## L3 Apprentissage

## Michèle Sebag - Benjamin Monmège LRI - LSV

27 février 2013

## Next course

## Tutorials/Videolectures

- http://www.iro.umontreal.ca/~bengioy/talks/icml2012-YBtutorial.pdf
- Part 1: 1-56; Part 2: 79-133
- Mar. 13th (oral participation)
- Some students present part 1-2
- Other students ask questions


## Hypothesis Space $\mathcal{H}$ / Navigation

This course

|  | $\mathcal{H}$ | navigation operators |
| :--- | ---: | ---: |
| Version Space | Logical | spec / gen |
| Decision Trees | Logical | specialisation |
| Neural Networks | Numerical | gradient |
| Support Vector Machines | Numerical | quadratic opt. |
| Ensemble Methods | - | adaptation $\mathcal{E}$ |

$$
h: \mathcal{X}=\mathbb{R}^{D} \mapsto \mathbb{R}
$$

Binary classification
$h(\mathbf{x})>0 \rightarrow \mathbf{x}$ classified as True else, classified as False

## Overview

Linear SVM, separable case
Linear SVM, non separable case
The kernel trick
The Kernel principle
Examples
Discussion
Extensions
Multi-class discrimination
Regression
Novelty detection
On the practitioner side
Improve precision
Reduce computational cost
Theory

## The separable case: <br> More than one separating hyperplane



## Linear Support Vector Machines

Linear Separators

$$
f(\mathbf{x})=\langle\mathbf{w}, \mathbf{x}\rangle+b
$$

Region $\hat{y}=1: f(\mathbf{x})>0$
Region $\hat{y}=-1$ : $f(\mathbf{x})<0$
Criterion

$$
\forall i, y_{i}\left(\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle+b\right)>0
$$

## Remark

Invariant by multiplication of $\mathbf{w}$ and $b$ by a positive value

## Canonical formulation

Fix the scale:

$$
\min _{i}\left\{y_{i}\left(\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle+b\right)\right\}=1
$$

$$
\forall i, y_{i}\left(\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle+b\right) \geq 1
$$

## Maximize the Margin

## Criterion

Maximize the minimal distance (points, hyperplane). Obtain the largest possible band
Margin

$$
\begin{gathered}
\left\langle\mathbf{w}, \mathbf{x}_{+}\right\rangle+b=1 \quad\left\langle\mathbf{w}, \mathbf{x}_{-}\right\rangle+b=-1 \\
\left\langle\mathbf{w}, \mathbf{x}_{+}-\mathbf{x}_{-}\right\rangle=2
\end{gathered}
$$

Margin $=$ projection of $\mathbf{x}_{+}-\mathbf{x}_{-}$on the normal vector of the hyperplane, $\frac{\mathbf{w}}{\|w\|_{2}}$

$$
\begin{aligned}
& \Rightarrow \text { Maximize } \frac{1}{\|\mathbf{W}\|} \\
& \Leftrightarrow \text { minimize }\|\mathbf{w}\|^{2}
\end{aligned}
$$

## Maximize the Margin (2)

Problem

$$
\begin{cases}\text { Minimize } & \frac{1}{2}\|\mathbf{w}\|^{2} \\ \text { with the constraints } & \forall i, y_{i}\left(\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle+b\right) \geq 1\end{cases}
$$

## Maximal Margin Hyperplane



## Quadratic Optimization (reminder)

Optimize $f$ with constraints $f_{i} \geq 0$
When $f$ and $f_{i}$ are convex Introduce the Lagrange multipliers $\alpha_{i}\left(\alpha_{i} \geq 0\right)$,
Consider
(penalization of the violated constraints)

$$
F(\mathbf{x}, \alpha)=f(\mathbf{x})-\sum_{i} \alpha_{i} f_{i}(\mathbf{x})
$$

Kuhn-Tucker principle (1951)

$$
F\left(\mathbf{x}_{0}, \alpha^{*}\right)=\min _{\alpha \geq 0} F\left(\mathbf{x}_{0}, \alpha\right)=\max _{x} F\left(\mathbf{x}, \alpha^{*}\right)
$$

## Primal Problem

$$
L(\mathbf{w}, b, \alpha)=\frac{1}{2}\|\mathbf{w}\|^{2}-\sum_{i} \alpha_{i}\left(y_{i}\left(\left\langle\mathbf{x}_{i}, \mathbf{w}\right\rangle+b\right)-1\right), \alpha_{i} \geq 0
$$

- Differentiate w.r.t. $b$ : at the optimum,

$$
\frac{\partial L}{\partial b}=0=\sum \alpha_{i} y_{i}
$$

- Differentiate w.r.t. w :

$$
\frac{\partial L}{\partial \mathbf{w}}=0=\mathbf{w}-\sum \alpha_{i} y_{i} \mathbf{x}_{i}
$$

- Replace in $L(\mathbf{w}, b, \alpha)$ :


## Dual problem (Wolfe)



$$
\begin{aligned}
& W(\alpha)=\sum_{i} \alpha_{i}-\frac{1}{2} \sum_{i, j} \alpha_{i} \alpha_{j} y_{i} y_{j}<\mathbf{x}_{i}, \mathbf{x}_{j}> \\
& \forall i, \alpha_{i} \geq 0 \\
& \sum_{i} \alpha_{i} y_{i}=0
\end{aligned}
$$

Quadratic form w.r.t. $\alpha$
quadratic optimization is easy
Solution: $\alpha_{i}^{*}$

- Compute $\mathbf{w}^{*}$ :

$$
\mathbf{w}^{*}=\sum_{i} \alpha_{i}^{*} y_{i} \mathbf{x}_{i}
$$

- If $\left(\left\langle\mathbf{x}_{i}, \mathbf{w}^{*}\right\rangle+b\right) y_{i}>1, \alpha_{i}^{*}=0$.
- IF $\left(\left\langle\mathbf{x}_{i}, \mathbf{w}^{*}\right\rangle+b\right) y_{i}=1, \alpha_{i}^{*}>0, \quad \mathbf{x}_{i} \quad$ support vector
- Compute $b^{*}$ :

$$
b^{*}=-\frac{1}{2}\left(\left\langle\mathbf{w}^{*}, \overline{\mathbf{x}}^{+}\right\rangle+\left\langle\mathbf{w}^{*}, \overline{\mathbf{x}}^{-}\right\rangle\right)
$$



