

Statistical Machine Learning
Methods for Bioinformatics
**V. Support Vector Machine
Theory**

Jianlin Cheng, PhD

Computer Science Department and Informatics Institute
University of Missouri, Columbia

2012

History of SVM

- Started in the late 1970s. (Vapnik, 1979)
- Hot since the middle of 1990s. (Vapnik, 1995 and Vapnik 1998), still hot now.
- There are a lot of applications in many areas such as bioinformatics, text classification, computer vision, handwriting recognition, object recognition, speaker identification, face detection, time series,

Theories Related to SVM

- Bias variance tradeoff (Geman and Bienenstock, 1992)
- Capacity control (Guyon et al., 1992, Vapnik, 1995 and Vapnik 1998)
- Overfitting (Montgomery and Peck, 1992)
- Basic idea: for a given learning task, with a given finite amount of training data, the best generalization performance will be achieved if the right balance is struck between the accuracy attained on the particular training set, and the capacity of the machine.

A Machine with Too Much / Little Capacity

- A machine with too much capacity is like a botanist with a photographic memory who, when presented with a new tree, concludes that it is not a tree because it has a different number of leaves from anything she has seen before;
- a machine with too little capacity is like the botanist's lazy brother, who declares that if it is green, it's a tree.

Notation

- l observations.
- Each observation consists of a pair: a vector X_i in \mathbb{R}^n , $i = 1, \dots, l$ and associated truth y_i .
- Tree recognition problem: X_i is a vector of pixel values (256, 16*16 image) and y_i is 1 if the image contains a tree and -1 otherwise.
- It is assumed that there exists some unknown probability distribution $P(X, y)$ from which these data are drawn, the data is assumed iid: independently drawn and identically distributed. P for cumulative probability distribution, p for their density.

Learning Machine

- Learning the mapping: $X_i \mapsto y_i$.
- The machine is defined by a set of mappings $X \mapsto f(X, a)$, where the functions $f(X, a)$ themselves are labeled by the adjustable parameters a . The machine is assumed to be deterministic.
- A particular choice of a generates a “trained” machine. (e.g. neural network with fixed architecture and trained weights)

Expectation of Test Error

$$R(a) = \int \frac{1}{2} |y - f(X, a)| dP(X, y)$$

$$dP(X, y) = p(X, y) dX dy \quad \text{if } p(X, y) \text{ exists.}$$

$R(a)$ is called expected risk or actual risk.

Empirical Risk $R_{emp}(a)$ is defined to be the measured mean error rate on the training set (for a fixed, finite number of observations).

$$R_{emp}(a) = \frac{1}{2l} \sum_{i=1}^l |y_i - f(X_i, a)|$$

Loss and Risk Bound

- $\frac{1}{2}|y_i - f(X_i, a)|$ is called the loss. It can only take the values 0 and 1.
- Choose η such that $0 \leq \eta \leq 1$. With probability $1 - \eta$, the following bound holds (Vapnik, 1995)

$$R(a) \leq R_{emp}(a) + \sqrt{\left(\frac{h(\log(2l/h) + 1) - \log(\eta/4)}{l}\right)}$$


Where h is a non-negative integer called the Vapnik Chevonenkis (VC) dimension, and is a measure of the notion of capacity. Second part of the right is called **VC confidence**.

Insights about Risk Bound


- It is independent of $P(X,y)$.
- It is usually not possible to compute the left hand side.
- If we know h , we can easily compute right hand side.
- **Structural Risk Minimization:** Given several learning machines $f(X,a)$, and choosing a fixed, sufficiently small η , by then taking the machine which minimize the right hand side, giving the lowest upper bound on the actual risk.
- **Question:** how does the bound change according to η ?

