

An Idiot's guide to Support vector machines (SVMs)

R. Berwick, Village Idiot

SVMs: A New Generation of Learning Algorithms

- **Pre 1980:**
 - Almost all learning methods learned linear decision surfaces.
 - Linear learning methods have nice theoretical properties
- **1980's**
 - Decision trees and NNs allowed efficient learning of non-linear decision surfaces
 - Little theoretical basis and all suffer from local minima
- **1990's**
 - Efficient learning algorithms for non-linear functions based on computational learning theory developed
 - Nice theoretical properties.

Key Ideas

- Two independent developments within last decade
 - Computational learning theory
 - New efficient separability of non-linear functions that use “kernel functions”
- The resulting learning algorithm is an optimization algorithm rather than a greedy search.

Statistical Learning Theory

- Systems can be mathematically described as a system that
 - Receives data (observations) as input and
 - Outputs a function that can be used to predict some features of future data.
- Statistical learning theory models this as a function estimation problem
- Generalization Performance (accuracy in labeling test data) is measured

Organization

- Basic idea of support vector machines
 - Optimal hyperplane for linearly separable patterns
 - Extend to patterns that are not linearly separable by transformations of original data to map into new space – Kernel function
- SVM algorithm for pattern recognition


Unique Features of SVM's and Kernel Methods

- Are explicitly based on a theoretical model of learning
- Come with theoretical guarantees about their performance
- Have a modular design that allows one to separately implement and design their components
- Are not affected by local minima
- Do not suffer from the curse of dimensionality

Support Vectors

- Support vectors are the data points that lie closest to the decision surface
- They are the most difficult to classify
- They have direct bearing on the optimum location of the decision surface
- We can show that the optimal hyperplane stems from the function class with the lowest "capacity" (VC dimension).


Recall: Which Hyperplane?



- In general, lots of possible solutions for a, b, c .
- Support Vector Machine (SVM) finds an optimal solution. (wrt what cost?)

Support Vector Machine (SVM)


- SVMs maximize the *margin* around the separating hyperplane.
- The decision function is fully specified by a subset of training samples, *the support vectors*.
- *Quadratic programming* problem
- Text classification method du jour



Separation by Hyperplanes


- Assume *linear separability* for now:
 - in 2 dimensions, can separate by a line
 - in higher dimensions, need hyperplanes
- Can find separating hyperplane by *linear programming* (e.g. perceptron):
 - separator can be expressed as $ax + by = c$

Linear Programming / Perceptron



Find a, b, c , such that
 $ax + by \geq c$ for red points
 $ax + by \leq c$ for green points.


Which Hyperplane?



In general, lots of possible solutions for a, b, c .

Which Hyperplane?

- Lots of possible solutions for a, b, c .
- Some methods find a separating hyperplane, but not the optimal one (e.g., perceptron)
- Most methods find an optimal separating hyperplane
- Which points should influence optimality?
 - All points
 - Linear regression
 - Naïve Bayes
 - Only “difficult points” close to decision boundary
 - Support vector machines
 - Logistic regression (kind of)




Support Vectors again for linearly separable case

- Support vectors are the elements of the training set that would change the position of the dividing hyper plane if removed.
- Support vectors are the critical elements of the training set
- The problem of finding the optimal hyper plane is an optimization problem and can be solved by optimization techniques (use Lagrange multipliers to get into a form that can be solved analytically).

Support Vectors: Input vectors for which

$$\mathbf{w}_0^T \mathbf{x} + \mathbf{b}_0 = 1 \quad \text{or} \quad \mathbf{w}_0^T \mathbf{x} + \mathbf{b}_0 = -1$$



Definitions

Define the hyperplane H such that:

$$\mathbf{x}_i \cdot \mathbf{w} + b \geq +1 \quad \text{when } y_i = +1$$

$$\mathbf{x}_i \cdot \mathbf{w} + b \leq -1 \quad \text{when } y_i = -1$$


$H1$ and $H2$ are the planes:

$$H1: \mathbf{x}_i \cdot \mathbf{w} + b = +1$$

$$H2: \mathbf{x}_i \cdot \mathbf{w} + b = -1$$

The points on the planes

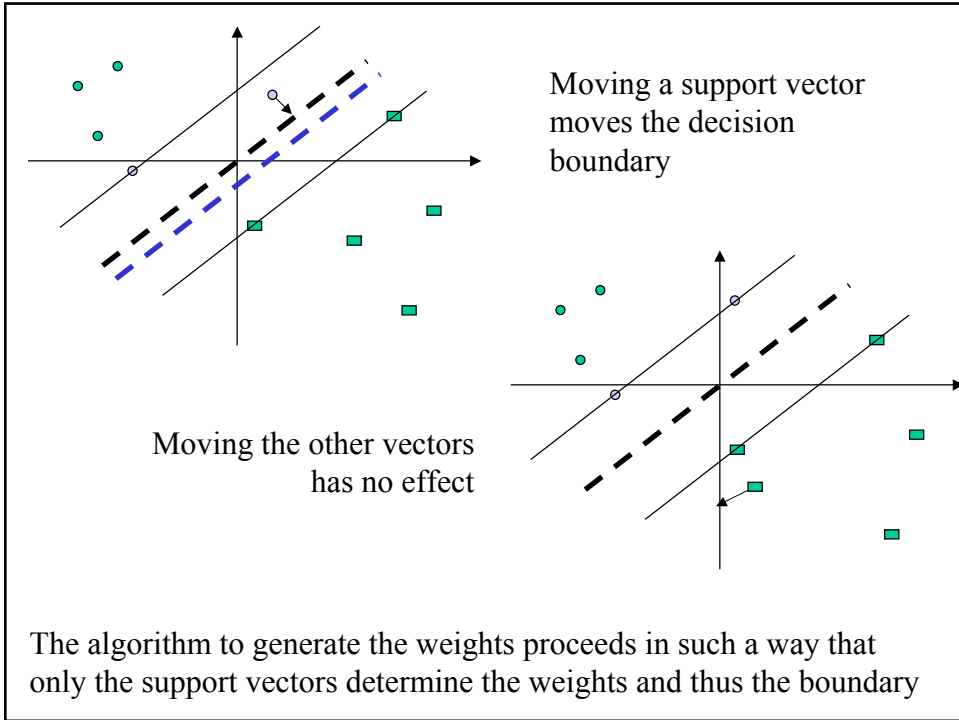
$H1$ and $H2$ are the
Support Vectors



$d+$ = the shortest distance to the closest positive point

$d-$ = the shortest distance to the closest negative point

The margin of a separating hyperplane is $d^+ + d^-$.



Maximizing the margin


We want a classifier with as big margin as possible.

Recall the distance from a point (x_0, y_0) to a line: $Ax + By + c = 0$ is $|Ax_0 + By_0 + c| / \sqrt{A^2 + B^2}$

The distance between H and H1 is:

$$|w \cdot x + b| / \|w\| = 1 / \|w\|$$

The distance between H1 and H2 is: $2 / \|w\|$



In order to maximize the margin, we need to minimize $\|w\|$. With the condition that there are no datapoints between H1 and H2:

$$\left. \begin{array}{l} x_i \cdot w + b \geq +1 \text{ when } y_i = +1 \\ x_i \cdot w + b \leq -1 \text{ when } y_i = -1 \end{array} \right\} \text{ Can be combined into } y_i(x_i \cdot w) \geq 1$$

We now must solve a quadratic programming problem


- Problem is: minimize $\|w\|$, **s.t.** discrimination boundary is obeyed, i.e., $\min f(x)$ s.t. $g(x)=0$, where

$$f: \frac{1}{2} \|w\|^2 \text{ and}$$

$$g: y_i(x_i \cdot w) - b = 1 \text{ or } [y_i(x_i \cdot w) - b] - 1 = 0$$


This is a **constrained optimization problem**

Solved by Lagrangian multiplier method




Intuition: intersection of two functions at a tangent point.

flattened paraboloid $2-x^2-2y^2$ with superimposed constraint
 $x^2+y^2 = 1$




flattened paraboloid $f: 2-x^2-2y^2=0$ with superimposed
constraint $g: x+y = 1$



Maximize when the constraint line g is tangent to the inner ellipse
contour line of f

flattened paraboloid $f: 2-x^2-2y^2=0$ with superimposed constraint $g: x+y=1$; at tangent solution p , gradient vectors of f, g are parallel (no possible move to incr f that also keeps you in region g)



Maximize when the constraint line g is tangent to the inner ellipse contour line of f

Two constraints

1. Parallel normal constraint (= gradient constraint on f, g solution is a max)
2. $G(x)=0$ (solution is on the constraint line)

