# CS60020: Foundations of <br> Algorithm Design and Machine Learning <br> <br> Sourangshu Bhattacharya 

 <br> <br> Sourangshu Bhattacharya}

## BOOSTING

## Boosting

- Train classifiers (e.g. decision trees) in a sequence.
- A new classifier should focus on those cases which were incorrectly classified in the last round.
- Combine the classifiers by letting them vote on the final prediction (like bagging).
- Each classifier is "weak" but the ensemble is "strong."
- AdaBoost is a specific boosting method.


## Boosting Intuition

- We adaptively weigh each data case.
- Data cases which are wrongly classified get high weight (the algorithm will focus on them)
- Each boosting round learns a new (simple) classifier on the weighed dataset.
- These classifiers are weighed to combine them into a single powerful classifier.
- Classifiers that that obtain low training error rate have high weight.
- We stop by using monitoring a hold out set (cross-validation).


## Boosting in a Picture

boosting rounds


## Boosting

- Combining multiple "base" classifiers to come up with a "good" classifier.
- Base classifiers have to be "weak learners", accuracy > 50\%
- Base classifiers are trained on a weighted training dataset.
- Boosting involves sequentially learning $\alpha_{m}$ and $y_{m}(x)$.


## Adaboost

1. Initialize the data weighting coefficients $\left\{w_{n}\right\}$ by setting $w_{n}^{(1)}=1 / N$ for $n=1, \ldots, N$.
2. For $m=1, \ldots, M$ :
(a) Fit a classifier $y_{m}(\mathrm{x})$ to the training data by minimizing the weighted error function

$$
J_{m}=\sum_{n=1}^{N} w_{n}^{(m)} I\left(y_{m}\left(\mathrm{x}_{n}\right) \neq t_{n}\right)
$$

where $I\left(y_{m}\left(\mathrm{x}_{n}\right) \neq t_{n}\right)$ is the indicator function and equals 1 when $y_{m}\left(\mathrm{x}_{n}\right) \neq t_{n}$ and 0 otherwise.
(b) Evaluate the quantities

$$
\epsilon_{m}=\frac{\sum_{n=1}^{N} w_{n}^{(m)} I\left(y_{m}\left(\mathrm{x}_{n}\right) \neq t_{n}\right)}{\sum_{n=1}^{N} w_{n}^{(m)}}
$$

and then use these to evaluate

$$
\alpha_{m}=\ln \left\{\frac{1-\epsilon_{m}}{\epsilon_{m}}\right\} .
$$

## Adaboost (contd..)

(c) Update the data weighting coefficients

$$
w_{n}^{(m+1)}=w_{n}^{(m)} \exp \left\{\alpha_{m} I\left(y_{m}\left(\mathrm{x}_{n}\right) \neq t_{n}\right)\right\}
$$

3. Make predictions using the final model, which is given by

$$
Y_{M}(\mathrm{x})=\operatorname{sign}\left(\sum_{m=1}^{M} \alpha_{m} y_{m}(\mathrm{x})\right) .
$$

## And in animation



Original training set: equal weights to all training samples

## AdaBoost example

$$
\begin{aligned}
& \varepsilon=\text { error rate of classifier } \\
& \alpha=\text { weight of classifier }
\end{aligned}
$$

## ROUND 1



## AdaBoost example

## ROUND 2



## AdaBoost example

ROUND 3


## AdaBoost example



## Adaboost illustration








## Adaboost - Observations

- $\epsilon_{m}:$ weighted error $\in[0,0.5)$
- $\alpha_{m} \geq 0$
- $w_{i}^{m+1}$ is higher than $w_{i}^{m}$ by a factor (1-
$\left.\epsilon_{m}\right) / \epsilon_{m}$, when $i$ is misclassified.


