can be factorized as

$$
L\left(\boldsymbol{\theta} \mid \mathcal{T}_{X Y}\right)=\log P\left(\mathcal{T}_{X Y} \mid \boldsymbol{\theta}\right)=\sum_{i=1}^{l} \log P_{X Y \mid \Theta}\left(\boldsymbol{x}_{i}, y_{i} \mid \boldsymbol{\theta}\right)
$$

The problem of Maximum-Likelihood (ML) estimation from the complete data $\mathcal{T}_{X Y}$ is defined as

$$
\begin{equation*}
\boldsymbol{\theta}^{*}=\underset{\boldsymbol{\theta} \in \Theta}{\operatorname{argmax}} L\left(\boldsymbol{\theta} \mid \mathcal{T}_{X Y}\right)=\underset{\boldsymbol{\theta} \in \Theta}{\operatorname{argmax}} \sum_{i=1}^{l} \log P_{X Y \mid \Theta}\left(\boldsymbol{x}_{i}, y_{i} \mid \boldsymbol{\theta}\right) . \tag{4.3}
\end{equation*}
$$

The ML estimate of the parameters $\boldsymbol{\theta}$ of the GMM from the complete data set $\mathcal{T}_{X Y}$ has an analytical solution. The computation of the ML estimate is implemented in the function mlcgmm.

## Example: Maximum-Likelihood estimate from complete data

The parameters of the Gaussian mixture model is estimated from the complete Riply's data set riply_trn. The data are binary labeled thus the GMM has two Gaussians components. The estimated GMM is visualized as: (i) shape of components of the GMM (see Figure 4.2a) and (ii) contours of the distribution function (see Figure 4.2b)

```
data = load('riply_trn'); % load labeled (complete) data
model = mlcgmm(data); % ML estimate of GMM
% visualization
figure; hold on;
ppatterns(data); pgauss(model);
figure; hold on;
ppatterns(data);
pgmm(model,struct('visual','contour'));
```


### 4.2.2 Incomplete data

The input is a set $\mathcal{T}_{X}=\left\{\boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{l}\right\}$ which contains only vectors of observations $\boldsymbol{x}_{i} \in \mathbb{R}^{n}$ without knowledge about the corresponding hidden states $y_{i} \in \mathcal{Y}$. The logarithm of probability $P\left(\mathcal{T}_{X} \mid \boldsymbol{\theta}\right)$ is referred to as the $\log$-likelihood $L\left(\boldsymbol{\theta} \mid \mathcal{T}_{X}\right)$ of $\boldsymbol{\theta}$ with respect to $\mathcal{T}_{X}$. Thanks to the i.i.d. assumption the log-likelihood can be factorized as

$$
L\left(\boldsymbol{\theta} \mid \mathcal{T}_{X}\right)=\log P\left(\mathcal{T}_{X} \mid \boldsymbol{\theta}\right)=\sum_{i=1}^{l} \sum_{y \in \mathcal{Y}} P_{X \mid Y \Theta}\left(\boldsymbol{x}_{i} \mid y_{i}, \boldsymbol{\theta}\right) P_{Y \mid \Theta}\left(y_{i} \mid \boldsymbol{\theta}\right) .
$$

The problem of Maximum-Likelihood (ML) estimation from the incomplete data $\mathcal{T}_{X}$ is defined as

$$
\begin{equation*}
\boldsymbol{\theta}^{*}=\underset{\boldsymbol{\theta} \in \Theta}{\operatorname{argmax}} L\left(\boldsymbol{\theta} \mid \mathcal{T}_{X}\right)=\underset{\boldsymbol{\theta} \in \Theta}{\operatorname{argmax}} \sum_{i=1}^{l} \sum_{y \in \mathcal{Y}} P_{X \mid Y \Theta}\left(\boldsymbol{x}_{i} \mid y_{i}, \boldsymbol{\theta}\right) P_{Y \mid \Theta}\left(y_{i} \mid \boldsymbol{\theta}\right) \tag{4.4}
\end{equation*}
$$

The ML estimate of the parameters $\boldsymbol{\theta}$ of the GMM from the incomplete data set $\mathcal{T}_{X Y}$ does not have an analytical solution. The numerical optimization has to be used instead.

The Expectation-Maximization (EM) algorithm is an iterative procedure for ML estimation from the incomplete data. The EM algorithm builds a sequence of parameter estimates

$$
\boldsymbol{\theta}^{(0)}, \boldsymbol{\theta}^{(1)}, \ldots, \boldsymbol{\theta}^{(t)}
$$

such that the $\log$-likelihood $L\left(\boldsymbol{\theta}^{(t)} \mid \mathcal{T}_{X}\right)$ monotonically increases, i.e.,

$$
L\left(\boldsymbol{\theta}^{(0)} \mid \mathcal{T}_{X}\right)<L\left(\boldsymbol{\theta}^{(1)} \mid \mathcal{T}_{X}\right)<\ldots<L\left(\boldsymbol{\theta}^{(t)} \mid \mathcal{T}_{X}\right)
$$

until a stationary point $L\left(\boldsymbol{\theta}^{(t-1)} \mid \mathcal{T}_{X}\right)=L\left(\boldsymbol{\theta}^{(t)} \mid \mathcal{T}_{X}\right)$ is achieved. The EM algorithm is a local optimization technique thus the estimated parameters depend on the initial guess $\boldsymbol{\theta}^{(0)}$. The random guess or $K$-means algorithm (see Section 4.5) are commonly used to select the initial $\boldsymbol{\theta}^{(0)}$. The EM algorithm for the GMM is implemented in the function emgmm. An interactive demo on the EM algorithm for the GMM is implemented in demo_emgmm.
References: The EM algorithm (including convergence proofs) is described in book [26]. A nice description can be found in book [3]. The monograph [18] is entirely devoted to the EM algorithm. The first papers describing the EM are [25,5].

## Example: Maximum-Likelihood estimate from incomplete data

The estimation of the parameters of the GMM is demonstrated on a synthetic data. The ground truth model is the GMM with two univariate Gaussian components. The ground truth model is sampled to obtain a training (incomplete) data. The EM algorithm is applied to estimate the GMM parameters. The random initialization is used by default. The ground truth model, sampled data and the estimate model are visualized in Figure 4.3a. The plot of the log-likelihood function $L\left(\boldsymbol{\theta}^{(t)} \mid \mathcal{T}_{X}\right)$ with respect to the number of iterations $t$ is visualized in Figure 4.3b.

```
% ground truth Gaussian mixture model
true_model.Mean = [-2 2];
true_model.Cov = [1 0.5];
true_model.Prior = [0.4 0.6];
sample = gmmsamp(true_model, 250);
```

```
% ML estimation by EM
options.ncomp = 2; % number of Gaussian components
options.verb = 1; % display progress info
estimated_model = emgmm(sample.X,options);
% visualization
figure;
ppatterns(sample.X);
h1 = pgmm(true_model,struct('color','r'));
h2 = pgmm(estimated_model,struct('color','b'));
legend([h1(1) h2(1)],'Ground truth', 'ML estimation');
figure; hold on;
xlabel('iterations'); ylabel('log-likelihood');
plot( estimated_model.logL );
```


### 4.3 Minimax estimation

The input is a set $\mathcal{T}_{X}=\left\{\boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{l}\right\}$ contains vectors $\boldsymbol{x}_{i} \in \mathbb{R}^{n}$ which are realizations of random variable distributed according to $P_{X \mid \boldsymbol{\Theta}}(\boldsymbol{x} \mid \boldsymbol{\theta})$. The vectors $\mathcal{T}_{X}$ are assumed to be a realizations with high value of the distribution function $P_{X \mid \boldsymbol{\Theta}}(\boldsymbol{x} \mid \boldsymbol{\theta})$. The $P_{X \mid \boldsymbol{\Theta}}(\boldsymbol{x} \mid \boldsymbol{\theta})$ is known up to the parameter $\boldsymbol{\theta} \in \Theta$. The minimax estimation is defined as

$$
\begin{equation*}
\boldsymbol{\theta}^{*}=\underset{\boldsymbol{\theta} \in \Theta}{\operatorname{argmax}} \min _{\boldsymbol{x} \in \mathcal{T}_{X}} \log P_{X \mid \Theta}(\boldsymbol{x} \mid \boldsymbol{\theta}) \tag{4.5}
\end{equation*}
$$

If a procedure for the Maximum-Likelihood estimate (global solution) of the distribution $P_{X \mid \Theta}(\boldsymbol{x} \mid \boldsymbol{\theta})$ exists then a numerical iterative algorithm which converges to the optimal estimate $\boldsymbol{\theta}^{*}$ can be constructed. The algorithm builds a series of estimates $\boldsymbol{\theta}^{(0)}, \boldsymbol{\theta}^{(1)}, \ldots, \boldsymbol{\theta}^{(t)}$ which converges to the optimal estimate $\boldsymbol{\theta}^{*}$. The algorithm converges in a finite number of iterations to such an estimate $\boldsymbol{\theta}$ that the condition

$$
\begin{equation*}
\min _{\boldsymbol{x} \in \mathcal{T}_{X}} \log P_{X \mid \Theta}\left(\boldsymbol{x} \mid \boldsymbol{\theta}^{*}\right)-\min _{\boldsymbol{x} \in \mathcal{T}_{X}} \log P_{X \mid \Theta}(\boldsymbol{x} \mid \boldsymbol{\theta})<\varepsilon \tag{4.6}
\end{equation*}
$$

holds for $\varepsilon>0$ which defines closeness to the optimal solution $\boldsymbol{\theta}^{*}$. The inequality (4.6) can be evaluated efficiently thus it is a natural choice for the stopping condition of the algorithm. The STPRtool contains an implementation of the minimax algorithm for the Gaussian distribution (4.1). The algorithm is implemented in the function mmgauss. An interactive demo on the minimax estimation for the Gaussian distribution is implemented in the function demo_mmgauss.
References: The minimax algorithm for probability estimation is described and analyzed in Chapter 8 of the book [26].

## Example: Minimax estimation of the Gaussian distribution

The bivariate Gaussian distribution is described by three samples which are known to have high value of the distribution function. The minimax algorithm is applied to estimate the parameters of the Gaussian. In this particular case, the minimax problem is equivalent to searching for the minimal enclosing ellipsoid. The found solution is visualized as the isoline of the distribution function at the point $P_{X \mid \Theta}(\boldsymbol{x} \mid \boldsymbol{\theta})$ (see Figure 4.4).

```
X = [[0;0] [1;0] [0;1]]; % points with high p.d.f.
model = mmgauss(X); % run minimax algorithm
% visualization
figure; ppatterns(X,'xr',13);
pgauss(model, struct('p',exp(model.lower_bound)));
```


### 4.4 Probabilistic output of classifier

In general, the discriminant function of an arbitrary classifier does not have meaning of probability (e.g., SVM, Perceptron). However, the probabilistic output of the classifier can help in post-processing, for instance, in combining more classifiers together. Fitting a sigmoid function to the classifier output is a way how to solve this problem.

Let $\mathcal{T}_{X Y}=\left\{\left(\boldsymbol{x}_{1}, y_{1}\right), \ldots,\left(\boldsymbol{x}_{l}, y_{l}\right)\right\}$ is the training set composed of the vectors $\boldsymbol{x}_{i} \in$ $\mathcal{X} \subseteq \mathbb{R}^{n}$ and the corresponding binary hidden states $y_{i} \in \mathcal{Y}=\{1,2\}$. The training set $\mathcal{T}_{X Y}$ is assumed to be identically and independently distributed from underlying distribution. Let $f: \mathcal{X} \subseteq \mathbb{R}^{n} \rightarrow \mathbb{R}$ be a discriminant function trained from the data $\mathcal{T}_{X Y}$ by an arbitrary learning method. The aim is to estimate parameters of a posteriori distribution $P_{Y \mid F \Theta}(y \mid f(\boldsymbol{x}), \boldsymbol{\theta})$ of the hidden state $y$ given the value of the discriminant function $f(\boldsymbol{x})$. The distribution is modeled by a sigmoid function

$$
P_{Y \mid F \Theta}(1 \mid f(\boldsymbol{x}), \boldsymbol{\theta})=\frac{1}{1+\exp \left(a_{1} f(\boldsymbol{x})+a_{2}\right)},
$$

which is determined by parameters $\boldsymbol{\theta}=\left[a_{1}, a_{2}\right]^{T}$. The log-likelihood function is defined as

$$
L\left(\boldsymbol{\theta} \mid \mathcal{T}_{X Y}\right)=\sum_{i=1}^{l} \log P_{Y \mid F \Theta}\left(y_{i} \mid f\left(\boldsymbol{x}_{i}\right), \boldsymbol{\theta}\right)
$$

The parameters $\boldsymbol{\theta}=\left[a_{1}, a_{2}\right]^{T}$ can be estimated by the maximum-likelihood method

$$
\begin{equation*}
\boldsymbol{\theta}=\underset{\boldsymbol{\theta}^{\prime}}{\operatorname{argmax}} L\left(\boldsymbol{\theta}^{\prime} \mid \mathcal{T}_{X Y}\right)=\underset{\boldsymbol{\theta}^{\prime}}{\operatorname{argmax}} \sum_{i=1}^{l} \log P_{Y \mid F \Theta}\left(y_{i} \mid f\left(\boldsymbol{x}_{i}\right), \boldsymbol{\theta}^{\prime}\right) . \tag{4.7}
\end{equation*}
$$

The STPRtool provides an implementation mlsigmoid which uses the Matlab Optimization toolbox function fminunc to solve the task (4.7). The function produces the parameters $\boldsymbol{\theta}=\left[a_{1}, a_{2}\right]^{T}$ which are stored in the data-type describing the sigmoid (see Table 4.6). The evaluation of the sigmoid is implemented in function sigmoid.

Table 4.6: Data-type used to describe sigmoid function.

| Sigmoid function | (structure array): |
| :--- | :--- |
| . $\mathrm{A}[2 \times 1]$ | Parameters $\boldsymbol{\theta}=\left[a_{1}, a_{2}\right]^{T}$ of the sigmoid function. |
| . fun $=$ 'sigmoid' | Identifies function associated with this data type. |

References: The method of fitting the sigmoid function is described in [22].

