

# Kernels

Machine Learning – CSE446

Carlos Guestrin

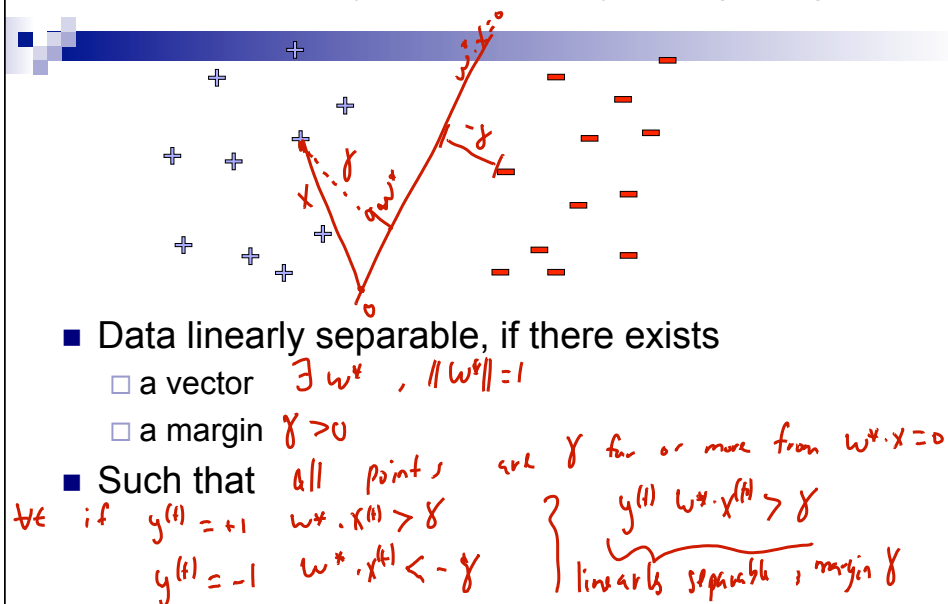
University of Washington

May 3, 2013

©Carlos Guestrin 2005-2013

1

## Linear Separability: More formally, Using Margin



©Carlos Guestrin 2005-2013

2

## Perceptron Analysis: Linearly Separable Case

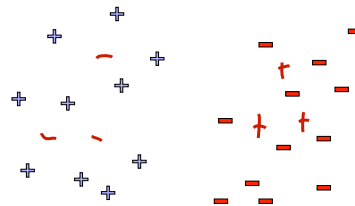
- Theorem [Block, Novikoff]:
  - Given a sequence of labeled examples:  $(x^{(1)}, y^{(1)}), (x^{(2)}, y^{(2)}), \dots, (x^{(T)}, y^{(T)})$   
*examples need not be iid. or random...*
  - Each feature vector has bounded norm:  
 $\forall t \quad \|x^{(t)}\| \leq R$   *$w^*$  is unknown!*
  - If dataset is linearly separable:  
 $\exists w^*, \|w^*\|=1 \quad \forall t \quad y^{(t)} w^* \cdot x^{(t)} \geq \gamma$ , for  $\gamma > 0$
- Then the number of mistakes made by the online perceptron on this sequence is bounded by  
 $\left(\frac{R}{\gamma}\right)^2$  *wow!!*  
 $\leftarrow$  *constant, doesn't depend on T*  
*dimensionality of X !!*

©Carlos Guestrin 2005-2013

3

## Beyond Linearly Separable Case

- Perceptron algorithm is super cool!
  - No assumption about data distribution!
    - Could be generated by an oblivious adversary, no need to be iid
  - Makes a fixed number of mistakes, and it's done for ever!  $\leftarrow \left(\frac{R}{\gamma}\right)^2$ 
    - Even if you see infinite data
- However, real world not linearly separable
  - Can't expect never to make mistakes again
  - Analysis extends to non-linearly separable case
  - Very similar bound, see Freund & Schapire
  - Converges, but ultimately may not give good accuracy (make many many many mistakes)

