

CS434a/541a: Pattern Recognition
Prof. Olga Veksler

Lecture 11
Support Vector Machines

SVM

- Said to start in 1979 with Vladimir Vapnik's paper
- Major developments throughout 1990's
- Elegant theory
 - Has good generalization properties
- Have been applied to diverse problems very successfully in the last 10-15 years
- One of the most important developments in pattern recognition in the last 10 years



Today


- Support Vector Machines (SVM)
 - Introduction
 - Linear Discriminant
 - Linearly Separable Case
 - Linearly Non Separable Case
 - Kernel Trick
 - Non Linear Discriminant

Linear Discriminant Functions

- A discriminant function is linear if it can be written as

$$g(x) = w^T x + w_0$$


$$\begin{aligned} g(x) > 0 &\Rightarrow x \in \text{class 1} \\ g(x) < 0 &\Rightarrow x \in \text{class 2} \end{aligned}$$



- which separating hyperplane should we choose?

Linear Discriminant Functions


- Training data is just a subset of all possible data
- Suppose hyperplane is close to sample x_i
- If we see new sample close to sample i , it is likely to be on the wrong side of the hyperplane



- Poor generalization (performance on unseen data)

SVM


- Idea: maximize distance to the closest example



- For the optimal hyperplane
 - distance to the closest negative example = distance to the closest positive example

Linear Discriminant Functions


- Hyperplane as far as possible from any sample



- New samples close to the old samples will be classified correctly
- Good generalization


SVM: Linearly Separable Case

- SVM: maximize the **margin**



- margin** is twice the absolute value of distance b of the closest example to the separating hyperplane
- Better generalization (performance on test data)
 - in practice
 - and in theory

SVM: Linearly Separable Case



- Support vectors are the samples closest to the separating hyperplane
 - they are the most difficult patterns to classify
 - Optimal hyperplane is completely defined by support vectors
 - of course, we do not know which samples are support vectors without finding the optimal hyperplane


SVM: Formula for the Margin

- For uniqueness, set $|w^T x_i + w_0| = 1$ for any example x_i closest to the boundary
- now distance from closest sample x_i to $g(x) = 0$ is

$$\frac{|w^T x_i + w_0|}{\|w\|} = \frac{1}{\|w\|}$$

- Thus the margin is


$$m = \frac{2}{\|w\|}$$



SVM: Formula for the Margin

- $g(x) = w^T x + w_0$
- absolute distance between x and the boundary $g(x) = 0$

$$\frac{|w^T x + w_0|}{\|w\|}$$



- distance is unchanged for hyperplane $g_1(x) = \alpha g(x)$

$$\frac{|\alpha w^T x + \alpha w_0|}{\|\alpha w\|} = \frac{|w^T x + w_0|}{\|w\|}$$

- Let x_i be an example closest to the boundary. Set $|w^T x_i + w_0| = 1$
- Now the largest margin hyperplane is unique

SVM: Optimal Hyperplane

- Maximize margin $m = \frac{2}{\|w\|}$
- subject to constraints

$$\begin{cases} w^T x_i + w_0 \geq 1 & \text{if } x_i \text{ is positive example} \\ w^T x_i + w_0 \leq -1 & \text{if } x_i \text{ is negative example} \end{cases}$$

- Let $\begin{cases} z_i = 1 & \text{if } x_i \text{ is positive example} \\ z_i = -1 & \text{if } x_i \text{ is negative example} \end{cases}$

- Can convert our problem to

$$\begin{aligned} &\text{minimize } J(w) = \frac{1}{2} \|w\|^2 \\ &\text{constrained to } z_i (w^T x_i + w_0) \geq 1 \quad \forall i \end{aligned}$$

- $J(w)$ is a quadratic function, thus there is a single global minimum

SVM: Optimal Hyperplane

- Use Kuhn-Tucker theorem to convert our problem to:

$$\begin{aligned} \text{maximize } L_D(\alpha) &= \sum_{i=1}^n \alpha_i - \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j z_i z_j x_i^t x_j \\ \text{constrained to } \alpha_i &\geq 0 \quad \forall i \quad \text{and} \quad \sum_{i=1}^n \alpha_i z_i = 0 \end{aligned}$$