## Adaboost - derivation

- Consider the error function:

$$
E=\sum_{n=1}^{N} \exp \left\{-t_{n} f_{m}\left(\mathrm{x}_{n}\right)\right\}
$$

- Where

$$
f_{m}(\mathrm{x})=\frac{1}{2} \sum_{l=1}^{m} \alpha_{l} y_{l}(\mathrm{x})
$$

- Goal: Minimize E w.r.t. $\alpha_{l}$ and $y_{l}(x)$, sequentially.


## Adaboost - derivation

- Minimize w.r.t. $\alpha_{m}$

$$
\begin{aligned}
E & =\sum_{n=1}^{N} \exp \left\{-t_{n} f_{m-1}\left(\mathrm{x}_{n}\right)-\frac{1}{2} t_{n} \alpha_{m} y_{m}\left(\mathrm{x}_{n}\right)\right\} \\
& =\sum_{n=1}^{N} w_{n}^{(m)} \exp \left\{-\frac{1}{2} t_{n} \alpha_{m} y_{m}\left(\mathrm{x}_{n}\right)\right\}
\end{aligned}
$$

- Let $\tau_{m}$ be the set of datapoints correctly classified by $y_{m}$.

$$
\begin{aligned}
E & =e^{-\alpha_{m} / 2} \sum_{n \in \mathcal{T}_{m}} w_{n}^{(m)}+e^{\alpha_{m} / 2} \sum_{n \in \mathcal{M}_{m}} w_{n}^{(m)} \\
& =\left(e^{\alpha_{m} / 2}-e^{-\alpha_{m} / 2}\right) \sum_{n=1}^{N} w_{n}^{(m)} I\left(y_{m}\left(\mathrm{x}_{n}\right) \neq t_{n}\right)+e^{-\alpha_{m} / 2} \sum_{n=1}^{N} w_{n}^{(m)} .
\end{aligned}
$$

## Adaboost - derivation

- Minimizing w.r.t. $y_{m}$ and $\alpha_{m}$, we get the updates 2(a) and 2(b).
- We can see that:
- Using:

$$
\begin{gathered}
w_{n}^{(m+1)}=w_{n}^{(m)} \exp \left\{-\frac{1}{2} t_{n} \alpha_{m} y_{m}\left(\mathrm{x}_{n}\right)\right\} . \\
t_{n} y_{m}\left(\mathrm{x}_{n}\right)=1-2 I\left(y_{m}\left(\mathrm{x}_{n}\right) \neq t_{n}\right)
\end{gathered}
$$

- We get:

$$
w_{n}^{(m+1)}=w_{n}^{(m)} \exp \left(-\alpha_{m} / 2\right) \exp \left\{\alpha_{m} I\left(y_{m}\left(\mathrm{x}_{n}\right) \neq t_{n}\right)\right\} .
$$

## Adaboost

1. Initialize the data weighting coefficients $\left\{w_{n}\right\}$ by setting $w_{n}^{(1)}=1 / N$ for $n=1, \ldots, N$.
2. For $m=1, \ldots, M$ :
(a) Fit a classifier $y_{m}(\mathrm{x})$ to the training data by minimizing the weighted error function

$$
J_{m}=\sum_{n=1}^{N} w_{n}^{(m)} I\left(y_{m}\left(\mathrm{x}_{n}\right) \neq t_{n}\right)
$$

where $I\left(y_{m}\left(\mathrm{x}_{n}\right) \neq t_{n}\right)$ is the indicator function and equals 1 when $y_{m}\left(\mathrm{x}_{n}\right) \neq t_{n}$ and 0 otherwise.
(b) Evaluate the quantities

$$
\epsilon_{m}=\frac{\sum_{n=1}^{N} w_{n}^{(m)} I\left(y_{m}\left(\mathrm{x}_{n}\right) \neq t_{n}\right)}{\sum_{n=1}^{N} w_{n}^{(m)}}
$$

and then use these to evaluate

$$
\alpha_{m}=\ln \left\{\frac{1-\epsilon_{m}}{\epsilon_{m}}\right\} .
$$

## Adaboost (contd..)

(c) Update the data weighting coefficients

$$
w_{n}^{(m+1)}=w_{n}^{(m)} \exp \left\{\alpha_{m} I\left(y_{m}\left(\mathrm{x}_{n}\right) \neq t_{n}\right)\right\}
$$