## Summary

$$
\left.\mathcal{E}=\left\{\left(\mathbf{x}_{i}, y_{i}\right)\right\}, \mathbf{x}_{i} \in \mathbb{R}^{d}, y_{i} \in\{-1,1\}, i=1 . . n\right\} \quad\left(\mathbf{x}_{i}, y_{i}\right) \sim P(\mathbf{x}, y)
$$



$$
h(\mathbf{x})=\langle\mathbf{w}, \mathbf{x}\rangle+b
$$

Two goals
Role

- Data fitting $\operatorname{sign}\left(y_{i}\right)=\operatorname{sign}\left(h\left(\mathbf{x}_{i}\right)\right) \rightarrow$ maximize margin $y_{i} . h\left(\mathbf{x}_{i}\right)$
achieve learning
- Regularization : minimize $\|\mathbf{w}\|$


## Support Vector Machines

## General scheme

- Minimize the regularization term
- ... subject to data constraints

$$
=\operatorname{margin} \geq 1(*)
$$

$$
\begin{cases}\text { Min. } & \frac{1}{2}\|\mathbf{w}\|^{2} \\ \text { s.t. } & y_{i}\left(\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle+b\right) \geq 1 \quad \forall i=1 \ldots n\end{cases}
$$

Constrained minimization of a convex function
$\rightarrow$ introduce Lagrange multipliers $\alpha_{i} \geq 0, i=1 \ldots n$

$$
\operatorname{Min} \mathcal{L}(\mathbf{w}, b, \alpha)=\frac{1}{2}\|\mathbf{w}\|^{2}+\sum_{i} \alpha_{i}\left(1-y_{i}\left(\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle+b\right)\right)
$$

## Primal problem

- $d+1$ variables ( $+n$ Lagrange multipliers)
$\left(^{*}\right)$ in the separable case; see later for non-separable case


## Support Vector Machines, 2

At the optimum

$$
\frac{\partial \mathcal{L}}{\partial \mathbf{w}}=\frac{\partial \mathcal{L}}{\partial b}=\frac{\partial \mathcal{L}}{\partial \alpha}=0
$$

Dual problem
Wolfe

$$
\begin{cases}\text { Max. } & \mathcal{Q}(\alpha)=\sum_{i} \alpha_{i}-\frac{1}{2} \sum_{i, j} \alpha_{i} \alpha_{j} y_{i} y_{j}\left\langle\mathbf{x}_{i}, \mathbf{x}_{j}\right\rangle \\ \text { s.t. } & \forall i, \alpha_{i} \geq 0 \\ & \sum_{i} \alpha_{i} y_{i}=0\end{cases}
$$

Support vectors

$$
\text { Examples }\left(\mathbf{x}_{i}, y_{i}\right) \text { s.t. } \alpha_{i}>0
$$

the only ones involved in the decision function

$$
\mathbf{w}=\sum \alpha_{i} y_{i} \mathbf{x}_{i}
$$

## Support vectors, examples



## Support vectors, examples

## MNIST data

## 7234567890 12 3 4 5 67890 1234567890

 Data13


Support vectors

## Remarks

- Support vectors are critical examples
near-miss
- Show that the Leave-One-Out error is less than \# sv.

LOO: iteratively, learn on all examples but one, and test on the remaining one

## Overview

## Linear SVM, separable case

Linear SVM, non separable case
The kernel trick
The Kernel principle
Examples
Discussion
Extensions
Multi-class discrimination
Regression
Novelty detection
On the practitioner side
Improve precision
Reduce computational cost
Theory

## Separable vs non-separable data

Training
Test




## Linear hypotheses, non separable data

Cortes \& Vapnik 95
Non-separable data $\Rightarrow$ not all constraints are satisfiable

$$
y_{i}\left(\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle+b\right) \geq 1-\xi_{\mathbf{i}}
$$

Formalization

- Introduce slack variables $\xi_{i}$
- And penalize them

$$
\left\{\begin{array}{lc}
\text { Minimize } & \frac{1}{2}\|\mathbf{w}\|^{2}+\mathbf{C} \sum_{\mathrm{i}} \xi_{\mathfrak{i}} \\
\text { Subject to } & \forall i, y_{i}\left(\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle+b\right) \geq 1-\xi_{\mathrm{i}} \\
& \xi_{i} \geq 0
\end{array}\right.
$$

Critical decision: adjust $C=$ error cost.

## Primal problem, non separable case

Same resolution: Lagrange Multipliers $\alpha_{i}$ and $\beta_{i}$, with $\alpha_{i} \geq 0, \beta_{i} \geq 0$

$$
\begin{aligned}
\mathcal{L}(\mathbf{w}, b, \xi, \alpha, \beta)= & \operatorname{Min} \frac{1}{2}\|\mathbf{w}\|^{2}+C \sum_{i} \xi_{i} \\
& -\sum_{i} \alpha_{i}\left(y_{i}\left(\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle+b\right)-1+\xi_{i}\right) \\
& -\sum_{i} \beta_{i} \xi_{i}
\end{aligned}
$$

At the optimum

$$
\begin{gathered}
\frac{\partial \mathcal{L}}{\partial \mathbf{w}}=\frac{\partial \mathcal{L}}{\partial b}=\frac{\partial \mathcal{L}}{\partial \xi_{i}}=0 \\
\mathbf{w}=\sum_{i} \alpha_{i} y_{i} \mathbf{x}_{i} \quad \sum_{i} \alpha_{i} y_{i}=0 \quad C-\alpha_{i}-\beta_{i}=0
\end{gathered}
$$

Likewise

- Convex (quadratic) optimization problem $\rightarrow$ it is equivalent to solve the primal and the dual problem (expressed with multipliers $\alpha, \beta$ )


## Dual problem, non separable case

$$
\operatorname{Min} \sum_{i} \alpha_{i}-\frac{1}{2} \sum_{i, j} \alpha_{i} \alpha_{j} y_{i} y_{j}\left\langle\mathbf{x}_{i}, \mathbf{x}_{j}\right\rangle, 0 \leq \alpha_{i} \leq C
$$

Mathematically nice problem

- $H=$ semi-definite positive $n \times n$ matrix

$$
H_{i, j}=y_{i} y_{j}\left\langle\mathbf{x}_{i}, \mathbf{x}_{j}\right\rangle
$$

- Dual problem quadratic form

Minimize $\langle\alpha, e\rangle-\alpha^{T} H \alpha$
with $e=(1, \ldots, 1) \in \mathbb{R}^{n}$.

