## Machine Learning - Lecture 7

## Statistical Learning Theory

### 23.05.2016

Bastian Leibe<br>RWTH Aachen<br>http://www.vision.rwth-aachen.de<br>leibe@vision.rwth-aachen.de<br>Many slides adapted from B. Schiele

## Course Outline

- Fundamentals (2 weeks)
, Bayes Decision Theory
, Probability Density Estimation

- Discriminative Approaches (5 weeks)
, Linear Discriminant Functions
, Statistical Learning Theory \& SVMs
, Ensemble Methods \& Boosting
> Randomized Trees, Forests \& Ferns
- Generative Models (4 weeks)
, Bayesian Networks
, Markov Random Fields




## Topics of This Lecture

- Recap: Generalized Linear Discriminants
- Logistic Regression
, Probabilistic discriminative models
, Logistic sigmoid (logit function)
, Cross-entropy error
, Gradient descent
, Iteratively Reweighted Least Squares
- Note on error functions
- Statistical Learning Theory
, Generalization and overfitting
, Empirical and actual risk
, VC dimension
, Empirical Risk Minimization
, Structural Risk Minimization


## Recap: Linear Discriminant Functions

- Basic idea
, Directly encode decision boundary
, Minimize misclassification probability directly.
- Linear discriminant functions

$$
y(\mathbf{x})=\mathbf{w}^{\mathrm{T}} \mathbf{x}+w_{\text {weight vector }}^{w_{0}}
$$

> $\mathbf{w}, w_{0}$ define a hyperplane in $\mathbb{R}^{D}$.

, If a data set can be perfectly classified by a linear discriminant, then we call it linearly separable.

## Recap: Extension to Nonlinear Basis Fcts.

- Generalization
, Transform vector $\mathbf{x}$ with $M$ nonlinear basis functions $\phi_{j}(\mathbf{x})$ :

$$
y_{k}(\mathbf{x})=\sum_{j=1}^{M} w_{k j} \phi_{j}(\mathbf{x})+w_{k 0}
$$

- Advantages
, Transformation allows non-linear decision boundaries.
> By choosing the right $\phi_{j}$, every continuous function can (in principle) be approximated with arbitrary accuracy.
- Disadvantage
, The error function can in general no longer be minimized in closed form.
$\Rightarrow$ Minimization with Gradient Descent


## Recap: Basis Functions

- Generally, we consider models of the following form

$$
y_{k}(\mathbf{x})=\sum_{j=0}^{M} w_{k j} \phi_{j}(\mathbf{x})=\mathbf{w}^{T} \phi(\mathbf{x})
$$

, where $\phi_{j}(\mathbf{x})$ are known as basis functions.
, In the simplest case, we use linear basis functions: $\phi_{d}(\mathbf{x})=x_{d}$.

- Other popular basis functions


B. Leibe


## Gradient Descent

- Iterative minimization
, Start with an initial guess for the parameter values $w_{k j}^{(0)}$.
, Move towards a (local) minimum by following the gradient.
- Basic strategies
, "Batch learning"

$$
w_{k j}^{(\tau+1)}=w_{k j}^{(\tau)}-\left.\eta \frac{\partial E(\mathbf{w})}{\partial w_{k j}}\right|_{\mathbf{w}^{(\tau)}}
$$

, "Sequential updating" $w_{k j}^{(\tau+1)}=w_{k j}^{(\tau)}-\left.\eta \frac{\partial E_{n}(\mathbf{w})}{\partial w_{k j}}\right|_{\mathbf{w}^{(\tau)}}$
where $\quad E(\mathbf{w})=\sum_{n=1}^{N} E_{n}(\mathbf{w})$

## Recap: Gradient Descent

- Example: Quadratic error function

$$
E(\mathbf{w})=\sum_{n=1}^{N}\left(y\left(\mathbf{x}_{n} ; \mathbf{w}\right)-\mathbf{t}_{n}\right)^{2}
$$

- Sequential updating leads to delta rule (=LMS rule)

$$
\begin{aligned}
w_{k j}^{(\tau+1)} & =w_{k j}^{(\tau)}-\eta\left(y_{k}\left(\mathbf{x}_{n} ; \mathbf{w}\right)-t_{k n}\right) \phi_{j}\left(\mathbf{x}_{n}\right) \\
& =w_{k j}^{(\tau)}-\eta \delta_{k n} \phi_{j}\left(\mathbf{x}_{n}\right)
\end{aligned}
$$

, where

$$
\delta_{k n}=y_{k}\left(\mathbf{x}_{n} ; \mathbf{w}\right)-t_{k n}
$$

$\Rightarrow$ Simply feed back the input data point, weighted by the classification error.