3. Make predictions using the final model, which is given by

$$
Y_{M}(\mathrm{x})=\operatorname{sign}\left(\sum_{m=1}^{M} \alpha_{m} y_{m}(\mathrm{x})\right) .
$$

## SUPPORT VECTOR MACHINES

## Support vector machines

- Let $\left\{x_{1}, \ldots, x_{n}\right\}$ be our data set and let $y_{i} \in\{1,-1\}$ be the class label of $x_{i}$

$$
\text { For } y_{i}=1 \quad w^{T} x_{i}+b \geq 1
$$

For $y_{i}=-1 \quad w^{T} x_{i}+b \leq-1$


## Large-margin Decision Boundary

- The decision boundary should be as far away from the data of both classes as possible



## Finding the Decision Boundary

- The decision boundary should classify all points correctly $\Rightarrow$

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

- The decision boundary can be found by solving the following constrained optimization problem

$$
\begin{aligned}
& \text { Minimize } \frac{1}{2}\|\mathrm{w}\|^{2} \\
& \text { subject to } y_{i}\left(\mathrm{w}^{T} \mathbf{x}_{i}+b\right) \geq 1
\end{aligned}
$$

- This is a constrained optimization problem. Solving it requires to use Lagrange multipliers


## KKT Conditions

- Problem:

$$
\min _{x} f(x) \text { sub. to: } \mathrm{g}_{\mathrm{i}}(\mathrm{x}) \leq 0 \forall i
$$

- Lagrangian: $L(x, \mu)=f(x)-\sum_{i} \mu_{i} g_{i}(x)$
- Conditions:
- Stationarity: $\nabla_{\mathrm{x}} \mathrm{L}(\mathrm{x}, \mu)=0$.
- Primal feasibility: $g_{i}(x) \leq 0 \quad \forall i$.
- Dual feasibility: $\mu_{i} \geq 0$.
- Complementary slackness: $\mu_{i} g_{i}(x)=0$.


## Finding the Decision Boundary

Minimize $\frac{1}{2}\|\mathrm{w}\|^{2}$
subject to $1-y_{i}\left(\mathrm{w}^{T} \mathbf{x}_{i}+b\right) \leq 0 \quad$ for $i=1, \ldots, n$

- The Lagrangian is

$$
\mathcal{L}=\frac{1}{2} \mathbf{w}^{T} \mathbf{w}+\sum_{i=1}^{n} \alpha_{i}\left(1-y_{i}\left(\mathbf{w}^{T} \mathbf{x}_{i}+b\right)\right)
$$

$-\alpha_{i} \geq 0$

- Note that $\|\mathbf{w}\|^{2}=\mathbf{w}^{\top} \mathbf{w}$


## The Dual Problem

- Setting the gradient $\mathcal{L}_{\text {if }}$ w.r.t. $\mathbf{w}$ and $b$ to zero, we have

$$
\begin{aligned}
& L=\frac{1}{2} w^{T} w+\sum_{i=1}^{n} \alpha_{i}\left(1-y_{i}\left(w^{T} x_{i}+b\right)\right)= \\
& =\frac{1}{2} \sum_{k=1}^{m} w^{k} w^{k}+\sum_{i=1}^{n} \alpha_{i}\left(1-y_{i}\left(\sum_{k=1}^{m} w^{k} x_{i}^{k}+b\right)\right)
\end{aligned}
$$

n : no of examples, m: dimension of the space

$$
\left\{\begin{aligned}
\frac{\partial L}{\partial w^{k}} & =0, \forall k \\
\frac{\partial L}{\partial b} & =0
\end{aligned}\right.
$$

