8 Generalized Support Vector Machines

Olvi L. Mangasarian
Computer Sciences Department
University of Wisconsin
1210 West Dayton Street
Madison, WI 53706, USA
olvi@cs.wisc.edu
http://www.cs.wisc.edu/∼olvi

By setting apart the two functions of a support vector machine: separation of points by a nonlinear surface in the original space of patterns, and maximizing the distance between separating planes in a higher dimensional space, we are able to define indefinite, possibly discontinuous, kernels, not necessarily inner product ones, that generate highly nonlinear separating surfaces.

Maximizing the distance between the separating planes in the higher dimensional space is surrogated in the present approach by support vector suppression, which is achieved by minimizing some desired norm of their Lagrange multipliers. The norm may be one induced by the separation kernel if it happens to be positive definite, or a Euclidean or a polyhedral norm (*i.e.* a norm induced by a bounded polyhedron such as the 1-norm or the ∞-norm). Polyhedral norms lead to linear programs whereas Euclidean norms lead to convex quadratic programs, all with an arbitrary separation kernel.

A standard support vector machine can be recovered by using the same kernel for separation and support vector suppression. On a simple test example, all models perform equally well when a positive definite kernel is used. When a negative definite kernel is used, we are unable to solve the nonconvex quadratic program associated with a conventional support vector machine, while all other proposed models remain convex and easily generate a surface that separates all given points.

8.1 Introduction

Support vector machines (Vapnik, 1995; Bennett and Blue, 1997; Girosi, 1998; Wahba, 1999b; Cherkassky and Mulier, 1998; Schölkopf, 1997; Smola, 1998) attempt to separate points belonging to two given sets in N-dimensional real (Euclidean) space $\mathbb{R}^N$ by a nonlinear surface, often only implicitly defined by a kernel function. In our approach here the nonlinear surface in the original input space $\mathbb{R}^N$ which is defined linearly in its parameters, can be represented as a linear function (plane) in a higher, often much higher dimensional feature space, say $\mathbb{R}^\ell$. Also, the original points of the two given sets can also be mapped into this higher dimensional space. If the two sets are linearly separable in $\mathbb{R}^\ell$, then it is intuitively plausible to generate a plane mid-way between the furthest parallel planes apart that bound the two sets. Using a distance induced by the kernel generating the nonlinear surface in $\mathbb{R}^N$, it can be shown (Vapnik and Lerner, 1963) that such a plane optimizes the generalization ability of the separating plane. If the two sets are not linearly separable, a similar approach can be used (Cortes and Vapnik, 1995; Vapnik, 1995) to maximize the distance between planes that bound each set with certain minimal error. Linear separation by planes with maximal 1-norm distance and by nonlinear surfaces were proposed earlier in (Mangasarian, 1965, 1968) as well as with soft margins (Bennett and Mangasarian, 1992)

In this paper we start with a nonlinear separating surface (8.1), defined by some arbitrary kernel k and by some linear parameters $u \in \mathbb{R}^m$, to be determined, that turn out to be closely related to some dual variables. Based on this surface we derive a general convex mathematical program (8.5) that attempts separation via the nonlinear surface (8.1) while minimizing some function θ of the parameters u. In our formulation here, the function θ which attempts to suppress u can be interpreted as minimizing the number of support vectors, or under more conventional assumptions as maximizing the distance between the separating planes in $\mathbb{R}^\ell$. The choice of θ leads to various support vector machines. We consider two classes of such machines based on whether θ is quadratic or piecewise linear. If we choose θ to be a quadratic function generated by the kernel defining the nonlinear surface (8.1), then we are led to the conventional dual quadratic program (8.9) associated with a support vector machine which requires positive definiteness of this kernel. However the quadratic function choice for θ can be divorced from the kernel defining the separating surface and this leads to other convex quadratic programs such as (8.10) *without* making any assumptions on the kernel. In (Smola and Schölkopf, 1998b) techniques for dealing with kernels that are not positive semidefinite were also presented. Another class of support vector machines are generated by choosing a piecewise linear convex function for θ and this leads to linear programs such as (8.11) and (8.12), both of which make no assumptions on the kernel. In Section 8.5 we give some simple applications of all four formulations to the Exclusive-Or (XOR) problem using first a positive definite second-order polynomial kernel and then a negative definite third-order polynomial kernel. For the positive definite kernel all four convex formulations

are easily solved and the resulting nonlinear surfaces separate all points in all cases. However, for the negative definite kernel, a powerful state-of-the-art package fails to solve the nonconvex quadratic program associated with the conventional support vector machine, whereas all other three convex formulations are easily solved and lead to complete separation of the data by the nonlinear surface.