## Support vectors



- Only support vectors $\left(\alpha_{i}>0\right)$ are involved in $h$

$$
\mathbf{w}=\sum \alpha_{i} y_{i} \mathbf{x}_{i}
$$

- Importance of support vector $\mathbf{x}_{i}$ : weight $\alpha_{i}$
- Difference with the separable case $0<\alpha_{i}<C$ bounded influence of examples


## The loss (error cost) function

Roles

- The goal is data fitting
loss function characterizes the learning goal
- while solving a convex optimization problem and makes it tractable/reproducible
The error cost
- Binary cost: $\ell(y, h(\mathbf{x}))=1$ iff $y \neq h(x)$
- Quadratic cost: $\ell(y, h(\mathbf{x}))=(y-h(x))^{2}$
- Hinge loss

$$
\ell(y, h(\mathbf{x}))=\max (0,1-y \cdot h(x))=(1-y \cdot h(x))_{+}=\xi
$$



## Complexity

## Learning complexity

- Worst case: $\mathcal{O}\left(n^{3}\right)$
- Empirical complexity: depends on $C$
- $\mathcal{O}\left(n^{2} n_{s v}\right)$ where $n_{s v}$ is the number of s.v.

Usage complexity

- $\mathcal{O}\left(n_{s v}\right)$


## Overview

Linear SVM, separable case
Linear SVM, non separable case
The kernel trick
The Kernel principle
Examples
Discussion
Extensions
Multi-class discrimination
Regression
Novelty detection
On the practitioner side
Improve precision
Reduce computational cost
Theory

## Non-separable data



Representation change

$$
\mathbf{x} \in \mathbb{R}^{2} \rightarrow \text { polar coordinates } \in \mathbb{R}^{2}
$$

## Principle



$$
\Phi: X \mapsto \Phi(X) \subset \mathbb{R}^{D}
$$

## Intuition

- In a high-dimensional space, every dataset is linearly separable $\rightarrow$ Map data onto $\Phi(X)$, and we are back to linear separation


## Glossary

- $X$ : input space
- $\Phi(X)$ : feature space


## The kernel trick

## Remark

- Generalization bounds do not depend on the dimension of input space $X$ but on the capacity of the hypothesis space $\mathcal{H}$.
- SVMs only involve scalar products $\left\langle\mathbf{x}_{i}, \mathbf{x}_{j}\right\rangle$.


## Intuition

- Representation change is only "virtual"
$\Phi: X \mapsto \Phi(X)$
- Consider scalar product in $\Phi(X)$
- ... and compute it in $X$

$$
K\left(\mathbf{x}_{i}, \mathbf{x}_{j}\right)=\left\langle\Phi\left(\mathbf{x}_{i}\right), \Phi\left(\mathbf{x}_{j}\right)\right\rangle
$$

## Example: polynomial kernel

## Principle

$$
\begin{gathered}
\mathbf{x} \in \mathbb{R}^{3} \mapsto \Phi(\mathbf{x}) \in \mathbb{R}^{10} \\
\mathbf{x}=\left(x_{1}, x_{2}, x_{3}\right) \\
\Phi(\mathbf{x})=\left(1, \sqrt{2} x_{1}, \sqrt{2} x_{2}, \sqrt{2} x_{3}, \sqrt{2} x_{1} x_{2}, \sqrt{2} x_{1} x_{3}, \sqrt{2} x_{2} x_{3}, x_{1}^{2}, x_{2}^{2}, x_{3}^{2}\right)
\end{gathered}
$$

Why $\sqrt{2}$ ?

## Example: polynomial kernel

## Principle

$$
\mathbf{x} \in \mathbb{R}^{3} \mapsto \Phi(\mathbf{x}) \in \mathbb{R}^{10}
$$

$$
\begin{aligned}
& \mathbf{x}=\left(x_{1}, x_{2}, x_{3}\right) \\
& \Phi(\mathbf{x})=\left(1, \sqrt{2} x_{1}, \sqrt{2} x_{2}, \sqrt{2} x_{3}, \sqrt{2} x_{1} x_{2}, \sqrt{2} x_{1} x_{3}, \sqrt{2} x_{2} x_{3}, x_{1}^{2}, x_{2}^{2}, x_{3}^{2}\right)
\end{aligned}
$$

Why $\sqrt{2}$ ?
because

$$
\left\langle\Phi(\mathbf{x}), \Phi\left(\mathbf{x}^{\prime}\right)\right\rangle=\left(1+\left\langle\mathbf{x}, \mathbf{x}^{\prime}\right\rangle\right)^{2}=K\left(\mathbf{x}, \mathbf{x}^{\prime}\right)
$$

## Primal and dual problems unchanged

Primal problem

$$
\begin{cases}\text { Min. } & \frac{1}{2}\|\mathbf{w}\|^{2} \\ \text { s.t. } & y_{i}\left(\left\langle\mathbf{w}, \Phi\left(\mathbf{x}_{i}\right)\right\rangle+b\right) \geq 1 \quad \forall i=1 \ldots n\end{cases}
$$

Dual problem

$$
\begin{cases}\text { Max. } & \mathcal{Q}(\alpha)=\sum_{i} \alpha_{i}-\frac{1}{2} \sum_{i, j} \alpha_{i} \alpha_{j} y_{i} y_{j} K\left(\mathbf{x}_{i}, \mathbf{x}_{j}\right) \\ \text { s.t. } & \forall i, \alpha_{i} \geq 0 \\ & \sum_{i} \alpha_{i} y_{i}=0\end{cases}
$$

Hypothesis

$$
h(\mathbf{x})=\sum_{i} \alpha_{i} y_{i} K\left(\mathbf{x}_{i}, \mathbf{x}\right)
$$

## Example, polynomial kernel

$$
K\left(\mathbf{x}, \mathbf{x}^{\prime}\right)=\left(a\left\langle\mathbf{x}, \mathbf{x}^{\prime}\right\rangle+1\right)^{b}
$$

- Choice of $a, b$ : cross validation
- Domination of high/low degree terms ?
- Importance of normalization



## Example, Radius-Based Function kernel (RBF)

$$
K\left(\mathbf{x}, \mathbf{x}^{\prime}\right)=\exp \left(-\gamma\left\|\mathbf{x}-\mathbf{x}^{\prime}\right\|^{2}\right)
$$

- No closed form $\Phi$
- $\Phi(X)$ of infinite dimension For $x$ in $\mathbb{R}$

$$
\left.\Phi(x)=\exp \left(-\gamma x^{2}\right)\right)\left[1, \sqrt{\frac{2 \gamma}{1!}} x, \sqrt{\frac{(2 \gamma)^{2}}{2!}} x^{2}, \sqrt{\frac{(2 \gamma)^{3}}{3!}} x^{3}, \ldots\right]
$$