- $\alpha = \{\alpha_1, \dots, \alpha_n\}$ are new variables, one for each sample
- Can rewrite $L_D(\alpha)$ using n by n matrix H :

$$L_D(\alpha) = \sum_{i=1}^n \alpha_i - \frac{1}{2} \begin{bmatrix} \alpha_1 \\ \vdots \\ \alpha_n \end{bmatrix}^t H \begin{bmatrix} \alpha_1 \\ \vdots \\ \alpha_n \end{bmatrix}$$

- where the value in the i th row and j th column of H is $H_{ij} = z_i z_j x_i^t x_j$

SVM: Optimal Hyperplane

- After finding the optimal $\alpha = \{\alpha_1, \dots, \alpha_n\}$
 - For every sample i , one of the following must hold
 - $\alpha_i = 0$ (sample i is not a support vector)
 - $\alpha_i \neq 0$ and $z_i(w^t x_i + w_0 - 1) = 0$ (sample i is support vector)
 - can find w using $w = \sum_{i=1}^n \alpha_i z_i x_i$
 - can solve for w_0 using any $\alpha_i > 0$ and $\alpha_i [z_i(w^t x_i + w_0) - 1] = 0$

$$w_0 = \frac{1}{z_i} - w^t x_i$$
- Final discriminant function:

$$g(x) = \left(\sum_{i \in S} \alpha_i z_i x_i \right)^t x + w_0$$

- where S is the set of support vectors

$$S = \{x_i \mid \alpha_i \neq 0\}$$

SVM: Optimal Hyperplane

- Use Kuhn-Tucker theorem to convert our problem to:

$$\begin{aligned} \text{maximize } L_D(\alpha) &= \sum_{i=1}^n \alpha_i - \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j z_i z_j x_i^t x_j \\ \text{constrained to } \alpha_i &\geq 0 \quad \forall i \quad \text{and} \quad \sum_{i=1}^n \alpha_i z_i = 0 \end{aligned}$$

- $\alpha = \{\alpha_1, \dots, \alpha_n\}$ are new variables, one for each sample
- $L_D(\alpha)$ can be optimized by quadratic programming
- $L_D(\alpha)$ formulated in terms of α
 - it depends on w and w_0 indirectly

SVM: Optimal Hyperplane

$$\begin{aligned} \text{maximize } L_D(\alpha) &= \sum_{i=1}^n \alpha_i - \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j z_i z_j x_i^t x_j \\ \text{constrained to } \alpha_i &\geq 0 \quad \forall i \quad \text{and} \quad \sum_{i=1}^n \alpha_i z_i = 0 \end{aligned}$$

- $L_D(\alpha)$ depends on the number of samples, not on dimension of samples
- samples appear only through the dot products $x_i^t x_j$
- This will become important when looking for a **nonlinear** discriminant function, as we will see soon

SVM: Example using Matlab


- Class 1: [1,6], [1,10], [4,11]
- Class 2: [5,2], [7,6], [10,4]
- Let's pile all data into array X

$$X = \begin{bmatrix} 1 & 6 \\ 1 & 10 \\ 4 & 11 \\ 5 & 2 \\ 7 & 6 \\ 10 & 4 \end{bmatrix}$$

- Pile z_i 's into vector $z = \begin{bmatrix} 1 \\ 1 \\ -1 \\ -1 \\ -1 \\ -1 \end{bmatrix}$

- Matrix H with $H_{ij} = z_i z_j x_i^t x_j$, in matlab use $H = (x * x') .* (z * z')$

$$H = \begin{bmatrix} 37 & 61 & 70 & -17 & -43 & -34 \\ 61 & 101 & 114 & -25 & -67 & -50 \\ 70 & 114 & 137 & -42 & -84 & -64 \\ -17 & -25 & -42 & 29 & 47 & 58 \\ -43 & -67 & -84 & 47 & 85 & 94 \\ -34 & -50 & -64 & 58 & 94 & 116 \end{bmatrix}$$