## Example: Probabilistic output for SVM

This example is implemented in the script demo_svmpout. The example shows how to train the probabilistic output for the SVM classifier. The Riply's data set riply_trn is used for training the SVM and the ML estimation of the sigmoid. The ML estimated of Gaussian mixture model (GMM) of the SVM output is used for comparison. The GMM allows to compute the a posteriori probability for comparison to the sigmoid estimate. The example is implemented in the script demo_svmpout. Figure 4.5a shows found probabilistic models of a posteriori probability of the SVM output. Figure 4.5b shows the decision boundary of SVM classifier for illustration.

```
% load training data
data = load('riply_trn');
% train SVM
options = struct('ker','rbf','arg',1,'C',10);
svm_model = smo(data,options);
% compute SVM output
[dummy,svm_output.X] = svmclass(data.X,svm_model);
svm_output.y = data.y;
% fit sigmoid function to svm output
sigmoid_model = mlsigmoid(svm_output);
% ML estimation of GMM model of SVM output
gmm_model = mlcgmm(svm_output);
% visulization
figure; hold on;
```

```
xlabel('svm output f(x)'); ylabel('p(y=1|f(x))');
% sigmoid model
fx = linspace(min(svm_output.X), max(svm_output.X), 200);
sigmoid_apost = sigmoid(fx,sigmoid_model);
hsigmoid = plot(fx,sigmoid_apost,'k');
ppatterns(svm_output);
% GMM model
pcond = pdfgauss( fx, gmm_model);
gmm_apost = (pcond(1,:)*gmm_model.Prior(1))./...
    (pcond(1,:)*gmm_model.Prior(1)+(pcond(2,:)*gmm_model.Prior(2)));
hgmm = plot(fx,gmm_apost,'g');
hcomp = pgauss(gmm_model);
legend([hsigmoid,hgmm,hcomp],'P(y=1|f(x)) ML-Sigmoid',...
    'P(y=1|f(x)) ML-GMM','P(f(x)|y=1) ML-GMM','P(f(x)|y=2) ML-GMM');
% SVM decision boundary
figure; ppatterns(data); psvm(svm_model);
```


## 4.5 $K$-means clustering

The input is a set $\mathcal{T}_{X}=\left\{\boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{l}\right\}$ which contains vectors $\boldsymbol{x}_{i} \in \mathbb{R}^{n}$. Let $\boldsymbol{\theta}=$ $\left\{\boldsymbol{\mu}_{1}, \ldots, \boldsymbol{\mu}_{c}\right\}, \boldsymbol{\mu}_{i} \in \mathbb{R}^{n}$ be a set of unknown cluster centers. Let the objective function $F(\boldsymbol{\theta})$ be defined as

$$
F(\boldsymbol{\theta})=\frac{1}{l} \sum_{i=1}^{l} \min _{y=1, \ldots, c}\left\|\boldsymbol{\mu}_{y}-\boldsymbol{x}_{i}\right\|^{2} .
$$

The small value of $F(\boldsymbol{\theta})$ indicates that the centers $\boldsymbol{\theta}$ well describe the input set $\mathcal{T}_{X}$. The task is to find such centers $\boldsymbol{\theta}$ which describe the set $\mathcal{T}_{X}$ best, i.e., the task is to solve

$$
\boldsymbol{\theta}^{*}=\underset{\boldsymbol{\theta}}{\operatorname{argmin}} F(\boldsymbol{\theta})=\underset{\boldsymbol{\theta}}{\operatorname{argmin}} \frac{1}{l} \sum_{i=1}^{l} \min _{y=1, \ldots, c}\left\|\boldsymbol{\mu}_{y}-\boldsymbol{x}_{i}\right\|^{2} .
$$

The $K$-means ${ }^{1}$ is an iterative procedure which iteratively builds a series $\boldsymbol{\theta}^{(0)}, \boldsymbol{\theta}^{(1)}, \ldots, \boldsymbol{\theta}^{(t)}$. The $K$-means algorithm converges in a finite number of steps to the local optimum of the objective function $F(\boldsymbol{\theta})$. The found solution depends on the initial guess $\boldsymbol{\theta}^{(0)}$ which

[^0]is usually selected randomly. The $K$-means algorithm is implemented in the function kmeans.
References: The $K$-means algorithm is described in the books $[26,6,3]$.

## Example: $K$-means clustering

The $K$-means algorithm is used to find 4 clusters in the Riply's training set. The solution is visualized as the cluster centers. The vectors classified to the corresponding centers are distinguished by color (see Figure 4.6a). The plot of the objective function $F\left(\boldsymbol{\theta}^{(t)}\right)$ with respect to the number of iterations is shown in Figure 4.6b.

```
data = load('riply_trn'); % load data
[model,data.y] = kmeans(data.X, 4 ); % run k-means
% visualization
figure; ppatterns(data);
ppatterns(model.X,'sk',14);
pboundary( model );
figure; hold on;
xlabel('t'); ylabel('F');
plot(model.MsErr);
```



Figure 4.2: Example of the Maximum-Likelihood of the Gaussian Mixture Model from the incomplete data. Figure (a) shows components of the GMM and Figure (b) shows contour plot of the distribution function.

(b)

Figure 4.3: Example of using the EM algorithm for ML estimation of the Gaussian mixture model. Figure (a) shows ground truth and the estimated GMM and Figure(b) shows the log-likelihood with respect to the number of iterations of the EM.


Figure 4.4: Example of the Minimax estimation of the Gaussian distribution.


Figure 4.5: Example of fitting a posteriori probability to the SVM output. Figure (a) shows the sigmoid model and GMM model both fitted by ML estimated. Figure (b) shows the corresponding SVM classifier.


Figure 4.6: Example of using the $K$-means algorithm to cluster the Riply's data set. Figure (a) shows found clusters and Figure (b) contains the plot of the objective function $F\left(\boldsymbol{\theta}^{(t)}\right)$ with respect to the number of iterations.

## Chapter 5

## Support vector and other kernel machines

### 5.1 Support vector classifiers

The binary support vector classifier uses the discriminant function $f: \mathcal{X} \subseteq \mathbb{R}^{n} \rightarrow \mathbb{R}$ of the following form

$$
\begin{equation*}
f(\boldsymbol{x})=\left\langle\boldsymbol{\alpha} \cdot \boldsymbol{k}_{S}(\boldsymbol{x})\right\rangle+b . \tag{5.1}
\end{equation*}
$$

The $\boldsymbol{k}_{S}(\boldsymbol{x})=\left[k\left(\boldsymbol{x}, \boldsymbol{s}_{1}\right), \ldots, k\left(\boldsymbol{x}, \boldsymbol{s}_{d}\right)\right]^{T}$ is the vector of evaluations of kernel functions centered at the support vectors $\mathcal{S}=\left\{s_{1}, \ldots, s_{d}\right\}, s_{i} \in \mathbb{R}^{n}$ which are usually subset of the training data. The $\boldsymbol{\alpha} \in \mathbb{R}^{l}$ is a weight vector and $b \in \mathbb{R}$ is a bias. The binary classification rule $q: \mathcal{X} \rightarrow \mathcal{Y}=\{1,2\}$ is defined as

The data-type used by the STPRtool to represent the binary SVM classifier is described in Table 5.2.

The multi-class generalization involves a set of discriminant functions $f_{y}: \mathcal{X} \subseteq \mathbb{R}^{n} \rightarrow$ $\mathbb{R}, y \in \mathcal{Y}=\{1,2, \ldots, c\}$ defined as

$$
f_{y}(\boldsymbol{x})=\left\langle\boldsymbol{\alpha}_{y} \cdot \boldsymbol{k}_{S}(\boldsymbol{x})\right\rangle+b_{y}, \quad y \in \mathcal{Y} .
$$

Let the matrix $\mathbf{A}=\left[\boldsymbol{\alpha}_{1}, \ldots, \boldsymbol{\alpha}_{c}\right]$ be composed of all weight vectors and $\boldsymbol{b}=\left[b_{1}, \ldots, b_{c}\right]^{T}$ be a vector of all biases. The multi-class classification rule $q: \mathcal{X} \rightarrow \mathcal{Y}=\{1,2, \ldots, c\}$ is defined as

$$
\begin{equation*}
q(\boldsymbol{x})=\underset{y \in \mathcal{Y}}{\operatorname{argmax}} f_{y}(\boldsymbol{x}) . \tag{5.3}
\end{equation*}
$$

The data-type used by the STPRtool to represent the multi-class SVM classifier is described in Table 5.3. Both the binary and the multi-class SVM classifiers are implemented in the function svmclass.

The majority voting strategy is other commonly used method to implement the multi-class SVM classifier. Let $q_{j}: \mathcal{X} \subseteq \mathbb{R}^{n} \rightarrow\left\{y_{j}^{1}, y_{j}^{2}\right\}, j=1, \ldots, g$ be a set of $g$ binary SVM rules (5.2). The $j$-th rule $q_{j}(\boldsymbol{x})$ classifies the inputs $\boldsymbol{x}$ into class $y_{j}^{1} \in \mathcal{Y}$ or $y_{j}^{2} \in \mathcal{Y}$. Let $\boldsymbol{v}(\boldsymbol{x})$ be a vector $[c \times 1]$ of votes for classes $\mathcal{Y}$ when the input $\boldsymbol{x}$ is to be classified. The vector $\boldsymbol{v}(\boldsymbol{x})=\left[v_{1}(\boldsymbol{x}), \ldots, v_{c}(\boldsymbol{x})\right]^{T}$ is computed as:

Set $v_{y}=0, \forall y \in \mathcal{Y}$.
For $j=1, \ldots, g$ do

$$
v_{y}(\boldsymbol{x})=v_{y}(\boldsymbol{x})+1 \text { where } y=q_{j}(\boldsymbol{x}) .
$$

end.
The majority votes based multi-class classifier assigns the input $\boldsymbol{x}$ into such class $y \in \mathcal{Y}$ having the majority of votes

$$
\begin{equation*}
y=\underset{y^{\prime}=1, \ldots, g}{\operatorname{argmax}} v_{y^{\prime}}(\boldsymbol{x}) . \tag{5.4}
\end{equation*}
$$

The data-type used to represent the majority voting strategy for the multi-class SVM classifier is described in Table 5.4. The strategy (5.4) is implemented in the function mvsvmclass. The implemented methods to deal with the SVM and related kernel classifiers are enlisted in Table 5.1
References: Books describing Support Vector Machines are for example [31, 4, 30]

### 5.2 Binary Support Vector Machines

The input is a set $\mathcal{T}_{X Y}=\left\{\left(\boldsymbol{x}_{1}, y_{1}\right), \ldots,\left(\boldsymbol{x}_{l}, y_{l}\right)\right\}$ of training vectors $\boldsymbol{x}_{i} \in \mathbb{R}^{n}$ and corresponding hidden states $y_{i} \in \mathcal{Y}=\{1,2\}$. The linear SVM aims to train the linear discriminant function $f(\boldsymbol{x})=\langle\boldsymbol{w} \cdot \boldsymbol{x}\rangle+b$ of the binary classifier

$$
q(\boldsymbol{x})=\left\{\begin{array}{lll}
1 & \text { for } & f(\boldsymbol{x}) \geq 0 \\
2 & \text { for } & f(\boldsymbol{x})<0
\end{array}\right.
$$

The training of the optimal parameters $\left(\boldsymbol{w}^{*}, b^{*}\right)$ is transformed to the following quadratic programming task

$$
\begin{align*}
&\left(\boldsymbol{w}^{*}, b^{*}\right)=\underset{\boldsymbol{w}, b}{\operatorname{argmin}} \frac{1}{2}\|\boldsymbol{w}\|^{2}+C \sum_{i=1}^{l} \xi_{i}^{p},  \tag{5.5}\\
&\left\langle\boldsymbol{w} \cdot \boldsymbol{x}_{i}\right\rangle+b \geq+1-\xi_{i}, \quad i \in \mathcal{I}_{1}, \\
&\left\langle\boldsymbol{w} \cdot \boldsymbol{x}_{i}\right\rangle+b \leq-1+\xi_{i}, \quad i \in \mathcal{I}_{2}, \\
& \xi_{i} \geq 0, \quad i \in \mathcal{I}_{1} \cup \mathcal{I}_{2}
\end{align*}
$$

Table 5.1: Implemented methods: Support vector and other kernel machines.

| svmclass | Support Vector Machines Classifier. |
| :--- | :--- |
| mvsvmclass | Majority voting multi-class SVM classifier. |
| evalsvm | Trains and evaluates Support Vector Machines classi- |
|  | fier. |
| bsvm2 | Multi-class BSVM with L2-soft margin. |
| oaasvm | Multi-class SVM using One-Against-All decomposi- |
|  | tion. |
| oaosvm | Multi-class SVM using One-Against-One decomposi- |
|  | tion. |
| smo | Sequential Minimal Optimization for binary SVM |
|  | with L1-soft margin. |
| svmlight | Interface to SVM Might software. |
| svmquadprog | SVM trained by Matlab Optimization Toolbox. |
| diagker | Returns diagonal of kernel matrix of given data. |
| kdist | Computes distance between vectors in kernel space. |
| kernel | Evaluates kernel function. |
| kfd | Kernel Fisher Discriminant. |
| kperceptr | Kernel Perceptron. |
| minball | Minimal enclosing ball in kernel feature space. |
| rsrbf | Reduced Set Method for RBF kernel expansion. |
| rbfpreimg | RBF pre-image by Schoelkopf's fixed-point algorithm. |
| rbfpreimg2 | RBF pre-image problem by Gradient optimization. |
| rbfpreimg3 | RBF pre-image problem by Kwok-Tsang's algorithm. |
| demo_svm | Demo on Support Vector Machines. |

The $\mathcal{I}_{1}=\left\{i: y_{i}=1\right\}$ and $\mathcal{I}_{2}=\left\{i: y_{i}=2\right\}$ are sets of indices. The $C>0$ stands for the regularization constant. The $\xi_{i} \geq 0, i \in \mathcal{I}_{1} \cup \mathcal{I}_{2}$ are slack variables used to relax the inequalities for the case of non-separable data. The linear $p=1$ and quadratic $p=2$ term for the slack variables $\xi_{i}$ is used. In the case of the parameter $p=1$, the $L 1$-soft margin SVM is trained and for $p=2$ it is denoted as the $L 2$-soft margin SVM.

