# Intelligent Systems: Reasoning and Recognition 

James L. Crowley

ENSIMAG 2 / MoSIG M1

Second Semester 2009/2010

Lesson 20
7 May 2010

## Linear Classification Methods

## Contents

Notation ..... 2
Learning Linear Classifiers. ..... 3
Pattern detectors as linear classifiers ..... 4
Least squares estimation of a hyperplane ..... 5
A Committee of Boosted Classifiers ..... 6
Learning a Committee of Classifiers with Boosting ..... 7
ROC Curve ..... 8
Learning a Multi-Stage Cascade of Classifiers ..... 9
Perceptrons ..... 10
Kernel Methods ..... 12

Sources Bibliographiques:
"Neural Networks for Pattern Recognition", C. M. Bishop, Oxford Univ. Press, 1995.
"Pattern Recognition and Scene Analysis", R. E. Duda and P. E. Hart, Wiley, 1973.

## Notation

x
X
$\vec{x}$
$\vec{X}$
D
E
k
K
$\omega_{k}$
$\mathrm{M}_{\mathrm{k}}$
M Total number of examples.
$M=\sum_{k=1}^{K} M_{k}$
$\left\{X_{m}^{k}\right\} \quad$ A set of $\mathrm{M}_{\mathrm{k}}$ examples for the class k .

$$
\left\{X_{m}\right\}=\bigcup_{k=1, K}\left\{X_{m}^{k}\right\}
$$

$\left\{y_{m}\right\} \quad$ A set of class labels (indicators) for the samples For detection $(\mathrm{K}=2), \mathrm{y} \in\{+1,-1\}$

## Learning Linear Classifiers.

$$
\hat{\omega}_{k}=d(\vec{g}(\vec{X}))
$$

In lesson 16 we saw that the classification function can be decomposed into two parts: $d\left(g_{k}(X)\right)$ and $g_{k}(X)$ :

$$
\begin{aligned}
& \vec{g}(\vec{X})=\left(\begin{array}{c}
g_{1}(\vec{X}) \\
g_{2}(\vec{X}) \\
\ldots \\
g_{K}(\vec{X})
\end{array}\right) \quad \text { A set of discriminant functions }: \mathrm{R}^{\mathrm{D}} \rightarrow \mathrm{R}^{\mathrm{K}} \\
& \mathrm{~d}(): \quad \text { a decision function } \quad \mathrm{R}^{\mathrm{K}} \rightarrow\left\{\omega_{\mathrm{K}}\right\}
\end{aligned}
$$

We derived the canonical form for the discriminant function.

$$
g_{k}(\vec{X})=\vec{X}^{T} D_{k} \vec{X}+\vec{W}_{k}^{T} \vec{X}+b_{k}
$$

where:

$$
\begin{aligned}
& D_{k}=-\frac{1}{2} C_{k}^{-1} \\
& \vec{W}_{k}=-2 C_{k}^{-1} \vec{\mu}_{k} \\
& b_{k}=-\frac{1}{2} \vec{\mu}_{k}^{T} C_{k}^{-1} \vec{\mu}_{k}-\log \left\{\operatorname{det}\left(C_{k}\right)\right\}+\log \left\{p\left(\omega_{k}\right)\right\}
\end{aligned}
$$

and

A set of K discrimination functions $\mathrm{g}_{\mathrm{k}}(\vec{X})$ partitions the space $\vec{X}$ into a disjoint set of regions with quadratic boundaries. At the boundaries between classes:

$$
g_{i}(\vec{X})-g_{j}(\vec{X})=0
$$

In our last lesson we saw that in many cases the quadratic term can be ignored and the partitions take on the form of hyper-surfaces. In this case, the discrimination function can be reduced to a linear equation.

$$
g_{k}(\vec{X})=\vec{W}_{k}^{T} \vec{X}+b_{k}
$$

This is very useful because there are simple powerful techniques to calculate the terms of such an equation.

## Pattern detectors as linear classifiers.

Linear classifiers are also widely used to define pattern "detectors". This is widely used in computer vision, for example to detect faces or publicity logos, or other patterns of interest.

In the case of pattern detectors, $\mathrm{K}=2$.

Class $\mathrm{k}=1$ : The target pattern.
Class k=2: Everything else.

