

Support Vector Machines for Regression

Provided with n training data points

$\{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_n, y_n)\} \subseteq \mathbb{R}^s \times \mathbb{R}$ we seek a function f for a fixed $\epsilon > 0$ such that:

$$|f(\mathbf{x}_i) - y_i| \leq \epsilon \quad (1)$$

Let us consider the space of linear functions for now and take $f(\mathbf{x}_i) = \langle \mathbf{w} \cdot \mathbf{x}_i \rangle + b$ with $\mathbf{w} \in \mathbb{R}^s$ and $b \in \mathbb{R}$. In order to avoid generalisation errors (over-fitting of the regression surface) we will require the function f to be as flat as possible so we want to minimise \mathbf{w} . We define a *convex quadratic* optimisation:

$$\begin{aligned} \min \quad & \frac{1}{2} \mathbf{w} \mathbf{w}^T \\ \text{subject to:} \quad & y_i - \langle \mathbf{w} \cdot \mathbf{x}_i \rangle - b \leq \epsilon \\ & \langle \mathbf{w} \cdot \mathbf{x}_i \rangle + b - y_i \leq \epsilon \end{aligned} \quad (2)$$

Feasibility

We assumed the existence of a function f satisfying the constraints. It is possible that for a given ϵ no function satisfying the constraint $|f(\mathbf{x}_i) - y_i| \leq \epsilon$ exists. So we define slack variables $\psi_i > 0$ and $\phi_i > 0$ and re-write the optimization as:

$$\text{minimise } \frac{1}{2} \mathbf{w} \mathbf{w}^T + \zeta \sum_i^n (\psi_i + \phi_i)$$

subject to:

$$\begin{aligned} y_i - \langle \mathbf{w} \cdot \mathbf{x}_i \rangle - b &\leq \epsilon + \psi_i \\ \langle \mathbf{w} \cdot \mathbf{x}_i \rangle + b - y_i &\leq \epsilon + \phi_i \\ \psi_i &\geq 0, \phi_i \geq 0, \forall n \end{aligned} \quad (3)$$

The constant ζ maintains the trade-off between how much deviation greater than ϵ is permitted versus the generalisation or in this case the flatness of the regression function.

We define the Lagrangian; the objective function plus a linear combination of the constraints:

$$\begin{aligned}
 L(\mathbf{w}, b, \phi, \psi | \alpha, \beta, \eta, \eta^*) &= \frac{1}{2} \|\mathbf{w}\|^2 + \zeta \sum_i^n (\psi_i + \phi_i) \\
 &- \sum_{i=1}^n \alpha_i (\epsilon + \psi_i - y_i + \langle \mathbf{w}, \mathbf{x}_i \rangle + b) \\
 &- \sum_{i=1}^n \beta_i (\epsilon + \phi_i + y_i - \langle \mathbf{w}, \mathbf{x}_i \rangle - b) \\
 &- \sum_{i=1}^n (\eta_i \psi_i + \eta_i^* \phi_i) \tag{4}
 \end{aligned}$$

where α_i , β_i , η_i and η_i^* are non-negative dual variables or lagrange multipliers.

Lagrangian Dual Problem

The dual of the lagrangian is defined as:

$$L_{dual}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\eta}, \boldsymbol{\eta}^*) = \inf_{\boldsymbol{w}, b, \boldsymbol{\phi}, \boldsymbol{\psi}} L(\boldsymbol{w}, b, \boldsymbol{\phi}, \boldsymbol{\psi} | \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\eta}, \boldsymbol{\eta}^*) \quad (5)$$

The Lagrangian Dual Problem is then:

$$\begin{aligned} & \max L_{dual}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\eta}, \boldsymbol{\eta}^*) \\ & \text{subject to: } \boldsymbol{\alpha} \geq 0 \text{ and } \boldsymbol{\beta} \geq 0 \text{ and } \boldsymbol{\eta} \geq 0 \text{ and } \boldsymbol{\eta}^* \geq 0 \end{aligned} \quad (6)$$

Lagrangian Weak Duality Theorem

The primal objective function f and its dual L_{dual} satisfy

$$f(\mathbf{w}, b, \psi, \phi) \geq L_{dual}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\eta}, \boldsymbol{\eta}^*) \quad (7)$$

Proof:

$$\begin{aligned} L_{dual}(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\eta}, \boldsymbol{\eta}^*) &= \inf_{\mathbf{w}, b, \phi, \psi} L(\mathbf{w}, b, \phi, \psi | \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\eta}, \boldsymbol{\eta}^*) \\ &\leq L(\mathbf{w}, b, \phi, \psi | \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\eta}, \boldsymbol{\eta}^*) \\ &\leq f(\mathbf{w}, b, \psi, \phi) \end{aligned} \quad (8)$$