The training of the non-linear (kernel) SVM with $L 1$-soft margin corresponds to solving the following QP task

$$
\begin{equation*}
\boldsymbol{\beta}^{*}=\underset{\boldsymbol{\beta}}{\operatorname{argmax}}\langle\boldsymbol{\beta} \cdot \mathbf{1}\rangle-\frac{1}{2}\langle\boldsymbol{\beta} \cdot \mathbf{H} \boldsymbol{\beta}\rangle, \tag{5.6}
\end{equation*}
$$

subject to

$$
\begin{aligned}
\langle\boldsymbol{\beta} \cdot \boldsymbol{\gamma}\rangle & =1 \\
\boldsymbol{\beta} & \geq 0 \\
\boldsymbol{\beta} & \leq 1 C .
\end{aligned}
$$

Table 5.2: Data-type used to describe binary SVM classifier.

| Binary SVM classifier | (structure array): |
| :--- | :--- |
| .Alpha $[d \times 1]$ | Weight vector $\boldsymbol{\alpha}$. |
| .b $[1 \times 1]$ | Bias $b$. |
| .sv.X $[n \times d]$ | Support vectors $\mathcal{S}=\left\{s_{1}, \ldots, \boldsymbol{s}_{d}\right\}$. |
| .options.ker $[$ string $]$ | Kernel identifier. |
| .options.arg $[1 \times p]$ | Kernel argument(s). |
| .fun $=$ 'svmclass' | Identifies function associated with this data type. |

Table 5.3: Data-type used to describe multi-class SVM classifier.

| Multi-class SVM classifier (structure array): |  |
| :--- | :--- |
| . Alpha $[d \times c]$ | Weight vectors $\mathbf{A}=\left[\boldsymbol{\alpha}_{1}, \ldots, \boldsymbol{\alpha}_{c}\right]$. |
| . $\mathrm{b}[c \times 1]$ | Vector of biased $\boldsymbol{b}=\left[b_{1}, \ldots, b_{c}\right]^{T}$. |
| . sv.X $[n \times d]$ | Support vectors $\mathcal{S}=\left\{s_{1}, \ldots, \boldsymbol{s}_{d}\right\}$. |
| .options.ker $[$ string $]$ | Kernel identifier. |
| .options.arg $[1 \times p]$ | Kernel argument(s). |
| . fun $=$ 'svmclass' | Identifies function associated with this data type. |

The $\mathbf{0}[l \times 1]$ is vector of zeros and the $\mathbf{1}[l \times 1]$ is vector of ones. The vector $\boldsymbol{\gamma}[l \times 1]$ and the matrix $\mathbf{H}[l \times l]$ are constructed as follows

$$
\gamma_{i}=\left\{\begin{array}{rll}
1 & \text { if } & y_{i}=1, \\
-1 & \text { if } & y_{i}=2,
\end{array} \quad H_{i, j}=\left\{\begin{array}{rll}
k\left(\boldsymbol{x}_{i}, \boldsymbol{x}_{j}\right) & \text { if } & y_{i}=y_{j} \\
-k\left(\boldsymbol{x}_{i}, \boldsymbol{x}_{j}\right) & \text { if } & y_{i} \neq y_{j},
\end{array}\right.\right.
$$

where $k: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is selected kernel function. The discriminant function (5.1) is determined by the weight vector $\boldsymbol{\alpha}$, bias $b$ and the set of support vectors $\mathcal{S}$. The set of support vectors is denoted as $\mathcal{S}=\left\{\boldsymbol{x}_{i}: i \in \mathcal{I}_{S V}\right\}$, where the set of indices is $\mathcal{I}_{S V}=\left\{i: \beta_{i}>0\right\}$. The weight vectors are compute as

$$
\alpha_{i}=\left\{\begin{array}{rll}
\beta_{i} & \text { if } & y_{i}=1,  \tag{5.7}\\
-\beta_{i} & \text { if } & y_{i}=2,
\end{array} \quad \text { for } \quad i \in \mathcal{I}_{S V} .\right.
$$

The bias $b$ can be computed from the Karush-Kuhn-Tucker (KKT) conditions as the following constrains

$$
\begin{aligned}
& f\left(\boldsymbol{x}_{i}\right)=\left\langle\boldsymbol{\alpha} \cdot \boldsymbol{k}_{S}\left(\boldsymbol{x}_{i}\right)\right\rangle+b=+1 \quad \text { for } \quad i \in\left\{j: y_{j}=1,0<\beta_{j}<C\right\} \\
& f\left(\boldsymbol{x}_{i}\right)=\left\langle\boldsymbol{\alpha} \cdot \boldsymbol{k}_{S}\left(\boldsymbol{x}_{i}\right)\right\rangle+b=-1 \quad \text { for } \quad i \in\left\{j: y_{j}=2,0<\beta_{j}<C\right\},
\end{aligned}
$$

must hold for the optimal solution. The bias $b$ is computed as an average over all the constrains such that

$$
b=\frac{1}{\left|\mathcal{I}_{b}\right|} \sum_{i \in \mathcal{I}_{b}} \gamma_{i}-\left\langle\boldsymbol{\alpha} \cdot \boldsymbol{k}_{S}\left(\boldsymbol{x}_{i}\right)\right\rangle
$$

Table 5.4: Data-type used to describe the majority voting strategy for multi-class SVM classifier.

| Majority voting SVM classifier (structure array): |  |
| :--- | :--- |
| .Alpha $[d \times g]$ | Weight vectors $\mathbf{A}=\left[\boldsymbol{\alpha}_{1}, \ldots, \boldsymbol{\alpha}_{g}\right]$. |
| .b $[g \times 1]$ | Vector of biased $\boldsymbol{b}=\left[b_{1}, \ldots, b_{g}\right]^{T}$. |
| .bin_y $[2 \times g]$ | Targets for the binary rules. The vector bin_y (1, : ) |
|  | contains $\left[y_{1}^{1}, y_{2}^{1}, \ldots, y_{g}^{1}\right]$ and the bin_y $(2,:$ ) contains |
|  | $\left[y_{1}^{2}, y_{2}^{2}, \ldots, y_{g}^{2}\right]$. |
| .sv.X $[n \times d]$ | Support vectors $\mathcal{S}=\left\{\boldsymbol{s}_{1}, \ldots, \boldsymbol{s}_{d}\right\}$. |
| .options.ker $[$ string $]$ | Kernel identifier. |
| .options.arg $[1 \times p]$ | Kernel argument $(\mathrm{s})$. |
| .fun $=$ 'mvsvmclass' | Identifies function associated with this data type. |

where $\mathcal{I}_{b}=\left\{i: 0<\beta_{i}<C\right\}$ are the indices of the vectors on the boundary.
The training of the non-linear (kernel) SVM with $L 2$-soft margin corresponds to solving the following QP task

$$
\begin{equation*}
\boldsymbol{\beta}^{*}=\underset{\boldsymbol{\beta}}{\operatorname{argmax}}\langle\boldsymbol{\beta} \cdot \mathbf{1}\rangle-\frac{1}{2}\left\langle\boldsymbol{\beta} \cdot\left(\mathbf{H}+\frac{1}{2 C} \mathbf{E}\right) \boldsymbol{\beta}\right\rangle, \tag{5.8}
\end{equation*}
$$

subject to

$$
\begin{aligned}
\langle\boldsymbol{\beta} \cdot \boldsymbol{\gamma}\rangle & =1 \\
\boldsymbol{\beta} & \geq \mathbf{0}
\end{aligned}
$$

where $\mathbf{E}$ is the identity matrix. The set of support vectors is set to $\mathcal{S}=\left\{\boldsymbol{x}_{i}: i \in \mathcal{I}_{S V}\right\}$ where the set of indices is $\mathcal{I}_{S V}=\left\{i: \beta_{i}>0\right\}$. The weight vector $\boldsymbol{\alpha}$ is given by (5.7). The bias $b$ can be computed from the KKT conditions as the following constrains

$$
\begin{aligned}
& f\left(\boldsymbol{x}_{i}\right)=\left\langle\boldsymbol{\alpha} \cdot \boldsymbol{k}_{S}\left(\boldsymbol{x}_{i}\right)\right\rangle+b=+1-\frac{\alpha_{i}}{2 C} \quad \text { for } \quad i \in\left\{i: y_{j}=1, \beta_{j}>0\right\} \\
& f\left(\boldsymbol{x}_{i}\right)=\left\langle\boldsymbol{\alpha} \cdot \boldsymbol{k}_{S}\left(\boldsymbol{x}_{i}\right)\right\rangle+b=-1+\frac{\alpha_{i}}{2 C} \quad \text { for } \quad i \in\left\{i: y_{j}=2, \beta_{j}>0\right\}
\end{aligned}
$$

must hold at the optimal solution. The bias $b$ is computed as an average over all the constrains thus

$$
b=\frac{1}{\left|\mathcal{I}_{S V}\right|} \sum_{i \in \mathcal{I}_{S V}} \gamma_{i}\left(1-\frac{\alpha_{i}}{2 C}\right)-\left\langle\boldsymbol{\alpha} \cdot \boldsymbol{k}_{S}\left(\boldsymbol{x}_{i}\right)\right\rangle .
$$

The binary SVM solvers implemented in the STPRtool are enlisted in Table 5.5. An interactive demo on the binary SVM is implemented in demo_svm.
References: The Sequential Minimal Optimizer (SMO) is described in the books [23, 27, 30]. The $S V M^{\text {light }}$ is described for instance in the book [27].
Example: Training binary SVM.

Table 5.5: Implemented binary SVM solvers.

| Function | description |
| :--- | :--- |
| smo | Implementation of the Sequential Minimal Optimizer (SMO) <br> used to train the binary SVM with L1-soft margin. It can <br> handle moderate size problems. |
| svmlight | Matlab interface to SV M <br> light <br> trains the binary SVM with L1-soft margin. The executable <br> file svm_learn for the Linux OS must be installed. It can be <br> downladed from http://svmlight. joachims.org/. The <br> $S V M^{\text {light }}$ is very powerful optimizer which can handle large <br> problems. |
| svmquadprog | SVM solver based on the QP solver quadprog of the Mat- <br> lab optimization toolbox. It trains both L1-soft and L2-soft <br> margin binary SVM. The function quadprog is a part of the <br> Matlab Optimization toolbox. It useful for small problems <br> only. |

The binary SVM is trained on the Riply's training set riply_trn. The SMO solver is used in the example but other solvers (svmlight, svmquadprog) can by used as well. The trained SVM classifier is evaluated on the testing data riply_tst. The decision boundary is visualized in Figure 5.1.

```
    trn = load('riply_trn'); % load training data
    options.ker = 'rbf'; % use RBF kernel
    options.arg = 1; % kernel argument
    options.C = 10; % regularization constant
    % train SVM classifier
    model = smo(trn,options);
% model = svmlight(trn,options);
% model = svmquadprog(trn,options);
% visualization
figure;
ppatterns(trn); psvm(model);
tst = load('riply_tst'); % load testing data
ypred = svmclass(tst.X,model); % classify data
cerror(ypred,tst.y) % compute error
```

0.0920


Figure 5.1: Example showing the binary SVM classifier trained on the Riply's data.

### 5.3 One-Against-All decomposition for SVM

The input is a set $\mathcal{T}_{X Y}=\left\{\left(\boldsymbol{x}_{1}, y_{1}\right), \ldots,\left(\boldsymbol{x}_{l}, y_{l}\right)\right\}$ of training vectors $\boldsymbol{x}_{i} \in \mathcal{X} \subseteq \mathbb{R}^{n}$ and corresponding hidden states $y_{i} \in \mathcal{Y}=\{1,2, \ldots, c\}$. The goal is to train the multi-class rule $q: \mathcal{X} \subseteq \mathbb{R}^{n} \rightarrow \mathcal{Y}$ defined by (5.3). The problem can be solved by the One-Against-All (OAA) decomposition. The OAA decomposition transforms the multi-class problem into a series of $c$ binary subtasks that can be trained by the binary SVM. Let the training set $\mathcal{T}_{X Y}^{y}=\left\{\left(\boldsymbol{x}_{1}, y_{1}^{\prime}\right), \ldots,\left(\boldsymbol{x}_{l}, y_{l}^{\prime}\right)\right\}$ contain the modified hidden states defined as

$$
y_{i}^{\prime}= \begin{cases}1 & \text { for } y=y_{i} \\ 2 & \text { for } y \neq y_{i}\end{cases}
$$

The discriminant functions

$$
f_{y}(\boldsymbol{x})=\left\langle\boldsymbol{\alpha}_{y} \cdot \boldsymbol{k}_{S}(\boldsymbol{x})\right\rangle+b_{y}, \quad y \in \mathcal{Y}
$$

are trained by the binary SVM solver from the set $\mathcal{T}_{X Y}^{y}, y \in \mathcal{Y}$. The OAA decomposition is implemented in the function oaasvm. The function allows to specify any binary SVM
solver (see Section 5.2) used to solve the binary problems. The result is the multi-class SVM classifier (5.3) represented by the data-type described in Table 5.3.
References: The OAA decomposition to train the multi-class SVM and comparison to other methods is described for instance in the paper [12].

