Computational Statistics
Linear Classification

Luca Bortolussi

Department of Mathematics and Geosciences
University of Trieste

Office 238, third floor, H2bis
luca@dmi.units.it

Trieste, Winter Semester 2016/2017

Outline

THE BAYESIAN WAY

- To recast logistic regression in a Bayesian framework, we need to put a prior on $p(\mathbf{w})$ of the coefficients $\mathbf{w}$ of $\sigma(\mathbf{w}^T\phi(\mathbf{x}))$ and compute the posterior distribution on $\mathbf{w}$ by Bayes theorem. Then we can make predictions by integrating out the parameters.
- Assume a Gaussian prior $p(\mathbf{w}) = \mathcal{N}(\mathbf{w}|\mathbf{m_0}, \mathbf{S_0})$. The posterior is $p(\mathbf{w}|\mathbf{t}) \propto p(\mathbf{w})p(\mathbf{t}|\mathbf{w})$, and the log-posterior is

$$\log p(\mathbf{w}|\mathbf{t}) = -\frac{1}{2}(\mathbf{w}-\mathbf{m_0})^T S_0^{-1}(\mathbf{w}-\mathbf{m_0}) + \sum_{i=1}^{N}[t_i \log y_i + (1-t_i)\log(1-y_i)] + c$$

  where $y_i = \sigma(\mathbf{w}\phi(\mathbf{x_i}))$.
- Computing the marginal likelihood and the normalisation constant is analytically intractable, due to quadratic and logistic terms. Hence we do a Laplace approximation of the posterior.

LAPLACE APPROXIMATION OF THE POSTERIOR

- Given $\log p(\mathbf{w}|\mathbf{t})$, we first find the maximum a-posteriori $\mathbf{w_{MAP}}$, by running a numerical optimisation, and then obtain the Laplace approximation computing the Hessian matrix at $\mathbf{w_{MAP}}$ and inverting it, obtaining

$$\mathbf{S_N} = -\nabla\nabla \log p(\mathbf{w}|\mathbf{t}) = \mathbf{S_0}^{-1} + \sum_{n=1}^{N} y_n(1-y_n)\phi(\mathbf{x_n})\phi(\mathbf{x_n})^T$$

  evaluated at $\mathbf{w} = \mathbf{w_{MAP}}$.
- Hence, the Laplace approximation of the posterior is

$$q(\mathbf{w}) = \mathcal{N}(\mathbf{w}|\mathbf{w_{MAP}}, \mathbf{S_N})$$

PREDICTIVE DISTRIBUTION

- The predictive distribution for class C_1 is given by

$$p(C_1|\phi, \mathbf{t}) = \int p(C_1|\phi, \mathbf{w}, \mathbf{t}) q(\mathbf{w}) d\mathbf{w} = \int \sigma(\mathbf{w}^T \phi(\mathbf{x})) q(\mathbf{w}) d\mathbf{w}$$

- This multi-dimensional integral can be simplified by noting that it depends on $\mathbf{w}$ only on the 1-dim projection $a = \mathbf{w}^T \phi(\mathbf{x})$, and that q restricted to this direction is still a Gaussian distribution $q(a)$ with mean and variance

$$\mu_a = \mathbf{w_{MAP}}^T \phi(\mathbf{x}) \quad \sigma_a^2 = \phi(\mathbf{x})^T \mathbf{S_N} \phi(\mathbf{x})$$

- Hence we have

$$p(C_1|\phi, \mathbf{t}) = \int \sigma(a) q(a) da$$

PROBIT APPROXIMATION

- The integral $p(C_1|\phi, \mathbf{t}) = \int \sigma(a)q(a)da$ can be approximated by approximating the logistic function by the probit: $\sigma(a) = \Psi(\lambda a)$, where λ is obtained by matching derivatives at zero and is $\lambda^2 = \pi/8$.
- We then use

$$\int \Psi(a)\mathcal{N}(a|\mu, \sigma^2) = \Psi\left(\frac{\mu}{(\lambda^{-2} + \sigma^2)^{1/2}}\right)$$