A word about our notation and background material. All vectors will be column vectors unless transposed to a row vector by a prime superscript $'$. For a vector s in the N-dimensional real space $\mathbb{R}^N$, the step function s_* of $s \in \mathbb{R}^N$ is defined as a vector of ones and zeros in $\mathbb{R}^N$, with ones corresponding to positive components of s and zeros corresponding to nonpositive components. The scalar (inner) product of two vectors s and y in the N-dimensional real space $\mathbb{R}^N$ will be denoted by $s \cdot y$. For an $m \times N$ matrix X, X_i will denote the ith row of X and $X_{\cdot j}$ will denote the jth column of X. The identity matrix in a real space of arbitrary dimension will be denoted by I, while a column vector of ones of arbitrary dimension will be denoted by e. We shall employ the MATLAB (1992) "dot" notation to signify application of a function to all components of a matrix or a vector. For example if $X \in \mathbb{R}^{m \times N}$, then $X_\bullet^2 \in \mathbb{R}^{m \times N}$ will denote the matrix of elements of X squared.

We begin by defining a general kernel function as follows.

general kernels with no assumptions

Definition 8.1
Let $X \in \mathbb{R}^{m \times N}$ and $B \in \mathbb{R}^{N \times \ell}$. The **kernel** $k(X, B)$ maps $\mathbb{R}^{m \times N} \times \mathbb{R}^{N \times \ell}$ into $\mathbb{R}^{m \times \ell}$.

In particular if s and t are column vectors in $\mathbb{R}^N$ then, $k(s', X')$ is a row vector in $\mathbb{R}^m$, $k(s', t)$ is a real number and $k(X, X')$ is an $m \times m$ matrix. Note that for our purposes here $k(X, X')$ need not be symmetric in general. Examples of kernels are given in the introduction to the book and below where $a \in \mathbb{R}^m$, $b \in \mathbb{R}^\ell$, $\mu \in R$ and d is an integer. For simplicity we restrict ourselves here to finite dimensional kernels although many of the results can be extended to infinite dimensional Hilbert spaces (Kimeldorf and Wahba, 1971).

Example 8.1
Polynomial Kernel $(XB + \mu ab')_\bullet^d$

Example 8.2
Neural Network Step Kernel $(XB + \mu ab')_{\bullet *}$

Example 8.3
Radial Basis Kernel $e^{-\mu \|X_i' - B_{\cdot j}\|^2}$, $i, j = 1, \ldots, m$, $\ell = m$,
where, here only, e is the base of the natural logarithm.

Note that our approach allows discontinuous kernels such as the neural network step kernel with a discontinuous step function without the need for a smoothing approximation such as the sigmoid or hyperbolic tangent approximation as is usually done (Vapnik, 1995; Cherkassky and Mulier, 1998).

8.2 GSVM: The General Support Vector Machine

We consider a given set $\mathcal{A}$ of m points in real N-dimensional space of features $\mathbb{R}^N$ represented by the matrix $X \in \mathbb{R}^{m\times N}$. Each point X_i, $i = 1,\ldots,m$, belongs to class 1 or class -1 depending on whether Y_{ii} is 1 or -1, where $Y \in \mathbb{R}^{m\times m}$ is a given diagonal matrix of plus or minus ones. We shall attempt to discriminate between the classes 1 and -1 by a nonlinear *separating surface*, induced by some kernel $k(X, X')$, as follows:

GSVM separating surface

$$k(s', X')Y \cdot u = b, \tag{8.1}$$

where $k(s', X') \in \mathbb{R}^m$, according to Definition 8.1. The parameters $u \in \mathbb{R}^m$ and $b \in R$ are determined by solving a mathematical program, typically quadratic or linear. A point $s \in \mathbb{R}^N$ is classified in class 1 or -1 according to whether the *decision function*

$$(k(s', X')Yu - b)_*, \tag{8.2}$$

yields 1 or 0 respectively. The kernel function $k(s', X')$ defines a nonlinear map from $s \in \mathbb{R}^N$ to some other space $s \in \mathbb{R}^\ell$ where ℓ may be much larger than N. In particular if the kernel k is an inner product kernel under Mercer's condition (Courant and Hilbert, 1953, pp 138-140),(Vapnik, 1995; Cherkassky and Mulier, 1998; Burges, 1998) (an assumption that we will not make in this paper) then for s and t in $\mathbb{R}^N$:

$$k(s, t) = \Phi(s) \cdot \Phi(t), \tag{8.3}$$

and the separating surface (8.1) becomes:

$$\Phi(s)'\Phi(X')Y \cdot u = b, \tag{8.4}$$

where Φ is a function, not easily computable, from $\mathbb{R}^N$ to $\mathbb{R}^\ell$, and $\Phi(X') \in \mathbb{R}^{\ell\times m}$ results from applying Φ to the m columns of X'. The difficulty in computing Φ and the possible high dimensionality of $\mathbb{R}^\ell$ have been important factors in using a kernel k as a generator of a nonlinear separating surface in the original feature space $\mathbb{R}^N$ but which is linear in the high dimensional space $\mathbb{R}^\ell$. Our separating surface (8.1) written in terms of a kernel function retains this advantage and is linear in its parameters, u, b. We now state a mathematical program that generates such a surface for a general kernel k as follows:

the GSVM

$$\begin{array}{lrcl} \min\limits_{u,b,\xi} & Ce \cdot \xi + \theta(u) & & \\ \text{s.t.} & Y(k(X, X')Yu - eb) + \xi & \geq & e \\ & \xi & \geq & 0. \end{array} \tag{8.5}$$

Here θ is some convex function on $\mathbb{R}^m$, typically some norm or seminorm, and C is some positive parameter that weights the separation error $e \cdot \xi$ versus suppression of the separating surface parameter u. Suppression of u, utilized in (Bradley and

Mangasarian, 1998) for feature selection, can be interpreted in one of two ways. We interpret it here as minimizing the number of support vectors, i.e. constraints of (8.5) with positive multipliers. A more conventional interpretation is that of maximizing some measure of the distance or margin between the bounding parallel planes in $\mathbb{R}^\ell$, under appropriate assumptions, such as θ being a quadratic function induced by a positive definite kernel k as in (8.9) below. As is well known, this leads to improved generalization by minimizing an upper bound on the VC dimension (Vapnik, 1995; Schölkopf, 1997). Girosi et al. (1993) used a quadratic function for the regularization term θ while Smola (1998) used linear and quadratic terms.

We term a solution of the mathematical program (8.5) and the resulting separating surface (8.1) and corresponding decision function (8.2) a *generalized support vector machine*, GSVM. In the following sections of the paper we derive a number of special cases, including the standard support vector machine. First, however, it is important to state under what conditions does the mathematical program (8.5) have a solution.

Proposition 8.2
Existence of a GSVM For any given $X \in \mathbb{R}^{m\times N}$, any $Y \in \mathbb{R}^{m\times m}$, $C > 0$ and any kernel k, the mathematical program (8.5) has a solution whenever θ is a piecewise-linear or quadratic function bounded below on $\mathbb{R}^m$.

Proof The feasible region of (8.5) is always nonempty. Just take: $u = 0$, $b = 0$ and $\xi = e$. When θ is piecewise-linear, existence follows from the standard linear programming result, that a feasible linear program with a bounded objective has a solution. Just apply this result to each piece of the objective on its polyhedral region. For a quadratic θ the result is a direct consequence of the Frank-Wolfe existence result for quadratic programming (Frank and Wolfe, 1956). ■

We note that no convexity of θ was needed for this existence result. However in our specific applications where duality theory will be invoked, θ will need to be convex.