## Example: Multi-class SVM using the OAA decomposition

The example shows the training of the multi-class SVM using the OAA decomposition. The SMO binary solver is used to train the binary SVM subtasks. The training data and the decision boundary is visualized in Figure 5.2.

```
data = load('pentagon'); % load data
options.solver = 'smo'; % use SMO solver
options.ker = 'rbf'; % use RBF kernel
options.arg = 1; % kernel argument
options.C = 10; % regularization constant
model = oaasvm(data,options ); % training
```

\% visualization
figure;
ppatterns(data);
ppatterns(model.sv.X, 'ko', 12);
pboundary (model) ;

### 5.4 One-Against-One decomposition for SVM

The input is a set $\mathcal{T}_{X Y}=\left\{\left(\boldsymbol{x}_{1}, y_{1}\right), \ldots,\left(\boldsymbol{x}_{l}, y_{l}\right)\right\}$ of training vectors $\boldsymbol{x}_{i} \in \mathcal{X} \subseteq \mathbb{R}^{n}$ and corresponding hidden states $y_{i} \in \mathcal{Y}=\{1,2, \ldots, c\}$. The goal is to train the multi-class rule $q: \mathcal{X} \subseteq \mathbb{R}^{n} \rightarrow \mathcal{Y}$ based on the majority voting strategy defined by (5.4). The one-against-one (OAO) decomposition strategy can be used train such a classifier. The OAO decomposition transforms the multi-class problem into a series of $g=c(c-1) / 2$ binary subtasks that can be trained by the binary SVM. Let the training set $\mathcal{T}_{X Y}^{j}=$ $\left\{\left(\boldsymbol{x}_{1}^{\prime}, y_{1}^{\prime}\right), \ldots,\left(\boldsymbol{x}_{l_{j}}^{\prime}, y_{l_{j}}^{\prime}\right)\right\}$ contain the training vectors $\boldsymbol{x}_{i} \in \mathcal{I}^{j}=\left\{i: y_{i}=y^{1} \bigvee y_{i}=y^{2}\right\}$ and the modified the hidden states defined as

$$
y_{i}^{\prime}=\left\{\begin{array}{ll}
1 & \text { for } y_{j}^{1}=y_{i}, \\
2 & \text { for } y_{j}^{2}=y_{i},
\end{array} \quad i \in \mathcal{I}^{j}\right.
$$

The training set $\mathcal{T}_{X Y}^{j}, j=1, \ldots, g$ is constructed for all $g=c(c-1) / 2$ combinations of classes $y_{j}^{1} \in \mathcal{Y}$ and $y_{j}^{2} \in \mathcal{Y} \backslash\left\{y_{j}^{1}\right\}$. The binary SVM rules $q_{j}, j=1, \ldots, g$ are trained on the data $\mathcal{T}_{X Y}^{j}$. The OAO decomposition for multi-class SVM is implemented in function


Figure 5.2: Example showing the multi-class SVM classifier build by one-against-all decomposition.
oaosvm. The function oaosvm allows to specify an arbitrary binary SVM solver to train the subtasks.
References: The OAO decomposition to train the multi-class SVM and comparison to other methods is described for instance in the paper [12].

## Example: Multi-class SVM using the OAO decomposition

The example shows the training of the multi-class SVM using the OAO decomposition. The SMO binary solver is used to train the binary SVM subtasks. The training data and the decision boundary is visualized in Figure 5.3.

```
data = load('pentagon'); % load data
options.solver = 'smo'; % use SMO solver
options.ker = 'rbf'; % use RBF kernel
options.arg = 1; % kernel argument
options.C = 10; % regularization constant
model = oaosvm(data,options); % training
% visualization
figure;
ppatterns(data);
ppatterns(model.sv.X,'ko',12);
pboundary(model);
```



Figure 5.3: Example showing the multi-class SVM classifier build by one-against-one decomposition.

### 5.5 Multi-class BSVM formulation

The input is a set $\mathcal{T}_{X Y}=\left\{\left(\boldsymbol{x}_{1}, y_{1}\right), \ldots,\left(\boldsymbol{x}_{l}, y_{l}\right)\right\}$ of training vectors $\boldsymbol{x}_{i} \in \mathcal{X} \subseteq \mathbb{R}^{n}$ and corresponding hidden states $y_{i} \in \mathcal{Y}=\{1,2, \ldots, c\}$. The goal is to train the multi-class rule $q: \mathcal{X} \subseteq \mathbb{R}^{n} \rightarrow \mathcal{Y}$ defined by (5.3). In the linear case, the parameters of the multiclass rule can be found by solving the multi-class BSVM formulation which is defined as

$$
\begin{equation*}
\left(\mathbf{W}^{*}, \boldsymbol{b}^{*}, \boldsymbol{\xi}^{*}\right)=\underset{\mathbf{W}, \boldsymbol{b}}{\operatorname{argmin}} \underbrace{\frac{1}{2} \sum_{y \in \mathcal{Y}}\left(\left\|\boldsymbol{w}_{y}\right\|^{2}+b_{y}^{2}\right)+C \sum_{i \in \mathcal{I}} \sum_{y \in \mathcal{Y} \backslash\{y\}}\left(\xi_{i}\right)^{p}}_{F(\mathbf{W}, \boldsymbol{b}, \boldsymbol{\xi})}, \tag{5.9}
\end{equation*}
$$

subject to

$$
\begin{array}{rlrl}
\left\langle\boldsymbol{w}_{y_{i}} \cdot \boldsymbol{x}_{i}\right\rangle+b_{y_{i}}-\left(\left\langle\boldsymbol{w}_{y} \cdot \boldsymbol{x}_{i}\right\rangle+b_{y}\right) & \geq 1-\xi_{i}^{y}, & & i \in \mathcal{I}, y \in \mathcal{Y} \backslash\left\{y_{i}\right\}, \\
\xi_{i}^{y} \geq 0, & i \in \mathcal{I}, y \in \mathcal{Y} \backslash\left\{y_{i}\right\},
\end{array}
$$

where $\mathbf{W}=\left[\boldsymbol{w}_{1}, \ldots, \boldsymbol{w}_{c}\right]$ is a matrix of normal vectors, $\boldsymbol{b}=\left[b_{1}, \ldots, b_{c}\right]^{T}$ is vector of biases, $\boldsymbol{\xi}$ is a vector of slack variables and $\mathcal{I}=[1, \ldots, l]$ is set of indices. The $L 1$-soft margin is used for $p=1$ and $L 2$-soft margin for $p=2$. The letter B in BSVM stands for the bias term added to the objective (5.9).

In the case of the $L 2$-soft margin $p=2$, the optimization problem (5.9) can be
expressed in its dual form

$$
\begin{equation*}
\mathbf{A}^{*}=\underset{\mathbf{A}}{\operatorname{argmax}} \sum_{i \in \mathcal{I}} \sum_{y \in \mathcal{Y} \backslash\left\{y_{i}\right\}} \alpha_{i}^{y}-\frac{1}{2} \sum_{i \in \mathcal{I}} \sum_{y \in \mathcal{Y} \backslash\left\{y_{i}\right\}} \sum_{j \in \mathcal{I}} \sum_{u \in \mathcal{Y} \backslash\left\{y_{j}\right\}} \alpha_{i}^{y} \alpha_{j}^{u} h(y, u, i, j), \tag{5.10}
\end{equation*}
$$

subject to

$$
\alpha_{i}^{y} \geq 0, \quad i \in \mathcal{I}, y \in \mathcal{Y} \backslash\left\{y_{i}\right\}
$$

where $\mathbf{A}=\left[\boldsymbol{\alpha}^{1}, \ldots, \boldsymbol{\alpha}^{c}\right]$ are optimized weight vectors. The function $h: \mathcal{Y} \times \mathcal{Y} \times \mathcal{I} \times \mathcal{I} \rightarrow \mathbb{R}$ is defined as

$$
\begin{equation*}
h(y, u, i, j)=\left(\left\langle\boldsymbol{x}_{i} \cdot \boldsymbol{x}_{j}\right\rangle+1\right)\left(\delta\left(y_{i}, y_{j}\right)+\delta(y, u)-\delta\left(y_{i}, u\right)-\delta\left(y_{j}, y\right)\right)+\frac{\delta(y, u) \delta(i, j)}{2 C} . \tag{5.11}
\end{equation*}
$$

The criterion (5.10) corresponds to the single-class SVM criterion which is simple to optimize. The multi-class rule (5.3) is composed of the set of discriminant functions $f_{y}: \mathcal{X} \rightarrow \mathbb{R}$ which are computed as

$$
\begin{equation*}
f_{y}(\boldsymbol{x})=\sum_{i \in \mathcal{I}}\left\langle\boldsymbol{x}_{i} \cdot \boldsymbol{x}\right\rangle \sum_{u \in \mathcal{Y} \backslash\left\{y_{i}\right\}} \alpha_{i}^{u}\left(\delta\left(u, y_{i}\right)-\delta(u, y)\right)+b_{y}, \quad y \in \mathcal{Y} \tag{5.12}
\end{equation*}
$$

where the bias $b_{y}, y \in \mathcal{Y}$ is given by

$$
b_{y}=\sum_{i \in \mathcal{I}} \sum_{y \in \mathcal{Y} \backslash\left\{y_{i}\right\}} \alpha_{i}^{u}\left(\delta\left(y, y_{i}\right)-\delta(y, u)\right), \quad y \in \mathcal{Y}
$$

The non-linear (kernel) classifier is obtained by substituting the selected kernel $k: \mathcal{X} \times$ $\mathcal{X} \rightarrow \mathbb{R}$ for the dot products $\left\langle\boldsymbol{x} \cdot \boldsymbol{x}^{\prime}\right\rangle$ to (5.11) and (5.12). The training of the multi-class BSVM classifier with $L 2$-soft margin is implemented in the function bsvm2. The optimization task (5.10) can be solved by most the optimizers. The optimizers implemented in the function bsvm2 are enlisted bellow:

Mitchell-Demyanov-Malozemov (solver = 'mdm').
Kozinec's algorithm (solver = 'kozinec').
Nearest Point Algorithm (solver = 'npa').
The above mentioned algorithms are of iterative nature and converge to the optimal solution of the criterion (5.9). The upper bound $F_{U B}$ and the lower bound $F_{L B}$ on the optimal value $F\left(\mathbf{W}^{*}, \boldsymbol{b}^{*}, \boldsymbol{\xi}^{*}\right)$ of the criterion (5.9) can be computed. The bounds are used in the algorithms to define the following two stopping conditions:
Relative tolerance: $\frac{F_{U B}-F_{L B}}{F_{L B}+1} \leq \varepsilon_{\text {rel }}$.

Absolute tolerance: $F_{U B} \leq \varepsilon_{a b s}$.
References: The multi-class BSVM formulation is described in paper [12]. The transformation of the multi-class BSVM criterion to the single-class criterion (5.10) implemented in the STPRtool is published in paper [9]. The Mitchell-Demyanov-Malozemov and Nearest Point Algorithm are described in paper [14]. The Kozinec's algorithm based SVM solver is described in paper [10].

## Example: Multi-class BSVM with $L 2$-soft margin.

The example shows the training of the multi-class BSVM using the NPA solver. The other solvers can be used instead by setting the variable options. solver (other options are 'kozinec', 'mdm'). The training data and the decision boundary is visualized in Figure 5.4.

```
data = load('pentagon'); % load data
options.solver = 'npa'; % use NPA solver
options.ker = 'rbf'; % use RBF kernel
options.arg = 1; % kernel argument
options.C = 10; % regularization constant
model = bsvm2(data,options); % training
% visualization
figure;
ppatterns(data);
ppatterns(model.sv.X,'ko',12);
pboundary(model);
```


### 5.6 Kernel Fisher Discriminant

The input is a set $\mathcal{T}_{X Y}=\left\{\left(\boldsymbol{x}_{1}, y_{1}\right), \ldots,\left(\boldsymbol{x}_{l}, y_{l}\right)\right\}$ of training vectors $\boldsymbol{x}_{i} \in \mathcal{X} \subseteq \mathbb{R}^{n}$ and corresponding values of binary state $y_{i} \in \mathcal{Y}=\{1,2\}$. The Kernel Fisher Discriminant is the non-linear extension of the linear FLD (see Section 2.3). The input training vectors are assumed to be mapped into a new feature space $\mathcal{F}$ by a mapping function $\phi: \mathcal{X} \rightarrow \mathcal{F}$. The ordinary FLD is applied on the mapped data training data $\mathcal{T}_{\Phi Y}=$ $\left\{\left(\boldsymbol{\phi}\left(\boldsymbol{x}_{1}\right), y_{1}\right), \ldots,\left(\boldsymbol{\phi}\left(\boldsymbol{x}_{l}\right), y_{l}\right)\right\}$. The mapping is performed efficiently by means of kernel functions $k: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ being the dot products of the mapped vectors $k\left(\boldsymbol{x}, \boldsymbol{x}^{\prime}\right)=$ $\left\langle\boldsymbol{\phi}(\boldsymbol{x}) \cdot \boldsymbol{\phi}\left(\boldsymbol{x}^{\prime}\right)\right\rangle$. The aim is to find a direction $\boldsymbol{\psi}=\sum_{i=1}^{l} \alpha_{i} \boldsymbol{\phi}\left(\boldsymbol{x}_{i}\right)$ in the feature space $\mathcal{F}$ given by weights $\boldsymbol{\alpha}=\left[\alpha_{1}, \ldots, \alpha_{l}\right]^{T}$ such that the criterion

$$
\begin{equation*}
F(\boldsymbol{\alpha})=\frac{\langle\boldsymbol{\alpha} \cdot \mathbf{M} \boldsymbol{\alpha}\rangle}{\langle\boldsymbol{\alpha} \cdot(\mathbf{N}+\mu \mathbf{E}) \boldsymbol{\alpha}\rangle}=\frac{\langle\boldsymbol{\alpha} \cdot \boldsymbol{m}\rangle^{2}}{\langle\boldsymbol{\alpha} \cdot(\mathbf{N}+\mu \mathbf{E}) \boldsymbol{\alpha}\rangle} \tag{5.13}
\end{equation*}
$$