and approximate back to the logistic to get

$$p(C_1|\phi, \mathbf{t}) \approx \sigma(\kappa(\sigma_a^2)\mu_a)$$

with $\kappa(\sigma_a^2) = (1 + \pi\sigma_a^2/8)^{-1/2}$

Outline

LAGRANGE MULTIPLIERS

- Suppose we want to maximise $f(\mathbf{x})$ subject to the constraint $g(\mathbf{x}) = 0$.
- $g(\mathbf{x}) = 0$ defines a surface and $\nabla g(\mathbf{x})$ is always orthogonal to it.
- In a point of this surface in which $f(\mathbf{x})$ is optimal, it must hold that $\nabla f(\mathbf{x}) = \lambda \nabla g(\mathbf{x})$, i.e. the projection of $\nabla f(\mathbf{x})$ on the tangent space of the surface is zero, otherwise we could increment the value of f by moving along the surface $g(\mathbf{x}) = 0$.

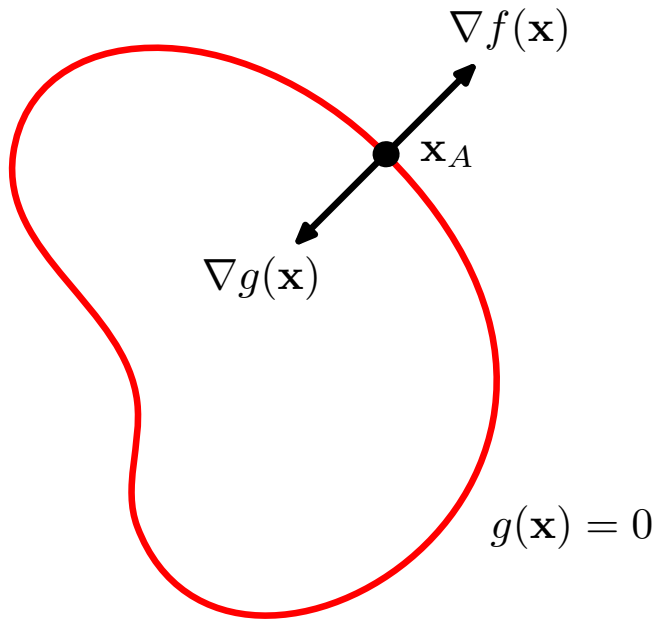

LAGRANGE MULTIPLIERS

- We can then do an unconstrained optimisation

$$\max_{\mathbf{x}} \inf_{\lambda} L(\mathbf{x}, \lambda)$$

of the Lagrangian function

$$L(\mathbf{x}, \lambda) = f(\mathbf{x}) + \lambda g(\mathbf{x})$$

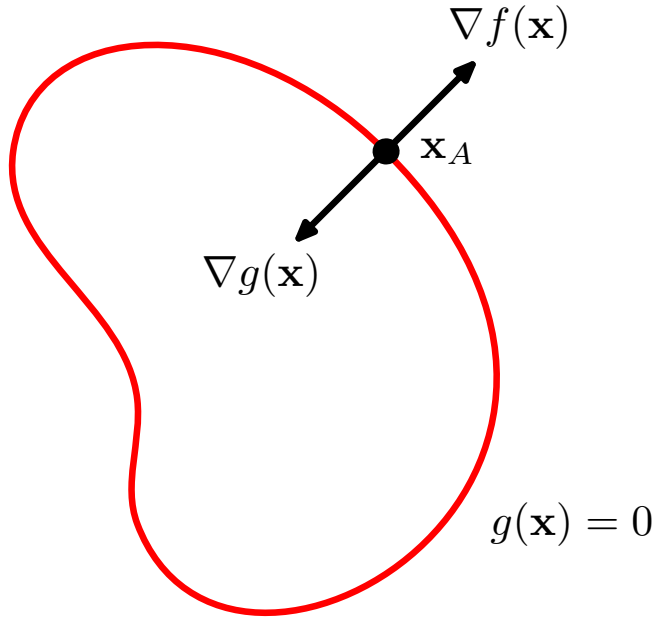

- In fact, if $g(\mathbf{x}) \neq 0$, then $\inf_\lambda L(\mathbf{x}, \lambda) = -\infty$, hence the Lagrangian optimization problem takes finite values only on $\{g(\mathbf{x}) = 0\}$.
- Deriving w.r.t $\mathbf{x}$ gives the condition on gradients, deriving w.r.t λ the constraint: setting the derivative to zero, we enforce the constraint and look for an optimal point.