SVM: Example using Matlab

- Multiply by -1 to convert to minimization:
minimize $L_D(\alpha) = -\sum_{i=1}^n \alpha_i + \frac{1}{2} \alpha' H \alpha$

- Let $f = \begin{bmatrix} -1 \\ \vdots \\ -1 \end{bmatrix} = -\mathbf{ones}(6,1)$, then can write

$$\text{minimize } L_D(\alpha) = f' \alpha + \frac{1}{2} \alpha' H \alpha$$

- First constraint is $\alpha_i \geq 0 \quad \forall i$

- Let $A = \begin{bmatrix} -1 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & -1 \end{bmatrix} = -\mathbf{eye}(6)$, $a = \begin{bmatrix} 0 \\ \vdots \\ 0 \end{bmatrix} = \mathbf{zeros}(6,1)$

- Rewrite the first constraint in canonical form:
 $A \alpha \leq a$

SVM: Example using Matlab

- Matlab expects quadratic programming to be stated in the *canonical* (standard) form which is

$$\begin{array}{l} \text{minimize } L_D(\alpha) = 0.5 \alpha' H \alpha + f' \alpha \\ \text{constrained to } A \alpha \leq a \text{ and } B \alpha = b \end{array}$$

- where A, B, H are matrices and f, a, b are vectors

- Need to convert our optimization problem to canonical form

$$\begin{array}{l} \text{maximize } L_D(\alpha) = \sum_{i=1}^n \alpha_i - \frac{1}{2} \begin{bmatrix} \alpha_1 \\ \vdots \\ \alpha_n \end{bmatrix}' H \begin{bmatrix} \alpha_1 \\ \vdots \\ \alpha_n \end{bmatrix} \\ \text{constrained to } \alpha_i \geq 0 \quad \forall i \text{ and } \sum_{i=1}^n \alpha_i z_i = 0 \end{array}$$

SVM: Example using Matlab

- Our second constraint is $\sum_{i=1}^n \alpha_i z_i = 0$

- Let $B = [z_1 \ z_2 \ z_3 \ z_4 \ z_5 \ z_6] = Z'$
and $b = 0$

- Second constraint in canonical form is:


$$B \alpha = b$$

- Thus our problem is in canonical form and can be solved by matlab:

$$\begin{array}{l} \text{minimize } L_D(\alpha) = 0.5 \alpha' H \alpha + f' \alpha \\ \text{constrained to } A \alpha \leq a \text{ and } B \alpha = b \end{array}$$

SVM: Example using Matlab

- $\alpha = \text{quadprog}(H + \text{eye}(6) * 0.001, f, A, a, B, b)$
for stability



- Solution $\alpha = \begin{bmatrix} 0.036 \\ 0 \\ 0.039 \\ 0 \\ 0.076 \\ 0 \end{bmatrix}$
- find w using $w = \sum_{i=1}^n \alpha_i z_i x_i = (\alpha * z)^T x = \begin{bmatrix} -0.33 \\ 0.20 \end{bmatrix}$
- since $\alpha_1 > 0$, can find w_0 using


$$w_0 = \frac{1}{z_1} - w^T x_1 = 0.13$$

SVM: Non Separable Case

- Use nonnegative "slack" variables ξ_1, \dots, ξ_n (one for each sample)
- Change constraints from $z_i(w^T x_i + w_0) \geq 1 \quad \forall i$ to $z_i(w^T x_i + w_0) \geq 1 - \xi_i \quad \forall i$


ξ_i is a measure of deviation from the ideal for sample i

- $\xi_i > 1$ sample i is on the wrong side of the separating hyperplane
- $0 < \xi_i < 1$ sample i is on the right side of separating hyperplane but within the region of maximum margin



SVM: Non Separable Case

- Data is most likely to be not linearly separable, but linear classifier may still be appropriate



- Can apply SVM in non linearly separable case
 - data should be "almost" linearly separable for good performance