VC Dimension

- VC dimension is a property of a set of functions $\{f(a)\}$. Here we consider functions that correspond to two-class pattern recognition case, so that $f(X,a) \in \{-1, +1\}$.
- If a given set of l points can be labeled in all possible 2^l ways, and **for each labeling**, a member of set $\{f(a)\}$ can be found to correctly assign those labels, we say that set of points is shattered by that set of functions.
- VC dimension for a set of functions $\{f(a)\}$ is defined as the maximum number of training points that can be shattered by $\{f(a)\}$.
- If the VC dimension is h , then there exists at least one set of h points that can be shattered. But not necessary for every set of h points.



8 possible labeling of 3 points can be separated by lines.



Simply can not separate the labeling of these four points using a line. So the VC dimension of a line is 3.


VC Dimension and the Number of Parameters

- Intuitively, more parameters \rightarrow higher VC dimension.
- However, 1 parameter function can have infinite VC dimension. (see Burge's tutorial)

$$f(x, \alpha) \equiv \theta(\sin(\alpha x)), \quad x, \alpha \in \mathbf{R}.$$

If $\sin(\alpha x) > 0$, $f(x, \alpha) = 1$, -1 otherwise

VC Confidence and VC Dimension h



VC confidence is monotonic in h . (here $l = 10,000$, $\eta = 0.05$ (95%))

Structural Risk Minimization

$$R(a) \leq R_{emp}(a) + \sqrt{\left(\frac{h(\log(2l/h) + 1) - \log(\eta/4)}{l}\right)}$$

Given some selection of learning machines whose empirical risk is zero, one wants to choose that learning machine whose associated set of functions has minimal VC dimension.

This is called **Occam's Razor**.

"All things being equal, the simplest solution tends to be the best one."

In general, for non-zero empirical risk, one wants to choose the learning machine which minimizes the Risk Bound.

Comments

- The risk bound equation gives (with some chosen **probability**) an upper bound on the actual risk. This does not prevent a particular machine with the same value for empirical risk, and whose function set has higher VC dimension from having better performance.
- For higher h value, the bound is guaranteed not tight.
- $h/l > 0.37$, VC confidence exceeds unity.

Example


- k^{th} nearest neighbor classifier, with $k = 1$, has infinite VC dimension and zero empirical risk, since any number of points, labeled arbitrarily, will be successfully learned by the algorithm (provided no two points of opposite class lie right on top of each other). Thus the bound provides no information.
- For any classifier with infinite VC dimension, the bound is not even valid.
- Nearest neighbor classifier can still perform well. Thus, infinite capacity does not guarantee poor performance.

Structure Risk Minimization

- We would like to find that subset of the chosen set of functions, such that the risk bound for that subset is minimized.

Linear Support Vector Machines

- Linear machine trained on separable data.
- Label training data $\{x_i, y_i\}$, $i = 1, \dots, l$, y_i in $\{-1, 1\}$. x_i in \mathbb{R}^d .
- A hyperplane separates the positive from negative examples. The points which lie on the hyperplane satisfy $w \cdot x + b = 0$, where w is normal to the hyperplane. $|b| / \|w\|$ is the perpendicular distance from the hyperplane to the origin, and $\|w\|$ is the Euclidean norm of w .




$$x_i \cdot w + b \geq +1 \text{ for } y_i = +1$$

$$x_i \cdot w + b \leq -1, \text{ for } y_i = -1$$

combined into: $y_i(x_i \cdot w + b) - 1 \geq 0$


- Distance from origin to $w\mathbf{x} + \mathbf{b} = 0$ is $|\mathbf{b}| / |\mathbf{w}|$.
- Choose a point \mathbf{x} on $w\mathbf{x} + \mathbf{b} = 0$ that vector $(0, \mathbf{x})$ is perpendicular to $w\mathbf{x} + \mathbf{b} = 0$. So \mathbf{x} is $\lambda\mathbf{w}$ because \mathbf{w} is norm of $w\mathbf{x} + \mathbf{b} = 0$.
- So $\lambda\mathbf{w}\mathbf{w} + \mathbf{b} = 0$. so $\lambda = -\mathbf{b} / \mathbf{w} \cdot \mathbf{w} = -\mathbf{b} / |\mathbf{w}|^2$
- So $\mathbf{x} = -\mathbf{b} / |\mathbf{w}|^2 * \mathbf{w}$
- So $|\mathbf{x}| = \text{sqrt} (\mathbf{b}^2 / |\mathbf{w}|^4 * \mathbf{w}^2) = |\mathbf{b}| / |\mathbf{w}|$.