KARUSH-KUHN-TUCKER CONDITIONS

- Suppose we want to optimise $f(\mathbf{x})$ subject to the constraint $g(\mathbf{x}) \geq 0$.
- If an optimum $\mathbf{x}$ satisfies $g(\mathbf{x}) > 0$ (inactive constraint), then $\nabla f(\mathbf{x}) = 0$ and $\lambda = 0$, if instead $g(\mathbf{x}) = 0$ (active constraint), then $\nabla f(\mathbf{x}) = -\lambda \nabla g(\mathbf{x})$, $\lambda > 0$ because an increase of f cannot bring inside the feasible region.

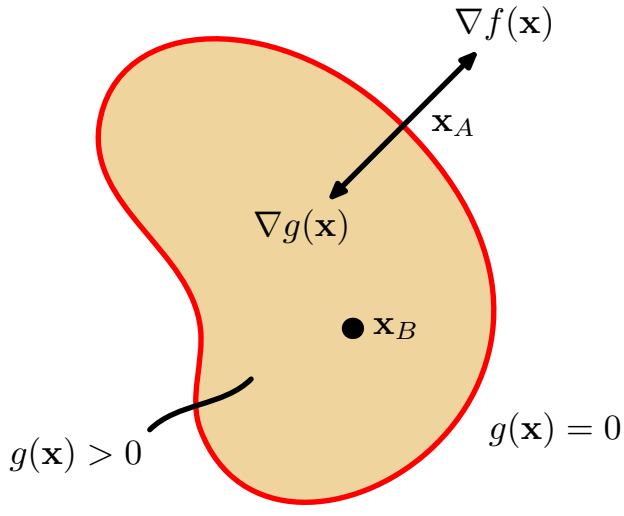

  - In any case $\lambda g(\mathbf{x}) = 0$ for an optimum point.
  - We can then optimise the Lagrangian function $L(\mathbf{x}, \lambda) = f(\mathbf{x}) + \lambda g(\mathbf{x})$ subject to $\lambda \geq 0$, $g(\mathbf{x}) \geq 0$, $\lambda g(\mathbf{x}) = 0$, known as the Karush-Kuhn-Tucker (KKT) conditions.

