## Support vector machines

## OPTIMAL SEPARATING HYPERPLANES

A main initiative in early computer science was to find separating hyperplanes among groups of data
(Rosenblatt (1958) with the perceptron algorithm)
The issue is that if there is a separating hyperplane, there is an infinite number

An optimal separating hyperplane can be generated by finding support points and bisecting them.
(Sometimes optimal separating hyperplanes are called maximum margin classifiers)

## BASIC LINEAR GEOMETRY

A hyperplane in $\mathbb{R}^{p}$ is given by

$$
\mathcal{H}=\left\{X \in \mathbb{R}^{p}: h(X)=\beta_{0}+\beta^{\top} X=0\right\}
$$

(Usually it is assumed that $\|\beta\|_{2}=1$ )

1. The vector $\beta$ is normal to $\mathcal{H}$
2. For any point $X \in \mathbb{R}^{p}$, the (signed) length of its orthogonal complement to $\mathcal{H}$ is $h(X)$

## Support vector machines (SVM)

Let $Y_{i} \in\{-1,1\}$
(It is common with SVMs to code $Y$ this way. With logistic regression, $Y$ is commonly phrased as $\{0,1\}$ due to the connection with Bernoulli trials)

We will generalize this to supervisors with more than 2 levels at the end

A classification rule induced by a hyperplane is

$$
g(X)=\operatorname{sgn}\left(X^{\top} \beta+\beta_{0}\right)
$$

## SEPARATING HYPERPLANES

Our classification rule is based on a hyperplane $\mathcal{H}$

$$
g(X)=\operatorname{sgn}\left(X^{\top} \beta+\beta_{0}\right)
$$

A correct classification is one such that $h(X) Y>0$ and $g(X) Y>0$

The larger the quantity $Y h(X)$, the more "sure" the classification

Under classical separability, we can find a function such that $Y_{i} h\left(X_{i}\right)>0$

## OPTIMAL SEPARATING HYPERPLANE

This idea can be encoded in the following convex program

$$
\begin{gathered}
M \rightarrow \max _{\beta_{0}, \beta}, \text { subject to } \\
Y_{i} h\left(X_{i}\right) \geq M \text { for each } i \text { and }\|\beta\|_{2}=1
\end{gathered}
$$

## Intuition:

- We know that $Y_{i} h\left(X_{i}\right)>0 \Rightarrow g\left(X_{i}\right)=Y_{i}$. Hence, larger $Y_{i} h\left(X_{i}\right) \Rightarrow$ "more" correct classification
- For "more" to have any meaning, we need to normalize $\beta$, thus the other constraint


## OPTIMAL SEPARATING HYPERPLANE

Let's take the original program:

$$
\begin{gathered}
M \rightarrow \max _{\beta_{0}, \beta} \text {, subject to } \\
Y_{i} h\left(X_{i}\right) \geq M \text { for each } i \text { and }\|\beta\|_{2}=1
\end{gathered}
$$

and rewrite it as

$$
\begin{aligned}
& \min _{\beta_{0}, \beta} \frac{1}{2}\|\beta\|_{2}^{2} \text { subject to } \\
& Y_{i} h\left(X_{i}\right) \geq 1 \text { for each } i
\end{aligned}
$$

This is still a convex optimization program: quadratic criterion, linear inequality constraints

## OPTIMAL SEPARATING HYPERPLANE

We can convert this constrained optimization problem into the Lagrangian (primal) form

$$
\min _{\beta_{0}, \beta} \frac{1}{2}\|\beta\|_{2}^{2}-\sum_{i=1}^{n} \alpha_{i}\left[Y_{i}\left(X_{i}^{\top} \beta+\beta_{0}\right)-1\right]
$$

Everything is nice and smooth, so we can take derivatives..