Figure 5.4: Example showing the multi-class BSVM classifier with $L 2$-soft margin.
is maximized. The criterion (5.13) is the feature space counterpart of the criterion (2.13) defined in the input space $\mathcal{X}$. The vector $\boldsymbol{m}[l \times 1]$, matrices $\mathbf{N}[l \times l]$ $\mathbf{M}[l \times l]$ are constructed as

$$
\begin{aligned}
\boldsymbol{m}_{y} & =\frac{1}{\left|\mathcal{I}_{y}\right|} \mathbf{K} \mathbf{1}_{y}, \quad y \in \mathcal{Y}, & \boldsymbol{m}=\boldsymbol{m}_{1}-\boldsymbol{m}_{2} \\
\mathbf{N} & =\mathbf{K K}^{T}-\sum_{y \in \mathcal{Y}}\left|\mathcal{I}_{y}\right| \boldsymbol{m}_{y} \boldsymbol{m}_{y}^{T}, & \mathbf{M}=\left(\boldsymbol{m}_{1}-\boldsymbol{m}_{2}\right)\left(\boldsymbol{m}_{1}-\boldsymbol{m}_{2}\right)^{T}
\end{aligned}
$$

where the vector $\mathbf{1}_{y}$ has entries $i \in \mathcal{I}_{y}$ equal to and remaining entries zeros. The kernel matrix $\mathbf{K}[l \times l]$ contains kernel evaluations $K_{i, j}=k\left(\boldsymbol{x}_{i}, \boldsymbol{x}_{j}\right)$. The diagonal matrix $\mu \mathbf{E}$ in the denominator of the criterion (5.13) serves as the regularization term. The discriminant function $f(\boldsymbol{x})$ of the binary classifier (5.2) is given by the found vector $\boldsymbol{\alpha}$ as

$$
\begin{equation*}
f(\boldsymbol{x})=\langle\boldsymbol{\psi} \cdot \boldsymbol{\phi}(\boldsymbol{x})\rangle+b=\langle\boldsymbol{\alpha} \cdot \boldsymbol{k}(\boldsymbol{x})\rangle+b, \tag{5.14}
\end{equation*}
$$

where the $\boldsymbol{k}(\boldsymbol{x})=\left[k\left(\boldsymbol{x}_{1}, \boldsymbol{x}\right), \ldots, k\left(\boldsymbol{x}_{l}, \boldsymbol{x}\right)\right]^{T}$ is vector of evaluations of the kernel. The discriminant function (5.14) is known up to the bias $b$ which is determined by the linear SVM with $L 1$-soft margin is applied on the projected data $\mathcal{T}_{Z Y}=\left\{\left(z_{1}, y_{1}\right), \ldots,\left(z_{l}, y_{l}\right)\right\}$. The projections $\left\{z_{1}, \ldots, z_{l}\right\}$ are computed as

$$
z_{i}=\left\langle\alpha \cdot \boldsymbol{k}\left(\boldsymbol{x}_{i}\right)\right\rangle .
$$

The KFD training algorithm is implemented in the function kfd.

References: The KFD formulation is described in paper [19]. The detailed analysis can be found in book [30].

## Example: Kernel Fisher Discriminant

The binary classifier is trained by the KFD on the Riply's training set riply_trn . mat. The found classifier is evaluated on the testing data riply_tst.mat. The training data and the found kernel classifier is visualized in Figure 5.5. The classifier is visualized as the SVM classifier, i.e., the isolines with $f(\boldsymbol{x})=+1$ and $f(\boldsymbol{x})=-1$ are visualized and the support vectors (whole training set) are marked as well.

```
trn = load('riply_trn'); % load training data
options.ker = 'rbf'; % use RBF kernel
options.arg = 1; % kernel argument
options.C = 10; % regularization of linear SVM
options.mu = 0.001; % regularization of KFD criterion
model = kfd(trn, options) % KFD training
% visualization
figure;
ppatterns(trn);
psvm(model);
% evaluation on testing data
tst = load('riply_tst');
ypred = svmclass( tst.X, model );
cerror( ypred, tst.y )
ans =
    0.0960
```


### 5.7 Pre-image problem

Let $\mathcal{T}_{X}=\left\{\boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{l}\right\}$ be a set of vector from the input space $\mathcal{X} \subseteq \mathbb{R}^{n}$. The kernel methods works with the data non-linearly mapped $\boldsymbol{\phi}: \mathcal{X} \rightarrow \mathcal{F}$ into a new feature space $\mathcal{F}$. The mapping is performed implicitly by using the kernel functions $k: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ which are dot products of the input data mapped into the feature space $\mathcal{F}$ such that $k\left(\boldsymbol{x}_{a}, \boldsymbol{x}_{b}\right)=\left\langle\boldsymbol{\phi}\left(\boldsymbol{x}_{a}\right) \cdot \boldsymbol{\phi}\left(\boldsymbol{x}_{b}\right)\right\rangle$. Let $\boldsymbol{\psi} \in \mathcal{F}$ be given by the following kernel expansion

$$
\begin{equation*}
\boldsymbol{\psi}=\sum_{i=1}^{l} \alpha_{i} \boldsymbol{\phi}\left(\boldsymbol{x}_{i}\right) \tag{5.15}
\end{equation*}
$$



Figure 5.5: Example shows the binary classifier trained based on the Kernel Fisher Discriminant.
where $\boldsymbol{\alpha}=\left[\alpha_{1}, \ldots, \alpha_{l}\right]^{T}$ is a weight vector of real coefficients. The pre-image problem aims to find an input vector $\boldsymbol{x} \in \mathcal{X}$, its feature space image $\phi(\boldsymbol{x}) \in \mathcal{F}$ would well approximate the expansion point $\Psi \in \mathcal{F}$ given by (5.15). In other words, the preimage problem solves mapping from the feature space $\mathcal{F}$ back to the input space $\mathcal{X}$. The pre-image problem can be stated as the following optimization task

$$
\begin{align*}
\boldsymbol{x} & =\underset{\boldsymbol{x}^{\prime}}{\operatorname{argmin}}\|\boldsymbol{\phi}(\boldsymbol{x})-\boldsymbol{\psi}\|^{2} \\
& =\underset{\boldsymbol{x}^{\prime}}{\operatorname{argmin}}\left\|\boldsymbol{\phi}\left(\boldsymbol{x}^{\prime}\right)-\sum_{i=1}^{l} \alpha_{i} \boldsymbol{\phi}\left(\boldsymbol{x}_{i}\right)\right\|^{2}  \tag{5.16}\\
& =\underset{\boldsymbol{x}^{\prime}}{\operatorname{argmin}} k\left(\boldsymbol{x}^{\prime}, \boldsymbol{x}^{\prime}\right)-2 \sum_{i=1}^{l} \alpha_{i} k\left(\boldsymbol{x}^{\prime}, \boldsymbol{x}_{i}\right)+\sum_{i=1}^{l} \sum_{j=1}^{l} \alpha_{i} \alpha_{j} k\left(\boldsymbol{x}_{i}, \boldsymbol{x}_{j}\right) .
\end{align*}
$$

Let the Radial Basis Function (RBF) $k\left(\boldsymbol{x}_{a}, \boldsymbol{x}_{b}\right)=\exp \left(-0.5\left\|\boldsymbol{x}_{a}-\boldsymbol{x}_{b}\right\|^{2} / \sigma^{2}\right)$ be assumed. The pre-image problem (5.16) for the particular case of the RBF kernel leads to the following task

$$
\begin{equation*}
\boldsymbol{x}=\underset{\boldsymbol{x}^{\prime}}{\operatorname{argmax}} \sum_{i=1}^{l} \alpha_{i} \exp \left(-0.5\left\|\boldsymbol{x}^{\prime}-\boldsymbol{x}_{i}\right\|^{2} / \sigma^{2}\right) . \tag{5.17}
\end{equation*}
$$

The STPRtool provides three optimization algorithms to solve the RBF pre-image problem (5.17) which are enlisted in Table 5.6. However, all the methods are guaranteed to find only a local optimum of the task (5.17).

Table 5.6: Methods to solve the RBF pre-image problem implemented in the STPRtool.

$$
\begin{array}{ll}
\hline \text { rbfpreimg } & \text { An iterative fixed point algorithm proposed in }[28] . \\
\text { rbfpreimg2 } & \text { A standard gradient optimization method using the } \\
& \begin{array}{l}
\text { Matlab optimization toolbox for } 1 D \text { search along the } \\
\text { gradient. }
\end{array} \\
\text { rbfpreimg3 } & \begin{array}{l}
\text { A method exploiting the correspondence between dis- } \\
\text { tances in the input space and the feature space pro- } \\
\text { posed in [16]. }
\end{array} \\
&
\end{array}
$$

References: The pre-image problem and the implemented algorithms for its solution are described in $[28,16]$.

## Example: Pre-image problem

The example shows how to visualize a pre-image of the data mean in the feature induced by the RBF kernel. The result is visualized in Figure 5.6.

```
% load input data
data = load('riply_trn');
% define the mean in the feature space
num_data = size(data,2);
expansion.Alpha = ones(num_data,1)/num_data;
expansion.sv.X = data.X;
expansion.options.ker = 'rbf';
expansion.options.arg = 1;
% compute pre-image problem
x = rbfpreimg(expansion);
% visualization
figure;
ppatterns(data.X);
ppatterns(x,'or',13);
```



Figure 5.6: Example shows training data and the pre-image of their mean in the feature space induced by the RBF kernel.

### 5.8 Reduced set method

The reduced set method is useful to simplify the kernel classifier trained by the SVM or other kernel machines. Let the function $f: \mathcal{X} \subseteq \mathbb{R}^{n} \rightarrow \mathbb{R}$ be given as

$$
\begin{equation*}
f(\boldsymbol{x})=\sum_{i=1}^{l} \alpha_{i} k\left(\boldsymbol{x}_{i}, \boldsymbol{x}\right)+b=\langle\boldsymbol{\psi} \cdot \boldsymbol{\phi}(\boldsymbol{x})\rangle+b, \tag{5.18}
\end{equation*}
$$

where $\mathcal{T}_{X}=\left\{\boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{l}\right\}$ are vectors from $\mathbb{R}^{n}, k: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is a kernel function, $\boldsymbol{\alpha}=\left[\alpha_{1}, \ldots, \alpha_{l}\right]^{T}$ is a weight vector of real coefficients and $b \in \mathbb{R}$ is a scalar. Let $\phi: \mathcal{X} \rightarrow \mathcal{F}$ be a mapping from the input space $\mathcal{X}$ to the feature space $\mathcal{F}$ associated with kernel function such that $k\left(\boldsymbol{x}_{a}, \boldsymbol{x}_{b}\right)=\left\langle\boldsymbol{\phi}\left(\boldsymbol{x}_{a}\right) \cdot \boldsymbol{\phi}\left(\boldsymbol{x}_{b}\right)\right\rangle$. The function $f(\boldsymbol{x})$ can be expressed in the feature space $\mathcal{F}$ as the linear function $f(\boldsymbol{x})=\langle\boldsymbol{\psi} \cdot \boldsymbol{\phi}(\boldsymbol{x})\rangle+b$. The reduced set method aims to find a new function

$$
\begin{equation*}
\tilde{f}(\boldsymbol{x})=\sum_{i=1}^{m} \beta_{i} k\left(\boldsymbol{s}_{i}, \boldsymbol{x}\right)+b=\langle\boldsymbol{\phi}(\boldsymbol{x}) \cdot \underbrace{\sum_{i=1}^{m} \beta_{i} \boldsymbol{\phi}\left(\boldsymbol{s}_{i}\right)}_{\tilde{\boldsymbol{\psi}}}\rangle+b, \tag{5.19}
\end{equation*}
$$

such that the expansion (5.19) is shorter, i.e., $m<l$, and well approximates the original one (5.18). The weight vector $\boldsymbol{\beta}=\left[\beta_{1}, \ldots, \beta_{m}\right]^{T}$ and the vectors $\mathcal{T}_{S}=\left\{\boldsymbol{s}_{1}, \ldots, \boldsymbol{s}_{m}\right\}$, $s_{i} \in \mathbb{R}^{n}$ determine the reduced kernel expansion (5.19). The problem of finding the
reduced kernel expansion (5.19) can be stated as the optimization task

$$
\begin{equation*}
\left(\boldsymbol{\beta}, \mathcal{T}_{S}\right)=\underset{\boldsymbol{\beta}^{\prime}, \mathcal{T}_{S}^{\prime}}{\operatorname{argmin}}\|\tilde{\boldsymbol{\psi}}-\boldsymbol{\psi}\|^{2}=\underset{\boldsymbol{\beta}^{\prime}, \mathcal{T}_{S}^{\prime}}{\operatorname{argmin}}\left\|\sum_{i=1}^{m} \beta_{i}^{\prime} \boldsymbol{\phi}\left(s_{i}^{\prime}\right)-\sum_{i=1}^{l} \alpha_{i} \boldsymbol{\phi}\left(\boldsymbol{x}_{i}\right)\right\|^{2} . \tag{5.20}
\end{equation*}
$$

Let the Radial Basis Function (RBF) $k\left(\boldsymbol{x}_{a}, \boldsymbol{x}_{b}\right)=\exp \left(-0.5\left\|\boldsymbol{x}_{a}-\boldsymbol{x}_{b}\right\|^{2} / \sigma^{2}\right)$ be assumed. In this case, the optimization task (5.20) can be solved by an iterative greedy algorithm. The algorithm converts the task (5.20) into a series of $m$ pre-image tasks (see Section 5.7). The reduced set method for the RBF kernel is implemented in function rsrbf.
References: The implemented reduced set method is described in [28].