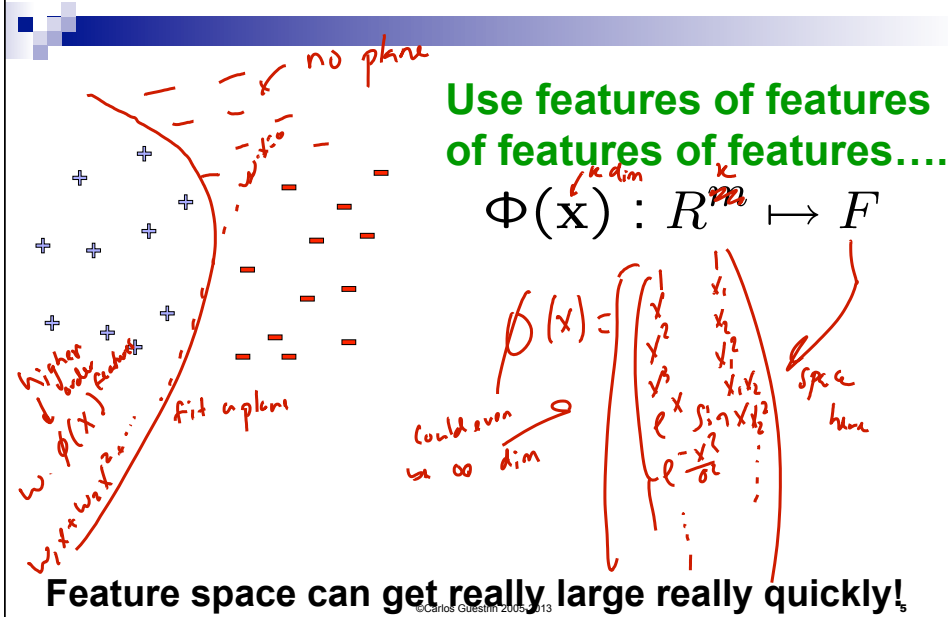


*we need features that make data as linearly separable as possible*

©Carlos Guestrin 2005-2013

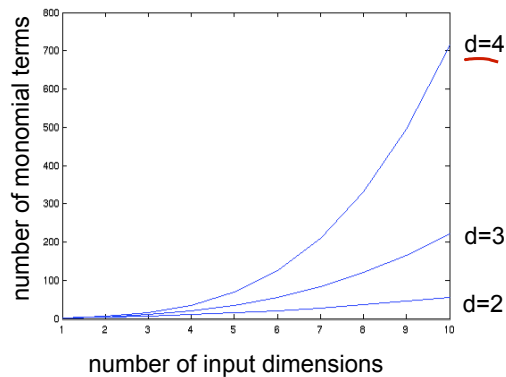
4

What if the data is not linearly separable?



## Higher order polynomials

$$\text{num. terms} = \binom{d+m-1}{d} = \frac{(d+m-1)!}{d!(m-1)!}$$



m – input features  
d – degree of polynomial

Even though  
dims of  $\phi(x)$   
are huge, AI model  
very ~~not~~ quickly  
grows fast!  
d = 6, m = 100  
about 1.6 billion terms

## Perceptron Revisited

- Given weight vector  $w^{(t)}$ , predict point  $\mathbf{x}$  by:
- Mistake at time  $t$ :  $w^{(t+1)} = w^{(t)} + y^{(t)} \mathbf{x}^{(t)}$
- Thus, write weight vector in terms of mistaken data points only:
  - Let  $M^{(t)}$  be time steps up to  $t$  when mistakes were made:
- Prediction rule now:
- When using high dimensional features:

©Carlos Guestrin 2005-2013

7

## Dot-product of polynomials

$\Phi(\mathbf{u}) \cdot \Phi(\mathbf{v}) = \text{polynomials of degree exactly } d$

©Carlos Guestrin 2005-2013

8

## Finally the Kernel Trick!!! (Kernelized Perceptron)

- Every time you make a mistake, remember  $(\mathbf{x}^{(t)}, y^{(t)})$

- Kernelized Perceptron prediction for  $\mathbf{x}$ :

$$\begin{aligned} \text{sign}(\mathbf{w}^{(t)} \cdot \phi(\mathbf{x})) &= \sum_{j \in M^{(t)}} y^{(j)} \phi(\mathbf{x}^{(j)}) \cdot \phi(\mathbf{x}) \\ &= \sum_{j \in M^{(t)}} y^{(j)} k(\mathbf{x}^{(j)}, \mathbf{x}) \end{aligned}$$

©Carlos Guestrin 2005-2013

9

## Polynomial kernels

- All monomials of degree  $d$  in  $O(d)$  operations:

$$\Phi(\mathbf{u}) \cdot \Phi(\mathbf{v}) = (\mathbf{u} \cdot \mathbf{v})^d = \text{polynomials of degree exactly } d$$

- How about all monomials of degree up to  $d$ ?