8.3 Quadratic Programming Support Vector Machines

We consider in this section support vector machines that include the standard ones (Vapnik, 1995; Cherkassky and Mulier, 1998; Burges, 1998) and which are obtained by setting θ of (8.5) to be a convex quadratic function $\theta(u) = \frac{1}{2}u \cdot Hu$, where $H \in \mathbb{R}^{m\times m}$ is some symmetric positive definite matrix. The mathematical program (8.5) becomes the following convex quadratic program:

$$\begin{array}{lll} \min\limits_{u,b,\xi} & Ce\cdot\xi + \frac{1}{2}u\cdot Hu & \\ \text{s.t.} & Y(k(X,X')Yu - eb) + \xi \;\geq\; e & \\ & \xi \;\geq\; 0. & \end{array} \tag{8.6}$$

dual QP GSVM

The Wolfe (1961) dual (Mangasarian, 1994) of this convex quadratic program is:

$$\begin{array}{lrcl} \min\limits_{\boldsymbol{\alpha}\in\mathbb{R}^m} & \frac{1}{2}\boldsymbol{\alpha}\cdot Yk(X,X')YH^{-1}Yk(X,X')'Y\boldsymbol{\alpha} - e\cdot\boldsymbol{\alpha} & & \\ \text{s.t.} & e\cdot Y\boldsymbol{\alpha} & = & 0 \\ & 0\le\boldsymbol{\alpha} & \le & Ce. \end{array} \tag{8.7}$$

Furthermore, the primal variable u is related to the dual variable $\boldsymbol{\alpha}$ by:

$$u = H^{-1}Yk(X,X')'Y\boldsymbol{\alpha}. \tag{8.8}$$

Equations (8.6)-(8.8) have also been given by Smola and Schölkopf (1998b). If we assume that the kernel $k(X,X')$ is symmetric positive definite and let $H = Yk(X,X')Y$, then our problem (8.6) becomes the standard SVM problem and our dual problem (8.7) degenerates to the dual problem of the standard support vector machine (Vapnik, 1995; Cherkassky and Mulier, 1998; Burges, 1998) with $u = \boldsymbol{\alpha}$:

$$\begin{array}{lrcl} \min\limits_{\boldsymbol{\alpha}\in\mathbb{R}^m} & \frac{1}{2}\boldsymbol{\alpha}\cdot Yk(X,X')Y\boldsymbol{\alpha} - e\cdot\boldsymbol{\alpha} & & \\ \text{s.t.} & e\cdot Y\boldsymbol{\alpha} & = & 0 \\ & 0\le\boldsymbol{\alpha} & \le & Ce. \end{array} \tag{8.9}$$

dual SVM

The positive definiteness assumption on $k(X,X')$ in (8.9) can be relaxed to positive *semi*definiteness while maintaining the convex quadratic program (8.6), with $H = Yk(X,X')Y$, as the direct dual of (8.9) without utilizing (8.7) and (8.8). The symmetry and positive semidefiniteness of the kernel $k(X,X')$ for this version of a support vector machine is consistent with the support vector machine literature. The fact that $\boldsymbol{\alpha} = u$ in the dual formulation (8.9), shows that the variable u appearing in the original formulation (8.6) is also the dual multiplier vector for the first set of constraints of (8.6). Hence the quadratic term in the objective function of (8.6) can be thought of as suppressing as many multipliers of support vectors as possible and thus minimizing the number of such support vectors. This is another interpretation of the standard support vector machine that is usually interpreted as maximizing the margin or distance between parallel separating planes.

This leads to the idea of using other values for the matrix H other than $Yk(X,X')Y$ that will also suppress u. One particular choice is interesting because it puts no restrictions on k: no symmetry, no positive definiteness or semidefiniteness and not even continuity. This is the choice $H = I$ in (8.6) which leads to a dual problem (8.7) with $H = I$ and $u = Yk(X,X')'Y\boldsymbol{\alpha}$ as follows:

arbitrary-kernel dual QP GSVM

$$\begin{array}{lrcl} \min\limits_{\boldsymbol{\alpha}\in\mathbb{R}^m} & \frac{1}{2}\boldsymbol{\alpha}\cdot Yk(X,X')k(X,X')'Y\boldsymbol{\alpha} - e'\boldsymbol{\alpha} & & \\ \text{s.t.} & e\cdot Y\boldsymbol{\alpha} & = & 0 \\ & 0\le\boldsymbol{\alpha} & \le & Ce. \end{array} \tag{8.10}$$

