## Machine Learning - Lecture 8

## Linear Support Vector Machines

$$
12.05 .2015
$$

Bastian Leibe
RWTH Aachen
http://www.vision.rwth-aachen.de/
leibe@vision.rwth-aachen.de

Recap: 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!


## 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



## Recap: 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")

RWITAACHER

## Recap: 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 $\{f(\alpha)\}$.

## Recap: 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 $P_{X, Y}(\mathbf{x}, y)!\quad \begin{gathered}\text { Guarated rike } \\ \text { (bound on generliation }\end{gathered}$
If we know $h$ (the VC dimension), we can easily compute the risk bound

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



## RWIHAACHE

## Recap: Structural Risk Minimization

- How can we implement Structural Risk Minimization?

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

- 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 $R_{\text {emp }}(\alpha)$ constant and minimize $\epsilon\left(N, p^{*}, h\right)$.
, In fact: $R_{\text {emp }}(\alpha)=0$ for separable data.
- Control $\epsilon\left(N, p^{*}, h\right)$ by adapting the VC dimension (controlling the "capacity" of the classifier).


## Topics of This Lecture

- Linear Support Vector Machines
, Lagrangian (primal) formulation
, Dual formulation
- Discussion
- Linearly non-separable case

Soft-margin classification

- Updated formulation
- Nonlinear Support Vector Machines
- Nonlinear basis functions
, The Kernel trick
- Mercer's condition

Popular kernels

- Applications


## Support Vector Machine (SVM)

- Let's first consider linearly separable data
, $N$ training data points $\left\{\left(\mathbf{x}_{i}, y_{i}\right)\right\}_{i=1}^{N} \quad \mathbf{x}_{i} \in \mathbb{R}^{d}$
. Target values $\quad t_{i} \in\{-1,1\}$
- Hyperplane separating the data



## Support Vector Machine (SVM)

- Since the data is linearly separable, there exists a hyperplane with

$$
\begin{aligned}
& \mathbf{w}^{\mathrm{T}} \mathbf{x}_{n}+b \geq+1 \text { for } t_{n}=+1 \\
& \mathbf{w}^{\mathrm{T}} \mathbf{x}_{n}+b \cdot-1 \text { for } t_{n}=-1
\end{aligned}
$$

- Combined in one equation, this can be written as

$$
t_{n}\left(\mathbf{w}^{\mathrm{T}} \mathbf{x}_{n}+b\right) \geq 1 \quad \forall n
$$

$\Rightarrow$ Canonical representation of the decision hyperplane.

- The equation will hold exactly for the points on the margin

$$
t_{n}\left(\mathbf{w}^{\mathrm{T}} \mathbf{x}_{n}+b\right)=1
$$

, By definition, there will always be at least one such point.


## Support Vector Machine (SVM)

- We can choose $w$ such that

$$
\begin{array}{lll}
\mathbf{w}^{\mathrm{T}} \mathbf{x}_{n}+b=+1 & \text { for one } & t_{n}=+1 \\
\mathbf{w}^{\mathrm{T}} \mathbf{x}_{n}+b=-1 & \text { for one } & t_{n}=-1
\end{array}
$$

- The distance between those two hyperplanes is then the margin

$$
\begin{aligned}
& d_{-}=d_{+}=\frac{1}{\|\mathbf{w}\|} \\
& d_{-}+d_{+}=\frac{2}{\|\mathbf{w}\|}
\end{aligned}
$$

$\Rightarrow$ We can find the hyperplane with maximal margin by minimizing $\|\mathrm{w}\|^{2}$.
B. Leibe

Support Vector Machine (SVM)

- Optimization problem

Find the hyperplane satisfying

$$
\underset{\mathbf{w}, b}{\arg \min } \frac{1}{2}\|\mathbf{w}\|^{2}
$$

under the constraints

$$
t_{n}\left(\mathbf{w}^{\mathrm{T}} \mathbf{x}_{n}+b\right) \geq 1 \quad \forall n
$$

- Quadratic programming problem with linear constraints.
- Can be formulated using Lagrange multipliers.
- Who is already familiar with Lagrange multipliers? - Let's look at a real-life example...


## Recap: Lagrange Multipliers

- Problem
, We want to maximize $K(\mathbf{x})$ subject to constraints $f(\mathbf{x})=0$.
, Example: we want to get as close as possible, but there is a fence.
, How should we move?
$f(\mathbf{x})=$


We want to maximize $\nabla K$
But we can only move parallel to the fence, i.e. along
$\nabla_{\|} K=\nabla K+\lambda \nabla f$ with $\lambda \neq 0$.

