

## Lagrange duality

Another way to arrive at the KKT conditions, and one which gives us some insight on solving constrained optimization problems, is through the *Lagrange dual*. The dual is a maximization program in  $\boldsymbol{\lambda}, \boldsymbol{\nu}$  — it is always concave (even when the original program is not convex), and gives us a systematic way to lower bound the optimal value.

### The Lagrangian

We consider an optimization program of the form

$$\begin{aligned} \underset{\boldsymbol{x} \in \mathbb{R}^N}{\text{minimize}} \quad & f_0(\boldsymbol{x}) \quad f_m(\boldsymbol{x}) \leq 0, \quad m = 1, \dots, M \\ & h_p(\boldsymbol{x}) = 0, \quad p = 1, \dots, P. \end{aligned} \quad (1)$$

Much of what we will say below applies equally well to nonconvex programs as well as convex programs, so we will make it clear when we are taking the  $f_m$  to be convex and the  $h_p$  to be affine. We will take the domain of all of the  $f_m$  and  $h_p$  to be all of  $\mathbb{R}^N$  below; this just simplifies the exposition, we can easily replace this with the intersections of the  $\text{dom } f_m$  and  $\text{dom } h_p$ . We will assume that the intersection of the feasible set,

$$\mathcal{C} = \{\boldsymbol{x} : f_m(\boldsymbol{x}) \leq 0, h_p(\boldsymbol{x}) = 0, m = 1, \dots, M, p = 1, \dots, P\}$$

is a non-empty and a subset  $\mathbb{R}^N$ .

The **Lagrangian** takes the constraints in the program above and integrates them into the objective function. The Lagrangian  $L : \mathbb{R}^N \times \mathbb{R}^M \times \mathbb{R}^P \rightarrow \mathbb{R}$  associated with this optimization program is

$$L(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\nu}) = f_0(\mathbf{x}) + \sum_{m=1}^M \lambda_m f_m(\mathbf{x}) + \sum_{p=1}^P \nu_p h_p(\mathbf{x})$$

The  $\mathbf{x}$  above are referred to as **primal variables**, and the  $\boldsymbol{\lambda}, \boldsymbol{\nu}$  as either **dual variables** or **Lagrange multipliers**.

The **Lagrange dual function**  $g(\boldsymbol{\lambda}, \boldsymbol{\nu}) : \mathbb{R}^M \times \mathbb{R}^P \rightarrow \mathbb{R}$  is the minimum of the Lagrangian over all values of  $\mathbf{x}$ :

$$g(\boldsymbol{\lambda}, \boldsymbol{\nu}) = \inf_{\mathbf{x} \in \mathbb{R}^N} \left( f_0(\mathbf{x}) + \sum_{m=1}^M \lambda_m f_m(\mathbf{x}) + \sum_{p=1}^P \nu_p h_p(\mathbf{x}) \right).$$

Since the dual is a pointwise infimum of a family of affine functions in  $\boldsymbol{\lambda}, \boldsymbol{\nu}$ ,  $g$  is **concave** regardless of whether or not the  $f_m, h_p$  are convex.

The key fact about the dual function is that it is everywhere a lower bound on the optimal value of the original program. If  $p^*$  is the optimal value for (1), then

$$g(\boldsymbol{\lambda}, \boldsymbol{\nu}) \leq p^*, \quad \text{for all } \boldsymbol{\lambda} \geq \mathbf{0}, \boldsymbol{\nu} \in \mathbb{R}^P.$$

This is (almost too) easy to see. For any feasible point  $\mathbf{x}_0$ ,

$$\sum_{m=1}^M \lambda_m f_m(\mathbf{x}_0) + \sum_{p=1}^P \nu_p h_p(\mathbf{x}_0) \leq 0,$$

and so

$$L(\mathbf{x}_0, \boldsymbol{\lambda}, \boldsymbol{\nu}) \leq f_0(\mathbf{x}_0), \quad \text{for all } \boldsymbol{\lambda} \geq \mathbf{0}, \boldsymbol{\nu} \in \mathbb{R}^P,$$

meaning

$$g(\boldsymbol{\lambda}, \boldsymbol{\nu}) = \inf_{\mathbf{x} \in \mathbb{R}^N} L(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\nu}) \leq L(\mathbf{x}_0, \boldsymbol{\lambda}, \boldsymbol{\nu}) \leq f_0(\mathbf{x}_0).$$

Since this holds for all feasible  $\mathbf{x}_0$ ,  $g(\boldsymbol{\lambda}, \boldsymbol{\nu}) \leq \inf_{\mathbf{x} \in \mathcal{C}} f_0(\mathbf{x}) = p^*$ .