KARUSH-KUHN-TUCKER CONDITIONS

- Also in this case, we can then solve the unconstrained optimisation

$$\max_{\mathbf{x}} \inf_{\lambda \geq 0} L(\mathbf{x}, \lambda)$$

of the Lagrangian function

$$L(\mathbf{x}, \lambda) = f(\mathbf{x}) + \lambda g(\mathbf{x})$$

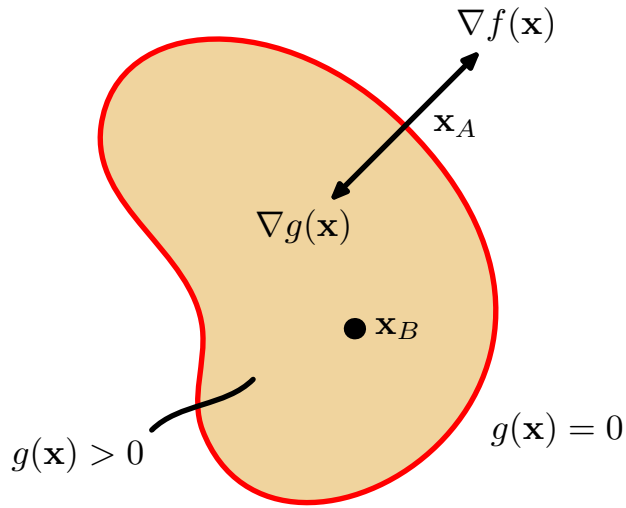

- In fact, if $g(\mathbf{x}) > 0$, then the inner optimization is solved by $\lambda = 0$, otherwise, if $g(\mathbf{x}) < 0$, it is solved by $\lambda = +\infty$ and the Lagrangian is $-\infty$. On the boundary $g(\mathbf{x}) = 0$, λ can take finite values.
- To minimise $f(\mathbf{x})$, we optimise $\min_{\mathbf{x}} \sup_{\lambda \geq 0} f(\mathbf{x}) - \lambda g(\mathbf{x})$
- Lagrange and KKT multipliers can be combined to solved constrained problems with both equalities and inequalities.

THE DUAL FORMULATION

- The dual formulation of the constrained minimisation problem with Lagrangian $L(\mathbf{x}, \lambda) = f(\mathbf{x}) - \sum_j \lambda_j g_j(\mathbf{x})$ is given by
$$\tilde{L}(\lambda) = \inf_{\mathbf{x} \in \mathcal{D}} L(\mathbf{x}, \lambda)$$
- $\tilde{L}(\lambda)$ is a lower bound on $f(\mathbf{x})$. The dual optimisation problem is to maximise $\tilde{L}(\lambda)$ subject to KKT conditions.
- If the original problem is convex (single global optimum), and under regularity conditions on the constraints (e.g. linear), then the solution of the dual gives exactly the minimum of the primal.
- For non-convex problems, there can be a duality gap.
- For quadratic objective functions and linear constraints, the dual objective can be computed easily, because $\partial L(\mathbf{x}, \lambda)/\partial \mathbf{x}$ gives a linear system that can be solved to express $\mathbf{x}$ as a function of λ's

Outline

KERNEL TRICK FOR CLASSIFICATION

- The trick works similarly as for regression. Consider class conditionals $p(C_1|\mathbf{x}) = \sigma(\mathbf{w}^T\phi(\mathbf{x}))$.
- We can make the assumption that $\mathbf{w} = \sum_{n=1}^{N} a_n\phi(\mathbf{x_n})$ (this is consistent, as the ML solution will belong to the space spanned by $\phi(\mathbf{x_n})$), thus getting

$$p(C_1|\mathbf{x}) = \sigma\left(\sum_{n=1}^{N} \alpha_n k(\mathbf{x}, \mathbf{x_n})\right)$$

  where we define the kernel function $k(\mathbf{x}, \mathbf{x}') = \phi(\mathbf{x})^T\phi(\mathbf{x}')$
- We can write also $p(C_1|\mathbf{x}) = \sigma(\mathbf{a}^T\mathbf{k}(\mathbf{x}))$. The maximum likelihood solution can be found using gradient based methods.

MAXIMUM MARGIN CLASSIFIERS

- We have 2-class data $\mathbf{x_n}, t_n$, with $t_n \in \{-1, 1\}$. We assume for the moment that the data is linearly separable in a feature space after applying the non-linear mapping $\phi(\mathbf{x})$.
- There may be many hyperplanes separating the data. An effective choice is to select the one maximising the margin, i.e. the smallest distance between the separating hyperplane and the data points.
- Only closest data points are needed to determine it.