Fence $f$
lide adanted from Mario Eritz

RWIIHAMEFE

## Recap: Lagrange Multipliers

- Problem
. Now let's look at constraints of the form $f(\mathbf{x}) \geq 0$
, Example: There might be a hill from which we can see better..
, Optimize $\max _{\mathbf{x}, \lambda} L(\mathbf{x}, \lambda)=K(\mathbf{x})+\lambda f(\mathbf{x})$

- Two cases

Solution lies on boundary $\Rightarrow f(\mathbf{x})=0$ for some $\lambda>0$

- Solution lies inside $f(\mathbf{x})>0$ $\Rightarrow$ Constraint inactive: $\lambda=0$
In both cases
$\Rightarrow \lambda f(\mathbf{x})=0$


## Recap: Lagrange Multipliers

- Problem
, We want to maximize $K(\mathbf{x})$ subject to constraints $f(\mathbf{x})=0$
Example: we want to get as close as possible, but there is a fence.
, How should we move?
$f(\mathbf{x})=0 \quad f(\mathbf{x})>0$ $\Rightarrow$ Optimize
$\max _{\mathbf{x}, \lambda} L(\mathbf{x}, \lambda)=K(\mathbf{x})+\lambda f(\mathbf{x})$
$\frac{\partial L}{\partial \mathbf{x}}=\nabla_{\|} K \stackrel{!}{=} 0$
$\frac{\partial L}{\partial \lambda}=f(x) \stackrel{!}{=} 0$
Fence $f$ 22


## Recap: Lagrange Multipliers

- Problem
- Now let's look at constraints of the form $f(\mathbf{x}) \geq 0$.
, Example: There might be a hill from which we can see better..
Optimize $\max _{\mathbf{x}, \lambda} L(\mathbf{x}, \lambda)=K(\mathbf{x})+\lambda f(\mathbf{x})$
- Two cases
, Solution lies on boundary $\Rightarrow f(\mathbf{x})=0$ for some $\lambda>0$
Solution lies inside $f(\mathbf{x})>0$
$\Rightarrow$ Constraint inactive: $\lambda=0$
In both cases
$\Rightarrow \lambda f(\mathbf{x})=0$
B. Leibe Fence $f \quad 24$


## SVM - Lagrangian Formulation

- Find hyperplane minimizing $\|\mathbf{w}\|^{2}$ under the constraints

$$
t_{n}\left(\mathbf{w}^{\mathrm{T}} \mathbf{x}_{n}+b\right)-1 \geq 0 \quad \forall n
$$

- Lagrangian formulation
- Introduce positive Lagrange multipliers: $a_{n} \geq 0 \quad \forall n$
, Minimize Lagrangian ("primal form")

$$
L(\mathbf{w}, b, \mathbf{a})=\frac{1}{2}\|\mathbf{w}\|^{2}-\sum_{n=1}^{N} a_{n}\left\{t_{n}\left(\mathbf{w}^{\mathrm{T}} \mathbf{x}_{n}+b\right)-1\right\}
$$

, I.e., find $\mathbf{w}, b$, and a such that

$$
\frac{\partial L}{\partial b}=0 \Rightarrow \sum_{n=1}^{N} a_{n} t_{n}=0 \quad \frac{\partial L}{\partial \mathbf{w}}=0 \Rightarrow \mathbf{w}=\sum_{n=1}^{N} a_{n} t_{n} \mathbf{x}_{n}
$$

## SVM - Lagrangian Formulation

- Lagrangian primal form

$$
\begin{aligned}
L_{p} & =\frac{1}{2}\|\mathbf{w}\|^{2}-\sum_{n=1}^{N} a_{n}\left\{t_{n}\left(\mathbf{w}^{\mathrm{T}} \mathbf{x}_{n}+b\right)-1\right\} \\
& =\frac{1}{2}\|\mathbf{w}\|^{2}-\sum_{n=1}^{N} a_{n}\left\{t_{n} y\left(\mathbf{x}_{n}\right)-1\right\}
\end{aligned}
$$

- The solution of $L_{p}$ needs to fulfill the KKT conditions
- Necessary and sufficient conditions

$$
\begin{array}{rl|}
a_{n} & \geq 0 \\
a_{n} y\left(\mathbf{x}_{n}\right)-1 & \geq 0 \\
a_{n}\left\{t_{n} y\left(\mathbf{x}_{n}\right)-1\right\} & =0 \\
\text { B. Leibe } & \\
\hline
\end{array}
$$

## SVM - Support Vectors

- The training points for which $a_{n}>0$ are called "support vectors".
- Graphical interpretation:
- Graphical interpretation:
, The support vectors are the
points on the margin.
, They define the margin
and thus the hyperplane.
$\Rightarrow$ Robustness to "too correct"
points!