## OPTIMAL SEPARATING HYPERPLANE

$$
\frac{1}{2}\|\beta\|_{2}^{2}-\sum_{i=1}^{n} \alpha_{i}\left[Y_{i}\left(X_{i}^{\top} \beta+\beta_{0}\right)-1\right]
$$

Derivatives with respect to $\beta$ and $\beta_{0}$ :

- $\beta=\sum_{i=1}^{n} \alpha_{i} Y_{i} X_{i}$
- $0=\sum_{i=1}^{n} \alpha_{i} Y_{i}$

Substituting into the Lagrangian:

$$
\text { Wolfe Dual }=\sum_{i=1}^{n} \alpha_{i}-\frac{1}{2} \sum_{i=1}^{n} \sum_{k=1}^{n} \alpha_{i} \alpha_{k} Y_{i} Y_{k} X_{i}^{\top} X_{k}
$$

(this is all subject to $\alpha_{i} \geq 0$ )
We want to maximize Wolfe Dual

## OPTIMAL SEPARATING HYPERPLANE

A side condition, known as complementary slackness states (or Karush-Kuhn-Tucker (KKT) conditions):

$$
\alpha_{i}\left[1-Y_{i} h\left(X_{i}\right)\right]=0 \text { for all } i
$$

(The product of Lagrangian parameters and inequalty constraint equals 0 )
This implies either:

- $\alpha_{i}=0$, which happens if the constraint $Y_{i} h\left(X_{i}\right)>1$
- $\alpha_{i}>0$, which happens if the constraint $Y_{i} h\left(X_{i}\right)=1$


## OPTIMAL SEPARATING HYPERPLANE

Taking this relationship

$$
\alpha_{i}\left[Y_{i} h\left(X_{i}\right)-1\right]=0
$$

we see that, for $i=1, \ldots, n$,

- The points $\left(X_{i}, Y_{i}\right)$ such that $\alpha_{i}>0$ are support vectors
- The points $\left(X_{i}, Y_{i}\right)$ such that $\alpha_{i}=0$ are irrelevant for classification

End Result: $\hat{g}(X)=\operatorname{sgn}\left(X^{\top} \hat{\beta}+\hat{\beta}_{0}\right)$

## Support vector classifier

## Support vector CLASSIFIER

Of course, we can't realistically assume that the data are linearly separated (even in a transformed space)

In this case, the previous program has no feasible solution
We need to introduce slack variables, $\xi$, that allow for overlap among the classes

These slack variables allow for us to encode training missclassifications into the optimization problem

## Support vector classifier

$$
\begin{gathered}
M \rightarrow \max _{\beta_{0}, \beta, \xi_{1}, \ldots, \xi_{n}}, \text { subject to } \\
Y_{i} h\left(X_{i}\right) \geq M \underbrace{\left(1-\xi_{i}\right), \quad \xi_{i} \geq 0, \quad \sum \xi_{i} \leq t}_{\text {new }}, \text { for each } i
\end{gathered}
$$

Note that

- $t$ is a tuning parameter. The literature usually refers to $t$ as a budget
- The separable case corresponds to $t=0$


## Support vector classifier

We can rewrite the problem again:

$$
\begin{gathered}
\min _{\beta_{0}, \beta, \xi} \frac{1}{2}\|\beta\|_{2}^{2}, \text { subject to } \\
Y_{i} h\left(X_{i}\right) \geq 1 \underbrace{1-\xi_{i}, \quad \xi_{i} \geq 0, \quad \sum \xi_{i} \leq t}_{\text {new }}, \text { for each } i
\end{gathered}
$$

Converting $\sum \xi_{i} \leq t$ to the Lagrangian (primal):

$$
\begin{aligned}
& \min _{\beta_{0}, \beta} \frac{1}{2}\|\beta\|_{2}^{2}+\lambda \sum \xi_{i} \text { subject to } \\
& Y_{i} h\left(X_{i}\right) \geq 1-\xi_{i}, \xi_{i} \geq 0, \text { for each } i
\end{aligned}
$$

## SVMS: SLACK VARIABLES

The slack variables give us insight into the problem

- If $\xi_{i}=0$, then that observation is on correct the side of the margin
- If $\xi_{i}=\in(0,1]$, then that observation is on the incorrect side of the margin, but still correctly classified
- If $\xi_{i}>1$, then that observation is incorrectly classified


## Support vector cLassifier

Continuing to convert constraints to Lagrangian

$$
\min _{\beta_{0}, \beta, \xi} \frac{1}{2}\|\beta\|_{2}^{2}+\lambda \sum \xi_{i} \underbrace{-\sum_{i=1}^{n} \alpha_{i}\left[Y_{i}\left(X_{i}^{\top} \beta+\beta_{0}\right)-\left(1-\xi_{i}\right)\right]-\sum_{i=1}^{n} \gamma_{i} \xi_{i}}_{\text {remaining constraints }}
$$

Necessary conditions (taking derivatives)