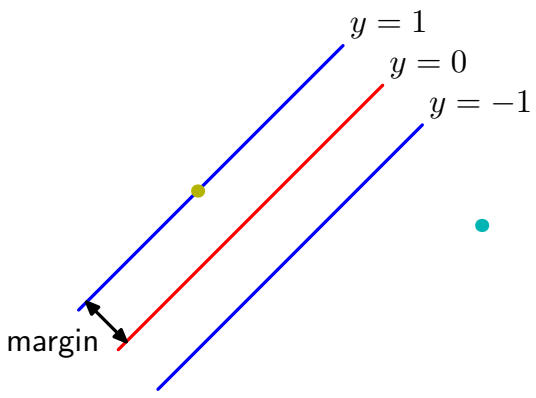

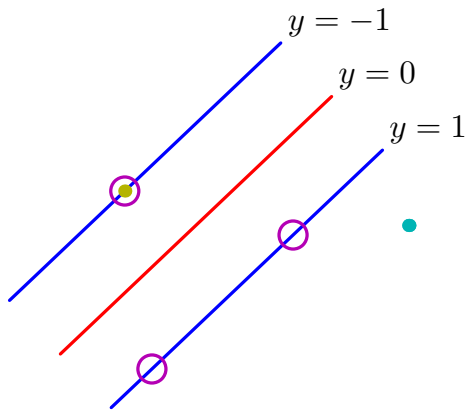

MAXIMUM MARGIN CLASSIFIERS

- Write $y(\mathbf{x}) = \mathbf{w}^T\phi(\mathbf{x}) + b$.
- The distance between a point and the separating hyperplane $\mathbf{w}^T\phi + b$ is $|y(\mathbf{x})|/\|\mathbf{w}\|$.
- As we want to classify correctly all points, it will hold that $t_n y(\mathbf{x_n}) \geq 0$, by the choice of t_n encoding.
- Hence, to find the maximum margin, we need to find $\mathbf{w}$ and b such that:

$$\max_{\mathbf{w},b}\left[\frac{1}{\|\mathbf{w}\|}\min_n\{t_n\mathbf{w}^T\phi(\mathbf{x_n}) + t_n b\}\right]$$

- The solution is defined up to an arbitrary rescaling of $\mathbf{w}$ and b, so we can set to 1 the margin, obtaining the constraint

$$t_n\mathbf{w}^T\phi(\mathbf{x_n}) + t_n b \geq 1, \quad n = 1, \ldots, N$$

MAXIMUM MARGIN CLASSIFIERS

- The constraints $t_n \mathbf{w}^T \phi(\mathbf{x_n}) + t_n b \geq 1$ known as the canonical representation. Points for which equality to 1 holds are called active, the others inactive.
- The maximisation above is equivalent to minimise $\|\mathbf{w}\|^2$:

$$\min_{\mathbf{w},b} \frac{1}{2}\|\mathbf{w}\|^2$$

subject to canonical constraints. b will be set via the constraints.

- To solve this quadratic program, we introduce a Langrange multiplier a_n for each constraint, resulting in the following Lagrangian

$$L(\mathbf{w}, b, \mathbf{a}) = \frac{1}{2}\|\mathbf{w}\|^2 - \sum_{n=1}^{N} a_n[t_n \mathbf{w}^T \phi(\mathbf{x_n}) + t_n b - 1]$$

which has to be minimised w.r.t $\mathbf{w}$ and b, and maximised w.r.t $\mathbf{a}$.

THE DUAL FORMULATION OF THE MAXIMUM MARGIN PROBLEM

- Starting from the Lagrangian $L(\mathbf{w}, b, \mathbf{a})$ we compute derivatives w.r.t. $\mathbf{w}$ and b and set them to zero, obtaining constraints

$$\mathbf{w} = \sum_n a_n t_n \phi(\mathbf{x_n}) \qquad 0 = \sum_n a_n t_n$$

- By substituting them in the Lagrangian, we obtain the dual representation

$$\tilde{L}(\mathbf{a}) = \sum_{n=1}^{N} a_n - \frac{1}{2} \sum_{n=1}^{N} \sum_{m=1}^{N} a_n a_m t_n t_m k(\mathbf{x_n}, \mathbf{x_m})$$

subject to the constraints

$$a_n \geq 0, \ n = 1, \ldots, N; \quad \sum_n a_n t_n = 0$$

- $k(\mathbf{x_n}, \mathbf{x_m}) = \phi(\mathbf{x_n})^T \phi(\mathbf{x_m})$ is the kernel function.