## SVM - Discussion (Part 1)

- Linear SVM
, Linear classifier
- Approximative implementation of the SRM principle.
- In case of separable data, the SVM produces an empirical risk of zero with minimal value of the VC confidence
(i.e. a classifier minimizing the upper bound on the actual risk).
, SVMs thus have a "guaranteed" generalization capability.
, Formulation as convex optimization problem.
$\Rightarrow$ Globally optimal solution!


## - Primal form formulation

, Solution to quadratic prog. problem in $M$ variables is in $\mathcal{O}\left(M^{3}\right)$.
, Here: $D$ variables $\Rightarrow \mathcal{O}\left(D^{3}\right)$
, Problem: scaling with high-dim. data ("curse of dimensionality")

## SVM - Dual Formulation

- Improving the scaling behavior: rewrite $L_{p}$ in a dual form

$$
\begin{aligned}
L_{p} & =\frac{1}{2}\|\mathbf{w}\|^{2}-\sum_{n=1}^{N} a_{n}\left\{t_{n}\left(\mathbf{w}^{\mathrm{T}} \mathbf{x}_{n}+b\right)-1\right\} \\
& =\frac{1}{2}\|\mathbf{w}\|^{2}-\sum_{n=1}^{N} a_{n} t_{n} \mathbf{w}^{\mathrm{T}} \mathbf{x}_{n}-b \sum_{\neq 1}^{N} a_{n} t_{n}+\sum_{n=1}^{N} a_{n}
\end{aligned}
$$

- Using the constraint $\sum_{n=1}^{N} a_{n} t_{n}=0$, we obtain $\quad \frac{\partial L_{p}}{\partial b}=0$

$$
L_{p}=\frac{1}{2}\|\mathbf{w}\|^{2}-\sum_{n=1}^{N} a_{n} t_{n} \mathbf{w}^{\mathrm{T}} \mathbf{x}_{n}+\sum_{n=1}^{N} a_{n}
$$

## SVM - Dual Formulation

$$
L_{p}=\frac{1}{2}\|\mathbf{w}\|^{2}-\sum_{n=1}^{N} a_{n} t_{n} \mathbf{w}^{\mathrm{T}} \mathbf{x}_{n}+\sum_{n=1}^{N} a_{n}
$$

Using the constraint $\mathbf{w}=\sum_{n=1}^{N} a_{n} t_{n} \mathbf{x}_{n}$, we obtain $\frac{\partial L_{p}}{\partial \mathbf{w}}=0$

$$
\begin{aligned}
L_{p} & =\frac{1}{2}\|\mathbf{w}\|^{2}-\sum_{n=1}^{N} a_{n} t_{n} \sum_{m=1}^{N} a_{m} t_{m} \mathbf{x}_{m}^{\mathrm{T}} \mathbf{x}_{n}+\sum_{n=1}^{N} a_{n} \\
& =\frac{1}{2}\|\mathbf{w}\|^{2}-\sum_{n=1}^{N} \sum_{m=1}^{N} a_{n} a_{m} t_{n} t_{m}\left(\mathbf{x}_{m}^{\mathrm{T}} \mathbf{x}_{n}\right)+\sum_{n=1}^{N} a_{n}
\end{aligned}
$$

## SVM - Dual Formulation

- Maximize

$$
L_{d}(\mathbf{a})=\sum_{n=1}^{N} a_{n}-\frac{1}{2} \sum_{n=1}^{N} \sum_{m=1}^{N} a_{n} a_{m} t_{n} t_{m}\left(\mathbf{x}_{m}^{\mathrm{T}} \mathbf{x}_{n}\right)
$$

under the conditions

$$
\begin{aligned}
a_{n} & \geq 0 \quad \forall n \\
\sum_{n=1}^{N} a_{n} t_{n} & =0
\end{aligned}
$$

- The hyperplane is given by the $N_{S}$ support vectors:

$$
\mathbf{w}=\sum_{n=1}^{N_{\mathcal{S}}} a_{n} t_{n} \mathbf{x}_{n}
$$

## SVM - Discussion (Part 2)

- Dual form formulation
- In going to the dual, we now have a problem in $N$ variables $\left(a_{n}\right)$.
, Isn't this worse??? We penalize large training sets!
- However...