## Example: Reduced set method for SVM classifier

The example shows how to use the reduced set method to simplify the SVM classifier. The SVM classifier is trained from the Riply's training set riply_trn/mat. The trained kernel expansion (discriminant function) is composed of 94 support vectors. The reduced expansion with 10 support vectors is found. Decision boundaries of the original and the reduced SVM are visualized in Figure 5.7. The original and reduced SVM are evaluated on the testing data riply_tst.mat.

```
% load training data
trn = load('riply_trn');
% train SVM classifier
model = smo(trn,struct('ker','rbf','arg',1,'C',10));
% compute the reduced rule with 10 support vectors
red_model = rsrbf(model,struct('nsv',10));
% visualize decision boundaries
figure; ppatterns(trn);
h1 = pboundary(model,struct('line_style','r'));
h2 = pboundary(red_model,struct('line_style','b'));
legend([h1(1) h2(1)],'Original SVM','Reduced SVM');
% evaluate SVM classifiers
tst = load('riply_tst');
ypred = svmclass(tst.X,model);
err = cerror(ypred,tst.y);
red_ypred = svmclass(tst.X,model);
```

```
red_err = cerror(red_ypred,tst.y);
fprintf(['Original SVM: nsv = %d, err = %.4f\n' ...
    'Reduced SVM: nsv = %d, err = %.4f\n'], ...
    model.nsv, err, red_model.nsv, red_err);
```

Original SVM: nsv $=94$, err $=0.0960$
Reduced SVM: nsv $=10$, err $=0.0960$


Figure 5.7: Example shows the decision boundaries of the SVM classifier trained by SMO and the approximated SVM classifier found by the reduced set method. The original decision rule involves 94 support vectors while the reduced one only 10 support vectors.

## Chapter 6

## Miscellaneous

### 6.1 Data sets

The STPRtool provides several datasets and functions to generate synthetic data. The datasets stored in the Matlab data file are enlisted in Table 6.1. The functions which generate synthetic data and scripts for converting external datasets to the Matlab are given in Table 6.2.

Table 6.1: Datasets stored in /data/ directory of the STPRtool.

| /riply_trn.mat | Training part of the Riply's data set [24]. |
| :--- | :--- |
| /riply_tst.mat | Testing part of the Riply's data set [24]. |
| /iris.mat | Well known Fisher's data set. |
| /anderson_task/*.mat | Input data for the demo on the Generalized Ander- <br> son's task demo_anderson. |
| /binary_separable/*.mat | Input data for the demo on training linear classifiers <br> from finite vector sets demo_linclass. |
| /gmm_samles/*.mat | Input data for the demo on the Expectation- <br>  <br> Maximization algorithm for the Gaussian mixture <br> model demo_emgmm. |
| /multi_separable/*.mat | Linearly separable multi-class data in 2D. <br> Lxamples of hand-written numerals. A description is <br> given in Section 7.1. |
| /svm_samples/*.mat | Input data for the demo on the Support Vector Ma- <br> chines demo_svm. |

Table 6.2: Data generation.

| createdata | GUI which allows to create interactively labeled vec- <br> tors sets and labeled sets of Gaussian distributions. It <br> operates in 2-dimensional space only. |
| :--- | :--- |
| genlsdata | Generates linearly separable binary data with pre- <br> scribed margin. |
| gmmsamp | Generates data from the Gaussian mixture model. <br> gsamp <br> Generates data from the multivariate Gaussian distri- <br> bution. <br> gencircledata <br> Generates data on circle corrupted by the Gaussian <br> noise. |
| usps2mat | Converts U.S. Postal Service (USPS) database of <br> hand-written numerals to the Matlab file. |

### 6.2 Visualization

The STPRtool provides tools for visualization of common models and data. The list of functions for visualization is given in Table 6.3. Each function contains an example of using.

Table 6.3: Functions for visualization.

| pandr | Visualizes solution of the Generalized Anderson's |
| :--- | :--- |
| pboundary | task. |
| Plots decision boundary of given classifier in 2D. |  |
| pgauss | Visualizes set of bivariate Gaussians. |
| pgmm | Visualizes the bivariate Gaussian mixture model. |
| pkernelproj | Plots isolines of kernel projection. |
| plane3 | Plots plane in 3D. |
| pline | Plots line in 2D. |
| ppatterns | Visualizes patterns as points in the feature space. It |
|  | works for 1D, 2D and 3D. <br> psvm |
|  | Plots decision boundary of binary SVM classifier in |
| showim | 2D. |
| Displays images entered as column vectors. |  |

### 6.3 Bayesian classifier

The object under study is assumed to be described by a vector of observations $\boldsymbol{x} \in \mathcal{X}$ and a hidden state $y \in \mathcal{Y}$. The $\boldsymbol{x}$ and $y$ are realizations of random variables with joint probability distribution $P_{X Y}(\boldsymbol{x}, y)$. A decision rule $q: \mathcal{X} \rightarrow \mathcal{D}$ takes a decision $d \in \mathcal{D}$ based on the observation $\boldsymbol{x} \in \mathcal{X}$. Let $W: \mathcal{D} \times \mathcal{Y} \rightarrow \mathbb{R}$ be a loss function which penalizes the decision $q(\boldsymbol{x}) \in \mathcal{D}$ when the true hidden state is $y \in \mathcal{Y}$. Let $\mathcal{X} \subseteq \mathbb{R}^{n}$ and the sets $\mathcal{Y}$ and $\mathcal{D}$ be finite. The Bayesian risk $R(q)$ is an expectation of the value of the loss function $W$ when the decision rule $q$ is applied, i.e.,

$$
\begin{equation*}
R(q)=\int_{\mathcal{X}} \sum_{y \in \mathcal{Y}} P_{X Y}(\boldsymbol{x}, y) W(q(\boldsymbol{x}), y) \boldsymbol{d} \boldsymbol{x} \tag{6.1}
\end{equation*}
$$

The optimal rule $q^{*}$ which minimizes the Bayesian risk (6.1) is referred to as the Bayesian rule

$$
q^{*}(\boldsymbol{x})=\underset{y}{\operatorname{argmin}} \sum_{y \in \mathcal{Y}} P_{X Y}(\boldsymbol{x}, y) W(q(\boldsymbol{x}), y), \quad \forall \boldsymbol{x} \in \mathcal{X} .
$$

The STPRtool implements the Bayesian rule for two particular cases:
Minimization of misclassification: The set of decisions $\mathcal{D}$ coincides to the set of hidden states $\mathcal{Y}=\{1, \ldots, c\}$. The $0 / 1$-loss function

$$
W_{0 / 1}(q(\boldsymbol{x}), y)=\left\{\begin{array}{lll}
0 & \text { for } & q(\boldsymbol{x})=y  \tag{6.2}\\
1 & \text { for } & q(\boldsymbol{x}) \neq y
\end{array}\right.
$$

is used. The Bayesian risk (6.1) with the $0 / 1$-loss function corresponds to the expectation of misclassification. The rule $q: \mathcal{X} \rightarrow \mathcal{Y}$ which minimizes the expectation of misclassification is defined as

$$
\begin{align*}
q(\boldsymbol{x}) & =\underset{y \in \mathcal{Y}}{\operatorname{argmax}} P_{Y \mid X}(y \mid \boldsymbol{x}), \\
& =\underset{y \in \mathcal{Y}}{\operatorname{argmax}} P_{X \mid Y}(\boldsymbol{x} \mid y) P_{Y}(y) . \tag{6.3}
\end{align*}
$$

Classification with reject-option: The set of decisions $\mathcal{D}$ is assumed to be $\mathcal{D}=$ $\mathcal{Y} \cup\{$ dont_know $\}$. The loss function is defined as

$$
W_{\varepsilon}(q(\boldsymbol{x}), y)=\left\{\begin{array}{lll}
0 & \text { for } & q(\boldsymbol{x})=y  \tag{6.4}\\
1 & \text { for } & q(\boldsymbol{x}) \neq y \\
\varepsilon & \text { for } & q(\boldsymbol{x})=\text { dont_know }
\end{array}\right.
$$

where $\varepsilon$ is penalty for the decision dont_know. The rule $q: \mathcal{X} \rightarrow \mathcal{Y}$ which minimizes the Bayesian risk with the loss function (6.4) is defined as

$$
q(\boldsymbol{x})= \begin{cases}\underset{y \in \mathcal{Y}}{\operatorname{argmax}} P_{X \mid Y}(\boldsymbol{x} \mid y) P_{Y}(y) & \text { if } 1-\max _{y \in \mathcal{Y}} P_{Y \mid X}(y \mid \boldsymbol{x})<\varepsilon  \tag{6.5}\\ \text { dont_know } & \text { if } 1-\max _{y \in \mathcal{Y}} P_{Y \mid X}(y \mid \boldsymbol{x}) \geq \varepsilon .\end{cases}
$$

To apply the optimal classification rules one has to know the class-conditional distributions $P_{X \mid Y}$ and a priory distribution $P_{Y}$ (or their estimates). The data-type used by the STPRtool is described in Table 6.4. The Bayesian classifiers (6.3) and (6.5) are implemented in function bayescls.

Table 6.4: Data-type used to describe the Bayesian classifier.

| Bayesian classifier (structure array): |  |
| :--- | :--- |
| .Pclass $[$ cell $1 \times c]$ | Class conditional distributions $P_{X \mid Y}$. The distribution |
|  | $P_{X \mid Y}(\boldsymbol{x} \mid y)$ is given by Pclassy which uses the data- |
|  | type described in Table 4.2. |
| . Prior $[1 \times c]$ | Discrete distribution $P_{Y}(y), y \in \mathcal{Y}$. |
| . fun $={ }^{\prime}$ 'bayescls' | Identifies function associated with this data type. |

References: The Bayesian decision making with applications to PR is treated in details in book [26].

## Example: Bayesian classifier

The example shows how to design the plug-in Bayesian classifier for the Riply's training set riply_trn.mat. The class-conditional distributions $P_{X \mid Y}$ are modeled by the Gaussian mixture models (GMM) with two components. The GMM are estimated by the EM algorithm emgmm. The a priori probabilities $P_{Y}(y), y \in\{1,2\}$ are estimated by the relative occurrences (it corresponds to the ML estimate). The designed Bayesian classifier is evaluated on the testing data riply_tst.mat.

```
% load input training data
trn = load('riply_trn');
inx1 = find(trn.y==1);
inx2 = find(trn.y==2);
% Estimation of class-conditional distributions by EM
bayes_model.Pclass{1} = emgmm(trn.X(:,inx1),struct('ncomp',2));
bayes_model.Pclass{2} = emgmm(trn.X(:,inx2),struct('ncomp',2));
% Estimation of priors
n1 = length(inx1); n2 = length(inx2);
bayes_model.Prior = [n1 n2]/(n1+n2);
% Evaluation on testing data
tst = load('riply_tst');
ypred = bayescls(tst.X,bayes_model);
cerror(ypred,tst.y)
```

ans $=$
0.0900

The following example shows how to visualize the decision boundary of the found Bayesian classifier minimizing the misclassification (solid black line). The Bayesian classifier splits the feature space into two regions corresponding to the first class and the second class. The decision boundary of the reject-options rule (dashed line) is displayed for comparison. The reject-option rule splits the feature space into three regions corresponding to the first class, the second class and the region in between corresponding to the dont_know decision. The visualization is given in Figure 6.1.

```
% Visualization
figure; hold on; ppatterns(trn);
bayes_model.fun = 'bayescls';
pboundary(bayes_model);
% Penalization for don't know decision
reject_model = bayes_model;
reject_model.eps = 0.1;
% Vislualization of rejet-option rule
pboundary(reject_model,struct('line_style','k--'));
```


### 6.4 Quadratic classifier

The quadratic classification rule $q: \mathcal{X} \subseteq \mathbb{R}^{n} \rightarrow \mathcal{Y}=\{1,2, \ldots, c\}$ is composed of a set of discriminant functions

$$
f_{y}(\boldsymbol{x})=\left\langle\boldsymbol{x} \cdot \mathbf{A}_{y} \boldsymbol{x}\right\rangle+\left\langle\boldsymbol{b}_{y} \cdot \boldsymbol{x}\right\rangle+c_{y}, \quad \forall y \in \mathcal{Y},
$$

which are quadratic with respect to the input vector $\boldsymbol{x} \in \mathbb{R}^{n}$. The quadratic discriminant function $f_{y}$ is determined by a matrix $\mathbf{A}_{y}[n \times n]$, a vector $\boldsymbol{b}_{y}[n \times 1]$ and a scalar $c_{y}[1 \times 1]$. The input vector $\boldsymbol{x} \in \mathbb{R}^{n}$ is assigned to the class $y \in \mathcal{Y}$ its corresponding discriminant function $f_{y}$ attains maximal value

$$
\begin{equation*}
y=\underset{y^{\prime} \in \mathcal{Y}}{\operatorname{argmax}} f_{y^{\prime}}(\boldsymbol{x})=\underset{y^{\prime} \in \mathcal{Y}}{\operatorname{argmax}}\left(\left\langle\boldsymbol{x} \cdot \mathbf{A}_{y} \boldsymbol{x}\right\rangle+\left\langle\boldsymbol{b}_{y} \cdot \boldsymbol{x}\right\rangle+c_{y}\right) . \tag{6.6}
\end{equation*}
$$

In the particular binary case $\mathcal{Y}=\{1,2\}$, the quadratic classifier is represented by a single discriminant function

$$
f(\boldsymbol{x})=\langle\boldsymbol{x} \cdot \mathbf{A} \boldsymbol{x}\rangle+\langle\boldsymbol{b} \cdot \boldsymbol{x}\rangle+c,
$$



Figure 6.1: Figure shows the decision boundary of the Bayesian classifier (solid line) and the decision boundary of the reject-option rule with $\varepsilon=0.1$ (dashed line). The class-conditional probabilities are modeled by the Gaussian mixture models with two components estimated by the EM algorithm.
given by a matrix $\mathbf{A}[n \times n]$, a vector $\boldsymbol{b}[n \times 1]$ and a scalar $c[1 \times 1]$. The input vector $\boldsymbol{x} \in \mathbb{R}^{n}$ is assigned to class $y \in\{1,2\}$ as follows

$$
q(\boldsymbol{x})=\left\{\begin{array}{lll}
1 & \text { if } & f(\boldsymbol{x})=\langle\boldsymbol{x} \cdot \mathbf{A} \boldsymbol{x}\rangle+\langle\boldsymbol{b} \cdot \boldsymbol{x}\rangle+c \geq 0,  \tag{6.7}\\
2 & \text { if } & f(\boldsymbol{x})=\langle\boldsymbol{x} \cdot \mathbf{A} \boldsymbol{x}\rangle+\langle\boldsymbol{b} \cdot \boldsymbol{x}\rangle+c<0 .
\end{array}\right.
$$

The STPRtool uses a specific structure array to describe the binary (see Table 6.5) and the multi-class (see Table 6.6) quadratic classifier. The implementation of the quadratic classifier itself provides function quadclass.