THE DUAL FORMULATION OF THE MAXIMUM MARGIN PROBLEM

- This optimisation problem can be solved in $O(N^3)$ time. Its main advantage is that it depends on the kernel, not on basis functions, hence it can be applied to more general kernels.
- The prediction for a new point $\mathbf{x}$ is obtained by using the dual formulation of $\mathbf{w}$, giving

$$y(\mathbf{x}) = \sum_n a_n t_n k(\mathbf{x}, \mathbf{x_n}) + b$$

SPARSITY OF THE SOLUTION

- The optimisation problem satisfies the KKT conditions:

$$a_n \geq 0; \quad t_n y(\mathbf{x_n}) - 1 \geq 0; \quad a_n[t_n y(\mathbf{x_n}) - 1] = 0$$

- This implies that either $t_n y(\mathbf{x_n}) = 1$ (the vector $\mathbf{x_n}$ is at minimum distance from the margin) or $a_n = 0$ (it does not contribute to the predictions).
- Let us indicate with $\mathcal{S}$ the set of support vectors.

DETERMINING b

- From any $\mathbf{x_n} \in \mathcal{S}$, by using $t_n y(\mathbf{x_n}) = 1$, we can determine b by solving

$$t_n \sum_{m \in \mathcal{S}} a_m t_m k(\mathbf{x_n}, \mathbf{x_m}) + t_n b = 1$$

- To have a more stable solution, one multiplies by t_n, uses $t_n^2 = 1$, and averages for the different support vectors:

$$b = \frac{1}{N_{\mathcal{S}}} \sum_{n \in \mathcal{S}} \left(t_n - \sum_{m \in \mathcal{S}} a_m t_m k(\mathbf{x_n}, \mathbf{x_m})) \right)$$

EXAMPLE OF SVM

- Example of data linearly separable in the space defined by the Gaussian kernel function.
- Sparsity: only support vectors define the maximum margin hyperplane: moving the other is irrelevant, as far as they remain on the same side.

Example of synthetic data from two classes in two dimensions showing contours of constant $y(\mathbf{x})$ obtained from a support vector machine having a Gaussian kernel function. Also shown are the decision boundary, the margin boundaries, and the support vectors.

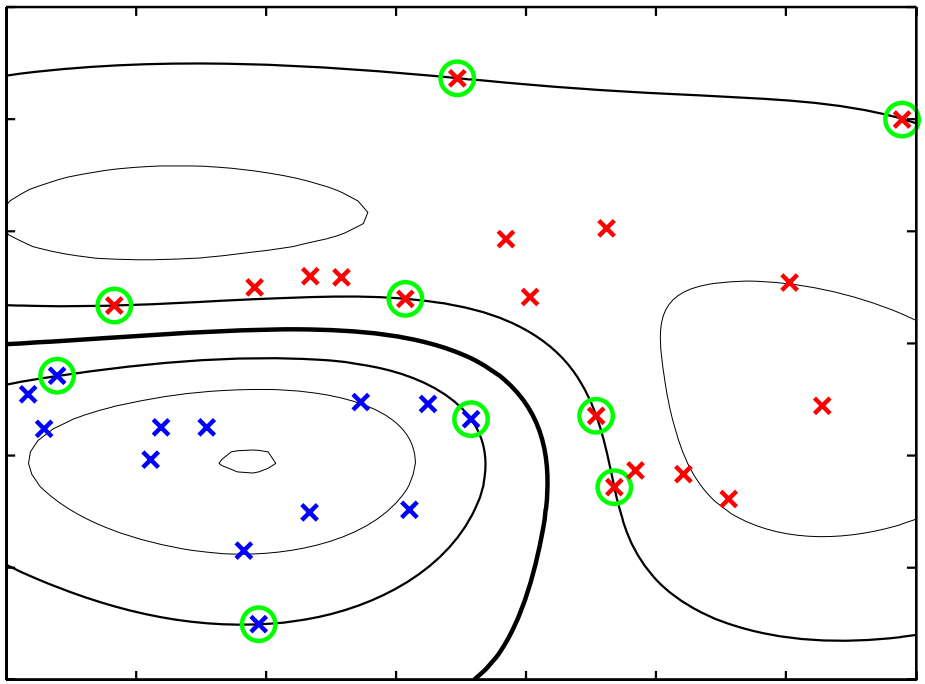

SOFT MARGIN SVM

- If class conditionals overlap, then an exact (non-linear) separation of training data may result in poor generalisation. It is better to allow some training points to be misclassified, by relaxing the constraint $t_n y(\mathbf{x_n}) \geq 1$
- We will do this by introducing N new slack variables $\xi_n \geq 0$, rewriting constraint as $t_n y(\mathbf{x_n}) \geq 1 - \xi_n$.