Q1: How perceptron finds a linear function?
Is the function found by perceptron unique?



Q2: From your intuition, why we choose this line 2?

Margin




Why Maximize Margin?

- Intuitively this feels safest.
- If we've made a small error in the location of the boundary (it's been jolted in its perpendicular direction) this gives us the least chance of causing a misclassification
- LOOCV is easy since the model is immune to removal of any non-support vector data points
- Related to VC dimension / structural risk minimization.
- Empirically it works very well.

How to Compute Margin?

Computing the margin width




How do we compute M in terms of w and b ?

- Plus-plane = $\{x : w \cdot x + b = +1\}$
- Minus-plane = $\{x : w \cdot x + b = -1\}$

Claim: The vector w is perpendicular to the Plus Plane. **Why?**

Computing the margin width



How do we compute M in terms of w and b ?


- Plus-plane = $\{x : w \cdot x + b = +1\}$
- Minus-plane = $\{x : w \cdot x + b = -1\}$

Claim: The vector w is perpendicular to the Plus Plane. Why?

Let u and v be two vectors on the Plus Plane. What is $w \cdot (u - v)$?

And so of course the vector w is also perpendicular to the Minus Plane

Computing the margin width




How do we compute M in terms of w and b ?

- Plus-plane = $\{ \mathbf{x} : \mathbf{w} \cdot \mathbf{x} + b = +1 \}$
- Minus-plane = $\{ \mathbf{x} : \mathbf{w} \cdot \mathbf{x} + b = -1 \}$
- The vector \mathbf{w} is perpendicular to the Plus Plane
- Let \mathbf{x}^- be any point on the minus plane
- Let \mathbf{x}^+ be the closest plus-plane-point to \mathbf{x}^- .


Any location in \mathbb{R}^m : not necessarily a datapoint

Computing the margin width




How do we compute M in terms of w and b ?

- Plus-plane = $\{ \mathbf{x} : \mathbf{w} \cdot \mathbf{x} + b = +1 \}$
- Minus-plane = $\{ \mathbf{x} : \mathbf{w} \cdot \mathbf{x} + b = -1 \}$
- The vector \mathbf{w} is perpendicular to the Plus Plane
- Let \mathbf{x}^- be any point on the minus plane
- Let \mathbf{x}^+ be the closest plus-plane-point to \mathbf{x}^- .
- **Claim:** $\mathbf{x}^+ = \mathbf{x}^- + \lambda \mathbf{w}$ for some value of λ . **Why?**



Computing the margin width




What we know:

- $w \cdot x^+ + b = +1$
- $w \cdot x^- + b = -1$
- $x^+ = x^- + \lambda w$
- $|x^+ - x^-| = M$

It's now easy to get M
in terms of w and b

Computing the margin width



What we know:

- $w \cdot x^+ + b = +1$
- $w \cdot x^- + b = -1$
- $x^+ = x^- + \lambda w$
- $|x^+ - x^-| = M$

It's now easy to get M
in terms of w and b

$$w \cdot (x^- + \lambda w) + b = 1$$

\Rightarrow

$$w \cdot x^- + b + \lambda w \cdot w = 1$$


\Rightarrow

$$-1 + \lambda w \cdot w = 1$$

\Rightarrow

$$? = \frac{2}{w \cdot w}$$

Computing the margin width



$$M = |x^+ - x^-| = |\lambda \mathbf{w}| =$$


$$= ? |\mathbf{w}| = ? \sqrt{\mathbf{w} \cdot \mathbf{w}}$$

$$= \frac{2\sqrt{\mathbf{w} \cdot \mathbf{w}}}{\mathbf{w} \cdot \mathbf{w}} = \frac{2}{\sqrt{\mathbf{w} \cdot \mathbf{w}}}$$