## Recap: Gradient Descent

- Cases with differentiable, non-linear activation function

$$
y_{k}(\mathbf{x})=g\left(a_{k}\right)=g\left(\sum_{j=0}^{M} w_{k i} \phi_{j}\left(\mathbf{x}_{n}\right)\right)
$$

- Gradient descent (again with quadratic error function)

$$
\begin{aligned}
\frac{\partial E_{n}(\mathbf{w})}{\partial w_{k j}} & =\frac{\partial g\left(a_{k}\right)}{\partial w_{k j}}\left(y_{k}\left(\mathbf{x}_{n} ; \mathbf{w}\right)-t_{k n}\right) \phi_{j}\left(\mathbf{x}_{n}\right) \\
w_{k j}^{(\tau+1)} & =w_{k j}^{(\tau)}-\eta \delta_{k n} \phi_{j}\left(\mathbf{x}_{n}\right) \\
\delta_{k n} & =\frac{\partial g\left(a_{k}\right)}{\partial w_{k j}}\left(y_{k}\left(\mathbf{x}_{n} ; \mathbf{w}\right)-t_{k n}\right)
\end{aligned}
$$

## Recap: Classification as Dim. Reduction


good separation


- Classification as dimensionality reduction
- Interpret linear classification as a projection onto a lower-dim. space.

$$
y=\mathbf{w}^{\mathrm{T}} \mathbf{x}
$$

$\Rightarrow$ Learning problem: Try to find the projection vector w that maximizes class separation.

## RWIHAACHEN <br> Recap: Fisher's Linear Discriminant Analysis



- Maximize distance between classes
- Minimize distance within a class
- Criterion: $J(\mathbf{w})=\frac{\mathbf{w}^{\mathrm{T}} \mathbf{S}_{B} \mathbf{w}}{\mathbf{w}^{\mathrm{T}} \mathbf{S}_{W} \mathbf{w}}$
$\mathrm{S}_{B} \ldots$ between-class scatter matrix $\mathbf{S}_{W} \ldots$ within-class scatter matrix
- The optimal solution for $\mathbf{w}$ can be obtained as:

$$
\mathbf{w} \propto \mathbf{S}_{W}^{-1}\left(\mathbf{m}_{2}-\mathbf{m}_{1}\right)
$$

- Classification function:

$$
\begin{aligned}
y(\mathbf{x}) & =\mathbf{w}^{T} \mathbf{x}+w_{0} \stackrel{\text { Class } 1}{\gtrless} 0 \\
\text { where } w_{0} & =-\mathbf{w}^{T} \mathbf{m}
\end{aligned}
$$

## Topics of This Lecture

- Recap: Generalized Linear Discriminants
- Logistic Regression
, Probabilistic discriminative models
, Logistic sigmoid (logit function)
, Cross-entropy error
, Gradient descent
, Iteratively Reweighted Least Squares
- Note on error functions
- Statistical Learning Theory
, Generalization and overfitting
, Empirical and actual risk
, VC dimension
, Empirical Risk Minimization
- Structural Risk Minimization


## R

## Probabilistic Discriminative Models

- We have seen that we can write

$$
p\left(\mathcal{C}_{1} \mid \mathbf{x}\right)=\sigma(a)
$$

logistic sigmoid function

- We can obtain the familiar probabilistic model by setting

$$
a=\ln \frac{p\left(\mathbf{x} \mid \mathcal{C}_{1}\right) p\left(\mathcal{C}_{1}\right)}{p\left(\mathbf{x} \mid \mathcal{C}_{2}\right) p\left(\mathcal{C}_{2}\right)}
$$

- Or we can use generalized linear discriminant models

$$
\begin{aligned}
& a=\mathbf{w}^{T} \mathbf{x} \\
& \text { or } \quad a=\mathbf{w}^{T} \boldsymbol{\phi}(\mathbf{x})
\end{aligned}
$$