- ☐ Solution 0:

- ☐ Better solution:

©Carlos Guestrin 2005-2013

10

## Common kernels

- Polynomials of degree exactly  $d$

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

- Polynomials of degree up to  $d$

$$K(\mathbf{u}, \mathbf{v}) = (\mathbf{u} \cdot \mathbf{v} + 1)^d$$

- Gaussian (squared exponential) kernel

$$K(\mathbf{u}, \mathbf{v}) = \exp\left(-\frac{\|\mathbf{u} - \mathbf{v}\|^2}{2\sigma^2}\right)$$

- Sigmoid

$$K(\mathbf{u}, \mathbf{v}) = \tanh(\eta \mathbf{u} \cdot \mathbf{v} + \nu)$$

©Carlos Guestrin 2005-2013

11

## What you need to know

- Notion of online learning
- Perceptron algorithm
- Mistake bounds and proofs
- The kernel trick
- Kernelized Perceptron
- Derive polynomial kernel
- Common kernels
- In online learning, report averaged weights at the end

©Carlos Guestrin 2005-2013

12

# Support Vector Machines

Machine Learning – CSE446

Carlos Guestrin

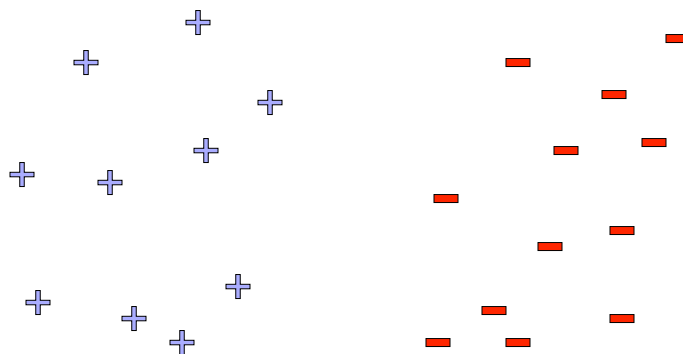
University of Washington

May 3, 2013

©Carlos Guestrin 2005-2013

13

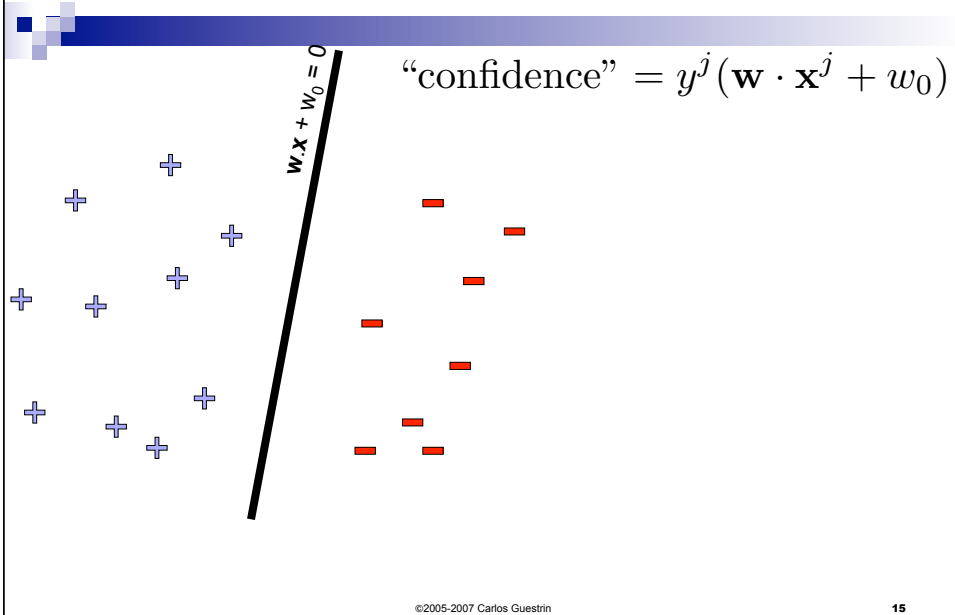
## Linear classifiers – Which line is better?