The (Lagrange) dual to the optimization program (1) is

$$\underset{\boldsymbol{\lambda} \in \mathbb{R}^M, \boldsymbol{\nu} \in \mathbb{R}^P}{\text{maximize}} \quad g(\boldsymbol{\lambda}, \boldsymbol{\nu}) \quad \text{subject to} \quad \boldsymbol{\lambda} \geq \mathbf{0}. \quad (2)$$

The dual optimal value  $d^*$  is

$$d^* = \sup_{\boldsymbol{\lambda} \geq \mathbf{0}, \boldsymbol{\nu}} g(\boldsymbol{\lambda}, \boldsymbol{\nu}) = \sup_{\boldsymbol{\lambda} \geq \mathbf{0}, \boldsymbol{\nu}} \inf_{\mathbf{x} \in \mathbb{R}^N} L(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\nu}).$$

Since  $g(\boldsymbol{\lambda}, \boldsymbol{\nu}) \leq p^*$ , we know that

$$d^* \leq p^*.$$

The quantity  $p^* - d^*$  is called the **duality gap**. If  $p^* = d^*$ , then we say that (1) and (2) exhibit **strong duality**.

## Certificates of (sub)optimality

Any dual feasible<sup>1</sup>  $(\boldsymbol{\lambda}, \boldsymbol{\nu})$  gives us a lower bound on  $p^*$ , since  $g(\boldsymbol{\lambda}, \boldsymbol{\nu}) \leq p^*$ . If we have a primal feasible  $\boldsymbol{x}$ , then we know that

$$f_0(\boldsymbol{x}) - p^* \leq f_0(\boldsymbol{x}) - g(\boldsymbol{\lambda}, \boldsymbol{\nu}).$$

We will refer to  $f_0(\boldsymbol{x}) - g(\boldsymbol{\lambda}, \boldsymbol{\nu})$  as the **duality gap** for primal/dual (feasible) pair  $\boldsymbol{x}, \boldsymbol{\lambda}, \boldsymbol{\nu}$ . We know that

$$p^* \in [g(\boldsymbol{\lambda}, \boldsymbol{\nu}), f_0(\boldsymbol{x})], \quad \text{and likewise} \quad d^* \in [g(\boldsymbol{\lambda}, \boldsymbol{\nu}), f_0(\boldsymbol{x})].$$

If we are ever able to reduce this gap to zero, then we know that  $\boldsymbol{x}$  is primal optimal, and  $\boldsymbol{\lambda}, \boldsymbol{\nu}$  are dual optimal.

There are certain kinds of “primal-dual” algorithms that produce a series of (feasible) points  $\boldsymbol{x}^{(k)}, \boldsymbol{\lambda}^{(k)}, \boldsymbol{\nu}^{(k)}$  at every iteration. We can then use

$$f_0(\boldsymbol{x}^{(k)}) - g(\boldsymbol{\lambda}^{(k)}, \boldsymbol{\nu}^{(k)}) \leq \epsilon,$$

as a stopping criteria, and know that our answer would yield an objective value no further than  $\epsilon$  from optimal.

## Strong duality and the KKT conditions

Suppose that for a convex program, the primal optimal value  $p^*$  and the dual optimal value  $d^*$  are equal

$$p^* = d^*.$$

---

<sup>1</sup>We simply need  $\boldsymbol{\lambda} \geq \mathbf{0}$  for  $(\boldsymbol{\lambda}, \boldsymbol{\nu})$  to be dual feasible.