- For points correctly classified and inside the margin, we have $\xi_n = 0$, while for other points we have $\xi_n = |t_n - y(\mathbf{x_n})|$. It follows that misclassified points will have $\xi_n > 1$, while $\xi_n = 1$ only if a point lies in the separating hyperplane.
- $\sum_n \xi_n$ is an upper bound on misclassified training points.

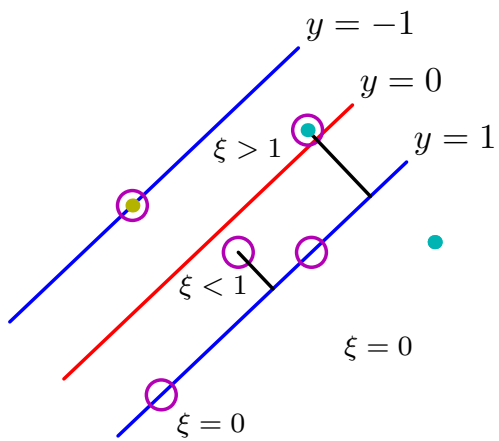

SOFT MARGIN SVM

- The primal objective function is modified to penalise the number of misclassified points:

$$C \sum_{n=1}^{N} \xi_n + \frac{1}{2}\|\mathbf{w}\|^2$$

- C is a regularisation term: it controls the trade-off between correct classification of training points and model complexity. For $C \to \infty$, we recover the previous SVM.
- The Lagrangian $L(\mathbf{w}, b, \mathbf{a}, \mu)$ is now given by

$$C \sum_{n-1}^{N} \xi_n + \frac{1}{2}\|\mathbf{w}\|^2 - \sum_{n=1}^{N} a_n[t_n \mathbf{w}^T \phi(\mathbf{x_n}) + t_n b - 1 + \xi_n] - \sum_{n=1}^{N} \mu_n \xi_n$$

with a_n, μ_n Lagrange multipliers. We omit the KKT conditions.

SOFT MARGIN SVM: DUAL FORMULATION

- By taking partial derivatives w.r.t $\mathbf{w}$, b, and ξ_n, we obtain the dual formulation:

$$\tilde{L}(\mathbf{a}) = \sum_{n=1}^{N} a_n - \frac{1}{2} \sum_{n=1}^{N} \sum_{m=1}^{N} a_n a_m t_n t_m k(\mathbf{x_n}, \mathbf{x_m})$$

  which has to satisfy the following box constraints

$$0 \leq a_n \leq C, \; n = 1, \ldots, N; \quad \sum_n a_n t_n = 0$$

- In the solution, we can have $a_n = 0$ (points inside the margin , for which $\xi_n = 0$), $0 < a_n < C$ (points on the margin, for which $\xi_n = 0$), or $a_n = C$ (points on the wrong side of the margin, $\xi_n > 0$).
- b can be determined as for the hard margin case, by restricting to support vectors on the margin.

SVM: COMMENTS

- The quadratic problem is convex, hence has a unique minimum, but a classic optimisation can be challenging for large problems (N large). Specialised methods have been developed, that try to decompose the problem into simpler pieces. E.g. Sequential minimal optimisation works by optimising two a_n's at time.
- SVM are hard to generalise to multi-class problems (one-versus-the-rest approach being the typical approach)
- SVM do not have a probabilistic interpretation, and some ad-hoc processing is required.
- SVM can be quite sensitive to outliers (misclassified points deeply inside the other's class region).