In the following examples, we will assume that our training data is composed of M sample observations $\left\{\vec{X}_{m}\right\}$ where each sample is labeled with an indicator $y_{m}$
$y_{m}=+1$ for examples of the target pattern (class 1 )
$y_{m}=-1$ for all other examples.

A variety of techniques exist to calculate the plane. The best choice can depend on the nature of the pattern class as well as the nature of the non-class data.

For example:

1) Vector between center of gravities.
2) Fisher linear discriminant analysis,
3) Least Squares estimation
4) Perceptrons

Wednesday we saw the first 2. Here we look at least-squares

## Least squares estimation of a hyperplane

Assume a training set of M labeled training samples $\left\{y_{m}, \vec{X}_{\mathrm{m}}\right\}$
such that $y_{m}=+1$ for class 1 and $y_{m}=-1$ for class 2 .
Our goal is to determine a function $g(\vec{X})=\vec{W}^{T} \vec{X}+b$

Such that we minimize a "Loss" function: $L(\hat{W})=\sum_{m=1}^{M}\left(y_{m}-\vec{W}^{T} \vec{X}_{m}\right)^{2}$
We will use the $M$ training samples to compose a matrix $\mathbf{X}$ and a vector $\mathbf{Y}$.

$$
\begin{aligned}
& \mathbf{X}=\left(\vec{X}_{1}, \overrightarrow{\mathrm{X}}_{2}, \ldots, \overrightarrow{\mathrm{X}}_{\mathrm{M}}\right)(\mathrm{D} \text { row by } \mathrm{M} \text { columns }) \\
& \mathbf{Y}=\left(\mathrm{y}_{1}, \mathrm{y}_{2}, \ldots, \mathrm{y}_{\mathrm{M}}\right)^{\mathrm{T}}(\mathrm{M} \text { coefficients }) .
\end{aligned}
$$

We write $\mathrm{L}(\mathrm{W})=\left(\mathbf{Y}-\mathbf{X}^{\mathrm{T}} \mathrm{W}\right)^{\mathrm{T}}\left(\mathbf{Y}-\mathbf{X}^{\mathrm{T}} \mathrm{W}\right)$
To minimize the loss function, we calculate the derivative and solve for W when the derivative is 0 .

$$
\frac{\partial L(W)}{\partial W}=-2 \mathbf{X}^{\mathrm{T}} \mathbf{Y}+\mathbf{2} \mathbf{X}^{\mathrm{T}} \mathbf{X} \mathbf{W}=0
$$

Thus: $\quad \mathbf{X}^{\mathrm{T}} \mathbf{Y}=\mathbf{X}^{\mathrm{T}} \mathbf{X} \mathbf{W}$

$$
\mathrm{W}=\left(\mathbf{X}^{\mathrm{T}} \mathbf{X}\right)^{-1} \mathbf{X}^{\mathrm{T}} \mathbf{Y}
$$

From this classifier an unknown event $\vec{X}$ as if $\vec{W}^{T} \vec{X}>0$ then $\hat{\omega}_{1}$ else $\hat{\omega}_{2}$

We can also trade False Positives for False negatives using a bias, B

$$
\text { if }\left(\vec{W}^{T} \vec{X}+B\right)>0 \text { then } \hat{\omega}_{1} \text { else } \hat{\omega}_{2}
$$

## A Committee of Boosted Classifiers

One of the more original ideas in machine learning the last decade is the discovery of a method by to learn a committee of classifiers by boosting. A boosted committee of classifiers can be made arbitrarily good: Adding a new classifier always improves performance.

A committee of classifiers decides by voting.


A feature vector is determined to be in the target class if the majority of classifiers vote $>0$.

$$
\text { if } \sum_{i=1}^{I}\left(\vec{W}_{i}^{T} \vec{X}_{m}\right)+b>0 \text { then } \hat{\omega}_{1} \text { else } \hat{\omega}_{2}
$$

To learn a boosted committee we iteratively add new classifiers to the committee. In each cycle we change the data set and learn a new classifier, $\mathrm{W}_{\mathrm{i}}$

The data set will be changed by giving additional weight to improperly classified samples. We learn the next class by multiplying the Y labels a weight vector, $\mathrm{A}_{\mathrm{i}}$.

$$
\vec{W}_{i}=\left(\boldsymbol{X}^{T} \boldsymbol{X}\right)^{-1} \boldsymbol{X}^{T}\left(\vec{A}_{i} \cdot \vec{Y}\right)
$$

## Learning a Committee of Classifiers with Boosting

We can iteratively apply the above procedure to learn a committee of classifiers using boosting. For this we will create a vector of "weights" $a_{m}$ for each training sample. Initially, all the weights are 1.