## R

## Probabilistic Discriminative Models

- In the following, we will consider models of the form

$$
\begin{aligned}
p\left(\mathcal{C}_{1} \mid \boldsymbol{\phi}\right) & =y(\boldsymbol{\phi})=\sigma\left(\mathrm{w}^{T} \boldsymbol{\phi}\right) \\
\text { with } & p\left(\mathcal{C}_{2} \mid \boldsymbol{\phi}\right)
\end{aligned}=1-p\left(\mathcal{C}_{1} \mid \boldsymbol{\phi}\right) \mathrm{l}
$$

- This model is called logistic regression.
- Why should we do this? What advantage does such a model have compared to modeling the probabilities?

$$
p\left(\mathcal{C}_{1} \mid \boldsymbol{\phi}\right)=\frac{p\left(\boldsymbol{\phi} \mid \mathcal{C}_{1}\right) p\left(\mathcal{C}_{1}\right)}{p\left(\boldsymbol{\phi} \mid \mathcal{C}_{1}\right) p\left(\mathcal{C}_{1}\right)+p\left(\boldsymbol{\phi} \mid \mathcal{C}_{2}\right) p\left(\mathcal{C}_{2}\right)}
$$

- Any ideas?


## Comparison

- Let's look at the number of parameters...
, Assume we have an $M$-dimensional feature space $\phi$.
, And assume we represent $p\left(\phi \mid \mathcal{C}_{k}\right)$ and $p\left(\mathcal{C}_{k}\right)$ by Gaussians.
, How many parameters do we need?
- For the means: $2 M$
- For the covariances: $\quad M(M+1) / 2$
- Together with the class priors, this gives $M(M+5) / 2+1$ parameters!
, How many parameters do we need for logistic regression?

$$
p\left(\mathcal{C}_{1} \mid \boldsymbol{\phi}\right)=y(\boldsymbol{\phi})=\sigma\left(\mathrm{w}^{T} \boldsymbol{\phi}\right)
$$

- Just the values of $\mathbf{w} \Rightarrow M$ parameters.
$\Rightarrow$ For large $M$, logistic regression has clear advantages!


## Logistic Sigmoid

- Properties
, Definition: $\sigma(a)=\frac{1}{1+\exp (-a)}$
, Inverse:

$$
a=\ln \left(\frac{\sigma}{1-\sigma}\right)
$$

"logit" function
, Symmetry property:

$$
\sigma(-a)=1-\sigma(a)
$$

, Derivative: $\frac{d \sigma}{d a}=\sigma(1-\sigma)$

## Logistic Regression

- Let's consider a data set $\left\{\phi_{n}, t_{n}\right\}$ with $n=1, \ldots, N$, where $\phi_{n}=\phi\left(\mathbf{x}_{n}\right)$ and $t_{n} \in\{0,1\}, \mathbf{t}=\left(t_{1}, \ldots, t_{N}\right)^{T}$.
- With $y_{n}=p\left(\mathcal{C}_{1} \mid \phi_{n}\right)$, we can write the likelihood as

$$
p(\mathbf{t} \mid \mathbf{w})=\prod_{n=1}^{N} y_{n}^{t_{n}}\left\{1-y_{n}\right\}^{1-t_{n}}
$$

- Define the error function as the negative log-likelihood

$$
\begin{aligned}
E(\mathbf{w}) & =-\ln p(\mathbf{t} \mid \mathbf{w}) \\
& =-\sum_{n=1}^{N}\left\{t_{n} \ln y_{n}+\left(1-t_{n}\right) \ln \left(1-y_{n}\right)\right\}
\end{aligned}
$$

, This is the so-called cross-entropy error function.

## Gradient of the Error Function

$$
y_{n}=\sigma\left(\mathbf{w}^{T} \boldsymbol{\phi}_{n}\right)
$$

- Error function

$$
E(\mathbf{w})=-\sum_{n=1}^{N}\left\{t_{n} \ln y_{n}+\left(1-t_{n}\right) \ln \left(1-y_{n}\right)\right\}
$$