If  $\mathbf{x}^*$  is a primal optimal point and  $\boldsymbol{\lambda}^*, \boldsymbol{\nu}^*$  is a dual optimal point, then we must have

$$\begin{aligned} f_0(\mathbf{x}^*) &= g(\boldsymbol{\lambda}^*, \boldsymbol{\nu}^*) \\ &= \inf_{\mathbf{x} \in \mathbb{R}^N} \left( f_0(\mathbf{x}) + \sum_{m=1}^M \lambda_m^* f_m(\mathbf{x}) + \sum_{p=1}^P \nu_p^* h_p(\mathbf{x}) \right) \\ &\leq f_0(\mathbf{x}^*) + \sum_{m=1}^M \lambda_m^* f_m(\mathbf{x}^*) + \sum_{p=1}^P \nu_p^* h_p(\mathbf{x}^*) \\ &\leq f_0(\mathbf{x}^*). \end{aligned}$$

The last inequality follows from the fact that  $\lambda_m^* \geq 0$  (dual feasibility),  $f_m(\mathbf{x}^*) \leq 0$ , and  $h_p(\mathbf{x}^*) = 0$  (primal feasibility). Since we started out and ended up with the same thing, all of the things above must be equal, and so

$$\lambda_m^* f_m(\mathbf{x}^*) = 0, \quad m = 1, \dots, M.$$

Also, since we know  $\mathbf{x}^*$  is a minimizer of  $L(\mathbf{x}, \boldsymbol{\lambda}^*, \boldsymbol{\nu}^*)$  (second equality above), which is an unconstrained convex function (with  $\boldsymbol{\lambda}, \boldsymbol{\nu}$  fixed), the gradient with respect to  $\mathbf{x}$  must be zero:

$$\nabla_{\mathbf{x}} L(\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\nu}^*) = \nabla f_0(\mathbf{x}^*) + \sum_{m=1}^M \lambda_m^* \nabla f_m(\mathbf{x}^*) + \sum_{p=1}^P \nu_p^* \nabla h_p(\mathbf{x}^*) = \mathbf{0}.$$

Thus strong duality immediately leads to the KKT conditions holding at the solution.

Also, if you can find  $\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\nu}^*$  that obey the KKT conditions, not only do you know that you have a primal optimal point on your hands, but also we have strong duality (and  $\boldsymbol{\lambda}^*, \boldsymbol{\nu}^*$  are dual optimal). For if KKT holds,

$$\nabla_{\mathbf{x}} L(\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\nu}^*) = \mathbf{0},$$

meaning that  $\mathbf{x}^*$  is a minimizer of  $L(\mathbf{x}, \boldsymbol{\lambda}^*, \boldsymbol{\nu}^*)$ , i.e.

$$L(\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\nu}^*) \leq L(\mathbf{x}, \boldsymbol{\lambda}^*, \boldsymbol{\nu}^*),$$

thus

$$\begin{aligned} g(\boldsymbol{\lambda}^*, \boldsymbol{\nu}^*) &= L(\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\nu}^*) \\ &= f_0(\mathbf{x}^*) + \sum_{m=1}^M \lambda_m^* f_m(\mathbf{x}^*) + \sum_{p=1}^P \nu_p^* h_p(\mathbf{x}^*) \\ &= f_0(\mathbf{x}^*), \quad (\text{by KKT}), \end{aligned}$$

and we have strong duality.

The upshot of this is that the conditions for strong duality are essentially the same as those under which KKT is necessary.

The program (1) and its dual (2) have strong duality if the  $f_m$  are affine inequality constraints, or there is an  $\mathbf{x} \in \mathbb{R}^N$  such that for all the  $f_i$  which are not affine we have  $f_i(\mathbf{x}) < 0$ .

## Examples

1. **Inequality LP.** Calculate the dual of

$$\underset{\mathbf{x} \in \mathbb{R}^N}{\text{minimize}} \langle \mathbf{x}, \mathbf{c} \rangle \quad \text{subject to} \quad \mathbf{Ax} \leq \mathbf{b}.$$

**Answer:** The Lagrangian is

$$\begin{aligned} L(\mathbf{x}, \boldsymbol{\lambda}) &= \langle \mathbf{x}, \mathbf{c} \rangle + \sum_{m=1}^M \lambda_m (\langle \mathbf{x}, \mathbf{a}_m \rangle - b_m) \\ &= \mathbf{c}^T \mathbf{x} - \boldsymbol{\lambda}^T \mathbf{b} + \boldsymbol{\lambda}^T \mathbf{Ax}. \end{aligned}$$