- Choice of $\gamma$ ? (intuition: think of $H, H_{i, j}=y_{i} y_{j} K\left(\mathbf{x}_{i}, \mathbf{x}_{j}\right)$ )


## String kernels

## Watkins 99, Lodhi 02

## Notations

- $s$ a string on alphabet $\Sigma$
- $\mathbf{i}=\left(i_{1}, i_{2}, \ldots, i_{n}\right)$ an ordered index sequence $\left(i_{j}<i_{j+1}\right)$, avec $\ell(\mathbf{i})=i_{n}-i_{1}+1$
- $s[\mathbf{i}]$ substring of $s$, extraction pattern is $\mathbf{i}$ $s=B I C Y C L E, \mathbf{i}=(1,3,6), s[\mathbf{i}]=B C L$

Definition

$$
K_{n}\left(s, s^{\prime}\right)=\sum_{u \in \Sigma^{n}} \sum_{\mathbf{i s . t . s}[\mathbf{i}]=u} \sum_{\mathbf{j} s . t . s^{\prime}[\mathbf{j}]=u} \varepsilon^{\ell(\mathbf{i})+\ell(\mathbf{j})}
$$

with $0<\varepsilon<1$ (discount)

## String kernels, followed

$\Phi$ : projection on $\mathbb{R}^{D}$ où $D=|\Sigma|^{n}$

$$
\begin{array}{|c|cccc|}
\hline & \text { CH } & \text { CA } & \text { CT } & \text { AT } \\
\text { CHAT } & \varepsilon^{2} & \varepsilon^{3} & \varepsilon^{4} & \varepsilon^{2} \\
\text { CARTOON } & 0 & \varepsilon^{2} & \varepsilon^{4} & \varepsilon^{3} \\
\hline
\end{array}
$$

Prefer the normalized version

$$
\left.\kappa_{( } s, s^{\prime}\right)=\frac{K\left(s, s^{\prime}\right)}{\sqrt{K(s, s) K\left(s^{\prime} s^{\prime}\right)}}
$$

## String kernels, followed

Application 1
Document mining

- Pre-processing matters a lot (stop-words, stemming)
- Multi-lingual aspects
- Document classification
- Information retrieval

Application 2, Bio-informatics

- Pre-processing matters a lot
- Classification (secondary structures)


Extension to graph kernels http://videolectures.net/gbr07_vert_ckac/

## Application to musical analysis

- Input: Midi files
- Pre-processing, rythm detection
- Representation: the musical worm (tempo, loudness)
- Output: Identification of performer styles


Using String Kernels to Identify Famous Performers from their Playing Style, Saunders et al., 2004

## Kernels: key features

Absolute $\rightarrow$ Relative representation

- $\left\langle\mathbf{x}, \mathbf{x}^{\prime}\right\rangle \propto$ angle of $\mathbf{x}$ and $\mathbf{x}^{\prime}$
- More generally $K\left(\mathbf{x}, \mathbf{x}^{\prime}\right)$ measures the (non-linear) similarity of $x$ and $x^{\prime}$
- $\mathbf{x}$ is described by its similarity to other examples

Necessary condition: the Mercer condition
$K$ must be positive semi-definite

$$
\forall g \in L_{2}, \int K\left(\mathbf{x}, \mathbf{x}^{\prime}\right) g(\mathbf{x}) g\left(\mathbf{x}^{\prime}\right) d \mathbf{x} \geq 0
$$

## Why ?

Related to $\Phi$ Mercer condition holds $\rightarrow \exists \phi_{1}, \phi_{2}, .$.

$$
k\left(\mathbf{x}, \mathbf{x}^{\prime}\right)=\sum_{i=1}^{\infty} \lambda_{i} \phi_{i}(\mathbf{x}) \phi_{i}\left(\mathbf{x}^{\prime}\right)
$$

with $\phi_{i}$ eigen functions, $\lambda_{i}>0$ eigen values
Kernel properties: let $K, K^{\prime}$ be p.d. kernels and $\alpha>0$, then

- $\alpha K$ is a p.d. kernel
- $K+K^{\prime}$ is a p.d. kernel
- $K . K^{\prime}$ is a p.d. kernel
- $K\left(\mathbf{x}, \mathbf{x}^{\prime}\right)=\operatorname{limit}_{p \rightarrow \infty} K_{p}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)$ is p.d. if it exists
- $K(A, B)=\sum_{\mathbf{x} \in A, \mathbf{x}^{\prime} \in B} K\left(x, x^{\prime}\right)$ is a p.d. kernel


## Overview

Linear SVM, separable case
Linear SV/M, non separable case
The kernel trick
The Kernel principle
Examples
Discussion
Extensions
Multi-class discrimination
Regression
Novelty detection
On the practitioner side
Improve precision
Reduce computational cost
Theory

## Multi-class discrimination

Input
Binary case
$\left.\mathcal{E}=\left\{\left(\mathbf{x}_{i}, y_{i}\right)\right\}, \mathbf{x}_{i} \in \mathbb{R}^{d}, y_{i} \in\{-1,1\}, i=1 . . n\right\} \quad\left(\mathbf{x}_{i}, y_{i}\right) \sim P(\mathbf{x}, y)$
Multi-class case
$\left.\mathcal{E}=\left\{\left(\mathbf{x}_{i}, y_{i}\right)\right\}, \mathbf{x}_{i} \in \mathbb{R}^{d}, y_{i} \in\{1 \ldots k\}, i=1 . . n\right\} \quad\left(\mathbf{x}_{i}, y_{i}\right) \sim P(\mathbf{x}, y)$
Output: $\hat{h}: \mathbb{R}^{d} \mapsto\{1 \ldots k\}$.

## Multi-class learning: one against all

First option: $k$ binary learning problems
Pb 1: class $1 \rightarrow+1$, classes $2 \ldots k \rightarrow-1$
Pb 2: class $2 \rightarrow+1$, classes $1,3, \ldots k \rightarrow-1$

Prediction

$$
h(\mathbf{x})=i \text { iff } h_{i}(\mathbf{x})=\operatorname{argmax}\left\{h_{j}(\mathbf{x}), j=1 \ldots k\right\}
$$

Justification
If $\mathbf{x}$ belongs to class 1 , one should have

$$
h_{1}(\mathbf{x}) \geq 1, h_{j}(\mathbf{x})<-1, j \neq 1
$$

## Where is the difficulty ?



What we get (one vs all)


What we want

## Multi-class learning: one vs one

Second option: $\frac{k(k-1)}{2}$ binary classification problems $\mathrm{Pb} i, j$ class $i \rightarrow+1$, class $j \rightarrow-1$

## Prediction

- Compute all $h_{i, j}(\mathbf{x})$
- Count the votes

| Classes | winner |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 12 | 1 |  |  |  |  |
| 13 | 1 |  |  |  |  |
| 14 | 1 |  |  |  |  |
| 23 | 2 |  |  |  |  |
| 24 | 4 | class | 12 | 3 | 4 |
| 34 | 3 | \# votes | 31 | 1 |  |

NB: One can also use the $h_{i, j}(\mathbf{x})$ values.