1. SVMs have sparse solutions: $a_{n} \neq 0$ only for support vectors!
$\Rightarrow$ This makes it possible to construct efficient algorithms

> e.g. Sequential Minimal Optimization (SMO)

Effective runtime between $\mathcal{O}(N)$ and $\mathcal{O}\left(N^{2}\right)$.
2. We have avoided the dependency on the dimensionality.
$\Rightarrow$ This makes it possible to work with infinite-dimensional feature spaces by using suitable basis functions $\phi(\mathbf{x})$.
$\Rightarrow$ We'll see that in a few minutes...


## SVM - Non-Separable Data

- Non-separable data
, I.e. the following inequalities cannot be satisfied for all data points

$$
\begin{array}{ll}
\mathbf{w}^{\mathrm{T}} \mathbf{x}_{n}+b \geq+1 & \text { for } t_{n}=+1 \\
\mathbf{w}^{\mathrm{T}} \mathbf{x}_{n}+b \cdot-1 & \text { for } \quad t_{n}=-1
\end{array}
$$

Instead use

$$
\begin{array}{ll}
\mathbf{w}^{\mathrm{T}} \mathbf{x}_{n}+b \geq+1-\xi_{n} & \text { for } \\
t_{n}=+1 \\
\mathbf{w}^{\mathrm{T}} \mathbf{x}_{n}+b \cdot-1+\xi_{n} & \text { for } \\
t_{n}=-1
\end{array}
$$

with "slack variables" $\xi_{n} \geq 0 \quad \forall n$

## SVM - Soft-Margin Classification

- Slack variables
- One slack variable $\xi_{n} \geq 0$ for each training data point.
- Interpretation
- $\xi_{n}=0$ for points that are on the correct side of the margin.
, $\xi_{n}=\left|t_{n}-y\left(\mathbf{x}_{n}\right)\right|$ for all other points (linear penalty).


Point on decision boundary: $\xi_{n}=1$

Misclassified point:
$\xi_{n}>1$

We do not have to set the slack variables ourselves!
$\Rightarrow$ They are jointly optimized together with w .
How that?

RWITMACHE

## SVM - Non-Separable Data

- Separable data
- Minimize
- Non-separable data
- Minimize

$\frac{1}{2}\|\mathbf{w}\|^{2} \quad$| Trade-off |
| :---: |
| parameter! |



RNIHAACHE
UNIVERSITY

## SVM - New Dual Formulation

- New SVM Dual: Maximize

$$
L_{d}(\mathbf{a})=\sum_{n=1}^{N} a_{n}-\frac{1}{2} \sum_{n=1}^{N} \sum_{m=1}^{N} a_{n} a_{m} t_{n} t_{m}\left(\mathbf{x}_{m}^{\mathrm{T}} \mathbf{x}_{n}\right)
$$

under the conditions
This is all that changed!

$$
\sum_{n=1}^{N} a_{n} t_{n}=0
$$

- This is again a quadratic programming problem $\Rightarrow$ Solve as before... (more on that later)


## SVM - New Solution

- Solution for the hyperplane
. Computed as a linear combination of the training examples

$$
\mathbf{w}=\sum_{n=1}^{N} a_{n} t_{n} \mathbf{x}_{n}
$$

- Again sparse solution: $a_{n}=0$ for points outside the margin.
$\Rightarrow$ The slack points with $\xi_{n}>0$ are now also support vectors!
, Compute $b$ by averaging over all $N_{\mathcal{M}}$ points with $0<a_{n}<C$ :

$$
b=\frac{1}{N_{\mathcal{M}}} \sum_{n \in \mathcal{M}}\left(t_{n}-\sum_{m \in \mathcal{M}} a_{m} t_{m} \mathbf{x}_{m}^{\mathrm{T}} \mathbf{x}_{n}\right)
$$

So Far...

- Only looked at linearly separable case Current problem formulation has no solution if the data are not linearly separable!
Need to introduce some tolerance to outlier data points.
$\Rightarrow$ Slack variables.

- Only looked at linear decision boundaries..
- This is not sufficient for many applications.

Want to generalize the ideas to non-linear boundaries.

## Interpretation of Support Vectors

- Those are the hard examples!