This is a linear functional in  $\mathbf{x}$  — it is unbounded below unless

$$\mathbf{c} + \mathbf{A}^T \boldsymbol{\lambda} = \mathbf{0}.$$

Thus

$$g(\boldsymbol{\lambda}) = \inf_{\mathbf{x}} \left( \mathbf{c}^T \mathbf{x} - \boldsymbol{\lambda}^T \mathbf{b} + \boldsymbol{\lambda}^T \mathbf{Ax} \right) = \begin{cases} -\langle \boldsymbol{\lambda}, \mathbf{b} \rangle, & \mathbf{c} + \mathbf{A}^T \boldsymbol{\lambda} = \mathbf{0} \\ -\infty, & \text{otherwise.} \end{cases}$$

So the Lagrange dual program is

$$\begin{aligned} \underset{\boldsymbol{\lambda} \in \mathbb{R}^M}{\text{maximize}} \quad & -\langle \boldsymbol{\lambda}, \mathbf{b} \rangle \quad \text{subject to} \quad \mathbf{A}^T \boldsymbol{\lambda} = -\mathbf{c} \\ & \boldsymbol{\lambda} \geq \mathbf{0}. \end{aligned}$$

2. **Standard form LP.** Calculate the dual of

$$\begin{aligned} \underset{\mathbf{x} \in \mathbb{R}^N}{\text{minimize}} \langle \mathbf{x}, \mathbf{c} \rangle \quad \text{subject to} \quad & \mathbf{Ax} = \mathbf{b} \\ & \boldsymbol{\lambda} \geq \mathbf{0}. \end{aligned}$$

**Least-squares.** Calculate the dual of

$$\underset{\mathbf{x} \in \mathbb{R}^N}{\text{minimize}} \|\mathbf{x}\|_2^2 \quad \text{subject to} \quad \mathbf{A}\mathbf{x} = \mathbf{b}.$$

Check that the duality gap is zero.

**Answer:**

$$\underset{\boldsymbol{\nu} \in \mathbb{R}^M}{\text{maximize}} -\frac{1}{4} \boldsymbol{\nu}^T \mathbf{A}\mathbf{A}^T \boldsymbol{\nu} - \mathbf{b}^T \boldsymbol{\nu}$$

3. **Minimum norm.** Calculate the dual of

$$\underset{\mathbf{x} \in \mathbb{R}^N}{\text{minimize}} \|\mathbf{x}\| \quad \text{subject to} \quad \mathbf{A}\mathbf{x} = \mathbf{b},$$

where  $\|\cdot\|$  is a general valid norm.

**Answer:** Use  $f_0(\mathbf{x}) = \|\mathbf{x}\|$  to ease notation below. We start with the Lagrangian:

$$\begin{aligned} L(\mathbf{x}, \boldsymbol{\nu}) &= f_0(\mathbf{x}) + \sum_{p=1}^P \nu_p (\langle \mathbf{x}, \mathbf{a}_m \rangle - b_m) \\ &= f_0(\mathbf{x}) - \langle \boldsymbol{\nu}, \mathbf{b} \rangle + (\mathbf{A}^T \boldsymbol{\nu})^T \mathbf{x} \end{aligned}$$

and so

$$\begin{aligned} g(\boldsymbol{\nu}) &= -\langle \boldsymbol{\nu}, \mathbf{b} \rangle + \inf_{\mathbf{x}} \left( f_0(\mathbf{x}) + (\mathbf{A}^T \boldsymbol{\nu})^T \mathbf{x} \right) \\ &= -\langle \boldsymbol{\nu}, \mathbf{b} \rangle - \sup_{\mathbf{x}} \left( -f_0(\mathbf{x}) - (\mathbf{A}^T \boldsymbol{\nu})^T \mathbf{x} \right) \\ &= -\langle \boldsymbol{\nu}, \mathbf{b} \rangle - f_0^*(-\mathbf{A}^T \boldsymbol{\nu}), \end{aligned}$$

where  $f_0^*$  is the Fenchel dual of  $f_0$ :

$$f_0^*(\mathbf{y}) = \sup_{\mathbf{x}} (\langle \mathbf{x}, \mathbf{y} \rangle - f_0(\mathbf{x})).$$

With  $f_0 = \|\cdot\|$ , we know already that

$$f_0^*(\mathbf{y}) = \begin{cases} 0, & \|\mathbf{y}\|_* \leq 1, \\ \infty, & \text{otherwise} \end{cases},$$

so

$$g(\boldsymbol{\nu}) = \begin{cases} -\langle \boldsymbol{\nu}, \mathbf{b} \rangle, & \|\mathbf{A}^T \boldsymbol{\nu}\|_* \leq 1 \\ -\infty, & \text{otherwise} \end{cases}.$$