Table 6.5: Data-type used to describe binary quadratic classifier.

| Binary linear quadratic | (structure array): |
| :--- | :--- |
| . $\mathrm{A}[n \times n]$ | Parameter matrix A. |
| . $\mathrm{b}[n \times 1]$ | Parameter vector $\boldsymbol{b}$. |
| . $\mathrm{c}[1 \times 1]$ | Parameter scalar $c$. |
| . fun $=$ 'quadclass' | Identifies function associated with this data type. |

## Example: Quadratic classifier

The example shows how to train the quadratic classifier using the Fisher Linear Discriminant and the quadratic data mapping. The classifier is trained from the Riply's

Table 6.6: Data-type used to describe multi-class quadratic classifier.

| Multi-class quadratic classifier (structure array): |  |
| :--- | :--- |
| . $\mathrm{A}[n \times n \times c]$ | Parameter matrices $\mathbf{A}_{y}, y \in \mathcal{Y}$. |
| . $\mathrm{b}[n \times c]$ | Parameter vectors $b_{y}, y \in \mathcal{Y}$. |
| . $\mathrm{c}[1 \times c]$ | Parameter scalars $c_{y}, y \in \mathcal{Y}$. |
| . fun $=$ 'quadclass' | Identifies function associated with this data type. |

training set riply_trn.mat. The found quadratic classifier is evaluated on the testing data riply_trn.mat. The boundary of the quadratic classifier is visualized (see Figure 6.2).

```
trn = load('riply_trn'); % load input data
quad_data = qmap(trn); % quadratic mapping
lin_model = fld(quad_data); % train FLD
% make quadratic classifier
quad_model = lin2quad(lin_model);
% visualization
figure;
ppatterns(trn); pboundary(quad_model);
% evaluation
tst = load('riply_tst');
ypred = quadclass(tst.X,quad_model);
cerror(ypred,tst.y)
```

ans $=$
0.0940

## 6.5 $K$-nearest neighbors classifier

Let $\mathcal{T}_{X Y}=\left\{\left(\boldsymbol{x}_{1}, y_{1}\right), \ldots,\left(\boldsymbol{x}_{l}, y_{l}\right)\right\}$ be a set of prototype vectors $\boldsymbol{x}_{i} \in \mathcal{X} \subseteq \mathbb{R}^{n}$ and corresponding hidden states $y_{i} \in \mathcal{Y}=\{1, \ldots, c\}$. Let $\mathbb{R}^{n}(\boldsymbol{x})=\left\{\boldsymbol{x}^{\prime}:\left\|\boldsymbol{x}-\boldsymbol{x}^{\prime}\right\| \leq r^{2}\right\}$ be a ball centered in the vector $\boldsymbol{x}$ in which lie $K$ prototype vectors $\boldsymbol{x}_{i}, i \in\{1, \ldots, l\}$, i.e., $\left|\left\{\boldsymbol{x}_{i}: \boldsymbol{x}_{i} \in \mathbb{R}^{n}(\boldsymbol{x})\right\}\right|=K$. The $K$-nearest neighbor (KNN) classification rule $q: \mathcal{X} \rightarrow \mathcal{Y}$ is defined as

$$
\begin{equation*}
q(\boldsymbol{x})=\underset{y \in \mathcal{Y}}{\operatorname{argmax}} v(\boldsymbol{x}, y), \tag{6.8}
\end{equation*}
$$



Figure 6.2: Figure shows the decision boundary of the quadratic classifier.
where $v(\boldsymbol{x}, y)$ is number of prototype vectors $\boldsymbol{x}_{i}$ with hidden state $y_{i}=y$ which lie in the ball $\boldsymbol{x}_{i} \in \mathbb{R}^{n}(\boldsymbol{x})$. The classification rule (6.8) is computationally demanding if the set of prototype vectors is large. The data-type used by the STPRtool to represent the $K$-nearest neighbor classifier is described in Table 6.7.

Table 6.7: Data-type used to describe the $K$-nearest neighbor classifier.

| $K$-NN classifier (structure array): |  |
| :--- | :--- |
| . $\mathrm{X}[n \times l]$ | Prototype vectors $\left\{\boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{l}\right\}$. |
| $\cdot \mathrm{y}[1 \times l]$ | Labels $\left\{y_{1}, \ldots, y_{l}\right\}$. |
| $. \mathrm{fun}=$ 'knnclass' | Identifies function associated with this data type. |

References: The $K$-nearest neighbors classifier is described in most books on PR, for instance, in book [6].

## Example: $K$-nearest neighbor classifier

The example shows how to create the $(K=8)$-NN classification rule from the Riply's training data riply_trn.mat. The classifier is evaluated on the testing data riply_tst.mat. The decision boundary is visualized in Figure 6.3.

```
% load training data and setup 8-NN rule
trn = load('riply_trn');
model = knnrule(trn,8);
```

\% visualize decision boundary and training data
figure; ppatterns(trn); pboundary(model);
\% evaluate classifier
tst $=$ load('riply_tst');
ypred = knnclass(tst.X,model);
cerror (ypred,tst.y)
ans =
0.1160


Figure 6.3: Figure shows the decision boundary of the $(K=8)$-nearest neighbor classifier.

## Chapter 7

## Examples of applications

We intend to demonstrate how methods implemented in the STPRtool can be used to solve a more complex applications. The applications selected are the optical character recognition (OCR) system based on the multi-class SVM (Section 7.1) and image denoising system using the kernel PCA (Section 7.2).

### 7.1 Optical Character Recognition system

The use of the STPRtool is demonstrated on a simple Optical Character Recognition (OCR) system for handwritten numerals from 0 to 9. The STPRtool provides GUI which allows to use standard computer mouse as an input device to draw images of numerals. The GUI is intended just for demonstration purposes and the scanner or other device would be used in practice instead. The multi-class SVM are used as the base classifier of the numerals. A numeral to be classified is represented as a vector $\boldsymbol{x}$ containing pixel values taken from the numeral image. The training stage of the OCR system involves (i) collection of training examples of numerals and (ii) training the multi-class SVM. To see the trained OCR system run the demo script demo_ocr.

The GUI for drawing numerals is implemented in a function mpaper. The function mpaper creates a form composed of 50 cells to which the numerals can be drawn (see Figure 7.1a). The drawn numerals are normalized to size $16 \times 16$ pixels (different size can be specified). The function mpaper is designed to invoke a specified function whenever the middle mouse button is pressed. In particular, the functions save_chars and ocr_fun are called from the mpaper to process drawn numerals. The function save_chars saves drawn numerals in the stage of collecting training examples and the function ocr_fun is used to recognize the numerals if the trained OCR is applied.

The set of training examples has to be collected before the OCR is trained. The function collect_chars invokes the form for drawing numerals and allows to save the numerals to a specified file. The STPRtool contains examples of numerals collected
from 6 different persons. Each person drew 50 examples of each of numerals from ' 0 ' to ' 9 '. The numerals are stored in the directory /data/ocr_numerals/. The file name format name_xx.mat (the character _ can be omitted from the name) is used. The string name is name of the person who drew the particular numeral and xx stands for a label assigned to the numeral. The labels from $y=1$ to $y=9$ corresponds to the numerals from ' 1 ' to ' 9 ' and the label $y=10$ is used for numeral ' 0 '. For example, to collect examples of numeral ' 1 ' from person honza_cech

```
collect_chars('honza_cech1')
```

was called. Figure 7.1a shows the form filled with the examples of the numeral '1'. The Figure 7.1b shows numerals after normalization which are saved to the specified file honza_cech.mat. The images of normalized numerals are stored as vectors to a matrix $\mathbf{X}$ of size $[256 \times 50]$. The vector dimension $n=256$ corresponds to $16 \times 16$ pixels of each image numeral and $l=50$ is the maximal number of numerals drawn to a single form. If the file name format name_xx.mat is met then the function mergesets(Directory, DataFile) can be used to merge all the file to a single one. The resulting file contains a matrix X of all examples and a vector y containing labels are assigned to examples according the number xx . The string Directory specifies the directory where the files name_xx.mat are stored and the string FileName is a file name to which the resulting training set is stored.

Having the labeled training set $\mathcal{T}_{X Y}=\left\{\left(\boldsymbol{x}_{1}, y_{1}\right), \ldots,\left(\boldsymbol{x}_{l}, y_{l}\right)\right\}$ created the multi-class SVM classifier can be trained. The methods to train the multi-class SVM are described in Chapter 5. In particular, the One-Against-One (OAO) decomposition was used as it requires the shortest training time. However, if a speed of classification is the major objective the multi-class BSVM formulation is often a better option as it produce less number of support vectors. The SVM has two free parameters which have to be tuned: the regularization constant $C$ and the kernel function $k\left(\boldsymbol{x}_{a}, \boldsymbol{x}_{b}\right)$ with its argument. The RBF kernel $k\left(\boldsymbol{x}_{a}, \boldsymbol{x}_{b}\right)=\exp \left(-0.5\left\|\boldsymbol{x}_{a}-\boldsymbol{x}_{b}\right\|^{2} / \sigma^{2}\right)$ determined by the parameter $\sigma$ was used in the experiment. Therefore, the free SVM parameters are $C$ and $\sigma$. A common method to tune the free parameters is to use the cross-validation to select the best parameters from a pre-selected set $\Theta=\left\{\left(C_{1}, \sigma_{1}\right), \ldots,\left(C_{k}, \sigma_{k}\right)\right\}$. The STPRtool contains a function evalsvm which evaluates the set of SVM parameters $\Theta$ based on the cross-validation estimation of the classification error. The function evalsvm returns the SVM model its parameters $\left(C^{*}, \sigma^{*}\right)$ turned out to be have the lowest cross-validation error over the given set $\Theta$. The parameter tuning is implemented in the script tune_ocr. The parameters $C=\infty$ (data are separable) and $\sigma=5$ were found to be the best.

After the parameters $\left(C^{*}, \sigma^{*}\right)$ are tuned the SVM classifier is trained using all training data available. The training procedure for OCR classifier is implemented in the function train_ocr. The directory /demo/ocr/models/ contains the SVM model trained by OAO decomposition strategy (function oaosvm) and model trained based on the multi-class BSVM formulation (function bsvm2).

The particular model used by OCR is specified in the function ocr_fun which implicitly uses the file ocrmodel.mat stored in the directory/demo/ocr/. To run the resulting OCR type demo_ocr or call

```
mpaper(struct('fun','ocr_fun'))
```

If one intends to modify the classifier or the way how the classification results are displayed then modify the function ocr_fun. A snapshot of running OCR is shown in Figure 7.2.

### 7.2 Image denoising

The aim is to train a model of a given class of images in order to restore images corrupted by noise. The U.S. Postal Service (USPS) database of images of handwritten numerals [17] was used as the class of images to model. The additive Gaussian noise was assumed as the source of image degradation. The kernel PCA was used to model the image class. This method was proposed in [20] and further developed in [15].

The STPRtool contains a script which converts the USPS database to the Matlab data file. It is implemented in the function usps2mat. The corruption of the USPS images by the Gaussian noise is implemented in script make_noisy_data. The script make_noisy_data produces the Matlab file usps_noisy.mat its contains is summarized in Table 7.1.

The aim is to train a model of images of all numerals at ones. Thus the unlabeled training data set $\mathcal{T}_{X}=\left\{\boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{l}\right\}$ is used. The model should describe the non-linear subspace $\mathcal{X}_{\text {img }} \subseteq \mathbb{R}^{n}$ where the images live. It is assumed that the non-linear subspace $\mathcal{X}_{\text {img }}$ forms a linear subspace $\mathcal{F}_{\text {img }}$ of a new feature space $\mathcal{F}$. A link between the input space $\mathcal{X}$ and the feature space $\mathcal{F}$ is given by a mapping $\phi: \mathcal{X} \rightarrow \mathcal{F}$ which is implicitly determined by a selected kernel function $k: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$. The kernel function corresponds to the dot product in the feature space, i.e., $k\left(\boldsymbol{x}_{a}, \boldsymbol{x}_{b}\right)=\left\langle\boldsymbol{\phi}\left(\boldsymbol{x}_{a}\right) \cdot \boldsymbol{\phi}\left(\boldsymbol{x}_{b}\right)\right\rangle$. The kernel PCA can be used to find the basis vectors of the subspace $\mathcal{F}_{\text {img }}$ from the input training data $\mathcal{T}_{X}$. Having the kernel PCA model of the subspace $\mathcal{F}_{\text {img }}$ it can be used to map the images $\boldsymbol{x} \in \mathcal{X}$ into the subspace $\mathcal{X}_{i m g}$. A noisy image $\hat{\boldsymbol{x}}$ is assumed to lie outside the subspace $\mathcal{X}_{\text {img }}$. The denoising procedure is based on mapping the noisy image $\hat{\boldsymbol{x}}$ into the subspace $\mathcal{X}_{\text {img }}$. The idea of using the kernel PCA for image denoising is demonstrated in Figure 7.3. The example which produced the figure is implemented in demo_kpcadenois.