We can visualize them, e.g. for face detection



RWIIHAMEHE
UNIVERSITY
Another Example

- Separable by a surface in 3D



## Nonlinear SVM - Feature Spaces

- General idea: The original input space can be mapped to some higher-dimensional feature space where the training set is separable:


## Nonlinear SVM

- General idea
, Nonlinear transformation $\phi$ of the data points $\mathbf{x}_{n}$ :

$$
\mathbf{x} \in \mathbb{R}^{D} \quad \phi: \mathbb{R}^{D} \rightarrow \mathcal{H}
$$

, Hyperplane in higher-dim. space $\mathcal{H}$ (linear classifier in $\mathcal{H}$ )

$$
\mathbf{w}^{\mathrm{T}} \phi(\mathbf{x})+b=0
$$

$\Rightarrow$ Nonlinear classifier in $\mathbb{R}^{D}$.

## What Could This Look Like?

- Example:
- Mapping to polynomial space, $\mathbf{x}, \mathbf{y} \in \mathbb{R}^{2}$ :
$\phi(\mathbf{x})=\left[\begin{array}{cc}x_{1}^{2} \\ \sqrt{2} x_{1} x_{2} \\ x_{2}^{2}\end{array}\right]$
. Motivation: Easier to separate data in higher-dimensional space.
, But wait - isn't there a big problem?
How should we evaluate the decision function?


## Solution: The Kernel Trick

- Important observation
- $\phi(\mathbf{x})$ only appears in the form of dot products $\phi(\mathbf{x})^{\top} \phi(\mathbf{y})$ :

$$
\begin{aligned}
y(\mathbf{x}) & =\mathbf{w}^{\mathrm{T}} \phi(\mathbf{x})+b \\
& =\sum_{n=1}^{N} a_{n} t_{n} \phi\left(\mathbf{x}_{n}\right)^{\mathrm{T}} \phi(\mathbf{x})+b
\end{aligned}
$$

- Trick: Define a so-called kernel function $k(\mathbf{x}, \mathbf{y})=\phi(\mathbf{x})^{\top} \phi(\mathbf{y})$.
, Now, in place of the dot product, use the kernel instead:

$$
y(\mathbf{x})=\sum_{n=1}^{N} a_{n} t_{n} k\left(\mathbf{x}_{n}, \mathbf{x}\right)+b
$$

- The kernel function implicitly maps the data to the higherdimensional space (without having to compute $\phi(\mathbf{x})$ explicitly)!


## Back to Our Previous Example...

## Problem with High-dim. Basis Functions

- Problem
- In order to apply the SVM, we need to evaluate the function

$$
y(\mathbf{x})=\mathbf{w}^{\mathrm{T}} \phi(\mathbf{x})+b
$$

, Using the hyperplane, which is itself defined as

$$
\mathbf{w}=\sum_{n=1}^{N} a_{n} t_{n} \phi\left(\mathbf{x}_{n}\right)
$$

$\Rightarrow$ What happens if we try this for a million-dimensional feature space $\phi(\mathbf{x})$ ?
, Oh-oh...

- $2^{\text {nd }}$ degree polynomial kernel:

$$
\begin{aligned}
\phi(\mathbf{x})^{\mathrm{T}} \phi(\mathbf{y}) & =\left[\begin{array}{c}
x_{1}^{2} \\
\sqrt{2} x_{1} x_{2} \\
x_{2}^{2}
\end{array}\right] \cdot\left[\begin{array}{c}
y_{1}^{2} \\
\sqrt{2} y_{1} y_{2} \\
y_{2}^{2}
\end{array}\right] \substack{\begin{subarray}{c}{1 \\
0.8 \\
0.6 \\
0.2} }} \\
{0} \\
& =x_{1}^{2} y_{1}^{2}+2 x_{1} x_{2} y_{1} y_{2}+x_{2}^{2} y_{2}^{2} \\
& =\left(\mathbf{x}^{\mathrm{T}} \mathbf{y}\right)^{2}=: k(\mathbf{x}, \mathbf{y})
\end{aligned}
$$



Whenever we evaluate the kernel function $k(\mathbf{x}, \mathbf{y})=\left(\mathbf{x}^{\top} \mathbf{y}\right)^{2}$, we implicitly compute the dot product in the higher-dimensional feature space.