We now recast these by combining f, g as the Lagrangian

Redescribing these conditions

- Want to look for solution point p where

$$\nabla f(p) = \nabla \lambda g(p)$$

$$g(x) = 0$$

- Or, combining these two as the *Langrangian* L & requiring derivative of L be zero:

$$L(x, \lambda) = f(x) - \lambda g(x)$$

$$\nabla(x, \lambda) = 0$$

How Langrangian solves constrained optimization

$$L(x, \lambda) = f(x) - \lambda g(x) \text{ where}$$

$$\nabla(x, \lambda) = 0$$

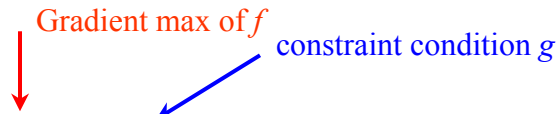
Partial derivatives wrt x recover the parallel normal constraint

Partial derivatives wrt λ recover the $g(x,y)=0$

In general,
$$L(x, \lambda) = f(x) + \sum_i \lambda_i g_i(x)$$

In general

Gradient max of f constraint condition g



$L(x, \alpha) = f(x) + \sum_i \alpha_i g_i(x)$ a function of $n + m$ variables

n for the x 's, m for the α . Differentiating gives $n + m$ equations, each set to 0. The n eqns differentiated wrt each x_i give the gradient conditions; the m eqns differentiated wrt each α_i recover the constraints g_i

In our case, $f(x): \frac{1}{2} \|\mathbf{w}\|^2$; $g(x): y_i(\mathbf{w} \cdot \mathbf{x}_i + b) - 1 = 0$ so Lagrangian is

$$L = \frac{1}{2} \|\mathbf{w}\|^2 - \sum \alpha_i [y_i(\mathbf{w} \cdot \mathbf{x}_i + b) - 1]$$

Lagrangian Formulation

- In the SVM problem the Lagrangian is

$$L_P \equiv \frac{1}{2} \|\mathbf{w}\|^2 - \sum_{i=1}^l \alpha_i y_i (\mathbf{x}_i \cdot \mathbf{w} + b) + \sum_{i=1}^l \alpha_i$$

$$\alpha_i \geq 0, \forall i$$

- From the derivatives = 0 we get

$$\mathbf{w} = \sum_{i=1}^l \alpha_i y_i \mathbf{x}_i, \quad \sum_{i=1}^l \alpha_i y_i = 0$$

The Lagrangian trick

Reformulate the optimization problem:

A "trick" often used in optimization is to do an Lagrangian formulation of the problem. The constraints will be replaced by constraints on the Lagrangian multipliers and the training data will occur only as dot products.

Gives us the task:

$$\text{Max } L = \sum \alpha_i - \frac{1}{2} \sum \alpha_i \alpha_j x_i \bullet x_j,$$

Subject to:

$$\mathbf{w} = \sum \alpha_i \gamma_i x_i$$

$$\sum \alpha_i \gamma_i = 0$$

What we need to see: x_i and x_j (input vectors) appear only in the form of dot product – we will soon see why that is important.

The Dual problem

- Original problem: fix value of f and find α
- New problem: Fix the values of α , and solve the (now unconstrained) problem $\max L(\alpha, x)$
- Ie, get a solution for each α , $f^*(\alpha)$
- Now minimize this over the space of α
- Kuhn-Tucker theorem: this is equivalent to original problem

At a solution p

- The the constraint line g and the contour lines of f must be tangent
- If they are tangent, their gradient vectors (perpendiculars) are parallel
- Gradient of g must be 0 – I.e., steepest ascent & so perpendicular to f
- Gradient of f must also be in the same direction as g

Inner products

The task:

$$\text{Max } L = \sum \alpha_i - \frac{1}{2} \sum \alpha_i \alpha_j \mathbf{x}_i \bullet \mathbf{x}_j$$

Subject to:

$$\mathbf{w} = \sum \alpha_i \gamma_i \mathbf{x}_i$$

$$\sum \alpha_i \gamma_i = 0$$

Inner product

Inner products

Why should inner product kernels be involved in pattern recognition?