What we know:

- $\mathbf{w} \cdot \mathbf{x}^+ + b = +1$
- $\mathbf{w} \cdot \mathbf{x}^- + b = -1$
- $\mathbf{x}^+ = \mathbf{x}^- + \lambda \mathbf{w}$
- $|\mathbf{x}^+ - \mathbf{x}^-| = M$
- $? = \frac{2}{\mathbf{w} \cdot \mathbf{w}}$

Learning the Maximum Margin Classifier




Given a guess of \mathbf{w} and b we can

- Compute whether all data points in the correct half-planes
- Compute the width of the margin

So now we just need to write a program to search the space of \mathbf{w} 's and b 's to find the widest margin that matches all the datapoints. *How?*

Gradient descent? Simulated Annealing? Matrix Inversion?
EM? Newton's Method?

Learning the Maximum Margin Classifier



Given guess of w , b we can

- Compute whether all data points are in the correct half-planes
- Compute the margin width


Assume R datapoints, each (\mathbf{x}_k, y_k) where $y_k = +/- 1$

What should our quadratic optimization criterion be?

How many constraints will we have?

What should they be?

Learning the Maximum Margin Classifier



Given guess of w , b we can

- Compute whether all data points are in the correct half-planes

- Compute the margin width

Assume R datapoints, each (x_k, y_k) where $y_k = +/- 1$

What should our quadratic optimization criterion be?

Minimize $w \cdot w$


How many constraints will we have? R

What should they be?

$$w \cdot x_k + b \geq 1 \text{ if } y_k = 1$$

$$w \cdot x_k + b \leq -1 \text{ if } y_k = -1$$

Margin



Objectives

- Maximize $2 / |w|$, subject to the linear constraints
- $x_i \cdot w + b \geq +1$, for $y_i = +1$
- $x_i \cdot w + b \leq -1$, for $y_i = -1$
- Combined into one set of inequalities
 $y_i(x_i \cdot w + b) - 1 \geq 0$, for all i .
- How many constraints are there?
- How to solve the constraint optimization problem?
(Lagrangian)

Optimization by Lagrange

- Maximize $2 / |w|$ (or minimize $|w|^2$), subject to constraints.
- Switch to a Lagrangian formulation.
- Inequality constraints will be replaced by constraints on Lagrangian multipliers
- Training data will only appear as dot products between vectors. This is a crucial property which will allow us to generalize the procedure to the nonlinear case.

Lagrange Multiplier

- An mathematical optimization technique named after Joseph Louis Lagrange
- A method for finding local minima of a function of several variables subject to one or more constraints
- The method reduces a problem in n variables with k constraints to a solvable problem in $n+k$ variables with no constraints.
- The method introduces a new unknown scalar variable, the Lagrange multiplier, for each constraint and forms a linear combination involving the multipliers as coefficients.

http://en.wikipedia.org/wiki/Lagrange_multipliers

Primal Optimization (L_p)

- Constraints: $y_i(x_i \cdot w + b) - 1 \geq 0$, for all i .
(constraints set C1)

- Introduce a Lagrange multiplier for each inequality: a_i .

$$L_P = \frac{1}{2} |w|^2 - \sum_{i=1}^l a_i y_i (x_i \cdot w + b) + \sum_{i=1}^l a_i$$

Dual Optimization Problem (L_D)

- Set the derivative of L_p with respect to w and b to 0.
- $w = \sum a_i y_i x_i$.
- $\sum a_i y_i = 0$
- Substitute the equality constraints above into the primal equation to get the dual equation.

$$L_D = \sum_i a_i - \frac{1}{2} \sum_{i,j} a_i a_j y_i y_j x_i \cdot x_j$$

L_p and L_D arise from the same objective function, but with different constraints ($\sum a_i y_i = 0$, and $a_i \geq 0$, constraint set C2); the solution is found by minimizing L_p or by maximizing L_D .