After each new classifier is added, we recalculate the weights to give more weight to improperly classified training samples.

As we add classifiers, whenever a sample is misclassified by the committee we will increase its weight so that it carries more weight in the next classifier added.

Recall the committee vote is $\sum_{i=1}^{l}\left(\vec{W}_{i}^{T} \vec{X}_{m}\right)>0$ for class 1 (positive detection).
For $\mathrm{m}=1$ to M: if $\left(\mathrm{y}_{\mathrm{m}} \cdot \sum_{i=1}^{I}\left(\vec{W}_{i}^{T} \vec{X}_{m}\right)\right)<0$ then $\mathrm{a}_{\mathrm{m}}=\mathrm{a}_{\mathrm{m}}+1$

The result is the $(i+1)^{\text {th }}$ weight vector $\mathrm{A}_{\mathrm{i}+1}$
We then learn the $i+1^{\text {th }}$ classifier from the re-weighted set by

$$
\vec{W}_{i+1}=\left(\boldsymbol{X}^{T} \boldsymbol{X}\right)^{-1} \boldsymbol{X}^{T}\left(\vec{A}_{i+1} \cdot \vec{Y}\right)
$$

## ROC Curve

As we saw in lesson 19, The ROC describes the True Positives (TP) and False Positives (FP) for a classifier as a function of the global bias B.

$$
\begin{aligned}
& \text { For } \mathrm{m}=1 \text { to } \mathrm{M} \text { : } \\
& \text { if } \sum_{i=1}^{I}\left(\vec{W}_{i}^{T} \vec{X}_{m}\right)+\mathrm{B}>0 \text { and } \mathrm{y}_{\mathrm{m}}>0 \text { then } \mathrm{TP}=\mathrm{TP}+1 \\
& \text { if } \quad \sum_{i=1}^{L}\left(\vec{W}_{i}^{T} \vec{X}_{m}\right)+\mathrm{B}>0 \text { and } \mathrm{y}_{\mathrm{m}}<0 \text { then } \mathrm{FP}=\mathrm{FP}+1
\end{aligned}
$$

The Boosting theorem states that adding a new boosted classifier to a committee always improves the committee's ROC curve. We can continue adding classifiers until we obtain a desired rate of false positives and false negatives.


However, in general, the improvement provided for each new classifier becomes progressively smaller. We can end with a very very large number of classifiers.

## Learning a Multi-Stage Cascade of Classifiers

We can optimize the computation time by using a multi-stage cascade.
With a multi-stage classifiers, only events labeled as positive are passed to the next stage.


Each stage is applied with a bias, so as to minimize False negatives.

Stages are organized so that each committee is successively more costly and more discriminant.

Assume a set of $M$ training samples $\left\{X_{m}\right\}$ with labels $\left\{y_{m}\right\}$.
Set a desired error rate for each stage $\mathrm{j}:\left(\mathrm{FP}_{\mathrm{j}}, \mathrm{FN}_{\mathrm{j}}\right)$.

For each stage, j , train the $\mathrm{j}+1$ stage with all positive samples from the previous stage.

Each stage acts as a filter, rejecting a grand number of easy cases, and passing the hard cases to the next stage. The stages become progressively more expensive, but are used progressively less often. Globally the computation cost decreases dramatically.

## Perceptrons

A perceptron is an incremental learning method for linear classifiers invented by Frank Rosenblatt in 1956. The perceptron is an on line learning method in which a linear classifier is improved by its own errors.

A perceptron learns a set of hyper-planes to separate training samples. When the training data are perfectly separated the data is said to be "separable". Otherwise, the data is said to be non-separable.

The "margin", $\gamma$, is the smallest separation between the two classes.

When are the training samples are separable, the algorithm uses the errors to update a plane until there are no more errors. When the training data is non-separable, the method may not converge, and must be arbitrarily stopped after a certain number of iterations.

Note that for all positive examples.

$$
\mathrm{ym}_{\mathrm{m}}\left(\mathrm{~W}^{\mathrm{T}} \overrightarrow{\mathrm{X}}_{\mathrm{m}}+\mathrm{B}\right)>0 \text { if the classification is correct. }
$$

The algorithm will apply a learning gain, $\eta$, to accelerate learning.