## The Dual Problem

- If we substitute ${ }_{\mathrm{w}}=\sum^{n} \alpha_{i} y_{i} \mathrm{x}_{i}{ }^{\text {to }}{ }_{\mathcal{L}}$, we have

$$
\begin{aligned}
\mathcal{L} & =\frac{1}{2} \sum_{i=1}^{n} \alpha_{i} y_{i} \mathbf{x}_{i}^{T} \sum_{j=1}^{n} \alpha_{j} y_{j} \mathbf{x}_{j}+\sum_{i=1}^{n} \alpha_{i}\left(1-y_{i}\left(\sum_{j=1}^{n} \alpha_{j} y_{j} \mathbf{x}_{j}^{T} \mathbf{x}_{i}+b\right)\right) \\
& =\frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} \alpha_{i} \alpha_{j} y_{i} y_{j} \mathbf{x}_{i}^{T} \mathbf{x}_{j}+\sum_{i=1}^{n} \alpha_{i}-\sum_{i=1}^{n} \alpha_{i} y_{i} \sum_{j=1}^{n} \alpha_{j} y_{j} \mathbf{x}_{j}^{T} \mathbf{x}_{i}-b \sum_{i=1}^{n} \alpha_{i} y_{i} \\
& =-\frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} \alpha_{i} \alpha_{j} y_{i} y_{j} \mathbf{x}_{i}^{T} \mathbf{x}_{j}+\sum_{i=1}^{n} \alpha_{i}
\end{aligned}
$$

Since

$$
\sum_{i=1}^{n} \alpha_{i} y_{i}=0
$$

- This is a function of $\alpha_{i}$ only


## The Dual Problem

- The new objective function is in terms of $\alpha_{i}$ only
- It is known as the dual problem: if we know $\mathbf{w}$, we know all $\alpha_{i}$; if we know all $\alpha_{i}$, we know w
- The original problem is known as the primal problem
- The objective function of the dual problem needs to be maximized (comes out from the KKT theory)
- The dual problem is therefore:

$$
\sum_{i=1}^{n} \alpha_{i} y_{i}=0
$$

Properties of $\alpha_{i}$ when we introduce the Lagrange multipliers

The result when we differentiate the original Lagrangian w.r.t. b

## The Dual Problem

$\max . W(\boldsymbol{\alpha})=\sum_{i=1}^{n} \alpha_{i}-\frac{1}{2} \sum_{i=1, j=1}^{n} \alpha_{i} \alpha_{j} y_{i} y_{j} \mathbf{x}_{i}^{T} \mathbf{x}_{j}$
subject to $\alpha_{i} \geq 0, \sum_{i=1}^{n} \alpha_{i} y_{i}=0$

- This is a quadratic programming (QP) problem
- A global maximum of $\alpha_{i}$ can always be found
- $\mathbf{w}$ can be recovered by $\mathbf{w}=\sum_{i=1}^{n} \alpha_{i} y_{i} \mathbf{x}_{i}$


## Characteristics of the Solution

- Many of the $\alpha_{i}$ are zero
- Complementary slackness: $\alpha_{i}\left(1-y_{i}\left(w^{T} x_{i}+\right.\right.$ b)) $=0$
- Sparse representation: wis a linear combination of a small number of data points
- $\mathbf{x}_{\mathrm{i}}$ with non-zero $\alpha_{i}$ are called support vectors (SV)
- The deci
SV
- Let $t_{\mathrm{j}}(j=1, \ldots, s)$ be the indices of the $s$ support vectors. We can write


## A Geometrical Interpretation



## Characteristics of the Solution

- For testing with a new data $\mathbf{z}$
- Computi ${ }^{\mathbf{w}}{ }^{T} \mathbf{z}+b=\sum_{j=1}^{s} \alpha_{t_{j}} y_{t_{j}}\left(\mathbf{x}_{t_{j}}^{T} \mathbf{z}\right)+b$
classify $\mathbf{z}$ as class 1 if the sum is positive, and class 2 otherwise
- Note: w need not be formed explicitly