- $\beta=\sum_{i=1}^{n} \alpha_{i} Y_{i} X_{i}$
- $0=\sum_{i=1}^{n} \alpha_{i} Y_{i}$
- $\alpha_{i}=\lambda-\gamma_{i}$
(As well as positivity constraints on Lagrangian parameters)


## Support vector CLASSIFIER

Substituting, we reaquire the Wolfe Dual
This, combined with the KKT conditions uniquely characterize the solution:
$\alpha$ subject to: $\max _{\text {KKT Wolfe Dual }} \sum_{i=1}^{n} \alpha_{i}-\frac{1}{2} \sum_{i=1}^{n} \sum_{i^{\prime}=1}^{n} \alpha_{i} \alpha_{i^{\prime}} Y_{i} Y_{i^{\prime}} X_{i}^{\top} X_{i^{\prime}}$

Note: the necessary conditions $\beta=\sum_{i=1}^{n} \alpha_{i} Y_{i} X_{i}$ imply estimators of the form

- $\hat{\beta}=\sum_{i=1}^{n} \hat{\alpha}_{i} Y_{i} X_{i}$
- $\hat{\beta}^{\top} X=\sum_{i=1}^{n} \hat{\alpha}_{i} Y_{i} X_{i}^{\top} X$


## SVMS: TUNING PARAMETER

We can think of $t$ as a budget for the problem
If $t=0$, then there is no budget and we won't tolerate any margin violations

If $t>0$, then no more than $\lfloor t\rfloor$ observations can be misclassified

A larger $t$ then leads to larger margins (we allow more margin violations)

## SVMs: TUNING PARAMETER

## Further intuition:

Like the optimal hyperplane, only observations that violate the margin determine $\mathcal{H}$

A large $t$ allows for many violations, hence many observations factor into the fit

A small $t$ means only a few observations do
Hence, $t$ calibrates a bias/variance trade-off, as expected
In practice, $t$ gets selected via cross-validation

## SVMs: TUNING PARAMETER




Figure: Figure 9.7 in ISL

## Kernel methods

Intuition: Many methods have linear decision boundaries
We know that sometimes this isn't sufficient to represent data

Example: Sometimes we need to included a polynomial effect or a log transform in multiple regression

Sometimes, a linear boundary, but in a different space makes all the difference..

## OPTIMAL SEPARATING HYPERPLANE

Reminder: The Wolfe dual, which gets maximized over $\alpha$, produces the optimal separating hyperplane

$$
\text { Wolf dual }=\sum_{i=1}^{n} \alpha_{i}-\frac{1}{2} \sum_{i=1}^{n} \sum_{k=1}^{n} \alpha_{i} \alpha_{k} Y_{i} Y_{k} X_{i}^{\top} X_{k}
$$

(this is all subject to $\alpha_{i} \geq 0$ )
A similar result holds after the introduction of slack variables (e.g. support vector classifiers)

Important: The features only enter via

$$
X^{\top} X^{\prime}=\left\langle X, X^{\prime}\right\rangle
$$

## Kernel Methods

## Nonnegative definite matrices

Let $A \in \mathbb{R}^{p \times p}$ be a symmetric, nonnegative definite matrix:

$$
z^{\top} A z \geq 0 \text { for all } z \text { and } A^{\top}=A
$$

Then, $A$ has an eigenvalue expansion

$$
A=U D U^{\top}=\sum_{j=1}^{p} d_{j} u_{j} u_{j}^{\top}
$$

where $d_{j} \geq 0$
Observation: Each such $A$, generates a new inner product

$$
\begin{gathered}
\left\langle z, z^{\prime}\right\rangle=z^{\top} z^{\prime}=z^{\top} \underbrace{1}_{\text {Identity }} z^{\prime} \\
\left\langle z, z^{\prime}\right\rangle_{A}=z^{\top} A z^{\prime}
\end{gathered}
$$

(If we enforce $A$ to be positive definite, then $\langle z, z\rangle_{A}=\|z\|_{A}^{2}$ is a norm)

## Nonnegative definite matrices

Suppose $A_{i}^{j}$ is the $(i, j)$ entry in $A$, and $A_{i}$ is the $i^{\text {th }}$ row

$$
A z=\left[\begin{array}{c}
A_{1}^{\top} \\
\vdots \\
A_{p}^{\top}
\end{array}\right] z=\left[\begin{array}{c}
A_{1}^{\top} z \\
\vdots \\
A_{p}^{\top} z
\end{array}\right]
$$

Note: Multiplication by $A$ is really taking inner products with its rows.