Note that setting $H = I$ corresponds to weight decay in neural networks and as shrinkage estimators in statistics (Smola and Schölkopf, 1998b). We also note that $k(X,X')k(X,X')'$ is positive semidefinite with no assumptions on $k(X,X')$, and hence the above problem is an always solvable convex quadratic program

for any kernel $k(X, X')$. In fact by Proposition 8.2 the quadratic program (8.6) is solvable for *any* symmetric positive semidefinite matrix H, and by quadratic programming duality so is its dual problem (8.7), the solution $\boldsymbol{\alpha}$ of which and $u = Yk(X, X')'Y\boldsymbol{\alpha}$ can be immediately used to generate a decision function (8.2). Thus we are free to choose any symmetric positive definite matrix H to generate a support vector machine. Experimentation will be needed to determine what are the most appropriate choices for H.

Note that even though (8.10) differs from (8.9) merely in the replacement of the kernel $k(X, X')$ by $k(X, X')k(X, X')'$, the two problems lead to distinct separating surfaces (8.1) because $u = \boldsymbol{\alpha}$ for (8.9) and $u = Yk(X, X')'Y\boldsymbol{\alpha}$ for (8.10). Note that some sparsity in the solution $\boldsymbol{\alpha}$ of (8.10) may be lost with the product kernel if it is not positive definite.

We turn our attention to linear programming support vector machines.

8.4 Linear Programming Support Vector Machines

In this section we consider problems generated from the mathematical program (8.5) by using a piecewise linear function θ in the objective function thus leading to linear programs.

LP1 GSVM with arbitrary kernel

The most obvious choice for θ is the 1-norm of u, which leads to the following linear programming formulation:

$$\begin{array}{rrcl} \min\limits_{u,b,\xi,t} & Ce \cdot \xi + e \cdot t & & \\ \text{s.t.} & Y(k(X, X')Yu - eb) + \xi & \geq & e \\ & t \geq u & \geq & -t \\ & \xi & \geq & 0. \end{array} \tag{8.11}$$

A solution (u, b, ξ, t) to this linear program for a chosen kernel $k(X, X')$ will provide a decision function as given by (8.2). This linear program parallels the quadratic programming formulation (8.10) that was obtained as the dual of (8.5) by setting $\theta(u)$ therein to half the 2-norm squared of u whereas $\theta(u)$ is set to the 1-norm of u in (8.11). Another linear programming formulation that somewhat parallels the quadratic programming formulation (8.9), which was obtained as the dual of (8.5) by setting $\theta(u)$ therein to half the 2-norm squared of $k(X, X')^{\frac{1}{2}}Yu$, is obtained setting θ to be the 1-norm of $k(X, X')Yu$. The motivation for this idea is to try capturing a norm induced by $k(X, X')$ even when the kernel is not positive semidefinite. This leads to the following linear program:

LP2 GSVM with arbitrary kernel

$$\begin{array}{rrcl} \min\limits_{u,b,\xi,t} & Ce \cdot \xi + e \cdot t & & \\ \text{s.t.} & Y(k(X, X')Yu - eb) + \xi & \geq & e \\ & t \geq k(X, X')Yu & \geq & -t \\ & \xi & \geq & 0. \end{array} \tag{8.12}$$

No assumptions of symmetry or positive definiteness on $k(X, X')$ are needed in either of the above linear programming formulations as was the case in the quadratic program (8.9).

It is interesting to note that if the linear kernel $k(X, X') = XX'$ is used in the linear program (8.11) we obtain the high-performing 1-norm linear SVM proposed by Bredensteiner and Bennett (1997) and utilized successfully in (Bredensteiner, 1997; Bennett et al., 1998; Bradley and Mangasarian, 1998). Hence, if we set $w = X'Yu$ in (8.11) we obtain (Bradley and Mangasarian, 1998, Equation (13)).