## SVMs with Kernels

- Using kernels
, Applying the kernel trick is easy. Just replace every dot product by a kernel function..

$$
\mathbf{x}^{\mathrm{T}} \mathbf{y} \quad \rightarrow \quad k(\mathbf{x}, \mathbf{y})
$$

- ...and we're done.
- Instead of the raw input space, we're now working in a higherdimensional (potentially infinite dimensional!) space, where the data is more easily separable.
"Sounds like magic..."
- Wait - does this always work?
- The kernel needs to define an implicit mapping to a higher-dimensional feature space $\phi(\mathbf{x})$.
When is this the case?
B. Leibe



## RWIHAACHE

## Recap: Kernels Fulfilling Mercer's Condition

- Polynomial kernel

$$
k(\mathbf{x}, \mathbf{y})=\left(\mathbf{x}^{\mathrm{T}} \mathbf{y}+1\right)^{p}
$$

- Radial Basis Function kernel

$$
k(\mathbf{x}, \mathbf{y})=\exp \left\{-\frac{(\mathbf{x}-\mathbf{y})^{2}}{2 \sigma^{2}}\right\}
$$

e.g. Gaussian

## - Hyperbolic tangent kernel



Actually, this was wrong in the original SVM paper...
(and many, many more...)
Slide credit: Bernt Schiele

## Which Functions are Valid Kernels?

- Mercer's theorem (modernized version):
, Every positive definite symmetric function is a kernel.
- Positive definite symmetric functions correspond to a positive definite symmetric Gram matrix:

$K=$| $k\left(\mathbf{x}_{1}, \mathbf{x}_{1}\right)$ | $k\left(\mathbf{x}_{1}, \mathbf{x}_{2}\right)$ | $k\left(\mathbf{x}_{1}, \mathbf{x}_{3}\right)$ | $\ldots$ | $k\left(\mathbf{x}_{1}, \mathbf{x}_{n}\right)$ |
| :---: | :---: | :---: | :---: | :---: |
| $k\left(\mathbf{x}_{2}, \mathbf{x}_{1}\right)$ | $k\left(\mathbf{x}_{2}, \mathbf{x}_{2}\right)$ | $k\left(\mathbf{x}_{2}, \mathbf{x}_{3}\right)$ |  | $k\left(\mathbf{x}_{2}, \mathbf{x}_{n}\right)$ |
|  |  |  |  |  |
| $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ |
| $k\left(\mathbf{x}_{n}, \mathbf{x}_{1}\right)$ | $k\left(\mathbf{x}_{n}, \mathbf{x}_{2}\right)$ | $k\left(\mathbf{x}_{n}, \mathbf{x}_{3}\right)$ | $\ldots$ | $k\left(\mathbf{x}_{n}, \mathbf{x}_{n}\right)$ |

(positive definite $=$ all eigenvalues are $>0$ )

Example: Bag of Visual Words Representation

- General framework in visual recognition
, Create a codebook (vocabulary) of prototypical image features
, Represent images as histograms over codebook activations
, Compare two images by any histogram kernel, e.g. $\chi^{2}$ kernel

$$
k_{\chi^{2}}\left(h, h^{\prime}\right)=\exp \left(-\frac{1}{\gamma} \sum_{j} \frac{\left(h_{j}-h_{j}^{\prime}\right)^{2}}{h_{j}+h_{j}^{\prime}}\right)
$$


$\downarrow$

$\downarrow$


Slide adapted from Christoon ampert
B. Leibe

Nonlinear SVM - Dual Formulation

- SVM Dual: Maximize

$$
L_{d}(\mathbf{a})=\sum_{n=1}^{N} a_{n}-\frac{1}{2} \sum_{n=1}^{N} \sum_{m=1}^{N} a_{n} a_{m} t_{n} t_{m} k\left(\mathbf{x}_{m}, \mathbf{x}_{n}\right)
$$

under the conditions

$$
\begin{gathered}
0 \cdot a_{n} \cdot C \\
\sum_{n=1}^{N} a_{n} t_{n}=0
\end{gathered}
$$

- Classify new data points using

$$
y(\mathbf{x})=\sum_{n=1}^{N} a_{n} t_{n} k\left(\mathbf{x}_{n}, \mathbf{x}\right)+b
$$

## VC Dimension for Gaussian RBF Kernel

- Radial Basis Function:

$$
k(\mathbf{x}, \mathbf{y})=\exp \left\{-\frac{(\mathbf{x}-\mathbf{y})^{2}}{2 \sigma^{2}}\right\}
$$

- In this case, $\mathcal{H}$ is infinite dimensional!

$$
\exp (\mathbf{x})=1+\frac{\mathbf{x}}{1!}+\frac{\mathbf{x}^{2}}{2!}+\ldots+\frac{\mathbf{x}^{n}}{n!}+\ldots
$$

- Since only the kernel function is used by the SVM, this is no problem.
, The hyperplane in $\mathcal{H}$ then has VC-dimension

$$
\operatorname{dim}(\mathcal{H})+1=\infty
$$

## VC Dimension for Gaussian RBF Kernel

- Intuitively
- If we make the radius of the RBF kernel sufficiently small, then each data point can be associated with its own kernel.
. However, this also means that we can get finite VC-dimension if we set a lower limit to the RBF radius.