The subspace $\mathcal{F}_{\text {img }}$ is modeled by the kernel PCA trained from the set $\mathcal{T}_{X}$ mapped into the feature space $\mathcal{F}$ by a function $\boldsymbol{\phi}: \mathcal{X} \rightarrow \mathcal{F}$. Let $\mathcal{T}_{\Phi}=\left\{\boldsymbol{\phi}\left(\boldsymbol{x}_{1}\right), \ldots, \boldsymbol{\phi}\left(\boldsymbol{x}_{l}\right)\right\}$ denote the feature space representation of the training images $\mathcal{T}_{X}$. The kernel PCA computes the lower dimensional representation $\mathcal{T}_{Z}=\left\{\boldsymbol{z}_{1}, \ldots, \boldsymbol{z}_{l}\right\}, \boldsymbol{z}_{i} \in \mathbb{R}^{m}$ of the mapped images

Table 7.1: The content of the file usps_noisy.mat produced by the script make_noisy_data.

| Structure array trn containing the training data: |  |
| :---: | :---: |
| gnd_X [256 × 7291] | Training images of numerals represented as column vectors. The input are 7291 gray-scale images of size $16 \times 16$. |
| $\mathrm{X}[256 \times 7291]$ | Training images corrupted by the Gaussian noise with signal-to-noise ration $\mathrm{SNR}=1$ (it can be changed). |
| $\mathrm{y}[1 \times 7291]$ | Labels $y \in \mathcal{Y}=\{1, \ldots, 10\}$ of training images. |


| Structure array tst containing the testing data: |  |  |
| :--- | :--- | :---: |
| gnd_X $[256 \times 2007]$ | Testing images of numerals represented as column vec- <br> tors. The input are 2007 gray-scale images of size |  |
| $\mathrm{X}[256 \times 2007]$ | $16 \times 16$. |  |
| Training images corrupted by the Gaussian noise with |  |  |
| signal-to-noise ration SNR $=1$ (it can be changed). |  |  |
| Labels $y \in \mathcal{Y}=\{1, \ldots, 10\}$ of training images. |  |  |

$\mathcal{T}_{\Phi}$ to minimize the mean square reconstruction error

$$
\varepsilon_{K M S}=\sum_{i=1}^{l}\left\|\phi\left(\boldsymbol{x}_{i}\right)-\tilde{\phi}\left(\boldsymbol{x}_{i}\right)\right\|^{2}
$$

where

$$
\begin{equation*}
\tilde{\phi}(\boldsymbol{x})=\sum_{i=1}^{l} \beta_{i} \boldsymbol{\phi}(\boldsymbol{x}), \quad \boldsymbol{\beta}=\mathbf{A}(\boldsymbol{z}-\boldsymbol{b}), \quad \boldsymbol{z}=\mathbf{A}^{T} \boldsymbol{k}(\boldsymbol{x})+\boldsymbol{b} \tag{7.1}
\end{equation*}
$$

The vector $k(\boldsymbol{x})=\left[k\left(\boldsymbol{x}_{1}, \boldsymbol{x}\right), \ldots, k\left(\boldsymbol{x}_{l}, \boldsymbol{x}\right)\right]^{T}$ contains kernel functions evaluated at the training data $\mathcal{T}_{X}$. The matrix $\mathbf{A}[l \times m]$ and the vector $\boldsymbol{b}[m \times 1]$ are trained by the kernel PCA algorithm. Given a projection $\tilde{\phi}(\boldsymbol{x})$ onto the kernel PCA subspace an image $\boldsymbol{x}^{\prime} \in \mathbb{R}^{n}$ of the input space can be recovered by solving the pre-image problem

$$
\begin{equation*}
x^{\prime}=\underset{x}{\operatorname{argmin}}\|\phi(x)-\tilde{\phi}(x)\|^{2} \tag{7.2}
\end{equation*}
$$

The Radial Basis Function (RBF) kernel $k\left(\boldsymbol{x}_{a}, \boldsymbol{x}_{b}\right)=\exp \left(-0.5\left\|\boldsymbol{x}_{a}-\boldsymbol{x}_{b}\right\|^{2} / \sigma^{2}\right)$ was used here as the pre-image problem (7.2) can be well solved for this case. The whole procedure of reconstructing the corrupted image $\hat{\boldsymbol{x}} \in \mathbb{R}^{n}$ involves (i) using (7.1) to compute $\boldsymbol{z}$ which determines $\tilde{\boldsymbol{\phi}}(\boldsymbol{x})$ and (ii) solving the pre-image problem (7.2) which yields the resulting $\boldsymbol{x}^{\prime}$. This procedure is implemented in function kpcarec.

In contrast to the linear PCA, the kernel PCA allows to model the non-linearity which is very likely in the case of images. However, the non-linearity hidden in
the parameter $\sigma$ of the used RBF kernel has to be tuned properly to prevent overfitting. Another parameter is the dimension $m$ of the subspace kernel PCA subspace to be trained. The best parameters $\left(\sigma^{*}, m^{*}\right)$ was selected out of a prescribed set $\Theta=\left\{\left(\sigma_{1}, m_{1}\right), \ldots,\left(\sigma_{k}, m_{k}\right)\right\}$ such that the cross-validation estimate of the mean square error

$$
\varepsilon_{M S}(\sigma, m)=\int_{\mathcal{X}} P_{X}(\boldsymbol{x})\left(\left\|\boldsymbol{x}-\boldsymbol{x}^{\prime}\right\|^{2}\right) \boldsymbol{d} \boldsymbol{x}
$$

was minimal. The $\boldsymbol{x}$ stands for the ground truth image and $\boldsymbol{x}^{\prime}$ denotes the image restored from noisy one $\hat{\boldsymbol{x}}$ using (7.1) and (7.2). The images are assumed to be generated from distribution $P_{X}(\boldsymbol{x})$. Figure 7.4 shows the idea of tuning the parameters of the kernel PCA. The linear PCA was used to solve the same task for comparison. The same idea was used to tune the output dimension of the linear PCA. Having the parameters tuned the resulting kernel PCA and linear PCA are trained on the whole training set. The whole training procedure of the kernel PCA is implemented in the script train_kpca_denois. The training for linear PCA is implemented in the script train_lin_denois. The greedy kernel PCA Algorithm 3.4 is used for training the model. It produces sparse model and can deal with large data sets. The results of tuning the parameters are displayed in Figure 7.5. The figure was produced by a script show_denois_tuning.

The function kpcarec is used to denoise images using the kernel PCA model. In the case of the linear PCA, the function pcarec was used. An example of denoising of the testing data of USPS is implemented in script show_denoising. Figure 7.6 shows examples of ground truth images, noisy images, images restored by the linear PCA and kernel PCA.


Figure 7.1: Figure (a) shows hand-written training examples of numeral '1' and Figure (b) shows corresponding normalized images of size $16 \times 16$ pixels.

(a)

| ${ }_{0.45}^{0.5}$ | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 0 | 0 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| ${ }_{0}^{0.45-}$ |  |  |  |  |  |  |  |  |  |  |
| ${ }_{0}^{0.35}$ |  | 4 |  |  |  | 0 |  |  |  |  |
| $\begin{aligned} & 0.2 \\ & 0.15 \end{aligned}$ | 5 |  | 3 |  | 1 |  | 9 |  |  |  |
| ${ }_{0}^{0.055}$ |  |  |  | 2 |  |  |  | 8 |  |  |

(b)

Figure 7.2: A snapshot of running OCR system. Figure (a) shows hand-written numerals and Figure (b) shows recognized numerals.


Figure 7.3: Demonstration of idea of image denoising by kernel PCA.


Figure 7.4: Schema of tuning parameters of kernel PCA model used for image restoration.


Figure 7.5: Figure (a) shows plot of $\varepsilon_{M S}$ with respect to output dimension of linear PCA. Figure(b) shows plot of $\varepsilon_{M S}$ with respect to output dimension and kernel argument of the kernel PCA.

Ground truth


Numerals with added Gaussian noise


Linear PCA reconstruction


Kernel PCA reconstruction


Figure 7.6: The results of denosing the handwritten images of USPS database.

## Chapter 8

## Conclusions and future planes

The STPRtool as well as its documentation is being updating. Therefore we are grateful for any comment, suggestion or other feedback which would help in the development of the STPRtool.

There are many relevant methods and proposals which have not been implemented yet. We welcome anyone who likes to contribute to the STPRtool in any way. Please send us an email and your potential contribution.

The future planes involve the extensions of the toolbox by the following methods:

- Feature selection methods.
- AdaBoost and related learning methods.
- Efficient optimizers to for large scale Support Vector Machines learning problems.
- Model selection methods for SVM.
- Probabilistic Principal Component Analysis (PPCA) and mixture of PPCA estimated by the Expectation-Maximization (EM) algorithm.
- Discrete probability models estimated by the EM algorithm.


## Bibliography

[1] T.W. Anderson and R.R. Bahadur. Classification into two multivariate normal distributions with differrentia covariance matrices. Anals of Mathematical Statistics, 33:420-431, June 1962.
[2] G. Baudat and F. Anouar. Generalized discriminant analysis using a kernel approach. Neural Computation, 12(10):2385-2404, 2000.
[3] C.M. Bishop. Neural Networks for Pattern Recognition. Clarendon Press, Oxford, Great Britain, 3th edition, 1997.
[4] N. Cristianini and J. Shawe-Taylor. Support Vector Machines. Cambridge University Press, 2000.
[5] A.P. Dempster, N.M. Laird, and D.B. Rubin. Maximum likelihood from incomplete data via the EM Algorithm. Journal of the Royal Statistical Society, 39:185197, 1977.
[6] R.O. Duda, P.E. Hart, and D.G. Stork. Pattern Classification. John Wiley \& Sons, 2nd. edition, 2001.
[7] R.P.W. Duin. Prtools: A matlab toolbox for pattern recognition, 2000.
[8] V. Franc. Programové nástroje pro rozpoznávání (Pattern Recognition Programming Tools, In Czech). Master's thesis, Ceské vysoké učení technické, Fakulta elektrotechnická, Katedra kybernetiky, February 2000.
[9] V. Franc and V. Hlaváč. Multi-class support vector machine. In R. Kasturi, D. Laurendeau, and Suen C., editors, 16th International Conference on Pattern Recognition, volume 2, pages 236-239. IEEE Computer Society, 2002.
[10] V. Franc and V. Hlaváč. An iterative algorithm learning the maximal margin classifier. Pattern Recognition, 36(9):1985-1996, 2003.
[11] V. Franc and V. Hlaváč. Greedy algorithm for a training set reduction in the kernel methods. In N. Petkov and M.A. Westenberg, editors, Computer Analysis of Images and Patterns, pages 426-433, Berlin, Germany, 2003. Springer.
[12] C.W. Hsu and C.J. Lin. A comparison of methods for multiclass support vector machins. IEEE Transactions on Neural Networks, 13(2), March 2002.
[13] I.T. Jollife. Principal Component Analysis. Springer-Verlag, New York, 1986.
[14] S.S. Keerthi, S.K. Shevade, C. Bhattacharya, and K.R.K. Murthy. A fast iterative nearest point algorithm for support vector machine classifier design. IEEE Transactions on Neural Networks, 11(1):124-136, January 2000.
[15] Kim Kwang, In, Franz Matthias, O., and Schölkopf Bernhard. Kernel hebbian algorithm for single-fram super-resolution. In Leonardis Aleš and Bischof Horst, editors, Statisical Learning in Computer Vision, ECCV Workshop. Springer, May 2004.
[16] J.T. Kwok and I.W. Tsang. The pre-image problem in kernel methods. In Proceedings of the Twentieth International Conference on Machine Learning (ICML2003), pages 408-415, Washington, D.C., USA, August 2003.
[17] Y. LeCun, B. Boser, J.S. Denker, D. Henderson, R.E. Howard, W. Hubbard, and L.J Jackel. Backpropagation applied to handwritten zip code recognition. Neural Computation, 1:541-551, 1989.
[18] G. McLachlan and T. Krishnan. The EM Algorithm and Extensions. John Wiley \& Sons, New York, 1997.
[19] S. Mika, G. Rätsch, J. Weston, B. Schölkpf, and K. Müller. Fisher discriminant analysis with kernel. In Y.H. Hu, J. Larsen, and S. Wilson, E. Douglas, editors, Neural Networks for Signal Processing, pages 41-48. IEEE, 1999.
[20] S. Mika, B. Schölkopf, A. Smola, K.R. Müller, M. Scholz, and G. Rätsch. Kernel pca and de-noising in feature spaces. In M.S. Kearns, S.A. Solla, and D.A. Cohn, editors, Advances in Neural Information Processing Systems 11, pages 536-542, Cambridge, MA, 1999. MIT Press.
[21] I.T. Nabney. NETLAB : algorithms for pattern recognition. Advances in pattern recognition. Springer, London, 2002.
[22] J. Platt. Probabilities for sv machines. In A.J. Smola, P.J. Bartlett, B. Scholkopf, and D. Schuurmans, editors, Advances in Large Margin Classifiers (Neural Information Processing Series). MIT Press, 2000.
[23] J.C. Platt. Sequential minimal optimizer: A fast algorithm for training support vector machines. Technical Report MSR-TR-98-14, Microsoft Research, Redmond, 1998.
[24] B.D. Riply. Neural networks and related methods for classification (with discusion). J. Royal Statistical Soc. Series B, 56:409-456, 1994.
[25] M.I. Schlesinger. A connection between learning and self-learning in the pattern recognition (in Russian). Kibernetika, 2:81-88, 1968.
[26] M.I. Schlesinger and V. Hlaváč. Ten lectures on statistical and structural pattern recognition. Kluwer Academic Publishers, 2002.
[27] B. Schölkopf, C. Burges, and A.J. Smola. Advances in Kernel Methods - Support Vector Learning. MIT Press, Cambridge, 1999.
[28] B. Schölkopf, P. Knirsch, and C. Smola, A. Burges. Fast approximation of support vector kernel expansions, and an interpretation of clustering as approximation in feature spaces. In P. Levi, M. Schanz, R.J. Ahler, and F. May, editors, Mustererkennung 1998-20. DAGM., pages 124-132, Berlin, Germany, 1998. SpringerVerlag.
[29] B. Scholkopf, A. Smola, and K.R. Muller. Nonlinear component analysis as a kernel eigenvalue problem. Neural Computation, 10:1299-1319, 1998.
[30] B. Schölkopf and A.J. Smola. Learning with Kernels. The MIT Press, MA, 2002.
[31] V. Vapnik. The nature of statistical learning theory. Springer Verlag, 1995.


[^0]:    ${ }^{1}$ The $K$ in the name $K$-means stands for the number of centers which is here denoted by $c$.