## Non-linearly Separable Problems

- We allow "error" $\xi_{i}$ in classification; it is based on the output of the discriminant function $\boldsymbol{w}^{\top} \boldsymbol{x}+b$
- $\xi_{\mathrm{i}}$ approximates the number of misclassified samples



## Soft Margin Hyperplane

- The new conditions become

$$
\begin{cases}\mathbf{w}^{T} \mathbf{x}_{i}+b \geq 1-\xi_{i} & y_{i}=1 \\ \mathbf{w}^{T} \mathbf{x}_{i}+b \leq-1+\xi_{i} & y_{i}=-1 \\ \xi_{i} \geq 0 & \forall i\end{cases}
$$

- $\xi_{\mathrm{i}}$ are "slack variables" in optimization
- Note that $\xi_{\mathrm{i}}=0$ if there is no error for $\mathbf{x}_{\mathrm{i}}$
- $\xi_{i}$ is an upper bound of the number of errors
- We want to minimize

$$
\frac{1}{2}\|w\|^{2}+C \sum_{i=1}^{n} \xi_{i}
$$

subject to $y_{i}\left(\mathbf{w}^{T} \mathbf{x}_{i}+b\right) \geq 1-\xi_{i}, \quad \xi_{i} \geq 0$

- $C$ : tradeott parameter between error āna margin


## The Optimization Problem

$$
L=\frac{1}{2} w^{T} w+C \sum_{i=1}^{n} \xi_{i}+\sum_{i=1}^{n} \alpha_{i}\left(1-\xi_{i}-y_{i}\left(w^{T} x_{i}+b\right)\right)-\sum_{i=1}^{n} \mu_{i} \xi_{i}
$$

With $a$ and $\mu$ Lagrange multipliers, POSITIVE

$$
\begin{aligned}
\frac{\partial L}{\partial w_{j}} & =w_{j}-\sum_{i=1}^{n} \alpha_{i} y_{i} x_{i j}=0 \\
\frac{\partial L}{\partial \xi_{j}} & =C-\alpha_{j}-\mu_{j}=0 \\
\frac{\partial L}{\partial b} & =\sum_{i=1}^{n} y_{i} \alpha_{i} y_{i} \vec{x}_{i}=0
\end{aligned}
$$

## The Dual Problem

$$
\begin{aligned}
& L=\frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} \alpha_{i} \alpha_{j} y_{i} y_{j} \vec{x}_{i}^{T} \vec{x}_{j}+C \sum_{i=1}^{n} \xi_{i}+ \\
& +\sum_{i=1}^{n} \alpha_{i}\left(1-\xi_{i}-y_{i}\left(\sum_{j=1}^{n} \alpha_{j} y_{j} x_{j}^{T} x_{i}+b\right)\right)-\sum_{i=1}^{n} \mu_{i} \xi_{i}
\end{aligned}
$$

With $\sum_{i=1}^{n} y_{i} \alpha_{i}=0 \quad$ and $\quad C=\alpha_{j}+\mu_{j}$

$$
L=-\frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} \alpha_{i} \alpha_{j} y_{i} y_{j} \vec{x}_{i}^{T} \vec{x}_{j}+\sum_{i=1}^{n} \alpha_{i}
$$

## The Optimization Problem

- The dual of this new constrained optimization problem is

$$
\begin{aligned}
& \max . W(\boldsymbol{\alpha})=\sum_{i=1}^{n} \alpha_{i}-\frac{1}{2} \sum_{i=1 . i=1}^{n} \alpha_{i} \alpha_{j} y_{i} y_{j} \mathbf{x}_{i}^{T} \mathbf{x}_{j} \\
& \text { subject to } C \geq \alpha_{i} \geq 0, \sum_{i=1}^{n} \alpha_{i} y_{i}=0
\end{aligned}
$$