## Example: RBF Kernels

- Decision boundary on toy problem


RBF Kernel width ( $\sigma$ )

- Vapnik has proven the following:
- The class of optimal linear separators has VC dimension $h$ bounded from above as

$$
\begin{aligned}
& \text { bove as } \\
& h \leq \min
\end{aligned}\left\{\left[\frac{D^{2}}{\rho^{2}}\right], m_{0}\right\}+1
$$

where $\rho$ is the margin, $D$ is the diameter of the smallest sphere that can enclose all of the training examples, and $m_{0}$ is the dimensionality.

- Intuitively, this implies that regardless of dimensionality $m_{0}$ we can minimize the VC dimension by maximizing the margin $\rho$.
- Thus, complexity of the classifier is kept small regardless of dimensionality.


## Summary: SVMs

## - Properties

- Empirically, SVMs work very, very well.
- SVMs are currently among the best performers for a number of classification tasks ranging from text to genomic data.
- SVMs can be applied to complex data types beyond feature vectors (e.g. graphs, sequences, relational data) by designing kernel functions for such data.
- SVM techniques have been applied to a variety of other tasks e.g. SV Regression, One-class SVMs, ...

The kernel trick has been used for a wide variety of applications. It can be applied wherever dot products are in use e.g. Kernel PCA, kernel FLD, ...

Good overview, software, and tutorials available on http://www.kernel-machines.org/

## Topics of This Lecture

- Linear Support Vector Machines (Recap)

Lagrangian (primal) formulation
Dual formulation
Discussion

- Linearly non-separable case

Soft-margin classification
Updated formulation

- Nonlinear Support Vector Machines

Nonlinear basis functions
The Kernel trick
Mercer's condition
Popular kernels

- Applications


## Summary: SVMs

- Limitations
- How to select the right kernel? Still something of a black art...
, How to select the kernel parameters?
(Massive) cross-validation.
Usually, several parameters are optimized together in a grid search.
. Solving the quadratic programming problem
Standard QP solvers do not perform too well on SVM task.
Dedicated methods have been developed for this, e.g. SMO.
- Speed of evaluation

Evaluating $y(\mathbf{x})$ scales linearly in the number of SVs.

- Too expensive if we have a large number of support vectors.
$\Rightarrow$ There are techniques to reduce the effective SV set.
, Training for very large datasets (millions of data points)
Stochastic gradient descent and other approximations can be used B. Leibe

- Results:

|  | Bayes | Rocchio | C4.5 | k-NN | SVM (poly)$\text { degree } d=$ |  |  |  |  | SVM (rbf) width $\gamma=$ |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  |  |  | 1 | 2 | 3 | 4 | 5 | 0.6 | 0.8 | 1.0 | 1.2 |
| earn | 95.9 | 96.1 | 96.1 | 97.3 | 98.2 | 98.4 | 98.5 | 98.4 | 98.3 | 98.5 | 98.5 | 98.4 | 98.3 |
| acq | 91.5 | 92.1 | 85.3 | 92.0 | 92.6 | 94.6 | 95.2 | 95.2 | 95.3 | 95.0 | 95.3 | 95.3 | 95.4 |
| money-fx | 62.9 | 67.6 | 69.4 | 78.2 | 66.9 | 72.5 | 75.4 | 74.9 | 76.2 | 74.0 | 75.4 | 76.3 | 75.9 |
| grain | 72.5 | 79.5 | 89.1 | 82.2 | 91.3 | 93.1 | 92.4 | 91.3 | 89.9 | 93.1 | 91.9 | 91.9 | 90.6 |
| crude | 81.0 | 81.5 | 75.5 | 85.7 | 86.0 | 87.3 | 88.6 | 88.9 | 87.8 | 88.9 | 89.0 | 88.9 | 88.2 |
| trade | 50.0 | 77.4 | 59.2 | 77.4 | 69.2 | 75.5 | 76.6 | 77.3 | 77.1 | 76.9 | 78.0 | 77.8 | 76.8 |
| interest | 58.0 | 72.5 | 49.1 | 74.0 | 69.8 | 63.3 | 67.9 | 73.1 | 76.2 | 74.4 | 75.0 | 76.2 | 76.1 |
| ship | 78.7 | 83.1 | 80.9 | 79.2 | 82.0 | 85.4 | 86.0 | 86.5 | 86.0 | 85.4 | 86.5 | 87.6 | 87.1 |
| wheat | 60.6 | 79.4 | 85.5 | 76.6 | 83.1 | 84.5 | 85.2 | 85.9 | 83.8 | 85.2 | 85.9 | 85.9 | 85.9 |
| corn | 47.3 | 62.2 | 87.7 | 77.9 | 86.0 | 86.5 | 85.3 | 85.7 | 83.9 | 85.1 | 85.7 | 85.7 | 84.5 |
| microavg. | 72.0 | 79.9 | 79.4 | 82.3 | $84.2$ | $\begin{aligned} & \|85.1\| \\ & \text { comb } \end{aligned}$ | $\begin{aligned} & 185.9 \\ & \text { bined: } \end{aligned}$ | $\begin{array}{\|l\|} \hline 86.2 \mid \\ \mathbf{8 6 . 0} \\ \hline \end{array}$ | 85.9 | $\begin{array}{\|r\|r\|} \hline 86.4 \\ \text { con } \end{array}$ | $\|86.5\|$ <br> mbine | $\begin{aligned} & \mid 86.3 \\ & \text { ed: } 86 \end{aligned}$ | $\begin{aligned} & \hline 86.2 \\ & 6.4 \end{aligned}$ |