- Gradient

$$
\begin{aligned}
& \nabla E(\mathbf{w})=-\sum_{n=1}^{N}\left\{t_{n} \frac{\frac{d}{d \mathbf{w}} y_{n}}{y_{n}}+\left(1-t_{n}\right) \frac{\frac{d}{d \mathbf{w}}\left(1-y_{n}\right)}{\left(1-y_{n}\right)}\right\} \\
&=-\sum_{n=1}^{N}\left\{t_{n} \frac{y_{n}\left(1-y_{n}\right)}{y_{n}} \phi_{n}-\left(1-t_{n}\right) \frac{y_{n}\left(1-y_{n}\right)}{\left(1-y_{n}\right)} \boldsymbol{\phi}_{n}\right\} \\
&=-\sum_{n=1}^{N}\left\{\left(t_{n}-t_{n} y_{n}-y_{n}+t_{n} \text { bn }_{n}\right) \phi_{n}\right\} \\
&=\sum_{n=1}^{N}\left(y_{n}-t_{n}\right) \phi_{n} \\
& \text { B. Leibe }
\end{aligned}
$$

## Gradient of the Error Function

- Gradient for logistic regression

$$
\nabla E(\mathbf{w})=\sum_{n=1}^{N}\left(y_{n}-t_{n}\right) \phi_{n}
$$

- Does this look familiar to you?
- This is the same result as for the Delta (=LMS) rule

$$
w_{k j}^{(\tau+1)}=w_{k j}^{(\tau)}-\eta\left(y_{k}\left(\mathbf{x}_{n} ; \mathbf{w}\right)-t_{k n}\right) \phi_{j}\left(\mathbf{x}_{n}\right)
$$

- We can use this to derive a sequential estimation algorithm.
. However, this will be quite slow...


## A More Efficient Iterative Method...

- Second-order Newton-Raphson gradient descent scheme

$$
\mathbf{w}^{(\tau+1)}=\mathbf{w}^{(\tau)}-\mathbf{H}^{-1} \nabla E(\mathbf{w})
$$

where $\mathbf{H}=\nabla \nabla E(\mathbf{w})$ is the Hessian matrix, i.e. the matrix of second derivatives.

- Properties
, Local quadratic approximation to the log-likelihood.
, Faster convergence.


## Newton-Raphson for Least-Squares Estimation

- Let's first apply Newton-Raphson to the least-squares error function:

$$
\begin{aligned}
& E(\mathbf{w})=\frac{1}{2} \sum_{n=1}^{N}\left(\mathbf{w}^{T} \boldsymbol{\phi}_{n}-t_{n}\right)^{2} \\
& \nabla E(\mathbf{w})=\sum_{n=1}^{N}\left(\mathbf{w}^{T} \boldsymbol{\phi}_{n}-t_{n}\right) \boldsymbol{\phi}_{n}=\boldsymbol{\Phi}^{T} \boldsymbol{\Phi} \mathbf{w}-\boldsymbol{\Phi}^{T} \mathbf{t} \\
& \mathbf{H}=\nabla \nabla E(\mathbf{w})=\sum_{n=1}^{N} \boldsymbol{\phi}_{n} \boldsymbol{\phi}_{n}^{T}=\boldsymbol{\Phi}^{T} \boldsymbol{\Phi} \quad \text { where } \boldsymbol{\Phi}=\left[\begin{array}{c}
\boldsymbol{\phi}_{1}^{T} \\
\vdots \\
\boldsymbol{\phi}_{N}^{T}
\end{array}\right]
\end{aligned}
$$

- Resulting update scheme:

$$
\begin{aligned}
\mathbf{w}^{(\tau+1)} & =\mathbf{w}^{(\tau)}-\left(\boldsymbol{\Phi}^{T} \boldsymbol{\Phi}\right)^{-1}\left(\boldsymbol{\Phi}^{T} \boldsymbol{\Phi} \mathbf{w}^{(\tau)}-\boldsymbol{\Phi}^{T} \mathbf{t}\right) \\
& =\left(\boldsymbol{\Phi}^{T} \boldsymbol{\Phi}\right)^{-1} \boldsymbol{\Phi}^{T} \mathbf{t} \quad \text { Closed-form solution! }
\end{aligned}
$$

## Newton-Raphson for Logistic Regression

- Now, let's try Newton-Raphson on the cross-entropy error function:

$$
\begin{aligned}
E(\mathbf{w}) & =-\sum_{n=1}^{N}\left\{t_{n} \ln y_{n}+\left(1-t_{n}\right) \ln \left(1-y_{n}\right)\right\} \\
\nabla E(\mathbf{w}) & =\sum_{n=1}^{N}\left(y_{n}-t_{n}\right) \boldsymbol{\phi}_{n}=\boldsymbol{\Phi}^{T}(\mathbf{y}-\mathbf{t}) \\
\mathbf{H}=\nabla \nabla E(\mathbf{w}) & =\sum_{n=1}^{N} y_{n}\left(1-y_{n}\right) \boldsymbol{\phi}_{n} \boldsymbol{\phi}_{n}^{T}=\boldsymbol{\Phi}^{T}\left(1-y_{n}\right) \boldsymbol{\phi}_{n}
\end{aligned}
$$

$$
\text { where } \mathbf{R} \text { is an } N \times N \text { diagonal matrix with } R_{n n}=y_{n}\left(1-y_{n}\right) \text {. }
$$

$\Rightarrow$ The Hessian is no longer constant, but depends on $w$ through the weighting matrix $\mathbf{R}$.