## Multi-class learning: additionnal constraints

Another option
Vapnik 98; Weston, Watkins 99

$$
\begin{cases}\text { Minimise } & \frac{1}{2} \sum_{j=1}^{k}\left\|\mathbf{w}_{j}\right\|^{2}+C \sum_{i=1}^{n} \sum_{\ell=1, \ell \neq y_{i}}^{k} \xi_{i, \ell} \\ \text { Subject to } & \forall i, \forall \ell \neq y_{i}, \\ & \left(\left\langle w_{y_{i}}, \mathbf{x}_{i}\right\rangle+b_{y_{i}}\right) \geq\left(\left\langle w_{\ell}, \mathbf{x}_{i}\right\rangle+b_{\ell}\right)+2-\xi_{i, \ell} \\ & \xi_{i, \ell} \geq 0\end{cases}
$$

Hum!

- $n \times k$ constraints: $n \times k$ dual variables


## Recommendations

In practice

- Results are in general (but not always !) similar
- 1 -vs- 1 is the fastest option


## Overview

Linear SVM, separable case
Linear SV/M, non separable case
The kernel trick
The Kernel principle
Examples
Discussion
Extensions
Multi-class discrimination
Regression
Novelty detection
On the practitioner side
Improve precision
Reduce computational cost
Theory

## Regression

## Input

$$
\left.\mathcal{E}=\left\{\left(x_{i}, y_{i}\right)\right\}, x_{i} \in \mathbb{R}^{d}, y_{i} \in \mathbb{R}, i=1 . . n\right\} \quad\left(x_{i}, y_{i}\right) \sim P(x, y)
$$

Output : $\hat{h}: \mathbb{R}^{d} \mapsto \mathbb{R}$.


## Regression with Support Vector Machines

## Intuition

- Find $h$ deviating by at most $\varepsilon$ from the data
- ... while being as flat as possible
loss function
regularization

Formulation

$$
\begin{cases}\text { Min. } & \frac{1}{2}\|\mathbf{w}\|^{2} \\ \text { s.t. } & \forall i=1 \ldots n \\ & \left(\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle+b\right) \geq y_{i}-\varepsilon \\ & \left(\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle+b\right) \leq y_{i}+\varepsilon\end{cases}
$$

Regression with Support Vector Machines, followed

Using slack variables

$$
\begin{cases}\text { Min. } & \frac{1}{2}\|\mathbf{w}\|^{2}+\mathbf{C} \sum_{i}\left(\xi_{i}^{+}+\xi_{i}^{-}\right) \\ \text {s.t. } & \forall i=1 \ldots n \\ & \left(\left\langle\mathbf{w}, \mathbf{x}_{\boldsymbol{i}}\right\rangle+b\right) \geq y_{i}-\varepsilon-\xi_{i^{-}}^{-} \\ & \left(\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle+b\right) \leq y_{i}+\varepsilon+\xi_{i}^{+}\end{cases}
$$



Regression with Support Vector Machines, followed
Primal problem

$$
\begin{aligned}
\mathcal{L}(\mathbf{w}, b, \xi, \alpha, \beta)= & \operatorname{Min} \frac{1}{2}\|\mathbf{w}\|^{2}+C \sum_{i}\left(\xi_{i}^{+}+\xi_{i}^{-}\right) \\
& -\sum_{i} \alpha_{i}^{+}\left(y_{i}+\varepsilon+\xi_{i}^{+}-\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle+b\right) \\
& -\sum_{i} \alpha_{i}^{-}\left(\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle+b-y_{i}+\varepsilon+\xi_{i}^{-}\right) \\
& -\sum_{i} \beta_{i}^{+} \xi_{i}^{+}-\sum_{i} \beta_{i}^{-} \xi_{i}^{-}
\end{aligned}
$$

Dual problem

$$
\begin{cases}\mathcal{Q}\left(\alpha^{+}, \alpha^{-}\right)= & \sum_{i} y_{i}\left(\alpha_{i}^{+}-\alpha_{i}^{-}\right)-\varepsilon \sum_{i}\left(\alpha_{i}^{+}+\alpha_{i}^{-}\right) \\ & +\sum_{i, j}\left(\alpha_{i}^{+}-\alpha_{i}^{-}\right)\left(\alpha_{j}^{+}-\alpha_{j}^{-}\right)\left\langle\mathbf{x}_{i}, \mathbf{x}_{j}\right\rangle \\ \text { s.t. } & \forall i=1 \ldots n \\ & \sum^{2}\left(\alpha_{i}^{+}-\alpha_{i}^{-}\right)=0 \\ & 0 \leq \alpha_{i}^{+} \leq C \\ & 0 \leq \alpha_{i}^{-} \leq C\end{cases}
$$

Regression with Support Vector Machines, followed Hypothesis

$$
h(\mathbf{x})=\sum\left(\alpha_{i}^{+}-\alpha_{i}^{-}\right)\left\langle\mathbf{x}_{i}, \mathbf{x}\right\rangle+b
$$

With no loss of generality you can replace everywhere

$$
\left\langle\mathbf{x}, \mathbf{x}^{\prime}\right\rangle \rightarrow K\left(\mathbf{x}, \mathbf{x}^{\prime}\right)
$$



## Beware

High-dimensional regression

$$
\left.\mathcal{E}=\left\{\left(\mathbf{x}_{i}, y_{i}\right)\right\}, \mathbf{x}_{i} \in \mathbb{R}^{D}, y_{i} \in \mathbb{R}, i=1 . . n\right\} \quad\left(\mathbf{x}_{i}, y_{i}\right) \sim P(\mathbf{x}, y)
$$

A very slippery game if $D \gg n$
curse of dimensionality
Dimensionality reduction mandatory

- Map x onto $\mathbb{R}^{d}$
- Central subspace:

$$
\pi: X \mapsto S \subset \mathbb{R}^{d}
$$

with $S$ minimal such that $y$ and $\mathbf{x}$ are independent conditionally to $\pi(x)$.

