

# Fast exact leave-one-out cross-validation of sparse least-squares support vector machines

Gavin C. Cawley\*, Nicola L.C. Talbot

*School of Computing Sciences, University of East Anglia, Norwich NR4 7TJ, UK*

Received 25 March 2002; accepted 29 July 2004

## Abstract

Leave-one-out cross-validation has been shown to give an almost unbiased estimator of the generalisation properties of statistical models, and therefore provides a sensible criterion for model selection and comparison. In this paper we show that exact leave-one-out cross-validation of sparse Least-Squares Support Vector Machines (LS-SVMs) can be implemented with a computational complexity of only  $\mathcal{O}(\ell n^2)$  floating point operations, rather than the  $\mathcal{O}(\ell^2 n^2)$  operations of a naïve implementation, where  $\ell$  is the number of training patterns and  $n$  is the number of basis vectors. As a result, leave-one-out cross-validation becomes a practical proposition for model selection in large scale applications. For clarity the exposition concentrates on sparse least-squares support vector machines in the context of non-linear regression, but is equally applicable in a pattern recognition setting.

© 2004 Elsevier Ltd. All rights reserved.

*Keywords:* Model selection; Cross-validation; Least-squares support vector machine

## 1. Introduction

The generalisation properties of statistical models are typically governed by a small set of parameters, for instance regularisation parameters controlling the bias-variance trade-off; finding the optimal values for these parameters is an activity known as ‘model selection’. Model selection frequently seeks to optimise a cross-validation estimate of an appropriate statistic measuring generalisation ability on unseen data, for instance misclassification rate in the case of pattern recognition or the sum-of-squares error in a regression setting. A  $k$ -fold cross-validation procedure partitions the available data into  $k$  disjoint subsets.  $k$  models are then trained, each model being trained on a different combination of  $k-1$  of the  $k$  subsets and the test statistic evaluated over the remaining partition. The mean of the test statistic for each of the  $k$  models is known as the cross-validation estimate of the test statistic. The most extreme form of cross-validation, such that  $k$  is given by the number of training

patterns,  $\ell$ , is known as leave-one-out cross-validation. Leave-one-out cross-validation provides a sensible model selection criterion as it has been shown to provide an almost unbiased estimate of the true generalisation ability of the model (Lemma 1). Empirical studies have shown that in some cases model selection based on  $k$ -fold cross-validation outperforms selection procedures based on the leave-one-out estimator as the latter is known to exhibit a comparatively high variance (e.g. Kohavi, 1995). However, bounds on, or approximations to the leave-one-out estimator have proved effective criteria for model selection for conventional support vector machines, (e.g. Chapelle, Vapnik, Bousquet, & Mukherjee, 2002; Keerthi, 2002; Vapnik & Chapelle, 2000).

**Lemma 1 (Luntz & Brailovsky, 1969).** *Given a sequence of observations  $\mathcal{Z} = \{\mathbf{z}_i\}_{i=1}^{\ell}$ , forming an independent and identically distributed (i.i.d.) sample from an underlying distribution  $P(\mathbf{z})$ , and a hypothesis class  $\mathcal{H}$ , the empirical risk can be written as*

$$R_{\text{emp}}(\mathcal{H}, \mathcal{Z}) = \frac{1}{\ell} \sum_{i=1}^{\ell} Q(\mathbf{z}_i, h | \mathcal{Z}), \quad h \in \mathcal{H}, \quad \mathbf{z}_i \in \mathcal{Z},$$

\* Corresponding author. Tel.: +44 1603 593258; fax: +44 1603 592245.  
E-mail address: [gcc@cmp.uea.ac.uk](mailto:gcc@cmp.uea.ac.uk) (G.C. Cawley).