## Iteratively Reweighted Least Squares

- Update equations

$$
\begin{aligned}
\mathbf{w}^{(\tau+1)}= & \mathbf{w}^{(\tau)}-\left(\boldsymbol{\Phi}^{T} \mathbf{R} \boldsymbol{\Phi}\right)^{-1} \boldsymbol{\Phi}^{T}(\mathbf{y}-\mathbf{t}) \\
= & \left(\boldsymbol{\Phi}^{T} \mathbf{R} \boldsymbol{\Phi}\right)^{-1}\left\{\boldsymbol{\Phi}^{T} \mathbf{R} \boldsymbol{\Phi} \mathbf{w}^{(\tau)}-\boldsymbol{\Phi}^{T}(\mathbf{y}-\mathbf{t})\right\} \\
= & \left(\boldsymbol{\Phi}^{T} \mathbf{R} \boldsymbol{\Phi}\right)^{-1} \boldsymbol{\Phi}^{T} \mathbf{R} \mathbf{z} \\
& \quad \text { with } \mathbf{z}=\boldsymbol{\Phi} \mathbf{w}^{(\tau)}-\mathbf{R}^{-1}(\mathbf{y}-\mathbf{t})
\end{aligned}
$$

- Again very similar form (normal equations)
, But now with non-constant weighing matrix $\mathbf{R}$ (depends on $\mathbf{w}$ ).
, Need to apply normal equations iteratively.
$\Rightarrow$ Iteratively Reweighted Least-Squares (IRLS)


## Summary: Logistic Regression

- Properties
, Directly represent posterior distribution $p\left(\phi \mid \mathcal{C}_{k}\right)$
> Requires fewer parameters than modeling the likelihood + prior.
, Very often used in statistics.
> It can be shown that the cross-entropy error function is concave
- Optimization leads to unique minimum
- But no closed-form solution exists
- Iterative optimization (IRLS)
> Both online and batch optimizations exist
, There is a multi-class version described in (Bishop Ch.4.3.4).
- Caveat
, Logistic regression tends to systematically overestimate odds ratios when the sample size is less than $\sim 500$.


## Topics of This Lecture

- Recap: Generalized Linear Discriminants
- Logistic Regression
, Probabilistic discriminative models
, Logistic sigmoid (logit function)
, Cross-entropy error
> Gradient descent
, Iteratively Reweighted Least Squares
- Note on error functions
- Statistical Learning Theory
, Generalization and overfitting
, Empirical and actual risk
, VC dimension
, Empirical Risk Minimization
, Structural Risk Minimization


## A Note on Error Functions

$$
t_{n} \in\{-1,1\}
$$



Ideal misclassification error

Not differentiable!


- Ideal misclassification error function (black)
, This is what we want to approximate,
, Unfortunately, it is not differentiable.
, The gradient is zero for misclassified points.
$\Rightarrow$ We cannot minimize it by gradient descent.


## A Note on Error Functions



- Squared error used in Least-Squares Classification
, Very popular, leads to closed-form solutions.
> However, sensitive to outliers due to squared penalty.
, Penalizes "too correct" data points
$\Rightarrow$ Generally does not lead to good classifiers.


## A Note ${ }^{\text {on }}$ Error Functions



- Squared error with sigmoid activation function (tanh)
, Fixes the problems with outliers and "too correct" data points.
, But: zero gradient for confidently misclassified data points.
$\Rightarrow$ Will give better performance than original squared error, but still does not fix all problems.


## A Note on Error Functions



- Cross-Entropy Error
, Minimizer of this error is given by posterior class probabilities.
, Concave error function, unique minimum exists.
, Robust to outliers, error increases only roughly linearly
. But no closed-form solution, requires iterative estimation.