Find $h, \mathbf{w}: y=h(\mathbf{w}, \mathbf{x})$

## Sliced Inverse Regression

Bernard-Michel et al, 09


More:
http://mistis.inrialpes.fr/learninria/ S. Girard

## Overview

Linear SVM, separable case
Linear SV/M, non separable case
The kernel trick
The Kernel principle
Examples
Discussion
Extensions
Multi-class discrimination
Regression
Novelty detection
On the practitioner side
Improve precision
Reduce computational cost
Theory

## Novelty Detection

## Input

$$
\left.\mathcal{E}=\left\{\left(x_{i}\right)\right\}, x_{i} \in X, i=1 . . n\right\} \quad\left(x_{i}\right) \sim P(x)
$$

Context

- Information retrieval

- Identification of the data support
estimation of distribution

Critical issue

- Classification approaches not efficient: too much noise


## One-class SVM

Formulation

$$
\begin{cases}\text { Min. } & \frac{1}{2}\|\mathbf{w}\|^{2}-\rho+\mathbf{C} \sum_{\mathrm{i}} \xi_{\mathrm{i}} \\ \text { s.t. } & \forall i=1 \ldots n \\ & \left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle \geq \rho-\xi_{\mathrm{i}}\end{cases}
$$



Dual problem

$$
\begin{cases}\text { Min. } & \sum_{i, j} \alpha_{i} \alpha_{j}\left\langle\mathbf{x}_{i}, \mathbf{x}_{j}\right\rangle \\ \text { s.t. } & \forall i=1 \ldots n \quad 0 \leq \alpha_{i} \leq C \\ & \sum_{i} \alpha_{i}=0\end{cases}
$$

## Implicit surface modelling

Schoelkopf et al, 04
Goal: find the surface formed by the data points

$$
\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle \geq \rho \text { becomes }-\varepsilon \leq\left(\left\langle\mathbf{w}, \mathbf{x}_{i}\right\rangle-\rho\right) \leq \varepsilon
$$



## Overview

Linear SVM, separable case
Linear SV/M, non separable case
The kernel trick
The Kernel principle
Examples
Discussion
Extensions
Multi-class discrimination
Regression
Novelty detection
On the practitioner side Improve precision
Reduce computational cost
Theory

## Normalisation / Scaling

Needed to prevent attributes to steal the game

|  | Height | Gender | Class |
| :---: | :---: | :---: | :---: |
| $\mathrm{x}_{1}$ | 150 | F | 1 |
| $\mathrm{x}_{2}$ | 180 | M | 0 |
| $\mathrm{x}_{3}$ | 185 | M | 0 |
| $\stackrel{\Delta}{\mathbf{x}_{1}}$ |  |  |  |
|  |  | $\stackrel{0}{x}^{0}{ }^{\circ}{ }_{3}$ |  |

$\Rightarrow$ Normalization

$$
\text { Height } \rightarrow \frac{\text { Height }-150}{180-150}
$$

## Beware

## Usual practice

- Normalize the whole dataset
- Learn on the training set
- Test on the test set


## Beware

## Usual practice

- Normalize the whole dataset
- Learn on the training set
- Test on the test set


## NO!

## Good practice

- Normalize the training set (Scale ${ }_{\text {train }}$ )
- Learn from the normalized training set
- Scale the test set according to Scale ${ }_{\text {train }}$ and test


## Imbalanced datasets

## Typically

- Normal transactions: 99.99\%
- Fraudulous transactions: not many


## Practice

- Define asymmetrical penalizations

$$
\begin{array}{ll}
\text { std penalization } & C \sum_{i} \xi_{i} \\
\text { asymmetrical penalizations } & C_{+} \sum_{i, y_{i}=1} \xi_{i}+C_{-} \sum_{i, y_{i}=-1} \xi_{i}
\end{array}
$$

Other options ?

## Overview

Linear SVM, separable case
Linear SV/M, non separable case
The kernel trick
The Kernel principle
Examples
Discussion
Extensions
Multi-class discrimination
Regression
Novelty detection
On the practitioner side Improve precision
Reduce computational cost
Theory

## Data sampling

Simple approaches

- Uniform sampling
often efficient
- Stratified sampling same distribution as in $\mathcal{E}$

Incremental approaches

- Partition $\mathcal{E} \rightarrow \mathcal{E}_{1}, \ldots \mathcal{E}_{N}$
- Learn from $\mathcal{E}_{1} \rightarrow$ support vectors $S V_{1}$
- Learn from $\mathcal{E}_{2} \cup S V_{1} \rightarrow$ support vectors $S V_{2}$
- etc.


## Data sampling, followed

Select examples

- Use $k$-nearest neighbors
- Train SVM on k-means (prototypes)
- Pb about distances

Hierarchical methods

- Use unsupervised learning and form clusters learning, J. Gama
- Learn a hypothesis on each cluster
- Aggregate hypotheses


## Reduce number of variables

Select candidate s.v. $\mathcal{F} \subset \mathcal{E}$

$$
w=\sum \alpha_{i} y_{i} \mathbf{x}_{i} \text { with }\left(\mathbf{x}_{i}, y_{i}\right) \in \mathcal{F}
$$

Optimize $\alpha_{i}$ on $\mathcal{E}$

$$
\begin{cases}\text { Min. } & \frac{1}{2} \sum_{i, j, \in \mathcal{F}} \alpha_{i} \alpha_{j} y_{i} y_{j}\left\langle x_{i}, x_{j}\right\rangle+C \sum_{\ell=1}^{n} \xi_{\ell} \\ \text { t.q. } & \forall \ell=1 \ldots n, \\ & \left(\left\langle w, \mathbf{x}_{\ell}\right\rangle+b\right) \geq 1-\xi_{\ell} \\ & \xi_{\ell} \geq 0\end{cases}
$$

## Sources

- Vapnik, The nature of statistical learning, Springer Verlag 1995; Statistical Learning Theory, Wiley 1998
- Cristianini \& Shawe Taylor, An introduction to Support Vector Machines, Cambridge University Press, 2000.
- http://www.kernel-machines.org/tutorials
- Videolectures + ML Summer Schools
- Large scale Machine Learning challenge, ICML 2008 wshop: http://largescale.ml.tu-berlin.de/workshop/


## Overview

Linear SVM, separable case
Linear SVM, non separable case
The kernel trick
The Kernel principle
Examples
Discussion
Extensions
Multi-class discrimination
Regression
Novelty detection
On the practitioner side
Improve precision
Reduce computational cost
Theory

## Reminder



Vapnik, 1995, 1998

## Input

$\left.\mathcal{E}=\left\{\left(x_{i}, y_{i}\right)\right\}, x_{i} \in \mathbb{R}^{m}, y_{i} \in\{-1,1\}, i=1 . . n\right\} \quad\left(x_{i}, y_{i}\right) \sim P(x, y)$
Output : $\hat{h}: \mathbb{R}^{m} \mapsto\{-1,1\}$ ou $\mathbb{R}$. $\quad \hat{h}$ approximates $y$
Criterion : ideally, minimize the generalization error

$$
\operatorname{Err}(h)=\int \ell(y, \hat{h}(x)) d P(x, y)
$$

$\ell=$ loss function: $1_{y \neq \hat{h}(x)},(y-\hat{h}(x))^{2}$
$P(x, y)=$ joint distribution of the data.