Primal Problem \Leftrightarrow Dual Problem

- We can minimize L_p with respect to w , b , and simultaneously require that the derivatives of L_p with respect to a_i vanish, all subject to the constraints $a_i \geq 0$ (Constraints C1 - original constraints)
- A convex quadratic programming problem.
- Equivalent to maximize L_D , subject to the constraints that the gradient of L_D with respect to w and b vanish, and subject also to the constraints that $a_i \geq 0$ (Constraints C2, new constraints). This dual formulation is called the Wolfe dual (Fletcher, 1987).
- It has the property that the maximum of L_D , subject to C2, occurs at the same values of the w , b , and a , as the minimum of L_p , subject to constraints C1.

Learning and Support Vectors

- Support vector training (for the separable, linear case) therefore amounts to maximizing L_d with respect to a_i , subject to $\sum a_i y_i = 0$ and $a_i \geq 0$.
- There is a Lagrange multiplier a_i for every training point. Those points for which $a_i > 0$ are called “support vectors”, and lie on one of the hyperplanes H_1, H_2 .
- All other training points have $a_i = 0$ and lie on that side of H_1 or H_2 such that strict inequality holds.

Margin



Questions

- Why the data points lying on the H1 and H2 are support vectors ($a_i > 0$)?
- Why the data points not lying on the H1 and H2 are not support vectors ($a_i = 0$)?
- What happen if a non-support vector is removed? Does the solution change?
- What happen if a support vector is removed? Does solution change?

Karush-Kuhn-Tucker (KKT) Condition

- $w = \sum a_i y_i x_i.$ (1)

- $\sum a_i y_i = 0$ (2)

- $y_i(x_i \cdot w + b) - 1 \geq 0, i = 1, \dots, l$ (3)

- $a_i \geq 0$ (4)

- **$a_i(y_i(w \cdot x_i + b) - 1) = 0$** (5)

(5) is called complementary slackness due to the Lagrange theory and can be explained in intuition.


How to Determine w and b

- Use quadratic programming to solve a_i and compute w is trivial. (use KKT condition (1))
- How to compute b ?
- Use KKT condition (5), for any support vector (point $a_i > 0$), $y_i(w \cdot x_i + b) - 1 = 0$.
- We compute b in terms of a support vector.
Better: we compute b in terms of all support vectors and take the average.

Test Phase


- We simply determine on which side of the decision boundary (that hyperplane lying half way between H_1 and H_2 and parallel to them) a given test pattern x lies and assign the corresponding class label
- Class of x is $sign(w \cdot x + b)$

Non-Separable Case



Can't satisfy the constraints $y_i(wx_i+b) \geq 1$ for some data points? What can we do?

Learning Maximum Margin with Noise



Given guess of w , b we can

- Compute sum of distances of points to their correct zones

- Compute the margin width

Assume R datapoints, each (\mathbf{x}_k, y_k) where $y_k = +/- 1$

What should our quadratic optimization criterion be?

Minimize $\frac{1}{2} \mathbf{w} \cdot \mathbf{w} + C \sum_{k=1}^R e_k$

How many constraints will we have? R

What should they be?

$\mathbf{w} \cdot \mathbf{x}_k + b \geq 1 - \epsilon_k$ if $y_k = 1$

$\mathbf{w} \cdot \mathbf{x}_k + b \leq -1 + \epsilon_k$ if $y_k = -1$

Relax Constraints – Soft Margin

- Introduce positive slack variables ξ_i , $i = 1, \dots, l$ to relax constraints. ($\xi_i \geq 0$)
- New constraints:
- $x_i \cdot w + b \geq +1 - \xi_i$ for $y_i = +1$
- $x_i \cdot w + b \leq -1 + \xi_i$ for $y_i = -1$
- Or $y_i(w x_i + b) \geq 1 - \xi_i$
- $\xi_i \geq 0$
- For an classification error to happen, the corresponding ξ_i must exceed unity, so $\sum \xi_i$ is an upper bound on the number of training errors.