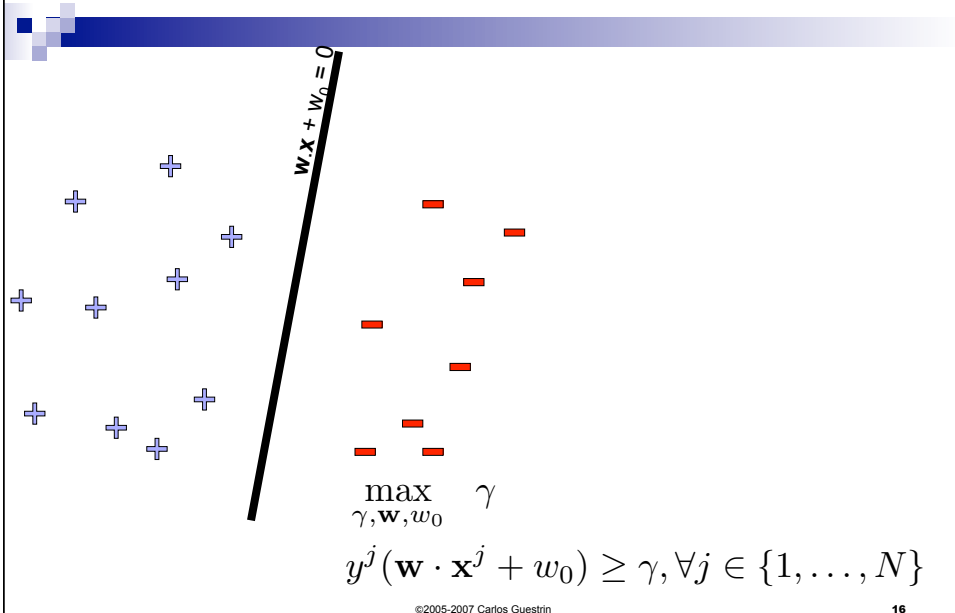
©2005-2007 Carlos Guestrin

14

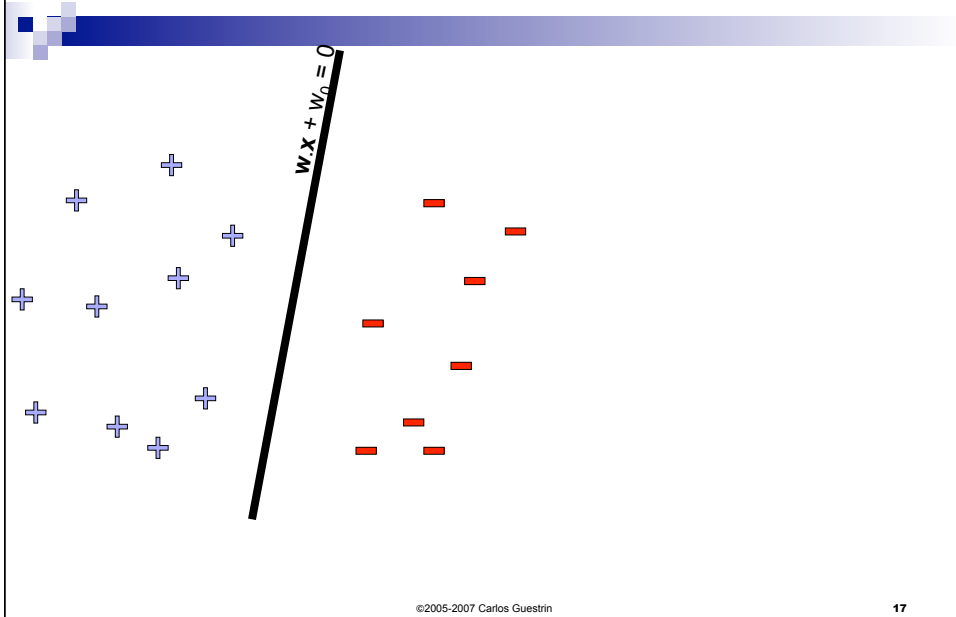
Pick the one with the largest margin!