## Topics of This Lecture

- Recap: Generalized Linear Discriminants
- Logistic Regression

Probabilistic discriminative models
. Logistic sigmoid (logit function)
, Cross-entropy error
, Gradient descent
, Iteratively Reweighted Least Squares

- Note on error functions
- Statistical Learning Theory
, Generalization and overfitting
, Empirical and actual risk
, VC dimension
, Empirical Risk Minimization
, Structural Risk Minimization


## Generalization and Overfitting



- Goal: predict class labels of new observations
, Train classification model on limited training set.
> The further we optimize the model parameters, the more the training error will decrease.
, However, at some point the test error will go up again.
$\Rightarrow$ Overfitting to the training set!


## Example: Linearly Separable Data

- Overfitting is often a problem with linearly separable data
, Which of the many possible decision boundaries is correct?
, All of them have zero error on the training set...

, However, they will most likely result in different predictions on novel test data.
$\Rightarrow$ Different generalization performance
- How to select the classifier with the best generalization performance?


## A Broader View on Statistical Learning

- Formal treatment: Statistical Learning Theory
- Supervised learning
, Environment: assumed stationary.
, I.e. the data x have an unknown but fixed probability density

$$
p_{X}(\mathbf{x})
$$

, Teacher: specifies for each data point $\mathbf{x}$ the desired classification $y$ (where $y$ may be subject to noise).

$$
y=g(\mathbf{x}, \nu) \quad \text { with noise } \nu
$$

, Learning machine: represented by class of functions, which produce for each $\mathbf{x}$ an output $y$ :

$$
y=f(\mathbf{x} ; \alpha) \quad \text { with parameters } \alpha
$$

## Statistical Learning Theory

- Supervised learning (from the learning machine's view)
, Selection of a specific function $f(\mathbf{x} ; \alpha)$
, Given: training examples $\quad\left\{\left(\mathbf{x}_{i}, y_{i}\right)\right\}_{i=1}^{N}$
, Goal: the desired response $y$ shall be approximated optimally.
- Measuring the optimality
. Loss function

$$
L(y, f(\mathbf{x} ; \alpha))
$$

, Example: quadratic loss

$$
L(y, f(\mathbf{x} ; \alpha))=(y-f(\mathbf{x} ; \alpha))^{2}
$$

## Risk

- Measuring the "optimality"
, Measure the optimality by the risk (= expected loss).
, Difficulty: how should the risk be estimated?
- Practical way
, Empirical risk (measured on the training/validation set)

$$
R_{e m p}(\alpha)=\frac{1}{N} \sum_{i=1}^{N} L\left(y_{i}, f\left(\mathbf{x}_{i} ; \alpha\right)\right)
$$

, Example: quadratic loss function

$$
R_{e m p}(\alpha)=\frac{1}{N} \sum_{i=1}^{N}\left(y_{i}-f\left(\mathbf{x}_{i} ; \alpha\right)\right)^{2}
$$

## Risk

- However, what we're really interested in is
, Actual risk (= Expected risk)

$$
R(\alpha)=\int L(y, f(\mathbf{x} ; \alpha)) d P_{X, Y}(\mathbf{x}, y)
$$

, $P_{X, Y}(\mathbf{x}, y)$ is the probability distribution of $(\mathbf{x}, y)$.
, $P_{X, Y}(\mathbf{x}, y)$ is fixed, but typically unknown.
$\Rightarrow$ In general, we can't compute the actual risk directly!
. The expected risk is the expectation of the error on all data.
> l.e., it is the expected value of the generalization error.

## Summary: Risk

- Actual risk
, Advantage: measure for the generalization ability
, Disadvantage: in general, we don't know $P_{X, Y}(\mathbf{x}, y)$
- Empirical risk
, Disadvantage: no direct measure of the generalization ability
, Advantage: does not depend on $P_{X, Y}(\mathbf{x}, y)$
, We typically know learning algorithms which minimize the empirical risk.
$\Rightarrow$ Strong interest in connection between both types of risk


## Statistical Learning Theory

- Idea
, Compute an upper bound on the actual risk based on the empirical risk

$$
R(\alpha) \cdot R_{e m p}(\alpha)+\epsilon\left(N, p^{*}, h\right)
$$

, where
$N$ : number of training examples
$p^{*}$ : probability that the bound is correct
$h$ : capacity of the learning machine ("VC-dimension")

- Side note:
. (This idea of specifying a bound that only holds with a certain probability is explored in a branch of learning theory called "Probably Approximately Correct" or PAC Learning).