## The Bias-Variance Tradeoff

Choice of a model: The space $\mathcal{H}$ where we are looking for $\hat{h}$.
Bias: Distance between $y$ and $h^{*}=\operatorname{argmin}\{\operatorname{Err}(h), h \in \mathcal{H}\}$.
the best we can hope for
Variance: Distance between $\hat{h}$ and $h^{*}$
between the best $h^{*}$ and the $\hat{h}$ we actually learn
Note:
Only the empirical risk (on the available data) is given

$$
E r r_{e m p, n}(\hat{h})=\frac{1}{n} \sum_{i=1}^{n} \ell\left(y_{i}, \hat{h}\left(x_{i}\right)\right)
$$

Principle:

$$
\operatorname{Err}(\hat{h})<\operatorname{Err}_{e m p, n}(\hat{h})+\mathcal{B}(n, \mathcal{H})
$$

If $\mathcal{H}$ is "reasonable", Err ${ }_{\text {emp }, n} \rightarrow$ Err when $n \rightarrow \infty$

## Statistical Learning

## Statistical Learning Theory

Learning from a statistical perspective.
Goal of the theory
in general
Model a real / artificial phenomenon, in order to:

* understand
* predict
* exploit


## General

A theory: hypotheses $\rightarrow$ predictions

- Hypotheses on the phenomenon
here, Learning
- Predictions about its behavior errors

Theory $\rightarrow$ algorithm

- Optimize the quantities allowing prediction
- Nothing practical like a good theory!

Vapnik

## General

A theory: hypotheses $\rightarrow$ predictions

- Hypotheses on the phenomenon
here, Learning
- Predictions about its behavior

Theory $\rightarrow$ algorithm

- Optimize the quantities allowing prediction
- Nothing practical like a good theory!

Vapnik

## Strength/Weaknesses

+ Stronger Hypotheses $\rightarrow$ more precise predictions
BUT if the hypotheses are wrong, nothing will work


## What Theory do we need?

Approach in expectation

- A set of data
one example
breast cancer
- $\bar{x}^{+}$: average of positive examples
- $\bar{x}^{-}$: average of negative examples
- $h(x)=+1$ iff $d\left(x, \bar{x}^{+}\right)<d\left(x, \bar{x}^{-}\right)$

Estimate the generalization error

- Data $\rightarrow$ Training set, test set
- Learn $\bar{x}^{+}$et $\bar{x}^{-}$on the training set, measure the errors on the test set


## Classical Statistics vs Statistical Learning

Classical Statistics

- Mean error

We want guarantees

- PAC Model Probably Approximately Correct
- What is the probability that the error is greater than a given threshold?


## Example

## Assume

$$
\operatorname{Err}(h)>\varepsilon
$$

What is the probability that $\operatorname{Err}_{\text {emp, } n}(h)=0$ ?

$$
\begin{aligned}
\operatorname{Pr}\left(E r r_{e m p, n}(h)=0, \operatorname{Err}(h)>\varepsilon\right) & =(1-\operatorname{Err}(h))^{n} \\
& <(1-\varepsilon)^{n} \\
& <\exp (-\varepsilon n)
\end{aligned}
$$

## Example

## Assume

$$
\operatorname{Err}(h)>\varepsilon
$$

What is the probability that $E r r_{e m p, n}(h)=0$ ?

$$
\begin{aligned}
\operatorname{Pr}\left(E r r_{e m p, n}(h)=0, \operatorname{Err}(h)>\varepsilon\right) & =(1-\operatorname{Err}(h))^{n} \\
& <(1-\varepsilon)^{n} \\
& <\exp (-\varepsilon n)
\end{aligned}
$$

Hence, in order to guarantee a risk $\delta$

$$
\operatorname{Pr}\left(E r r_{e m p, n}(h)=0, \operatorname{Err}(h)>\varepsilon\right)<\delta
$$

## Example

## Assume

$$
\operatorname{Err}(h)>\varepsilon
$$

What is the probability that $E r r_{e m p, n}(h)=0$ ?

$$
\begin{aligned}
\operatorname{Pr}\left(E r r_{e m p, n}(h)=0, \operatorname{Err}(h)>\varepsilon\right) & =(1-\operatorname{Err}(h))^{n} \\
& <(1-\varepsilon)^{n} \\
& <\exp (-\varepsilon n)
\end{aligned}
$$

Hence, in order to guarantee a risk $\delta$

$$
\operatorname{Pr}\left(\operatorname{Err}_{\text {emp }, n}(h)=0, \operatorname{Err}(h)>\varepsilon\right)<\delta
$$

The error should not be greater than

$$
\varepsilon<\frac{1}{n} \ln \frac{1}{\delta}
$$

## Statistical Learning

## Principle

- Find a bound on the generalization error
- Minimize the bound.


## Note

$\hat{h}$ should be considered as a random variable, depending on the training set $\mathcal{E}$ and the number of examples $n$.