New Objective Function

- Minimize $|w|^2/2 + C(\sum \xi_i)^k$.
- C is parameter to be chosen by the user, a larger C corresponding to assigning a higher penalty to errors.
- This is a convex programming problem for any positive integer k . The choice $k=1$ has the further advantage that neither ξ_i and their multipliers appear in the Wolfe dual problem.

Primal Optimization (Lp)

$$L_P = \frac{1}{2} |w|^2 + C \sum_i \xi_i - \sum_{i=1}^l a_i (y_i (x_i \cdot w + b) - 1 + \xi_i) - \sum_{i=1}^N u_i \xi_i$$

$$\frac{\partial L_P}{\partial \xi_i} = C - a_i - u_i = 0$$

$$\Rightarrow a_i \leq C$$

u_i is the Lagrange multipliers introduced to enforce
Non-negativity of ξ_i

KKT Conditions

$$1. \partial_{\mathbf{w}} \mathcal{L}_P = 0 \rightarrow \mathbf{w} - \sum_i \alpha_i y_i \mathbf{x}_i = 0$$

$$2. \partial_b \mathcal{L}_P = 0 \rightarrow \sum_i \alpha_i y_i = 0$$

$$3. \partial_{\xi} \mathcal{L}_P = 0 \rightarrow C - \alpha_i - \mu_i = 0$$

$$4. \text{constraint-1} \quad y_i (\mathbf{w}^T \mathbf{x}_i - b) - 1 + \xi_i \geq 0$$

$$5. \text{constraint-2} \quad \xi_i \geq 0$$

$$6. \text{multiplier condition-1} \quad \alpha_i \geq 0$$

$$7. \text{multiplier condition-2} \quad \mu_i \geq 0$$

$$8. \text{complementary slackness-1} \quad \alpha_i [y_i (\mathbf{w}^T \mathbf{x}_i - b) - 1 + \xi_i] = 0$$

$$9. \text{complementary slackness-1} \quad \mu_i \xi_i = 0$$

Dual Optimization (\mathcal{L}_D)

$$\text{maximize } \mathcal{L}_D = \sum_{i=1}^N \alpha_i - \frac{1}{2} \sum_{ij} \alpha_i \alpha_j y_i y_j \mathbf{x}_i^T \mathbf{x}_j$$


$$\text{subject to } \sum_i \alpha_i y_i = 0$$

$$0 \leq \alpha_i \leq C \quad \forall i$$

We can get LD by substitute w by

What's the only difference from the linearly separable cases?

Values of Multipliers



Solution of w and b

$$w = \sum_{i=1}^{N_s} a_i y_i x_i$$

Use complementary slackness to compute b (choose a support vector ($0 < a_i < C$) to compute b , where $\xi_i = 0$. $\xi_i = 0$ is derived by combining equations 3 and 9.

Non-Linear Classification

SVM Demo

- <http://www.youtube.com/watch?v=3liCbRZPrZA>

Nonlinear Support Vector Machines

- In the L_D function, what really matters is dot products: $x_i \cdot x_j$.
- Idea: map the data to some other (possibly infinite dimensional) Euclidean space H , using a mapping.

$$\Phi : R^d \mapsto H$$

Then the training algorithm would only depend on the data through dot products in H , i.e. $\Phi(x_i) \cdot \Phi(x_j)$.