Hence, $A_{i}$ is called the (multiplication) kernel of matrix $A$

## Kernel methods

$k: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is a symmetric, nonnegative definite kernel Write the eigenvalue expansion of $k$ as

$$
k\left(X, X^{\prime}\right)=\sum_{j=1}^{\infty} \theta_{j} \phi_{j}(X) \phi_{j}\left(X^{\prime}\right)
$$

with

- $\theta_{j} \geq 0 \quad$ (nonnegative definite)
- $\left\|\left(\theta_{j}\right)_{j=1}^{\infty}\right\|_{2}=\sum_{j=1}^{\infty} \theta_{j}^{2}<\infty$
- The $\phi_{j}$ are orthogonal eigenfunctions: $\int \phi_{j} \phi_{j^{\prime}}=\delta_{j, j^{\prime}}$

We can write any $f \in \mathcal{H}_{k}$ with two constraints

- $f(x)=\sum_{j=1}^{\infty} f_{j} \phi_{j}(x)$
- $\langle f, f\rangle_{\mathcal{H}_{k}}=\|f\|_{\mathcal{H}_{k}}^{2}=\sum_{j=1}^{\infty} f_{j}^{2} / \theta_{j}<\infty$


## Kernel: Example

Back to polynomial terms/interactions:
Form

$$
k_{d}\left(X, X^{\prime}\right)=\left(X^{\top} X^{\prime}+1\right)^{d}
$$

$k_{d}$ has $M=\binom{p+d}{d}$ eigenfunctions
These span the space of polynomials in $\mathbb{R}^{p}$ with degree $d$

## Kernel: Example

Example: Let $d=p=2 \Rightarrow M=6$ and

$$
\begin{aligned}
k(u, v) & =1+2 u_{1} v_{1}+2 u_{2} v_{2}+u_{1}^{2} v_{1}^{2}+u_{2}^{2} v_{2}^{2}+2 u_{1} u_{2} v_{1} v_{2} \\
& =\sum_{k=1}^{M} \Phi_{k}(u) \Phi_{k}(v) \\
& =\Phi(u)^{\top} \Phi(v) \\
& =\langle\Phi(u), \Phi(v)\rangle
\end{aligned}
$$

where

$$
\Phi(v)^{\top}=\left(1, \sqrt{2} v_{1}, \sqrt{2} v_{2}, v_{1}^{2}, v_{2}^{2}, \sqrt{2} v_{1} v_{2}\right)
$$

## Kernel: Conclusion

Let's recap:

$$
\begin{aligned}
k(u, v) & =1+2 u_{1} v_{1}+2 u_{2} v_{2}+u_{1}^{2} v_{1}^{2}+u_{2}^{2} v_{2}^{2}+2 u_{1} u_{2} v_{1} v_{2} \\
& =\langle\Phi(u), \Phi(v)\rangle
\end{aligned}
$$

- Some methods only involve features via inner products $X^{\top} X^{\prime}=\left\langle X, X^{\prime}\right\rangle$
(We've explicitly seen two: ridge regression and support vector classifiers)
- If we make transformations of $X$ to $\Phi(X)$, the procedure depends on $\Phi(X)^{\top} \Phi\left(X^{\prime}\right)=\left\langle\Phi(X), \Phi\left(X^{\prime}\right)\right\rangle$
- We can compute this inner product via the kernel:

$$
k\left(X, X^{\prime}\right)=\left\langle\Phi(X), \Phi\left(X^{\prime}\right)\right\rangle
$$

## (Kernel) SVMs

## Kernel SVM

Recall:

$$
\frac{1}{2}\|\beta\|_{2}^{2}-\sum_{i=1}^{n} \alpha_{i}\left[Y_{i}\left(X_{i}^{\top} \beta+\beta_{0}\right)-1\right]
$$

Derivatives with respect to $\beta$ and $\beta_{0}$ imply:

- $\beta=\sum_{i=1}^{n} \alpha_{i} Y_{i} X_{i}$
- $0=\sum_{i=1}^{n} \alpha_{i} Y_{i}$

Write the solution function

$$
h(X)=\beta_{0}+\beta^{\top} X=\beta_{0}+\sum_{i=1}^{n} \alpha_{i} Y_{i} X_{i}^{\top} X
$$