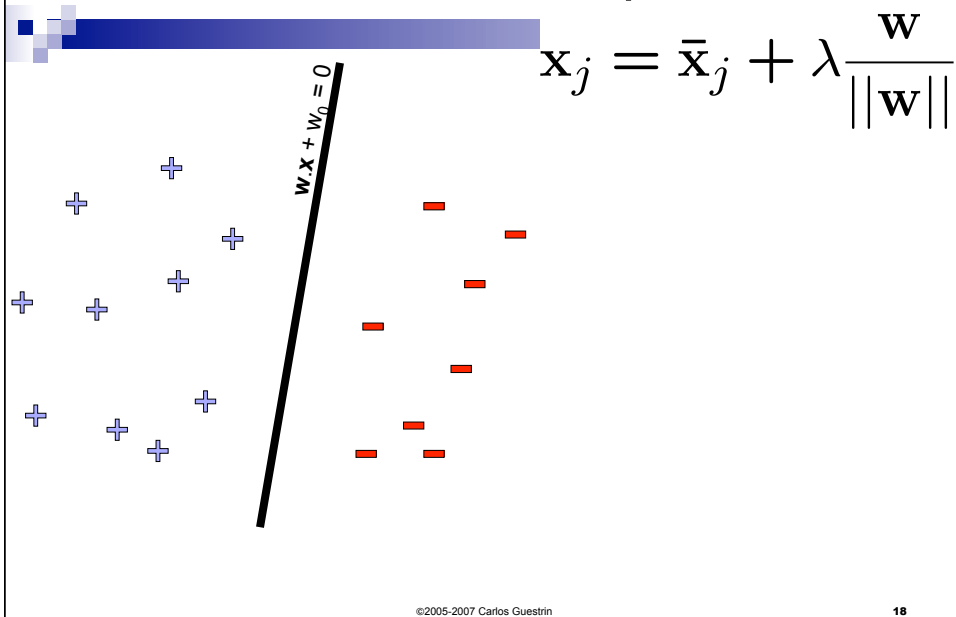
Maximize the margin



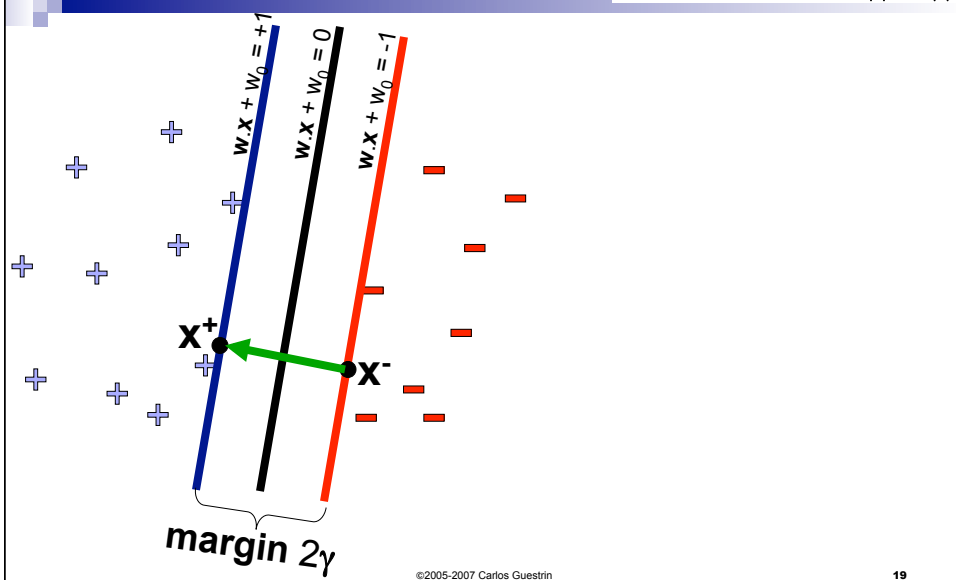
# But there are many planes...