8.5 A Simple Illustrative Example

We first demonstrate the workings and sometimes different, yet equally effective, decision surfaces obtained by the various proposed mathematical programming formulations, for a positive definite symmetric kernel. We then show that for a negative definite symmetric kernel, the conventional support vector machine fails to generate a decision function that correctly separates the given points, whereas all the new formulations do succeed in generating a decision surface that correctly separates all the given points.

For our positive definite kernel we use a polynomial kernel of order 2, based on Example 8.1 with $B = X'$, $\mu = 1$, $a = b = e$ and $d = 2$, and apply it to the classical Exclusive-Or (XOR) problem. We thus have:

$$X = \begin{bmatrix} 1 & 1 \\ 1 & -1 \\ -1 & -1 \\ -1 & 1 \end{bmatrix}, \; Y = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}. \tag{8.13}$$

Hence with the MATLAB (1992) "dot" notation signifying componentwise exponentiation we get that:

$$k(X, X') = (XX' + ee')^2_\bullet = \begin{bmatrix} 9 & 1 & 1 & 1 \\ 1 & 9 & 1 & 1 \\ 1 & 1 & 9 & 1 \\ 1 & 1 & 1 & 9 \end{bmatrix}, \tag{8.14}$$

and

$$k(s', X') = (s'X' + e')^2_\bullet = [s_1 + s_2 + 1 \;\; s_1 - s_2 + 1 \;\; -s_1 - s_2 + 1 \;\; -s_1 + s_2 + 1]^2_\bullet. \tag{8.15}$$

Solution of the linear program (8.11) with $C \geq 1$ gives:

$$u = \frac{1}{8}e, \; b = 0, \; \xi = 0, \; t = \frac{1}{8}e, \; Ce \cdot \xi + e \cdot t = \frac{1}{2}. \tag{8.16}$$

Note that the $\xi = 0$ in the above solution means that the decision surface correctly classifies all given points represented by X and Y. Solution of either quadratic program (8.9) or (8.10) with the same kernel and for $C \geq 1$ also yields $u = \frac{1}{8}e$. Substitution of this u in (8.6) and solving the resulting linear program gives the same b, ξ as in (8.16). Thus all mathematical programs (8.9), (8.10) and (8.11) yield exactly the same decision surface (8.2):

$$(k(s', X')Yu - b)_* = ((s'X' + e')^2_\bullet Yu - b)_* = (s_1 s_2)_*, \tag{8.17}$$

a step function of the quadratic function $s_1 s_2$, which correctly classifies the two categories class 1 and class -1 and is in agreement with the solution obtained in (Cherkassky and Mulier, 1998, pages 372-375) for the conventional support vector machine (8.9). Note that neither mathematical program (8.10) or (8.11) required positive definiteness of $k(X, X')$, whereas (8.9) does.

However, it is rather interesting to observe that the linear programming solution (8.16) is not unique. In fact another solution is the following:

$$u = \begin{bmatrix} 0 \\ \frac{1}{4} \\ 0 \\ \frac{1}{4} \end{bmatrix}, \; b = -\frac{3}{2}, \; \xi = 0, \; t = \begin{bmatrix} 0 \\ \frac{1}{4} \\ 0 \\ \frac{1}{4} \end{bmatrix}, \; Ce \cdot \xi + e \cdot t = \frac{1}{2}. \tag{8.18}$$

For this solution the decision surface (8.2) turns out to be:

$$(k(z', X')Yu - b)_* = ((s'X' + e')^2_\bullet Yu - b)_* = \frac{1}{2}(2 - (s_1 - s_2)^2)_*. \tag{8.19}$$

This decision surface is rather different from that of (8.17), but it does separate the two classes correctly and in fact it consists of two parallel lines separating $\mathbb{R}^2$ into 3 regions, whereas (8.17) separates $\mathbb{R}^2$ into four quadrants each pair of which contains one class. Both of these decision functions are depicted in Figure 8.1 and Figure 8.2.