## Example Application: OCR

- Handwritten digit recognition
, US Postal Service Database
Standard benchmark task for many learning algorithms

2601496757146371037314497 11027120 330193301029603510029012 $940525067240124 \leq 50299855$ $51012401832-70124 \times 24064$ $1161,1685712960015870189$. 11575 7.212579648327499516 99505200453622203242320 351211273133905388311 $13191419129192 \$ 1912014$ $1011915457368.2326414,8,4$

 (0, 4, $1.103047,3,2009979965$ $891042985 \geq 101422955460$
 01097075233197201351985 1075318518254389096 $12,25 \times 655605.54035405$ 1425510

## Historical Importance

- USPS benchmark
, 2.5\% error: human performance
- Different learning algorithms
, 16.2\% error: Decision tree (C4.5)
. 5.9\% error: (best) 2-layer Neural Network
- 5.1\% error: LeNet 1 - (massively hand-tuned) 5-layer network


## - Different SVMs

. $4.0 \%$ error: Polynomial kernel ( $p=3,274$ support vectors)
, 4.1\% error: Gaussian kernel ( $\sigma=0.3,291$ support vectors)

Example Application: OCR

- Results
. Almost no overfitting with higher-degree kernels.

| degree of <br> polynomial | dimensionality of <br> feature space | support <br> vectors | raw <br> error |
| :---: | :---: | :---: | :---: |
| 1 | 256 | 282 | 8.9 |
| 2 | $\approx 33000$ | 227 | 4.7 |
| 3 | $\approx 1 \times 10^{6}$ | 274 | 4.0 |
| 4 | $\approx 1 \times 10^{9}$ | 321 | 4.2 |
| 5 | $\approx 1 \times 10^{12}$ | 374 | 4.3 |
| 6 | $\approx 1 \times 10^{14}$ | 377 | 4.5 |
| 7 | $\approx 1 \times 10^{16}$ | 422 | 4.5 |

Example Application: Pedestrian Detection

N. Dalal, B. Triggs, Histograms of Oriented Gradients for Human Detection, CVPR 2005

## Many Other Applications

- Lots of other applications in all fields of technology
- OCR
, Text classification
Computer vision



## You Can Try It At Home...

- Lots of SVM software available, e.g.
, svmlight (http://svmlight.joachims.org/)
Command-line based interface
Source code available (in C)
Interfaces to Python, MATLAB, Perl, Java, DLL,...
- libsvm (http://www.csie.ntu.edu.tw/~cjlin/libsvm/)

Library for inclusion with own code
C++ and Java sources
Interfaces to Python, R, MATLAB, Perl, Ruby, Weka, C+ .NET,...

- Both include fast training and evaluation algorithms, support for multi-class SVMs, automated training and cross-validation, ..
$\Rightarrow$ Easy to apply to your own problems!

RWIHAACHE
UNIVERSIT

## References and Further Reading

- More information on SVMs can be found in Chapter 7.1 of Bishop's book. You can also look at Schölkopf \& Smola (some chapters available online).

|  | Christopher M. Bishop <br> Pattern Recognition and Machine Learning <br> Springer, 2006 |  |
| :---: | :---: | :---: |
|  | B. Schölkopf, A. Smola |  |
|  | Learning with Kernels MIT Press, 2002 |  |
|  | http://www.learning-with-kernels.org/ |  |

- A more in-depth introduction to SVMs is 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.