Kernelize the support vector classifier $\Rightarrow$ support vector machine (SVM):

$$
h(X)=\beta_{0}+\sum_{i=1}^{n} \alpha_{i} Y_{i} k\left(X_{i}, X\right)
$$

## General kernel machines

After specifying a kernel function, it can be shown that many procedures have a solution of the form

$$
\hat{f}(X)=\sum_{i=1}^{n} \gamma_{i} k\left(X, X_{i}\right)
$$

For some $\gamma_{1}, \ldots, \gamma_{n}$
Also, this is equivalent to performing the method in the space given by the eigenfunctions of $k$

$$
k(u, v)=\sum_{j=1}^{\infty} \theta_{j} \phi_{j}(u) \phi_{j}(v)
$$

Also, (the) feature map is

$$
\Phi=\left[\phi_{1}, \ldots, \phi_{p}, \ldots\right]
$$

## Kernel SVM: A REMInder

The dual Lagrangian is:

$$
\ell_{D}(\gamma)=\sum_{i} \gamma_{i}-\frac{1}{2} \sum_{i} \sum_{i^{\prime}} \gamma_{i} \gamma_{i^{\prime}} Y_{i} Y_{i^{\prime}} X_{i}^{\top} X_{i^{\prime}}
$$

with side conditions: $\gamma_{i} \in[0, C]$ and $\gamma^{\top} Y=0$
Let's replace the term $X_{i}^{\top} X_{i^{\prime}}=\left\langle X_{i}, X_{i^{\prime}}\right\rangle$ with $\left\langle\Phi\left(X_{i}\right), \Phi\left(X_{i^{\prime}}\right)\right\rangle$

## Kernel SVMs

Hence (and luckily) specifying $\Phi$ itself unnecessary,
(Luckily, as many kernels have difficult to compute eigenfunctions)
We need only define the kernel that is symmetric, positive definite

Some common choices for SVMs:

- Polynomial: $k(x, y)=\left(1+x^{\top} y\right)^{d}$
- RADIAL BASIS: $k(x, y)=e^{-\tau\|x-y\|_{b}^{b}}$
(For example, $b=2$ and $\tau=1 /\left(2 \sigma^{2}\right)$ is (proportional to) the Gaussian density)


## Kernel SVMs: Summary

Reminder: the solution form for SVM is

$$
\beta=\sum_{i=1}^{n} \alpha_{i} Y_{i} X_{i}
$$

Kernelized, this is

$$
\beta=\sum_{i=1}^{n} \alpha_{i} Y_{i} \Phi\left(X_{i}\right)
$$

Therefore, the induced hyperplane is:

$$
\begin{aligned}
h(X)=\Phi(X)^{\top} \beta+\beta_{0} & =\sum_{i=1}^{n} \alpha_{i} Y_{i}\left\langle\Phi(X), \Phi\left(X_{i}\right)\right\rangle+\beta_{0} \\
& =\sum_{i=1}^{n} \alpha_{i} Y_{i} k\left(X, X_{i}\right)+\beta_{0}
\end{aligned}
$$

The final classification is still $\hat{g}(X)=\operatorname{sgn}(\hat{h}(X))$

SVMs via penalization

## SVMs Via penalization

Note: SVMs can be derived from penalized loss methods
The support vector classifier optimization problem:

$$
\begin{gathered}
\min _{\beta_{0}, \beta} \frac{1}{2}\|\beta\|_{2}^{2}+\lambda \sum \xi_{i} \text { subject to } \\
Y_{i} h\left(X_{i}\right) \geq 1-\xi_{i}, \xi_{i} \geq 0, \text { for each } i
\end{gathered}
$$

Writing $h(X)=\Phi(X)^{\top} \beta+\beta_{0}$, consider

$$
\min _{\beta, \beta_{0}} \sum_{i=1}^{n}\left[1-Y_{i} h\left(X_{i}\right)\right]_{+}+\tau\|\beta\|_{2}^{2}
$$

These optimization problems are the same!
(With the relation: $2 \lambda=1 / \tau$ )

## SVMS VIA PENALIZATION

The loss part is the hinge loss function

$$
\ell(X, Y)=[1-Y h(X)]_{+}
$$

The hinge loss approximates the zero-one loss function underlying classification

It has one major advantage, however: convexity

## Surrogate losses: Convex relaxation

Looking at

$$
\min _{\beta, \beta_{0}} \sum_{i=1}^{n}\left[1-Y_{i} h\left(X_{i}\right)\right]_{+}+\tau\|\beta\|_{2}^{2}
$$