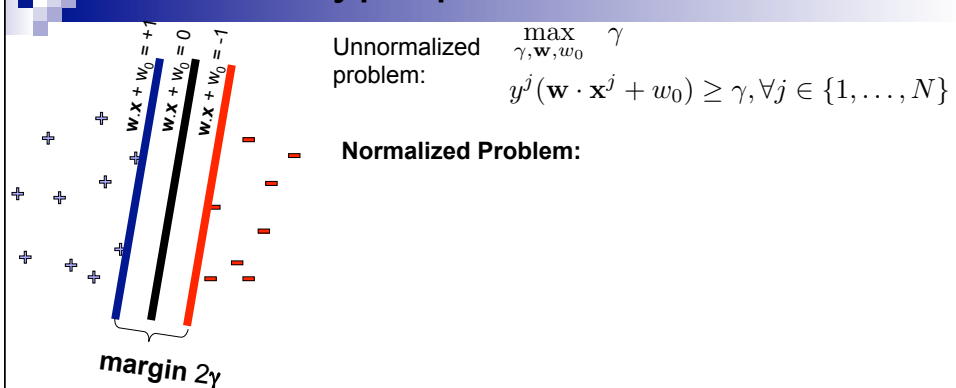
## Review: Normal to a plane



A Convention: Normalized margin – Canonical hyperplanes  $\mathbf{x}_j = \bar{\mathbf{x}}_j + \lambda \frac{\mathbf{w}}{\|\mathbf{w}\|}$



## Margin maximization using canonical hyperplanes



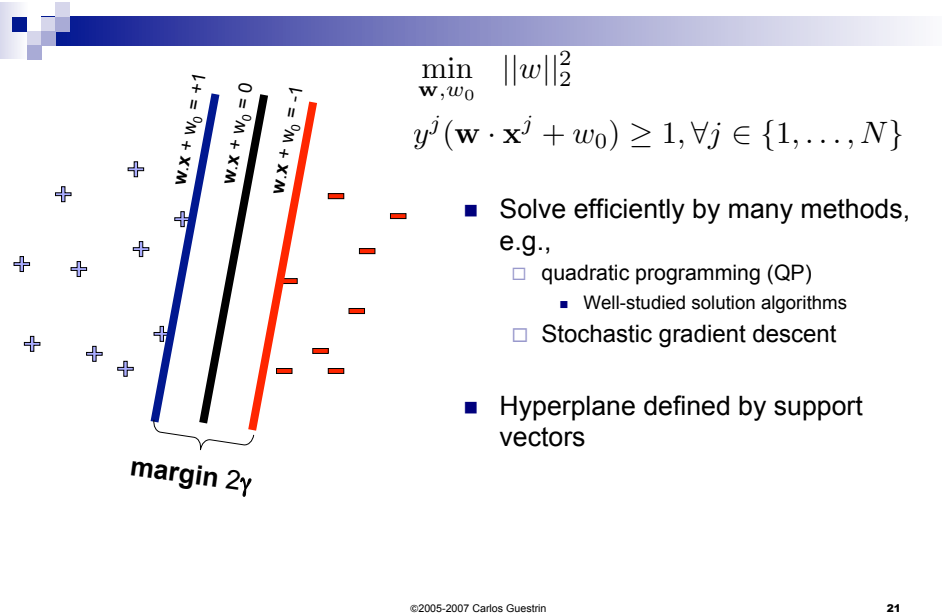
Unnormalized problem:  $\max_{\gamma, \mathbf{w}, w_0} \gamma$   
 $y^j (\mathbf{w} \cdot \mathbf{x}^j + w_0) \geq \gamma, \forall j \in \{1, \dots, N\}$

Normalized Problem:

$$\min_{\mathbf{w}, w_0} \|\mathbf{w}\|_2^2$$

$$y^j (\mathbf{w} \cdot \mathbf{x}^j + w_0) \geq 1, \forall j \in \{1, \dots, N\}$$