- New constraints derived from $C=\alpha_{j}+\mu_{j}$ since $\mu$ and $\alpha$ are positive.
- $\mathbf{w}$ is recovered as $\mathbf{w}=\sum_{j=1}^{s} \alpha_{t_{j}} y_{t_{j}} \mathbf{x}_{t_{j}}$
- This is very similar to the optimization problem in the linear separable case, except that there is an upper bound $C$ on $\alpha_{i}$ now
- Once again, a QP solver can be used to find $\alpha_{i}$

$$
\frac{1}{2}\|w\|^{2}+C \sum_{i=1}^{n} \xi_{i}
$$

- The algorithm try to keep $\xi$ low, maximizing the margin
- The algorithm does not minimize the number of error. Instead, it minimizes the sum of distances from the hyperplane.
- When C increases the number of errors tend to lower. At the limit of $C$ tending to infinite, the solution tend to that given by the hard margin formulation, with 0 errors


## Soft margin is more robust to outliers




Soft Margin SVM
Hard Margin SVM

## Extension to Non-linear Decision <br> Boundary

- So far, we have only considered large-margin classifier with a linear decision boundary
- How to generalize it to become nonlinear?
- Key idea: transform $\mathbf{x}_{i}$ to a higher dimensional space to "make life easier"
- Input space: the space the point $\mathbf{x}_{i}$ are located
- Feature space: the space of $\phi\left(\mathbf{x}_{\mathrm{i}}\right)$ after transformation
- Why transform?
- Linear operation in the feature space is equivalent to non-linear operation in input space
- Classification can become easier with a proper transformation. In the XOR problem, for example, adding a new feature of $x_{1} x_{2}$ make the problem linearly separable


## Extension to Non-linear Decision <br> Boundary

- So far, we have only considered large-margin classifier with a linear decision boundary
- How to generalize it to become nonlinear?
- Key idea: transform $\mathbf{x}_{i}$ to a higher dimensional space to "make life easier"
- Input space: the space the point $\mathbf{x}_{i}$ are located
- Feature space: the space of $\phi\left(\mathbf{x}_{\mathrm{i}}\right)$ after transformation
- Why transform?
- Linear operation in the feature space is equivalent to non-linear operation in input space
- Classification can become easier with a proper transformation. In the XOR problem, for example, adding a new feature of $x_{1} x_{2}$ make the problem linearly separable


## XOR



| $\mathbf{X}$ | $\mathbf{Y}$ | $\mathbf{X Y}$ |  |
| :--- | :--- | :--- | :--- |
| 0 | 0 | 0 | 0 |
| 0 | 1 | 0 | 1 |
| 1 | 0 | 0 | 1 |
| 1 | 1 | 1 | 0 |

## Find a feature space




Note: feature space is of higher dimension than the input space in practice

- Computation in the feature space can be costly because it is high dimensional
- The feature space is typically infinite-dimensional!
- The kernel trick comes to rescue


## The Kernel Trick

- Recall the SVM optimization problem

$$
\begin{aligned}
& \text { max. } W(\boldsymbol{\alpha})=\sum_{i=1}^{n} \alpha_{i}-\frac{1}{2} \sum_{i=1, j=1}^{n} \alpha_{i} \alpha_{j} y_{i} y_{j} \mathbf{x}_{i}^{T} \mathbf{x}_{j} \\
& \text { subject to } C \geq \alpha_{i} \geq 0, \sum_{i=1}^{n} \alpha_{i} y_{i}=0
\end{aligned}
$$

- The data points only appear as inner product
- As long as we can calculate the inner product in the feature space, we do not need the mapping explicitly
- Many common geometric operations (angles, distances) can be expressed by inner products
- Define the kernel function $K$ by

$$
K\left(\mathbf{x}_{i}, \mathbf{x}_{j}\right)=\phi\left(\mathbf{x}_{i}\right)^{T} \phi\left(\mathbf{x}_{j}\right)
$$