Solution of the linear program (8.12) with $C > 1$ yields:

$$u = \begin{bmatrix} \frac{1}{24} \\ \frac{5}{24} \\ \frac{1}{24} \\ \frac{5}{24} \end{bmatrix}, \; b = -1, \; \xi = 0, \; t = \begin{bmatrix} 0 \\ 2 \\ 0 \\ 2 \end{bmatrix}, \tag{8.20}$$

which gives the decision surface:

$$((s'X' + e')^2_\bullet Yu - b)_* = \frac{1}{3}(2 + s_1 s_2 - (s_1 - s_2)^2)_*. \tag{8.21}$$

This decision function divides $\mathbb{R}^2$ into three regions by a pair of "square root" curves that correctly classify the two classes as depicted in Figure 8.3.

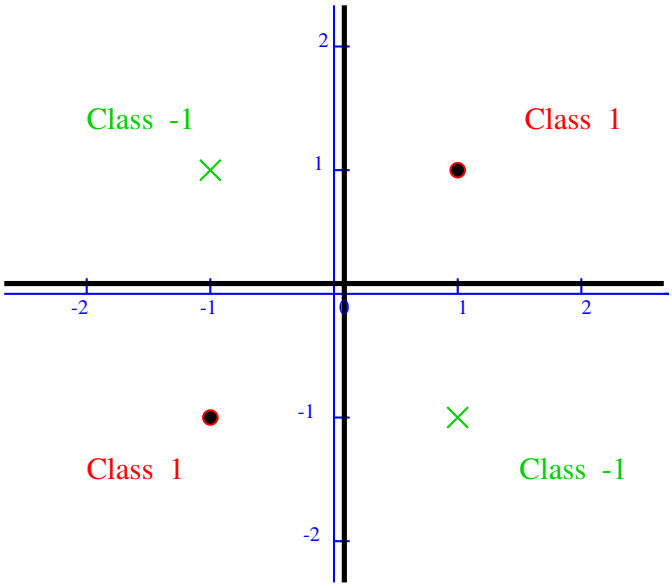

Figure 8.1 XOR discrimination by a step function of a quadratic function: $(s_1 s_2)_*$ obtained by the linear program (8.11) and the quadratic programs (8.9) and (8.10).

Finally in order to show that positive definiteness of the kernel $k(X, X')$ is not essential in any of our new mathematical programming formulations (8.10), (8.11) or (8.12), and whereas it is in the conventional quadratic programming formulation (8.9), we consider the following negative definite kernel:

$$k(X, X') = (-XX' - ee')^3_\bullet, \tag{8.22}$$

and attempt to solve the mathematical programs (8.9),(8.10), (8.11) and (8.12) with this kernel and with $C = 1$. The powerful PATH mathematical programming package (Ferris and Munson, 1999; Dirkse and Ferris, 1995) failed to solve the nonconvex quadratic programming formulation (8.9) for the conventional support vector machine. (PATH is a Newton-based method for solving a nonsmooth systems of equations that subsume the Karush Kuhn Tucker conditions of quadratic and nonlinear programming.) In contrast, the same package solved the quadratic program (8.10) giving $\boldsymbol{\alpha} = \frac{1}{576}e$ and a corresponding $u = Yk(X, X')'Y\boldsymbol{\alpha} = -\frac{1}{24}e$. Substitution of this u in the quadratic program (8.6) and solving the resulting linear program gives: $b = 0,\ \xi = 0$. The solution $\xi = 0$ indicates that all points represented by X have been correctly classified, which is corroborated by the resulting decision surface $(s_1 s_2)_*$, the same as that of (8.17). This indicates the effectiveness of the quadratic program (8.10) in its ability to extract from the negative definite cubic kernel just the required quadratic term to achieve correct separation. Similarly both linear programs (8.11) and (8.12) gave $\xi = 0$ thus also achieving complete separation with this negative definite kernel.