Lagrangian Strong Duality Theorem

1. A real valued function is said to be convex if $\forall \mathbf{w}, \mathbf{u} \in \mathbb{R}^s$ and for any $\theta \in (0, 1)$ we have: $f(\theta \mathbf{w} + (1 - \theta) \mathbf{u}) \leq \theta f(\mathbf{w}) + (1 - \theta) f(\mathbf{u})$
2. Every strictly convex constrained optimisation problem has a unique solution.
3. An affine function can be expressed in the form $f(\mathbf{w}) = \mathbf{A}\mathbf{w} + \mathbf{b}$.
4. Affine functions are convex.
5. A convex optimisation problem has a convex objective function and affine constraints.

Strong Duality Theorem: For a convex optimisation problem *the duality gap is zero* at primal optimality.

Karush-Kuhn-Tucker Complementary Conditions

Let $(\alpha_o, \beta_o, \eta_o, \eta^*_o)$ and $(\mathbf{w}_o, b, \psi_o, \phi_o)$ be optimal solutions of the dual and primal respectively, then $L_{dual}(\alpha_o, \beta_o, \eta_o, \eta^*_o) = f(\mathbf{w}_o, b, \psi_o, \phi_o)$. From (4) and (8) we then derive the KKT conditions:

$$\begin{aligned}
 \alpha_i(\epsilon + \psi_i - y_i + \langle \mathbf{w}, \mathbf{x}_i \rangle + b) &= 0 \\
 \beta_i(\epsilon + \phi_i + y_i - \langle \mathbf{w}, \mathbf{x}_i \rangle - b) &= 0 \\
 \eta_i \psi_i &= 0 \\
 \eta_i^* \phi_i &= 0 \\
 i &= 1, \dots, n
 \end{aligned} \tag{9}$$

A constraint $c_t(\mathbf{w})$ is said to be tight or active if the solution \mathbf{w}_o satisfies $c_t(\mathbf{w}_o) = 0$ and is otherwise said to be inactive.

Lagrangian Saddlepoint Equivalence Theorem

When we have an optimal primal solution (so the duality gap is zero) it is a saddle point of the lagrangian function of the primal problem and so we have:

$$\frac{\partial L}{\partial b} = \sum_{i=1}^n (\beta_i - \alpha_i) = 0 \quad (10)$$

$$\frac{\partial L}{\partial \mathbf{w}} = \mathbf{w}^* - \sum_{i=1}^n (\alpha_i - \beta_i) \mathbf{x}_i = 0 \quad (11)$$

$$\frac{\partial L}{\partial \phi_i} = \zeta - \beta_i - \eta_i^* = 0 \quad (12)$$

$$\frac{\partial L}{\partial \psi_i} = \zeta - \alpha_i - \eta_i = 0 \quad (13)$$

(11) can be written as $\mathbf{w} = \sum_{i=1}^n (\alpha_i - \beta_i) \mathbf{x}_i$ and so
 $f(\mathbf{x}_a) = \langle \mathbf{w} \cdot \mathbf{x}_a \rangle + b = \sum_{i=1}^n (\alpha_i - \beta_i) \langle \mathbf{x}_i \cdot \mathbf{x}_a \rangle + b$

Dual Formulation

To remove the dependence on the primal variables we explicitly compute (5) by differentiating $L(\mathbf{w}, b, \phi, \psi | \alpha, \beta, \eta, \eta^*)$ with respect to the primal variables which leaves us with (10) to (13) which we substitute into (4):

$$\begin{aligned}
 &\text{maximise} && -\frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n (\alpha_i - \beta_i)(\alpha_j - \beta_j) \langle \mathbf{x}_i \cdot \mathbf{x}_j \rangle \\
 &&& -\epsilon \sum_{i=1}^n (\alpha_i + \beta_i) + \sum_{i=1}^n y_i (\alpha_i - \beta_i) \\
 &\text{subject to:} && \sum_{i=1}^n (\alpha_i - \beta_i) = 0 \\
 &&& \alpha_i, \beta_i \in [0, \zeta] \tag{14}
 \end{aligned}$$

So the dual of our quadratic program (3) is another quadratic program but with simpler constraints.