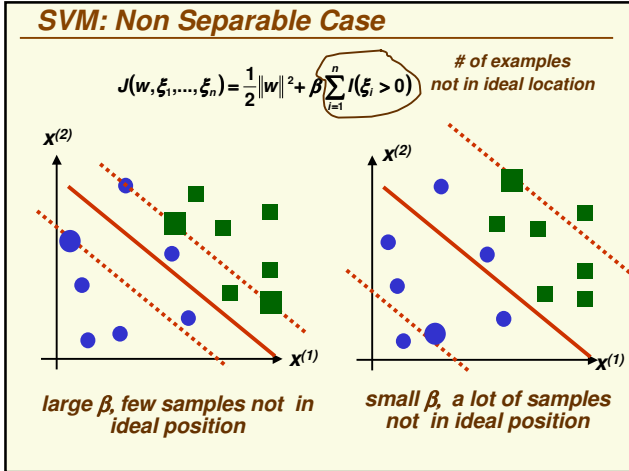
SVM: Non Separable Case

- Would like to minimize

$$J(w, \xi_1, \dots, \xi_n) = \frac{1}{2} \|w\|^2 + \beta \sum_{i=1}^n I(\xi_i > 0)$$

of samples not in ideal location

- where $I(\xi_i > 0) = \begin{cases} 1 & \text{if } \xi_i > 0 \\ 0 & \text{if } \xi_i \leq 0 \end{cases}$
- constrained to $z_i(w^T x_i + w_0) \geq 1 - \xi_i$ and $\xi_i \geq 0 \quad \forall i$
- β is a constant which measures relative weight of the first and second terms
 - if β is small, we allow a lot of samples not in ideal position
 - if β is large, we want to have very few samples not in ideal position



SVM: Non Separable Case

- Instead we minimize

$$J(w, \xi_1, \dots, \xi_n) = \frac{1}{2} \|w\|^2 + \beta \sum_{i=1}^n \xi_i$$

a measure of # of misclassified examples
- constrained to

$$\begin{cases} z_i(w'x_i + w_0) \geq 1 - \xi_i & \forall i \\ \xi_i \geq 0 & \forall i \end{cases}$$
- Can use Kuhn-Tucker theorem to converted to

$$\begin{aligned} &\text{maximize} && L_0(\alpha) = \sum_{i=1}^n \alpha_i - \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j z_i z_j x_i' x_j \\ &\text{constrained to} && 0 \leq \alpha_i \leq \beta \quad \forall i \quad \text{and} \quad \sum_{i=1}^n \alpha_i z_i = 0 \end{aligned}$$
- find w using

$$w = \sum_{i=1}^n \alpha_i z_i x_i$$
- solve for w_0 using any $0 < \alpha_i < \beta$ and $\alpha_i [z_i(w'x_i + w_0) - 1] = 0$

SVM: Non Separable Case

- Unfortunately this minimization problem is NP-hard due to discontinuity of functions $I(\xi_i)$

$$J(w, \xi_1, \dots, \xi_n) = \frac{1}{2} \|w\|^2 + \beta \sum_{i=1}^n I(\xi_i > 0)$$

of examples not in ideal location


- where $I(\xi_i > 0) = \begin{cases} 1 & \text{if } \xi_i > 0 \\ 0 & \text{if } \xi_i \leq 0 \end{cases}$
- constrained to $z_i(w'x_i + w_0) \geq 1 - \xi_i$ and $\xi_i \geq 0 \quad \forall i$

Non Linear Mapping

- Cover's theorem:
 - "pattern-classification problem cast in a high dimensional space non-linearly is more likely to be linearly separable than in a low-dimensional space"
- One dimensional space, not linearly separable
- Lift to two dimensional space with $\phi(x) = (x, x^2)$

Non Linear Mapping

- To solve a non linear classification problem with a linear classifier
 - Project data \mathbf{x} to high dimension using function $\phi(\mathbf{x})$
 - Find a linear discriminant function for transformed data $\phi(\mathbf{x})$
 - Final nonlinear discriminant function is $g(\mathbf{x}) = \mathbf{w}^t \phi(\mathbf{x}) + w_0$



In 2D, discriminant function is linear


$$g\left(\begin{bmatrix} x^{(1)} \\ x^{(2)} \end{bmatrix}\right) = [w_1 \ w_2] \begin{bmatrix} x^{(1)} \\ x^{(2)} \end{bmatrix} + w_0$$