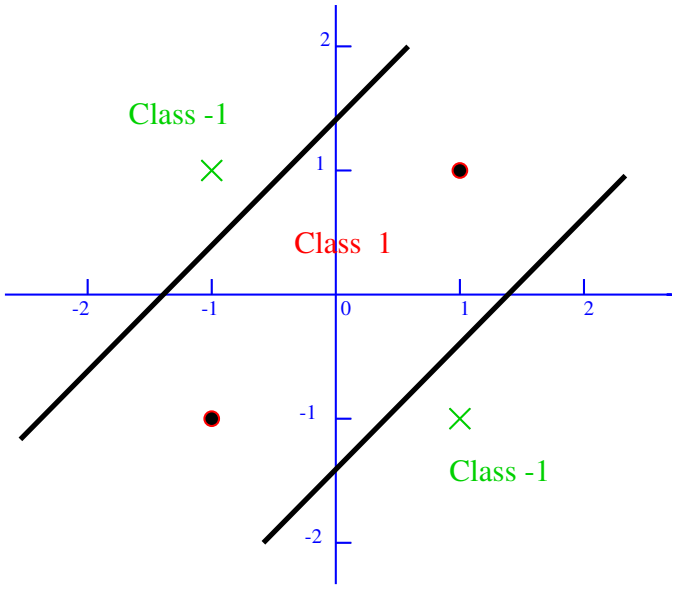

Figure 8.2 XOR discrimination by a step function of a quadratic function: $(2 - (s_1 - s_2)^2)_*$ obtained by another solution of the linear program (8.11).

8.6 Conclusion

We have proposed a direct mathematical programming framework for general support vector machines that makes essentially no or few assumptions on the kernel employed. We have derived new kernel-based linear programming formulations (8.11) and (8.12), and a new quadratic programming formulation (8.10) that require no assumptions on the kernel k. These formulations can lead to effective but different decision functions from those obtained by the quadratic programming formulation (8.9) for a conventional support vector machine that requires symmetry and positive definiteness of the kernel. Even for negative definite kernels these new formulations can generate decision functions that separate the given points whereas the conventional support vector machine does not. This leads us to suggest that further testing and experimentation with mathematical programming formulations such as (8.11), (8.12) and (8.10) and others are worthwhile. These formulations may open the way for a variety of support vector machines that could be tested computationally against each other and against existing ones. Furthermore, broad classes of serial and parallel optimization algorithms such as (Bennett and Mangasarian, 1994) can be brought to bear on these different formulations exploiting their structure in order to permit the processing of massive databases.

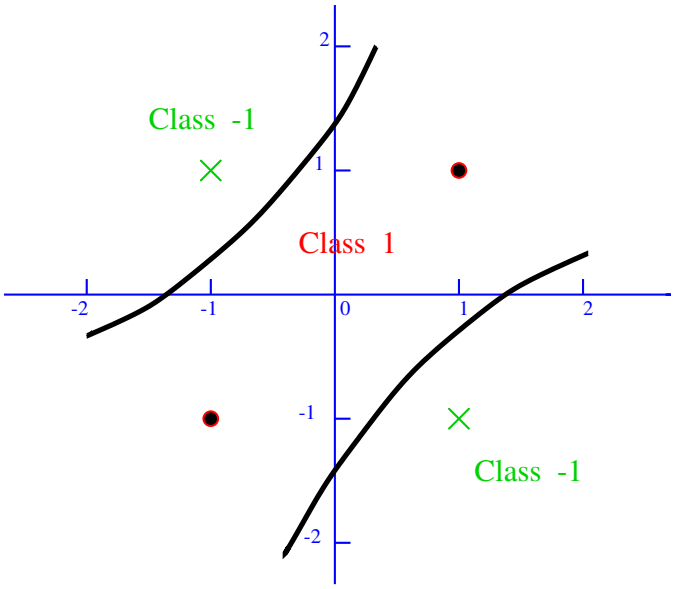

Figure 8.3 XOR discrimination by a step function of a quadratic function: $(2 + s_1 s_2 - (s_1 - s_2)^2)_*$ obtained by the linear program (8.12).

Acknowledgments

I am grateful to my colleague Grace Wahba for very helpful discussions on kernel theory, to my PhD student David Musicant and to Paul Bradley for important comments on the paper.

This research is supported by National Science Foundation Grants CCR-9322479, CCR-9729842 and CDA-9623632, and by Air Force Office of Scientific Research Grant F49620-97-1-0326 as Mathematical Programming Technical Report 98-14, October 1998.