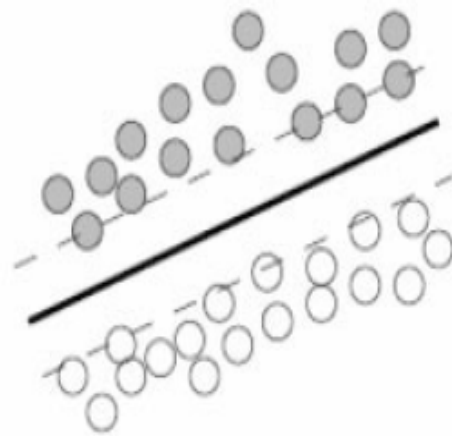
Kernel Trick

- If there were a kernel function K such that $K(x_i, x_j) = \Phi(x_i) \cdot \Phi(x_j)$, we would only need to use K in the training algorithm and would never need to explicitly do the mapping Φ .
- One example Gaussian kernel: $K(x_i, x_j) = e^{-|x_i - x_j|^2 / 2\sigma^2}$. In this example, H is infinite dimensional.
- So we simply replace $x_i \cdot x_j$ with $K(x_i, x_j)$ in the training algorithm, the algorithm will happily produce a support vector machine which lives in an infinite dimensional space, and furthermore do so in roughly the same amount of time it would take to train on the un-mapped data.

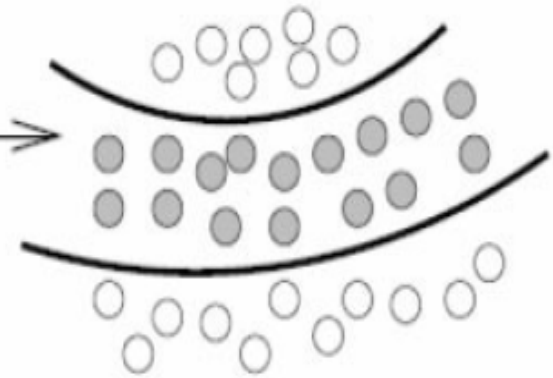
What's the effect of a mapping?



(a) no linear separating hyperplane in input space X



(b) linear separating hyperplane in feature space H



(c) Corresponding non-linear separating surface in input space X

How to Use the Machine?

- We can't get w if we do not do explicit mapping.
- Once again we use kernel trick.

$$f(x) = \left(\sum_{i=1}^{N_s} a_i y_i \Phi(s_i) \right) \Phi(x) + b = \sum_{i=1}^{N_s} a_i y_i K(s_i, x) + b$$

What's the problem from a computational point of view?

Speedup SVM Prediction

- Remove some redundant support vectors.
- Burges C.J.C. Simplified support vector decision rules. ICML, 1996
- Osuna E and Girosi F. Reducing the runtime complexity of support vector machines. International Conference on Pattern Recognition, 1998.

A Simple Kernel Example

- $K(\mathbf{x}_i, \mathbf{x}_j) = (\mathbf{x}_i \cdot \mathbf{x}_j)^2$. $\mathbf{x}_i, \mathbf{x}_j$ live in \mathbb{R}^2 space.
- Try to find a Φ from \mathbb{R}^2 to H , such that $(\mathbf{x} \cdot \mathbf{y})^2 = \Phi(\mathbf{x}) \cdot \Phi(\mathbf{y})$.
- Here are two possible mappings Φ (map \mathbb{R}^2 to \mathbb{R}^3 and \mathbb{R}^4 spaces)

$$\Phi(\mathbf{x}) = \begin{pmatrix} x_1^2 \\ \sqrt{2} x_1 x_2 \\ x_2^2 \end{pmatrix}$$

$$\Phi(\mathbf{x}) = \begin{pmatrix} x_1^2 \\ x_1 x_2 \\ x_1 x_2 \\ x_2^2 \end{pmatrix}$$

Kernel Function and Hilbert Space

- Hilbert space is a generalization of Euclidean space.
- It is a linear space, with an inner product define.
- Its inner product can be any inner product, not just scalar dot.

What Conditions Make a Function a Kernel?

- Mercer's condition

$$K(x, y) = \sum_i \Phi(x)_i \Phi(y)_i \Leftrightarrow \int K(x, y) g(x) g(y) dx dy \geq 0$$

where $\int g(x)^2 dx$ is finite. $g(x)$ is any function.

It is hard to check Mercer's condition because it must hold for every g with finite L2 norm.

What happens if one uses a kernel which does not satisfy Mercer's condition? Some time QP has no solution. Sometime, there is a solution, but the geometrical interpretation is lacking.