Algorithm:
$\vec{W}_{\mathrm{o}} \leftarrow 0 ; \mathrm{b}_{\mathrm{o}} \leftarrow 0 ; \mathrm{i}=0 ;$
$\mathrm{R} \leftarrow \max \left\{\left\|\vec{X}_{\mathrm{m}}\right\|\right\}$
REPEAT
FOR m $=1$ TO MDO
IF $\mathrm{ym}_{\mathrm{m}}\left(\overrightarrow{\mathrm{W}}_{\mathrm{i}}{ }^{\mathrm{T}} \overrightarrow{\mathrm{X}}_{\mathrm{m}}+\mathrm{b}_{\mathrm{i}}\right) \leq 0$ THEN
$\vec{W}_{i+1} \leftarrow \vec{W}_{i}+\eta y_{m} \vec{X}_{m}$;
$\mathrm{b}_{\mathrm{i}+1} \leftarrow \mathrm{~b}_{\mathrm{i}}+\eta \mathrm{ym}_{\mathrm{m}} \mathrm{R}^{2}$;
$\mathrm{i} \leftarrow \mathrm{i}+1$;
END IF
END FOR
UNTIL no mistakes in FOR loop.

After each stage the margin for each sample, $m$, is

$$
\gamma_{\mathrm{m}}=y_{\mathrm{m}}\left(\overrightarrow{\mathrm{~W}}_{\mathrm{i}}^{\mathrm{T}} \overrightarrow{\mathrm{X}}_{\mathrm{m}}+\mathrm{b}_{\mathrm{i}}\right)
$$

The coefficients must be normalised to compute the margin.

$$
W_{i}^{\prime}=\frac{W_{i}}{\left\|W_{i}\right\|} \quad b_{i}^{\prime}=\frac{b_{i}}{\left\|W_{i}\right\|}
$$

The decision rule is as before :

$$
\text { if } \left.\left(\vec{W}^{T} \vec{X}+B\right)\right)>0 \text { then } \hat{\omega}_{1} \text { else } \hat{\omega}_{2}
$$

The quality of the perceptron is give by the histogram of the margins.

## Kernel Methods

Linear methods are very well suited for use with very high dimensional feature space. We can map a quadratic decision space into a linear space by addiing additional dimensions.


A quadratic suface in a D dimensional space can be transformed into a linear surface in a $\mathrm{D}(\mathrm{D}+1) / 2$ space by from the D dimensional space to a space $\mathrm{P}>\mathrm{D}$. using a kernel function, K() .

For example a $\mathrm{D}=2$ quadratic space is linear in 5 dimensions:
$\vec{X}=\left(x_{1}, x_{2}, \ldots, x_{D}\right)$ can be projected into a $P=\frac{D(D+1)}{2}$ dimension defined by $\mathrm{K}(\mathrm{X})=\mathrm{W}=\left(\mathrm{x}_{1}, \mathrm{x}_{2}, \mathrm{x}_{1}{ }^{2}, \mathrm{x}_{1} \mathrm{x}_{2}, \mathrm{x}_{2}{ }^{2}\right)$

$$
\vec{W}=\left(x_{1}, x_{2}, \ldots, x_{D}, x_{1}^{2}, x_{1} x_{2}, x_{1} x_{3}, \ldots x_{D}-1 x_{D}, x_{D}^{2}\right)
$$

Ainsi, une fonction quadratique en $\mathrm{D}=2$ dimensions est linéaire en $\mathrm{P}=5$ dimensions

$$
\begin{aligned}
& \vec{X}=\left(x_{1}, x_{2}\right) \quad K(\vec{X})=\vec{W}=\left(x_{1}, x_{2}, x_{1}^{2}, x_{1} x_{2}, x_{2}^{2}\right) \\
& g(K(\vec{X}))=g(\vec{W})=a w_{1}+b w_{2}+c w_{3}+d w_{4}+e w_{5} \\
& \quad=a x_{1}+b x_{2}+c x_{1}^{2}+d x_{1} x_{2}+e x_{2}^{2}
\end{aligned}
$$

The Normal is often used as a Kernel:
$K(\vec{X})=\mathcal{N}\left(\vec{X} ; \vec{\mu}, \sigma^{2}\right)$