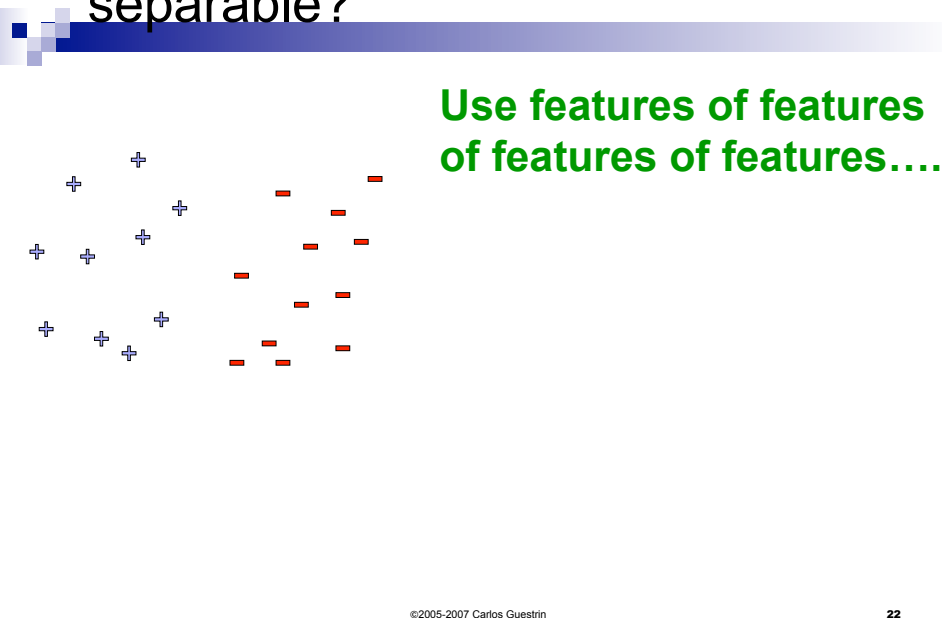
# Support vector machines (SVMs)



©2005-2007 Carlos Guestrin

21

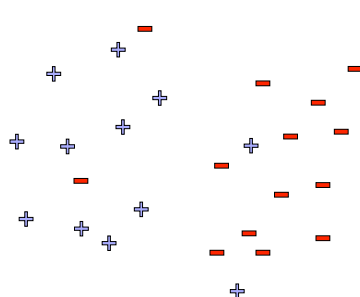
## What if the data is not linearly separable?



©2005-2007 Carlos Guestrin

22

## What if the data is still not linearly separable?



$$\min_{\mathbf{w}, w_0} \|\mathbf{w}\|_2^2$$

$$y^j(\mathbf{w} \cdot \mathbf{x}^j + w_0) \geq 1, \forall j$$

- If data is not linearly separable, some points don't satisfy margin constraint:
- How bad is the violation?
- Tradeoff margin violation with  $\|\mathbf{w}\|$ :

©2005-2007 Carlos Guestrin 23

## SVMs for Non-Linearly Separable meet my friend the Perceptron...

- Perceptron was minimizing the hinge loss:
 
$$\sum_{j=1}^N (-y^j(\mathbf{w} \cdot \mathbf{x}^j + w_0))_+$$
- SVMs minimizes the regularized hinge loss!!
 
$$\|\mathbf{w}\|_2^2 + C \sum_{j=1}^N (1 - y^j(\mathbf{w} \cdot \mathbf{x}^j + w_0))_+$$

©Carlos Guestrin 2005-2013 24

## Stochastic Gradient Descent for SVMs

- Perceptron minimization:

$$\sum_{j=1}^N (-y^j (\mathbf{w} \cdot \mathbf{x}^j + w_0))_+$$

- SGD for Perceptron:

$$\mathbf{w}^{(t+1)} \leftarrow \mathbf{w}^{(t)} + \mathbb{1} [y^{(t)} (\mathbf{w}^{(t)} \cdot \mathbf{x}^{(t)}) \leq 0] y^{(t)} \mathbf{x}^{(t)}$$

- SVMs minimization:

$$\|\mathbf{w}\|_2^2 + C \sum_{j=1}^N (1 - y^j (\mathbf{w} \cdot \mathbf{x}^j + w_0))_+$$

- SGD for SVMs:

©Carlos Guestrin 2005-2013

25

## What you need to know

- Maximizing margin
- Derivation of SVM formulation
- Non-linearly separable case
  - Hinge loss
  - A.K.A. adding slack variables
- SVMs = Perceptron + L2 regularization
- Can optimize SVMs with SGD
  - Many other approaches possible

©2005-2007 Carlos Guestrin

26