- Intuition is that they provide some measure of similarity
- cf Inner product in 2D between 2 vectors of unit length returns the cosine of the angle between them.

e.g. $\underline{x} = [1, 0]^T$, $\underline{y} = [0, 1]^T$

I.e. if they are parallel inner product is 1


$$\underline{x}^T \underline{x} = \underline{x} \cdot \underline{x} = 1$$

If they are perpendicular inner product is 0

$$\underline{x}^T \underline{y} = \underline{x} \cdot \underline{y} = 0$$


But...are we done???

Not Linearly Separable




Find a line that penalizes points on “the wrong side”.

Transformation to separate




Non Linear SVMs


- The idea is to gain linear separation by mapping the data to a higher dimensional space
 - The following set can't be separated by a linear function, but can be separated by a quadratic one



- So if we map $x \mapsto \{x^2, x\}$ we gain linear separation



Problems with linear SVM




What if the decision function is not linear? What transform would separate these?

Ans: polar coordinates!

Non-linear SVM 1

The Kernel trick

Imagine a function ϕ that maps the data into another space:



Remember the function we want to optimize: $L_{\text{dual}} = \sum \alpha_i - \frac{1}{2} \sum \alpha_i \alpha_j \mathbf{x}_i \bullet \mathbf{x}_j$,
 \mathbf{x}_i and \mathbf{x}_j as a dot product. We will have $\phi(\mathbf{x}_i) \bullet \phi(\mathbf{x}_j)$ in the non-linear case.


If there is a "kernel function" \mathbf{K} such as $\mathbf{K}(\mathbf{x}_i, \mathbf{x}_j) = \phi(\mathbf{x}_i) \bullet \phi(\mathbf{x}_j)$, we do not need to know ϕ explicitly. One example:

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

We've already seen a nonlinear transform...

- What is it???

- $\tanh(\beta_0 x^T \mathbf{x}_i + \beta_1)$



Examples for Non Linear SVMs

$$K(\mathbf{x}, \mathbf{y}) = (\mathbf{x} \cdot \mathbf{y} + 1)^p$$

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

$$K(\mathbf{x}, \mathbf{y}) = \tanh(\kappa \mathbf{x} \cdot \mathbf{y} - \delta)$$

1st is polynomial (includes $\mathbf{x} \cdot \mathbf{x}$ as special case)

2nd is radial basis function (gaussians)

3rd is sigmoid (neural net activation function)

Inner Product Kernels

Type of Support Vector Machine	Inner Product Kernel $K(x, x_i), i = 1, 2, \dots, N$	Comments
Polynomial learning machine	$(\mathbf{x}^T \mathbf{x}_i + 1)^p$	Power p is specified apriori by the user
Radial-basis function network	$\exp(1/(2\sigma^2) \ x-x_i\ ^2)$	The width σ^2 is specified apriori
Two layer perceptron	$\tanh(\beta_0 \mathbf{x}^T \mathbf{x}_i + \beta_1)$	Mercer's theorem is satisfied only for some values of β_0 and β_1

Non-linear svm2

The function we end up optimizing is:

$$\text{Max } L_d = \sum \alpha_i - \frac{1}{2} \sum \alpha_i \alpha_j K(x_i \bullet x_j),$$

Subject to:

$$\mathbf{w} = \sum \alpha_i y_i x_i$$


$$\sum \alpha_i y_i = 0$$

Another kernel example: The polynomial kernel


$$K(x_i, x_j) = (x_i \bullet x_j + 1)^p, \text{ where } p \text{ is a tunable parameter.}$$

Evaluating K only require one addition and one exponentiation more than the original dot product.


Examples for Non Linear SVMs 2 – Gaussian Kernel




Nonlinear rbf kernel



Admiral's delight w/ diff kernel functions



Overfitting by SVM



Building an SVM Classifier

- Now we know how to build a separator for two linearly separable classes
- What about classes whose exemplary examples are not linearly separable?

40004 4310
 37879 05153
~~3502~~ 75216
 35460 44209

101191548572630324414184
 4359720299299722510046701
 3084114591010615406103631
 1064111030775242001979966
 8912054768559131427955460
 2019730187112991099970984
 0109707597331972015519055
 1075318255182314338010963
 1787521655460354603546055
 18255108503047520739401

FIGURE 10.8

Examples of ZIP code image, and segmented and normalized numerals from the testing set. (Source: Reprinted with permission from Y. Le Cun, et al., "Backpropagation Applied to Handwritten Zip Code Recognition," *Neural Computation*, 1:541-551, 1989. ©1989 The MIT Press.)