It is tempting to minimize (analogous to linear regression)

$$
\sum_{i=1}^{n} \mathbf{1}\left(Y_{i} \neq \hat{g}\left(X_{i}\right)\right)+\tau\|\beta\|_{2}^{2}
$$

However, this is nonconvex (in $u=h(X) Y$ )

A common trick is to approximate the nonconvex objective with a convex one
(This is known as convex relaxation with a surrogate loss function)

## Surrogate losses

Idea: We can use a surrogate loss that mimics this function while still being convex

It turns out we have already done that! (twice)

- Hinge: [1 - Yh( $X$ ) $]_{+}$
- Logistic: $\log \left(1+e^{-Y h(X)}\right)$

Multiclass classification

## Multiclass SVMs

Sometimes, it becomes necessary to do multiclass classification
There are two main approaches:

- One-versus-one
- One-vesus-all


## Multiclass SVMs: One-VERSus-one

Here, for $G$ possible classes, we run $G(G-1) / 2$ possible pairwise classifications

For a given test point $X$, we find $\hat{g}_{k}(X)$ for $k=1, \ldots, G(G-1) / 2$ fits

The result is a vector $\hat{G} \in \mathbb{R}^{G}$ with the total number of times $X$ was assigned to each class

We report $\hat{g}(X)=\arg \max _{g} \hat{G}$
This approach uses all the class information, but can be slow

## Multiclass SVMs: One-vesus-all

Here, we fit only $G$ SVMs by respectively collapsing over all size $G-1$ subsets of $\{1, \ldots, G\}$
(This is compared with $G(G-1) / 2$ comparisons for one-versus-one)
Take all $\hat{h}_{g}(X)$ for $g=1, \ldots, G$, where class $g$ is coded 1 and "the rest" is coded -1

Assign $\hat{g}(X)=\arg \max _{g} \hat{h}_{g}(X)$

# Background: Structural Risk Minimization 

## Capacity and Generalization

- Generalization: Figure out similarities between already-seen data and new data
- Too much: "Square piece of paper? That's a $\$ 100$ bill"
- Capacity: Ability to allocate new categories for data
- Too much: "\#L26118670? It's a fake; all \$100 bills I've seen had other serial numbers"
- They are competitive with one another
- How to strike the right balance?


## Empirical Risk

- We are given $n$ observations $\left(\mathbf{x}_{i}, y_{i}\right)$
- $\mathbf{x}_{i} \in \mathbb{R}^{p}$
- $y_{i} \in\{-1,1\}$
- Learn $y=f(\mathbf{x}, \alpha)$ by tuning $\alpha$
- Expected test error (risk) and empirical risk:

$$
\begin{gathered}
R(\alpha)=\frac{1}{2} \int|y-f(\mathbf{x}, \alpha)| d P(\mathbf{x}, y) \\
R_{e m p}(\alpha)=\frac{1}{2 l} \sum\left|y_{i}-f\left(\mathbf{x}_{i}, \alpha\right)\right|
\end{gathered}
$$

## Risk Bound

- For $0 / 1$ loss and with probability $1-\eta, 0<\eta<1$ :

$$
R(\alpha) \leq R_{e m p}(\alpha)+\sqrt{\frac{h\left(1+\log \frac{2 n}{h}\right)-\log \frac{\eta}{4}}{n}}
$$

where $h \in \mathbb{N}$ is the Vapnik-Chervonenkis (VC) dimension

- Second term: "VC confidence"


## Importance of Risk Bound

1. Not dependent on $P(\mathbf{x}, y)$
2. Ihs not computable
3. rhs computable if we know $h$

- For a given task, choose the machine that minimizes the risk bound!
- Even when bound not tight, we can contrast "tightness" of various families of machines


## The VC DIMENSION

- For a family of functions $f(\alpha)$ :
- Choose a set of $n$ points
- Label them in any way
- $\exists \alpha$ s.t. $f(\alpha)$ can recognize ("shatter") them
- Then $f(\alpha)$ has VC at least $n$


## Example: Hyperplanes in $\mathbb{R}^{n}$

- Choosing 4 planar points:
- they can't be separated by one line for all of their possible labelings (one labeling will be inseparable)
- Similarly, $p+1$ points in $\mathbb{R}^{p}$ can't be separated for all labelings
- So the VC dimension of hyperplanes in $\mathbb{R}^{p}$ is $p+1$