## An Example for $\phi($.$) and K(.,$.

- Suppose $\phi($.$) is given as follows$

$$
\phi\left(\left[\begin{array}{l}
x_{1} \\
x_{2}
\end{array}\right]\right)=\left(1, \sqrt{2} x_{1}, \sqrt{2} x_{2}, x_{1}^{2}, x_{2}^{2}, \sqrt{2} x_{1} x_{2}\right)
$$

- An inner product in the feature space is

$$
\left\langle\phi\left(\left[\begin{array}{l}
x_{1} \\
x_{2}
\end{array}\right]\right), \phi\left(\left[\begin{array}{l}
y_{1} \\
y_{2}
\end{array}\right]\right)\right\rangle=\left(1+x_{1} y_{1}+x_{2} y_{2}\right)^{2}
$$

- So, if we define the kernel function as follows, there is no need to carry out $\phi($.$) explicitly$

$$
K(\mathbf{x}, \mathbf{y})=\left(1+x_{1} y_{1}+x_{2} y_{2}\right)^{2}
$$

- This use of kernel function to avoid carrying out $\phi$ (.) explicitly is known as the kernel trick


## Kernels

- Given a mapping: $\mathbf{x} \rightarrow \boldsymbol{\varphi}(\mathbf{x})$
a kernel is represented as the inner product

$$
K(\mathbf{x}, \mathbf{y}) \rightarrow \sum_{i} \varphi_{i}(\mathbf{x}) \varphi_{i}(\mathbf{y})
$$

A kernel must satisfy the Mercer's condition:

$$
\forall g(\mathbf{x}) \int K(\mathbf{x}, \mathbf{y}) g(\mathbf{x}) g(\mathbf{y}) d \mathbf{x} d \mathbf{y} \geq 0
$$

## Modification Due to Kernel Function

- Change all inner products to kernel functions
- For training,

Original

$$
\begin{aligned}
& \max . W(\boldsymbol{\alpha})=\sum_{i=1}^{n} \alpha_{i}-\frac{1}{2} \sum_{i=1, j=1}^{n} \alpha_{i} \alpha_{j} y_{i} y_{j} \mathbf{x}_{i}^{T} \mathbf{x}_{j} \\
& \text { subject to } C \geq \alpha_{i} \geq 0, \sum_{i=1}^{n} \alpha_{i} y_{i}=0
\end{aligned}
$$

With kernel max. $W(\boldsymbol{\alpha})=\sum_{i=1}^{n} \alpha_{i}-\frac{1}{2} \sum_{i=1, j=1}^{n} \alpha_{i} \alpha_{j} y_{i} y_{j} K\left(\mathbf{x}_{i}, \mathbf{x}_{j}\right)$
function
subject to $C \geq \alpha_{i} \geq 0, \sum_{i=1}^{n} \alpha_{i} y_{i}=0$

## Modification Due to Kernel Function

- For testing, the new data $\mathbf{z}$ is classified as class 1 if $f \geq 0$, and as class 2 if $f<0$

Original

$$
\begin{aligned}
\mathbf{w} & =\sum_{j=1}^{s} \alpha_{t_{j}} y_{t_{j}} \mathbf{x}_{t_{j}} \\
f & =\mathbf{w}^{\prime} \mathbf{z}+b=\sum_{j=1}^{s} \alpha_{t_{j}} y_{t_{j}} \mathbf{x}_{t_{j}}^{T} \mathbf{z}+b
\end{aligned}
$$

$$
\mathrm{w}=\sum_{j=1}^{s} \alpha_{t_{j}} y_{t_{j}} \phi\left(\mathrm{x}_{t_{j}}\right)
$$

$$
f=\langle\mathbf{w}, \phi(\mathbf{z})\rangle+b=\sum_{j=1}^{s} \alpha_{t_{j}} y_{t_{j}} K\left(\mathbf{x}_{t_{j}}, \mathbf{z}\right)+b
$$