In 1D, discriminant function is not linear $g(\mathbf{x}) = w_1 x + w_2 x^2 + w_0$

Non Linear SVM

- Can use any linear classifier after lifting data into a higher dimensional space. However we will have to deal with the “curse of dimensionality”
 - poor generalization to test data
 - computationally expensive
- SVM avoids the “curse of dimensionality” problems by
 - enforcing largest margin permits good generalization
 - It can be shown that generalization in SVM is a function of the margin, independent of the dimensionality
 - computation in the higher dimensional case is performed only implicitly through the use of **kernel** functions

Non Linear Mapping: Another Example



Non Linear SVM: Kernels

- Recall SVM optimization

$$\text{maximize } L_0(\alpha) = \sum_{i=1}^n \alpha_i - \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j z_i z_j x_i^t x_j$$
- Note this optimization depends on samples \mathbf{x}_i only through the dot product $\mathbf{x}_i^t \mathbf{x}_j$
- If we lift \mathbf{x}_i to high dimension using $\phi(\mathbf{x})$, need to compute high dimensional product $\phi(\mathbf{x}_i)^t \phi(\mathbf{x}_j)$

$$\text{maximize } L_0(\alpha) = \sum_{i=1}^n \alpha_i - \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j z_i z_j \phi(\mathbf{x}_i)^t \phi(\mathbf{x}_j)$$

$K(\mathbf{x}_i, \mathbf{x}_j)$

- Idea: find **kernel** function $K(\mathbf{x}_i, \mathbf{x}_j)$ s.t.

$$K(\mathbf{x}_i, \mathbf{x}_j) = \phi(\mathbf{x}_i)^t \phi(\mathbf{x}_j)$$

Non Linear SVM: Kernels

$$\text{maximize } L_D(\alpha) = \sum_{i=1}^n \alpha_i - \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j z_i z_j \phi(x_i)^T \phi(x_j)$$

$K(x_i, x_j)$

- Then we only need to compute $K(x_i, x_j)$ instead of $\phi(x_i)^T \phi(x_j)$
- "kernel trick": do not need to perform operations in high dimensional space explicitly

Non Linear SVM: Kernels

- How to choose kernel function $K(x_i, x_j)$?
 - $K(x_i, x_j)$ should correspond to product $\phi(x_i)^T \phi(x_j)$ in a higher dimensional space
 - Mercer's condition tells us which kernel function can be expressed as dot product of two vectors

- Some common choices:

- Polynomial kernel

$$K(x_i, x_j) = (x_i^T x_j + 1)^p$$

- Gaussian radial Basis kernel (data is lifted in infinite dimension)

$$K(x_i, x_j) = \exp\left(-\frac{1}{2\sigma^2} \|x_i - x_j\|^2\right)$$

Non Linear SVM: Kernels

- Suppose we have 2 features and $K(x, y) = (x^T y)^2$
- Which mapping $\phi(x)$ does it correspond to?

$$\begin{aligned} K(x, y) &= (x^T y)^2 = \left(\begin{bmatrix} x^{(1)} & x^{(2)} \end{bmatrix} \begin{bmatrix} y^{(1)} \\ y^{(2)} \end{bmatrix} \right)^2 = (x^{(1)} y^{(1)} + x^{(2)} y^{(2)})^2 \\ &= (x^{(1)} y^{(1)})^2 + 2(x^{(1)} y^{(1)})(x^{(2)} y^{(2)}) + (x^{(2)} y^{(2)})^2 \\ &= \left[(x^{(1)})^2 \quad \sqrt{2} x^{(1)} x^{(2)} \quad (x^{(2)})^2 \right] \begin{bmatrix} (y^{(1)})^2 \\ \sqrt{2} y^{(1)} y^{(2)} \\ (y^{(2)})^2 \end{bmatrix} \end{aligned}$$

- Thus $\phi(x) = \left[(x^{(1)})^2 \quad \sqrt{2} x^{(1)} x^{(2)} \quad (x^{(2)})^2 \right]^T$