## Results

- deviation of the empirical error

$$
\operatorname{Err}\left(\widehat{h_{n}}\right) \leq \operatorname{Err}_{e m p, n}\left(\widehat{h_{n}}\right)+\mathcal{B}_{1}(n, \mathcal{H})
$$

- bias-variance

$$
\operatorname{Err}\left(\widehat{h_{n}}\right) \leq \operatorname{Err}\left(h^{*}\right)+\mathcal{B}_{2}(n, \mathcal{H})
$$

## Approaches

Minimization of the empirical risk

- Model selection: Choose hypothesis space $\mathcal{H}$
- Choose $\widehat{h}_{n}=\operatorname{argmin}\left\{\operatorname{Err}_{n}(h), h \in \mathcal{H}\right\}$
beware of overfitting
Minimization of the structual risk
Given $\mathcal{H}_{1} \subset \mathcal{H}_{2} \subset \ldots \subset \mathcal{H}_{k}$,

$$
\text { Find } \widehat{h}_{n}=\operatorname{argmin}\left\{\operatorname{Err}_{n}(h)+\operatorname{pen}(n, k), h \in \mathcal{H}_{k}\right\}
$$

Which penalization?
Regularization

$$
\text { Find } \widehat{h}_{n}=\operatorname{argmin}\left\{\operatorname{Err}_{n}(h)+\lambda\|h\|, h \in \mathcal{H}\right\}
$$

$\lambda$ is identified by cross-validation

## Structural Risk Minimization



## Tool 1. Hoeffding bound

Hoeffing 1963
Let $X_{1} \ldots, X_{n}$ be independent random variables, and assume $X_{i}$ takes values in [ $a_{i}, b_{i}$ ]
Let $\bar{X}=\left(X_{1}+\cdots+X_{n}\right) / n$ be their empirical mean.

## Theorem

$$
\operatorname{Pr}(|\bar{X}-\mathrm{E}[\bar{X}]| \geq \varepsilon) \leq 2 \exp \left(-\frac{2 \varepsilon^{2} n^{2}}{\sum_{i=1}^{n}\left(b_{i}-a_{i}\right)^{2}}\right)
$$

where $\mathrm{E}[\bar{X}]$ is the expectation of $\bar{X}$.

## Hoeffding Bound (2)

Application: if

$$
\operatorname{Pr}\left(\left|\operatorname{Err}(g)-\operatorname{Err}_{n}(g)\right|>\varepsilon\right)<2 e^{-2 n \varepsilon^{2}}
$$

then with probability at least $1-\delta$

$$
\operatorname{Err}(g) \leq \operatorname{Err}_{n}(g)+\sqrt{\frac{\log 2 / \delta}{2 n}}
$$

but this does not say anything about $\hat{h}_{n} \ldots$
Uniform deviations

$$
\left|\operatorname{Err}\left(\hat{h}_{n}\right)-\operatorname{Err}_{n}\left(\hat{h}_{n}\right)\right| \leq \sup _{h \in H}\left|\operatorname{Err}(h)-\operatorname{Err}_{n}(h)\right|
$$

- if $\mathcal{H}$ is finite, consider the sum of $\left|\operatorname{Err}(h)-\operatorname{Err}_{n}(h)\right|$
- sif $\mathcal{H}$ is infinite, consider its trace on the data


## Statistical Learning. Definitions

Vapnik 92, 95, 98
Trace of $\mathcal{H}$ on $\left\{x_{1}, \ldots x_{n}\right\}$

$$
\operatorname{Tr}_{x_{1}, . . x_{n}}(\mathcal{H})=\left\{\left(h\left(x_{1}\right), . . h\left(x_{n}\right)\right), h \in \mathcal{H}\right\}
$$

Growth Function

$$
S(\mathcal{H}, n)=\sup _{\left(x_{1}, . . . x_{n}\right)}\left|\operatorname{Tr}_{x_{1}, . . x_{n}}(\mathcal{H})\right|
$$

## Statistical Learning. Definitions (2)

Capacity of an hypothesis space $\mathcal{H}$
If the training set is of size $n$, and some function of $\mathcal{H}$ can have "any behavior" on $n$ examples, nothing can be said!
$\mathcal{H}$ shatters $\left(x_{1}, \ldots x_{n}\right)$ iff

$$
\forall\left(y_{1}, \ldots y_{n}\right) \in\{1,-1\}^{n}, \exists h \in \mathcal{H} \text { s.t. } \forall i=1 \ldots n, h\left(x_{i}\right)=y_{i}
$$

Vapnik Cervonenkis Dimension $\mathrm{VC}(\mathcal{H})=\max \left\{n ;\left(x_{1}, \ldots x_{n}\right)\right.$ shattered by $\left.\mathcal{H}\right\}$

$$
V C(\mathcal{H})=\max \left\{n / S(\mathcal{H}, n)=2^{n}\right\}
$$

## A shattered set

3 points in $\mathbb{R}^{2}$
$\mathcal{H}=$ lines of the plane


## Growth Function of linear functions over $\mathbb{R}^{20}$



THe growth function is exponental w.r.t. $n$ for $n<d=V C(\mathcal{H})$, then polynomial (in $n^{d}$ ).

## Theorem, separable case

$\forall \delta>0$, with probability at least $1-\delta$

$$
\operatorname{Err}(h) \leq \operatorname{Err}_{n}(h)+\sqrt{2 \frac{\log (S(H, 2 n))+\log (2 / \delta)}{n}}
$$

Idea 1: Double sample trick
Consider a second sample $\mathcal{E}^{\prime}$

$$
\begin{aligned}
\operatorname{Pr}\left(\sup _{h}\left(\operatorname{Err}(h)-\operatorname{Err}_{n}(h)\right) \geq \varepsilon\right) \leq & \\
& 2 \operatorname{Pr}\left(\sup _{h}\left(\operatorname{Err}_{n}^{\prime}(h)-\operatorname{Err}_{n}(h)\right) \geq \varepsilon / 2\right)
\end{aligned}
$$

where $\operatorname{Err}_{n}^{\prime}(h)$ is the empirical error on $\mathcal{E}^{\prime}$.

## Double sample trick

- There exists $h$ s.t.
- A: $\operatorname{Err}(h)=0$
- B: $\operatorname{Err}(h) \geq \varepsilon$
- C: $E r r_{\mathcal{E}^{\prime}} \geq \frac{\varepsilon}{2}$

$$
\begin{aligned}
P(A(h) \& C(h)) & \geq P(A(h) \& B(h) \& C(h)) \\
& =P(A(h) \& B(h)) . P(C(h) \mid A(h) \& B(h)) \\
& \geq \frac{1}{2} P(A(h) \& B(h))
\end{aligned}
$$

## Tool 2. Sauer Lemma

Sauer Lemma
If $d=V C(\mathcal{H})$

$$
S(\mathcal{H}, n)=\sum_{i=1}^{d}\binom{n}{i}
$$

For $n>d$,

$$
S(H, n) \leq\left(\frac{e n}{d}\right)^{d}
$$

Idea 2: Symmetrization
Count the permutations that swap $\mathcal{E}$ et $\mathcal{E}^{\prime}$.
Summary

$$
\operatorname{Err}(h) \leq E r r_{n}(h)+\mathcal{O}\left(\sqrt{\frac{d \log n}{n}}\right)
$$