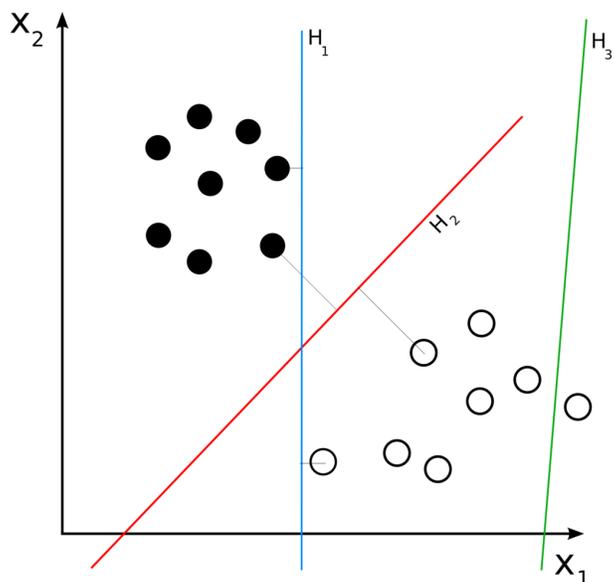
Thus the dual program is

$$\underset{\boldsymbol{\nu} \in \mathbb{R}^P}{\text{maximize}} \quad - \langle \boldsymbol{\nu}, \mathbf{b} \rangle \quad \text{subject to} \quad \|\mathbf{A}^T \boldsymbol{\nu}\|_* \leq 1,$$

where  $\|\cdot\|_*$  is the dual norm of  $\|\cdot\|$ .

# Support vector machines

Consider the following fundamental binary classification problem. We are given points  $\mathbf{x}_1, \dots, \mathbf{x}_M \in \mathbb{R}^N$  with labels  $y_1, \dots, y_M$ , where  $y_m \in \{-1, +1\}$ . We would like to find a hyperplane (i.e. affine functional) which *separates* the points<sup>1</sup>:



$\mathcal{H}_1$  and  $\mathcal{H}_2$  above separate the points in  $\mathbb{R}^2$ , but  $\mathcal{H}_3$  does not. To choose among the hyperplanes which separate the points, we will take the one with maximum margin (maximize the distance to the closest point in either class).

To restate this, we want to find a  $\mathbf{w} \in \mathbb{R}^N$  and  $b \in \mathbb{R}$  such that

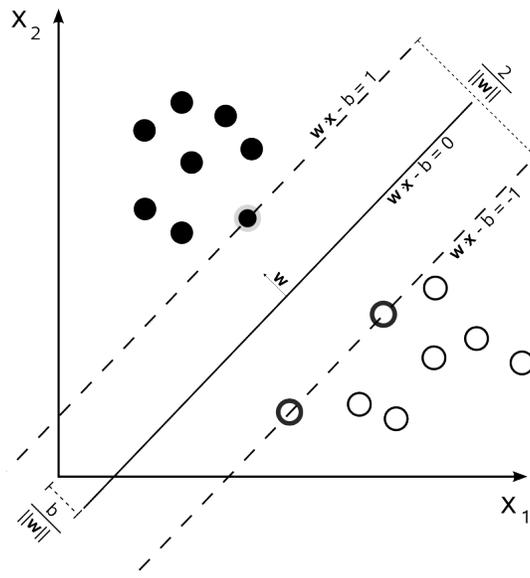
$$\begin{aligned} \langle \mathbf{x}_m, \mathbf{w} \rangle - b &\geq 1, & \text{when } y_m = 1, \\ \langle \mathbf{x}_m, \mathbf{w} \rangle - b &\leq -1, & \text{when } y_m = -1. \end{aligned}$$

---

<sup>1</sup>From Wikipedia: “Svm separating hyperplanes (SVG)” by User:ZackWeinberg, based on PNG version by User:Cyc.

Of course, it is possible that no separating hyperplane exists; in this case, there will be no feasible points in the program above. It is straightforward, though, to modify this discussion to allow “misclassified” points.