Common Kernels

(1) $K(x,y) = (x \cdot y + 1)^p$.
 p is degree. $p = 1$,
linear kernel.


(2) Gaussian radial basis kernel $K(x, y) = e^{-|x-y|^2 / 2\sigma^2}$

(3) Hyperbolic Tanh kernel $K(x, y) = \tanh(kx \cdot y - \delta)$

Note: RBF kernel, the weights (a_i) and centers (S_i) are automatically Learned.

Tanh kernel is equivalent to two-layer neural network, where Number of hidden units is number of support vectors. a_i corresponds To the weights of the second layer.


Connection between NN and SVM



Multi-Class SVM

- Most widely used method: one versus all
- Also direct multi-classification using SVM.

(K. Crammer and Y. Singer. On the Algorithmic Implementation of Multi-class SVMs, JMLR, 2001)



Global Solutions and Uniqueness

- For SVM optimization, every local solution is global due to the property to the convex objective function.
- The solution is guaranteed to be unique.
- SVM training always finds a global solution is in contrast to the case of neural networks, where many local minima usually exist.

Method of Solution

- The support vector optimization problem can be solved analytically only when the number of training data is very small, or for the separable case when it is known beforehand which of the training data become support vectors.
- For the general analytic case, the worst case computational complexity is of order N_s^3 (inversion of Hessian), where N_s is the number of support vectors.

Method of Solution

- In most real world cases, the quadratic optimization problem must be solved numerically.
- For small problems, any general purpose optimization package that solves linearly constrained convex quadratic programs will do. (a good survey: More and Wright, 1993)
- For large problems, divide and conquer technique is usually used, e.g. Sequential Minimal Optimization algorithm (J. Platt, 1998, <http://research.microsoft.com/users/jplatt/smo.html>)

Complexity, Scalability, and Parallelizability

- Striking property: both training and test functions depend on the data only through the kernel functions $K(x_i, x_j)$. Even though it corresponds to a dot product in a space of high dimension d_H . Only $O(d_L)$ operations is required to compute the dot product. d_L is the dimension of the original data.
- Training time (Bunch-Kaufman) (Kaufman, 1998). L is the number of training points, N_s the number of support vectors, d_L the dimension of the input data. In the case where most SVs are not at the upper bound, and $N_s / L \ll L$, the number of operations C is $(N_s^3 + (N_s^2)L + N_s d_L L)$. If Instead $N_s / L \approx 1$, then C is $(N_s^3 + N_s^2 L + N_s d_L L)$. For the case where most SVs are at the upper bound, and $N_s / L \ll 1$, then C is $O(N_s^2 + N_s d_L L)$. Finally if most SVs are at the upper bound, and $N_s / L \approx 1$, we have C of $O(d_L L^2)$.

Time Complexity of Testing

- $O(MN_s)$. M is the number of operations required to evaluate the kernel. For RBF kernel, M is $O(d_L)$. N_s is the number of support vectors.

A Bound from Leave-One-Out

- $E[P(\text{error})] = N_s / \text{number of training samples}$, where N_s is the number of support vectors
- What does this tells us?

Limitations and Extensions

- Choice of kernel. Once the kernel is fixed, SVM classifiers have only one user-chosen parameter (the error penalty) and kernel parameters
- Speed in test phase (Burges 96, Burges and Scholkopf 97, speed up 50 times)
- Challenge: Training for very large datasets (millions of data points)

SVM Tools

- SVM-light: <http://svmlight.joachims.org/>
- LIBSVM:
<http://www.csie.ntu.edu.tw/~cjlin/libsvm/>
- Gist: <http://bioinformatics.ubc.ca/gist/>
- More:
[http://www.kernel-machines.org/
software.html](http://www.kernel-machines.org/software.html)

Acknowledgements

- **Chris Burges**'s excellent tutorial. A Tutorial on Support Vector Machines for Pattern Recognition. *Data Mining and Knowledge Discovery*, 1998.
- **Andrew Moore**'s SVM Slides.