Support Vectors

The Karush-Kuhn-Tucker complementary conditions state that only the active constraints will have non-zero dual variables.

$$\epsilon + \phi_i - y_i + f(\mathbf{x}_i) > 0 \text{ implies } \alpha_i = 0$$

$$\epsilon + \phi_i - y_i + f(\mathbf{x}_i) = 0 \text{ implies } \alpha_i \neq 0$$

$$\epsilon + \psi_i + y_i - f(\mathbf{x}_i) > 0 \text{ implies } \beta_i = 0$$

$$\epsilon + \psi_i + y_i - f(\mathbf{x}_i) = 0 \text{ implies } \beta_i \neq 0$$

$$\forall i : 1 \leq i \leq n \quad (15)$$

The \mathbf{x}_i with non-zero α_i **or** β_i are the *support vectors*; if we were to train the SVM on only these \mathbf{x}_i , ignoring all the examples for which $\alpha_i = 0$ **and** $\beta_i = 0$, we would still induce the same regression surface.

Sparse Support Vector Expansion

We see that the vector \mathbf{w} can be written as a linear combination of the input training data points

$\{(\mathbf{x}_1, y_1), (\mathbf{x}_2, y_2), \dots, (\mathbf{x}_n, y_n)\} \subseteq \mathbb{R}^s \times \mathbb{R}$ and so:

$$f(\mathbf{x}_a) = \langle \mathbf{w} \cdot \mathbf{x}_a \rangle + b = \sum_{i=1}^n (\alpha_i - \beta_i) \langle \mathbf{x}_i \cdot \mathbf{x}_a \rangle + b \quad (16)$$

Let $\Lambda \subseteq \{1, 2, \dots, n\}$ such that $\forall i \in \Lambda$ we have *both* $\alpha_i \neq 0$ and $\beta_i \neq 0$. Then we can rewrite our regression function as:

$$f(\mathbf{x}_a) = \sum_{i \in \Lambda} (\alpha_i - \beta_i) \langle \mathbf{x}_i \cdot \mathbf{x}_a \rangle + b \quad (17)$$

We have a sparse expansion of $f(\mathbf{x}_a)$.

Non-linear SVM Regression - $\Phi : \mathbb{R}^s \rightarrow \mathcal{F}$

Our machine is linear. Our data might be non-linear.


We will apply a mapping function Φ to our input data, essentially projecting it into a higher dimensional feature space \mathcal{F} .

$$f(\mathbf{x}_i) = \langle \mathbf{w} \cdot \Phi(\mathbf{x}_i) \rangle + b \quad (18)$$

Then apply our linear machinery to find a linear regression in this new feature space. The corresponding regression surface in the input space will be non-linear, specifically $\Phi^{-1}(f(x))$.


Example: Quadratic Map

$$(u_1, u_2) \rightarrow \phi(u_1, u_2) = (u_1^2, u_2^2, \sqrt{2}u_1u_2) \quad (19)$$



Input Space: circular data in \mathbb{R}^2

Input data mapped into Feature Space



Data is linearly separable in feature space

Computational Feasibility

What computations are we performing in the feature space?

Since we are projecting only our input data \mathbf{x}_i - the dual form (14) has *slightly* changed; instead of computing the dot product between training examples in the input space $\langle \mathbf{x}_i \cdot \mathbf{x}_j \rangle$, we compute a dot product in the feature space \mathcal{F} :

$$\langle \Phi(\mathbf{x}_i) \cdot \Phi(\mathbf{x}_j) \rangle \quad (20)$$

If the dimension of \mathcal{F} is large then our problem might be computationally infeasible. We need to avoid computing in the feature space and hence projecting data points into the feature space.

But we still have to compute the dot product of the features efficiently.

Kernel Functions

What is a kernel: a function K such that for all \mathbf{x}, \mathbf{z} in some input space:

$$K(\mathbf{x}, \mathbf{z}) = \langle \phi(\mathbf{x}) \cdot \phi(\mathbf{z}) \rangle \quad (21)$$



where ϕ is a mapping from the input space to a feature space.

Mercer's Theorem: characterises what constitutes a valid kernel and how they can be built.

So the dimension of the feature space does *not* affect the computational complexity!

Inverse Multiquadric Kernel

$$k(\mathbf{x}, \mathbf{y}) = \frac{1}{\sqrt{\|\mathbf{x} - \mathbf{y}\|^2 + c^2}} \quad (22)$$



Quadratic Optimization using a Kernel Function

$$\text{maximise} \quad -\frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n (\alpha_i - \beta_i)(\alpha_j - \beta_j) \mathbf{K}(\mathbf{x}_i, \mathbf{x}_j)$$

$$-\epsilon \sum_{i=1}^n (\alpha_i + \beta_i) + \sum_{i=1}^n y_i (\alpha_i - \beta_i)$$

$$\text{subject to:} \quad \sum_{i=1}^n (\alpha_i - \beta_i) = 0$$

$$\alpha_i, \beta_i \in [0, \zeta] \quad (23)$$

So using the dual representation helps us in two ways:

1. we can operate in high dimensional spaces with the help of kernel functions which replace the dot product in the dual
2. lets us make use of optimisation algorithms designed for the dual form