In the formulation above, the distance between the two (parallel) hyperplanes<sup>2</sup> is  $2/\|\mathbf{w}\|_2$ :



Thus maximizing this distance is the same as minimizing  $\|\mathbf{w}\|_2$ .

We have the program

$$\underset{\mathbf{w} \in \mathbb{R}^N, b \in \mathbb{R}}{\text{minimize}} \quad \frac{1}{2} \|\mathbf{w}\|_2^2 \quad \text{subject to} \quad y_m(b - \langle \mathbf{x}_m, \mathbf{w} \rangle) + 1 \leq 0, \quad m = 1, \dots, M.$$

This is a linearly constrained quadratic program, and is clearly con-

---

<sup>2</sup>From Wikipedia: “Svm max sep hyperplane with margin” by Cyc - Own work. Licensed under Public Domain via Wikimedia Commons.

vex. The Lagrangian is

$$\begin{aligned} L(\mathbf{w}, b, \boldsymbol{\lambda}) &= \frac{1}{2} \|\mathbf{w}\|_2^2 + \sum_{m=1}^M \lambda_m [y_m(b - \langle \mathbf{x}_m, \mathbf{w} \rangle) + 1] \\ &= \frac{1}{2} \|\mathbf{w}\|_2^2 + b \boldsymbol{\lambda}^T \mathbf{y} - \boldsymbol{\lambda}^T \mathbf{X}^T \mathbf{w} + \boldsymbol{\lambda}^T \mathbf{1}, \end{aligned}$$

where  $\mathbf{X}$  is the  $N \times M$  matrix

$$\mathbf{X} = \begin{bmatrix} y_1 \mathbf{x}_1 & y_2 \mathbf{x}_2 & \cdots & y_M \mathbf{x}_M \end{bmatrix}.$$

The dual function is

$$g(\boldsymbol{\lambda}) = \inf_{\mathbf{w}, b} \left( \frac{1}{2} \|\mathbf{w}\|_2^2 + b \boldsymbol{\lambda}^T \mathbf{y} - \boldsymbol{\lambda}^T \mathbf{X}^T \mathbf{w} + \boldsymbol{\lambda}^T \mathbf{1} \right).$$

Since  $b$  is unconstrained above, we see that the presence of  $b \boldsymbol{\lambda}^T \mathbf{y}$  means that the dual will be  $-\infty$  unless  $\langle \boldsymbol{\lambda}, \mathbf{y} \rangle = 0$ . Minimizing over  $\mathbf{w}$ , we need the gradient equal to zero,

$$\nabla_{\mathbf{w}} L(\mathbf{w}, b, \boldsymbol{\lambda}) = \mathbf{0}, \quad \Rightarrow \quad \mathbf{w} - \mathbf{X} \boldsymbol{\lambda} = \mathbf{0}.$$

This means that we must have  $\mathbf{w} = \mathbf{X} \boldsymbol{\lambda}$ , which itself is a very handy fact as it gives us a direct passage from the dual solution to the primal solution. With these substitutions, the dual function is

$$g(\boldsymbol{\lambda}) = \begin{cases} \frac{1}{2} \|\mathbf{X} \boldsymbol{\lambda}\|_2^2 - \boldsymbol{\lambda}^T \mathbf{X}^T \mathbf{X} \boldsymbol{\lambda} + \boldsymbol{\lambda}^T \mathbf{1}, & \langle \boldsymbol{\lambda}, \mathbf{y} \rangle = 0, \\ -\infty, & \text{otherwise.} \end{cases}$$

The dual SVM program is then

$$\begin{aligned} \underset{\boldsymbol{\lambda}}{\text{maximize}} \quad & -\frac{1}{2} \|\mathbf{X} \boldsymbol{\lambda}\|_2^2 + \sum_{m=1}^M \lambda_m \quad \text{subject to} \quad \langle \boldsymbol{\lambda}, \mathbf{y} \rangle = 0 \\ & \boldsymbol{\lambda} \geq \mathbf{0}. \end{aligned}$$

Given the solution  $\boldsymbol{\lambda}^*$  above, we can take  $\boldsymbol{w}^* = \mathbf{X}\boldsymbol{\lambda}^*$ , and the classifier is

$$\begin{aligned} f(\boldsymbol{x}) &= \langle \boldsymbol{x}, \boldsymbol{w}^* \rangle - b^* \\ &= \langle \boldsymbol{x}, \mathbf{X}\boldsymbol{\lambda}^* \rangle - b^* \\ &= \sum_{m=1}^M \lambda_m^* y_m \langle \boldsymbol{x}, \boldsymbol{x}_m \rangle - b^*. \end{aligned}$$

Notice that the data  $\boldsymbol{x}_m$  appear only as linear functionals (i.e. inner products with)  $\boldsymbol{x}$ .