## VC Dimension

- Vapnik-Chervonenkis dimension
- Measure for the capacity of a learning machine.
- Formal definition:
, If a given set of $\ell$ points can be labeled in all possible $2^{\ell}$ ways, and for each labeling, a member of the set $\{f(\alpha)\}$ can be found which correctly assigns those labels, we say that the set of points is shattered by the set of functions.
, The VC dimension for the set of functions $\{f(\alpha)\}$ is defined as the maximum number of training points that can be shattered by $\{f(\alpha)\}$.


## VC Dimension

- Interpretation as a two-player game
, Opponent's turn: He says a number $N$.
, Our turn: We specify a set of $N$ points $\left\{\mathbf{x}_{1}, \ldots, \mathbf{x}_{N}\right\}$.
, Opponent's turn: He gives us a labeling $\left\{\mathbf{x}_{1}, \ldots, \mathbf{x}_{N}\right\} \in\{0,1\}^{N}$
, Our turn:
We specify a function $f(\alpha)$ which correctly classifies all $N$ points.
$\Rightarrow$ If we can do that for all $2^{N}$ possible labelings, then the VC dimension is at least $N$.


## VC Dimension

- Example
, The VC dimension of all oriented lines in $\mathbb{R}^{2}$ is 3.

1. Shattering 3 points with an oriented line:

2. More difficult to show: it is not possible to shatter 4 points (XOR)...
, More general: the VC dimension of all hyperplanes in $\mathbb{R}^{n}$ is $n+1$.

## VC Dimension

- Intuitive feeling (unfortunately wrong)
, The VC dimension has a direct connection with the number of parameters.
- Counterexample

$$
\begin{aligned}
f(x ; \alpha) & =g(\sin (\alpha x)) \\
g(x) & = \begin{cases}1, & x>0 \\
-1, & x \cdot\end{cases}
\end{aligned}
$$

, Just a single parameter $\alpha$.
, Infinite VC dimension

- Proof: Choose $\quad x_{i}=10^{-i}, \quad i=1, \ldots, \ell$

Slide adapted from Bernt Schiele

$$
\alpha=\pi\left(1+\sum_{i=1}^{\ell} \frac{\left(1-y_{i}\right) 10^{i}}{2}\right)
$$

## Upper Bound on the Risk

- Important result (Vapnik 1979, 1995)
, With probability $(1-\eta)$, the following bound holds

$$
R(\alpha) \cdot R_{e m p}(\alpha)+\underbrace{\sqrt{\frac{h(\log (2 N / h)+1)-\log (\eta / 4)}{N}}}_{\text {"VC confidence" }}
$$

, This bound is independent of $P_{X, Y}(\mathbf{x}, y)$ !
, Typically, we cannot compute the left-hand side (the actual risk)

- If we know $h$ (the VC dimension), we can however easily compute the risk bound

$$
R(\alpha) \cdot R_{e m p}(\alpha)+\epsilon\left(N, p^{*}, h\right)
$$

## Upper Bound on the Risk


B. Leibe

## Structural Risk Minimization

- How can we implement this?

$$
R(\alpha) \cdot \quad R_{e m p}(\alpha)+\epsilon\left(N, p^{*}, h\right)
$$

- Classic approach
, Keep $\epsilon\left(N, p^{*}, h\right)$ constant and minimize $R_{\text {emp }}(\alpha)$.
> $\epsilon\left(N, p^{*}, h\right)$ can be kept constant by controlling the model parameters.
- Support Vector Machines (SVMs)
, Keep $R_{e m p}(\alpha)$ constant and minimize $\epsilon\left(N, p^{*}, h\right)$.
, In fact: $R_{e m p}(\alpha)=0$ for separable data.
, Control $\epsilon\left(N, p^{*}, h\right)$ by adapting the VC dimension (controlling the "capacity" of the classifier).


## References and Further Reading

- More information on SVMs can be found in Chapter 7.1 of Bishop's book.

Christopher M. Bishop
Pattern Recognition and Machine Learning Springer, 2006

- Additional information about Statistical Learning Theory and a more in-depth introduction to SVMs are available in the following tutorial:
- C. Burges, A Tutorial on Support Vector Machines for Pattern Recognition, Data Mining and Knowledge Discovery, Vol. 2(2), pp. 121-167 1998.