Non Linear SVM

- search for separating hyperplane in high dimension
- $$w \phi(x) + w_0 = 0$$

- Choose $\phi(x)$ so that the first ("0"th) dimension is the augmented dimension with feature value fixed to 1

$$\phi(x) = \begin{bmatrix} 1 & x^{(1)} & x^{(2)} & x^{(1)} x^{(2)} \end{bmatrix}^T$$

- Threshold parameter w_0 gets folded into the weight vector w

$$\begin{bmatrix} w_0 & w \end{bmatrix} \begin{bmatrix} 1 \\ \phi(x) \end{bmatrix} = 0$$

Non Linear SVM

- Will not use notation $\mathbf{a} = [w_0 \ \mathbf{w}]$, we'll use old notation \mathbf{w} and seek hyperplane through the origin

$$\mathbf{w}\phi(\mathbf{x}) = 0$$
- If the first component of $\phi(\mathbf{x})$ is not 1, the above is equivalent to saying that the hyperplane has to go through the origin in high dimension
 - removes only one degree of freedom
 - But we have introduced many new degrees when we lifted the data in high dimension

Non Linear SVM Recipe

- Weight vector \mathbf{w} in the high dimensional space:

$$\mathbf{w} = \sum_{x_i \in S} \alpha_i z_i \phi(x_i)$$
 - where S is the set of support vectors $S = \{x_i \mid \alpha_i \neq 0\}$
- Linear discriminant function of largest margin in the high dimensional space:

$$g(\phi(\mathbf{x})) = \mathbf{w}'\phi(\mathbf{x}) = \left(\sum_{x_i \in S} \alpha_i z_i \phi(x_i) \right)' \phi(\mathbf{x})$$
- Non linear discriminant function in the original space

$$g(\mathbf{x}) = \left(\sum_{x_i \in S} \alpha_i z_i \phi(x_i) \right)' \phi(\mathbf{x}) = \sum_{x_i \in S} \alpha_i z_i \phi'(x_i) \phi(\mathbf{x}) = \sum_{x_i \in S} \alpha_i z_i K(x_i, \mathbf{x})$$
- decide class 1 if $g(\mathbf{x}) > 0$, otherwise decide class 2

Non Linear SVM Receptie

- Start with data $\mathbf{x}_1, \dots, \mathbf{x}_n$ which lives in feature space of dimension d
- Choose kernel $K(\mathbf{x}_i, \mathbf{x}_j)$ or function $\phi(\mathbf{x}_i)$ which takes sample \mathbf{x}_i to a higher dimensional space
- Find the largest margin linear discriminant function in the higher dimensional space by using quadratic programming package to solve:

$$\begin{aligned} &\text{maximize } L_D(\alpha) = \sum_{i=1}^n \alpha_i - \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \alpha_i \alpha_j z_i z_j K(\mathbf{x}_i, \mathbf{x}_j) \\ &\text{constrained to } 0 \leq \alpha_i \leq \beta \quad \forall i \quad \text{and} \quad \sum_{i=1}^n \alpha_i z_i = 0 \end{aligned}$$

Non Linear SVM

- Nonlinear discriminant function

$$g(\mathbf{x}) = \sum_{x_i \in S} \alpha_i z_i K(\mathbf{x}_i, \mathbf{x})$$


$$g(\mathbf{x}) = \sum \begin{matrix} \text{weight of support} \\ \text{vector } \mathbf{x}_i \end{matrix} \begin{matrix} \mp 1 \end{matrix} \begin{matrix} \text{"inverse distance"} \\ \text{from } \mathbf{x} \text{ to} \\ \text{support vector } \mathbf{x}_i \end{matrix}$$

most important training samples, i.e. support vectors

$$K(\mathbf{x}_i, \mathbf{x}) = \exp\left(-\frac{1}{2\sigma^2} \|\mathbf{x}_i - \mathbf{x}\|^2\right)$$

SVM Example: XOR Problem

- Class 1: $\mathbf{x}_1 = [1, -1]$, $\mathbf{x}_2 = [-1, 1]$
- Class 2: $\mathbf{x}_3 = [1, 1]$, $\mathbf{x}_4 = [-1, -1]$
- Use polynomial kernel of degree 2:



- $K(\mathbf{x}_i, \mathbf{x}_j) = (\mathbf{x}_i^T \mathbf{x}_j + 1)^2$
- This kernel corresponds to mapping

$$\phi(\mathbf{x}) = \begin{bmatrix} 1 \\ \sqrt{2}x^{(1)} \\ \sqrt{2}x^{(2)} \\ \sqrt{2}x^{(1)}x^{(2)} \\ (x^{(1)})^2 \\ (x^{(2)})^2 \end{bmatrix}$$
- Need to maximize

$$L_D(\alpha) = \sum_{i=1}^4 \alpha_i - \frac{1}{2} \sum_{i=1}^4 \sum_{j=1}^4 \alpha_i \alpha_j z_i z_j (\mathbf{x}_i^T \mathbf{x}_j + 1)^2$$
 constrained to $0 \leq \alpha_i, \forall i$ and $\alpha_1 + \alpha_2 - \alpha_3 - \alpha_4 = 0$

SVM Example: XOR Problem

$$\phi(\mathbf{x}) = \begin{bmatrix} 1 \\ \sqrt{2}x^{(1)} \\ \sqrt{2}x^{(2)} \\ \sqrt{2}x^{(1)}x^{(2)} \\ (x^{(1)})^2 \\ (x^{(2)})^2 \end{bmatrix}$$

- Class 1: $\mathbf{x}_1 = [1, -1]$, $\mathbf{x}_2 = [-1, 1]$
- Class 2: $\mathbf{x}_3 = [1, 1]$, $\mathbf{x}_4 = [-1, -1]$
- Weight vector \mathbf{w} is:

$$\mathbf{w} = \sum_{i=1}^4 \alpha_i z_i \phi(\mathbf{x}_i) = 0.25(\phi(\mathbf{x}_1) + \phi(\mathbf{x}_2) - \phi(\mathbf{x}_3) - \phi(\mathbf{x}_4))$$

$$= [0 \ 0 \ 0 \ -\sqrt{2} \ 0 \ 0]$$
- Thus the nonlinear discriminant function is:

$$g(\mathbf{x}) = \mathbf{w} \phi(\mathbf{x}) = \sum_{i=1}^6 w_i \phi_i(\mathbf{x}) = -\sqrt{2}(\sqrt{2}x^{(1)}x^{(2)}) = -2x^{(1)}x^{(2)}$$

SVM Example: XOR Problem

- Can rewrite $L_D(\alpha) = \sum_{i=1}^4 \alpha_i - \frac{1}{2} \alpha^T H \alpha$
- where $\alpha = [\alpha_1 \ \alpha_2 \ \alpha_3 \ \alpha_4]^T$ and $H = \begin{bmatrix} 9 & 1 & -1 & -1 \\ 1 & 9 & -1 & -1 \\ -1 & -1 & 9 & 1 \\ -1 & -1 & 1 & 9 \end{bmatrix}$


- Take derivative with respect to α and set it to 0

$$\frac{d}{d\alpha} L_D(\alpha) = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} - \begin{bmatrix} 9 & 1 & -1 & -1 \\ 1 & 9 & -1 & -1 \\ -1 & -1 & 9 & 1 \\ -1 & -1 & 1 & 9 \end{bmatrix} \alpha = 0$$


- Solution to the above is $\alpha_1 = \alpha_2 = \alpha_3 = \alpha_4 = 0.25$
- satisfies the constraints $\forall i, 0 \leq \alpha_i$ and $\alpha_1 + \alpha_2 - \alpha_3 - \alpha_4 = 0$
- all samples are support vectors

SVM Example: XOR Problem

$$g(\mathbf{x}) = -2x^{(1)}x^{(2)}$$



decision boundaries nonlinear



decision boundary is linear

SVM Summary

- **Advantages:**
 - Based on nice theory
 - excellent generalization properties
 - objective function has no local minima
 - can be used to find non linear discriminant functions
 - Complexity of the classifier is characterized by the number of support vectors rather than the dimensionality of the transformed space
- **Disadvantages:**
 - tends to be slower than other methods
 - quadratic programming is computationally expensive