The key realization is that for the dual program, the functional depends on the data  $\boldsymbol{x}_m$  only through inner products, as

$$\|\mathbf{X}\boldsymbol{\lambda}\|_2^2 = \sum_{\ell=1}^M \sum_{m=1}^M y_\ell y_m \langle \boldsymbol{x}_\ell, \boldsymbol{x}_m \rangle.$$

This means we can replace  $\langle \boldsymbol{x}_\ell, \boldsymbol{x}_m \rangle$  with any “positive kernel function”  $K(\boldsymbol{x}_\ell, \boldsymbol{x}_m) : \mathbb{R}^N \otimes \mathbb{R}^N \rightarrow \mathbb{R}$  — a positive kernel just means that the  $M \times M$  matrix  $K(\boldsymbol{x}_\ell, \boldsymbol{x}_m)$  is in  $S_+^M$  for all choices of  $\boldsymbol{x}_1, \dots, \boldsymbol{x}_M$ .

For example: you might take

$$K(\boldsymbol{x}_\ell, \boldsymbol{x}_m) = (1 + \langle \boldsymbol{x}_\ell, \boldsymbol{x}_m \rangle)^2 = 1 + 2\langle \boldsymbol{x}_\ell, \boldsymbol{x}_m \rangle + \langle \boldsymbol{x}_\ell, \boldsymbol{x}_m \rangle^2.$$

This means we have replaced the inner product of two vectors with the inner product between two vectors which have been mapped into

a higher dimensional space:

$$\begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_N \end{bmatrix} \rightarrow \begin{bmatrix} 1 \\ x_1 \\ \vdots \\ x_N \\ x_1^2 \\ x_2^2 \\ \vdots \\ x_N^2 \\ \sqrt{2}x_1x_2 \\ \vdots \\ \sqrt{2}x_{N-1}x_N \end{bmatrix}$$

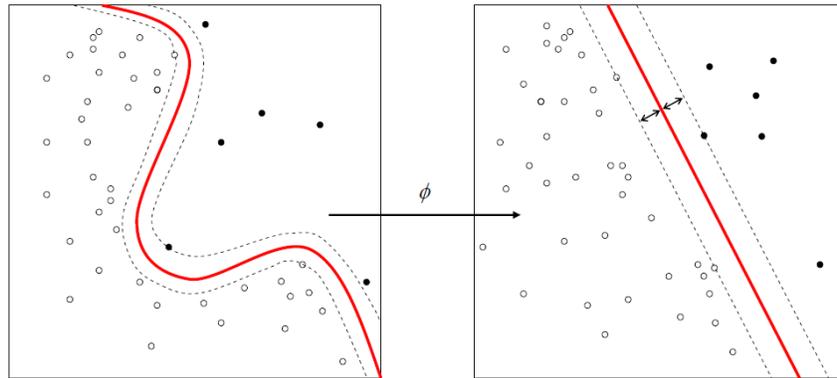
A set of linear constraints on the coordinates on the right, then, corresponds to a second order curve constraint (parabola, ellipse, hyperbola) on the coordinate on the left.

Many kernels are possible. The advantage is that to train and use the classifier, you never have to explicitly move to the higher dimensional space — you just need to be able to compute  $K(\mathbf{x}_\ell, \mathbf{x}_m)$  for any pair of inputs in  $\mathbb{R}^N$ . A popular choice of kernel is

$$K(\mathbf{x}_\ell, \mathbf{x}_m) = \exp(-\gamma \|\mathbf{x}_\ell - \mathbf{x}_m\|_2^2).$$

This is a perfectly valid positive kernel, and it is straightforward to compute it for any pair of inputs. But it corresponds to mapping the  $\mathbf{x}_m$  into an infinite dimensional space, then finding a hyperplane.

Here is an example from Wikipedia<sup>3</sup>:



---

<sup>3</sup>“Kernel Machine” by Alisneaky — Own work. Licensed under CC0 via Wikimedia Commons.