## More on Kernel Functions

- Since the training of SVM only requires the value of $K\left(\mathbf{x}_{i}, \mathbf{x}_{\mathrm{j}}\right)$, there is no restriction of the form of $\mathbf{x}_{i}$ and $\mathbf{x}_{\mathrm{j}}$
$-x_{i}$ can be a sequence or a tree, instead of a feature vector
- $K\left(\mathbf{x}_{\mathrm{i}}, \mathbf{x}_{\mathrm{j}}\right)$ is just a similarity measure comparing $\mathbf{x}_{\mathrm{i}}$ and $\mathbf{x}_{\mathrm{j}}$
- For a test object $\mathbf{z}$, the discriminant function essentially is a weighted sum of the similarity between $z$ and a pre-selected set of objects (the support vectors)

$$
f(\mathbf{z})=\sum_{\mathbf{x}_{i} \in \mathcal{S}} \alpha_{i} y_{i} K\left(\mathbf{z}, \mathbf{x}_{i}\right)+b
$$

$\mathcal{S}$ : the set of support vectors

## Kernel Functions

- In practical use of SVM, the user specifies the kernel function; the transformation $\phi($.$) is not explicitly stated$
- Given a kernel function $K\left(\mathbf{x}_{\mathrm{i}}, \mathbf{x}_{\mathrm{j}}\right)$, the transformation $\phi($. is given by its eigenfunctions (a concept in functional analysis)
- Eigenfunctions can be difficult to construct explicitly
- This is why people only specify the kernel function without worrying about the exact transformation
- Another view: kernel function, being an inner product, is really a similarity measure between the objects


## A kernel is associated to a transformation

-Given a kernel, in principle it should be recovered the transformation in the feature space that originates it.
$-K(x, y)=(x y+1)^{2}=x^{2} y^{2}+2 x y+1$

It corresponds the transformation

$$
x \rightarrow\left(\begin{array}{c}
x^{2} \\
\sqrt{2} x \\
1
\end{array}\right)
$$

## Examples of Kernel Functions

- Polynomial kernel of degree $d$

$$
K(\mathbf{u}, \mathbf{v})=(\mathbf{u} \cdot \mathbf{v})^{d}
$$

- Polynomial kernel up to degree $d$

$$
K(\mathbf{x}, \mathbf{y})=\left(\mathbf{x}^{T} \mathbf{y}+1\right)^{d}
$$

- Radial basis function kernel with width $\sigma$

$$
K(\mathbf{x}, \mathbf{y})=\exp \left(-\|\mathbf{x}-\mathbf{y}\|^{2} /\left(2 \sigma^{2}\right)\right)
$$

- The feature space is infinite-dimensional
- Sigmoid with parameter $\kappa$ and $\theta$

$$
K(\mathbf{x}, \mathbf{y})=\tanh \left(\kappa \mathbf{x}^{T} \mathbf{y}+\theta\right)
$$

- It does not satisfy the Mercer condition on all $\kappa$ and $\theta$


## Building new kernels

- If $k_{1}(x, y)$ and $k_{2}(x, y)$ are two valid kernels then the following kernels are valid
- Linear Combination

$$
k(x, y)=c_{1} k_{1}(x, y)+c_{2} k_{2}(x, y)
$$

- Exponential

$$
k(x, y)=\exp \left[k_{1}(x, y)\right]
$$

- Product

$$
k(x, y)=k_{1}(x, y) \cdot k_{2}(x, y)
$$

- Polynomial transformation ( $Q$ : polynomial with non negative coeffcients)

$$
k(x, y)=Q\left[k_{1}(x, y)\right]
$$

- Function product (f: any function)

$$
k(x, y)=f(x) k_{1}(x, y) f(y)
$$

## Polynomial kernel



Ben-Hur et al, PLOS computational Biology 4 (2008)

## Gaussian RBF kernel



Ben-Hur et al, PLOS computational Biology 4 (2008)