where  $Q(\mathbf{z}, h|\mathcal{Z})$  measures the loss for pattern  $z$  of the optimal hypothesis  $h \in \mathcal{H}$  formed on the basis on the set of observations  $\mathcal{Z}$ , denoted by  $h|\mathcal{Z}$ . Let  $\mathcal{L}(\mathcal{H}, \mathcal{Z})$  be the corresponding leave-one-out estimator

$$\mathcal{L}(\mathcal{H}, \mathcal{Z}) = \frac{1}{\ell} \sum_{i=1}^{\ell} Q(\mathbf{z}_i, h_i|\mathcal{Z} \setminus \{\mathbf{z}_i\}), \quad h_i \in \mathcal{H}, \quad \mathbf{z}_i \in \mathcal{Z},$$

then the leave-one-out cross-validation estimate of the risk of the hypothesis class  $\mathcal{H}$  on data  $\mathcal{Z}$ , with respect to  $P(\mathbf{z})$

$$R(\mathcal{H}, \mathcal{Z}) = \int Q(\mathbf{z}, h|\mathcal{Z}) dP(\mathbf{z}), \quad h \in \mathcal{H},$$

is almost unbiased in the sense that

$$E\{\mathcal{L}(\mathcal{H}, \mathcal{Z})\} = E\{R(\mathcal{H}, \mathcal{Z}')\},$$

where the expectations,  $E\{\cdot\}$ , are taken over i.i.d. samples  $\mathcal{Z}, \mathcal{Z}' \sim P(\mathbf{z})$  of size  $\ell$  on the left hand side and  $\ell - 1$  on the right (Luntz & Brailovsky, 1969; Vapnik, 1998).

Regularisation Machines (Poggio & Girosi, 1990), Kernel Ridge-Regression (Saunders, Gammernan, & Vovk, 1998) and the Least-Squares Support Vector Machine (LS-SVM) (Suykens et al., 2000; Suykens, De Brabanter, Lukas, & Vandewalle, 2002) comprise a family of closely related techniques constructing a linear model in a kernel-induced feature space minimising a regularised sum-of-squares loss functional. These methods provide an attractive approach for non-linear pattern recognition and regression problems as the optimal model parameters are given by the solution of a simple system of linear equations. The use of formal regularisation (Tikhonov & Arsenin, 1977) also provides convenient control of the bias-variance trade-off (Geman, Bienenstock, & Doursat, 1992) necessary for adequate generalisation in non-trivial applications. In this paper, we demonstrate that the computational complexity of leave-one-out cross-validation of a sparse formulation of the least squares support vector machine is of the same order as that of the basic training algorithm, i.e.  $\mathcal{O}(\ell^3)$ . Exact leave-one-out cross-validation therefore provides a feasible criterion for model selection in large scale applications.

The remainder of this paper is structured as follows. Section 2 introduces the sparse least-squares support vector machine and describes a computationally efficient method for performing leave-one-out cross-validation for this class of kernel machines. Section 3 presents results obtained using the proposed method on two real-world non-linear regression tasks, based on well known, public domain benchmark datasets. Section 4 discusses the origins of the proposed procedure and suggests some possible extensions. Finally, the work is summarised in Section 5.

## 2. Method

In this section, we first describe the least-squares support vector machine as a kernel form of ridge regression

(Hoerl & Kennard, 1970), before introducing a sparse formulation with an efficient training algorithm for large scale applications. A computationally efficient algorithm for leave-one-out cross-validation is then proposed, based on a corollary of the familiar Sherman–Woodbury–Morrison formula.

### 2.1. Ridge regression

Ridge regression (Hoerl & Kennard, 1970) is a method from classical statistics that implements a regularised form of least-squares regression. Given training data

$$\mathcal{D} = \{\mathbf{x}_i, t_i\}_{i=1}^{\ell}, \quad \mathbf{x} \in \chi \subset \mathbb{R}^d, \quad t_i \in \mathcal{T} \subset \mathbb{R},$$

ridge regression determines the weight vector,  $\mathbf{w} \in \mathbb{R}^d$ , and bias,  $b$ , of a linear model,  $f(\mathbf{x}) = \mathbf{w} \cdot \mathbf{x} + b$ , via minimisation of the following objective function:

$$L(\mathbf{w}, b) = \frac{1}{2} \|\mathbf{w}\|^2 + \gamma \sum_{i=1}^{\ell} (t_i - \mathbf{w} \cdot \mathbf{x}_i - b)^2. \quad (1)$$

The objective function used in ridge regression (1) implements a form of Tikhonov regularisation (Tikhonov & Arsenin, 1977) of a sum-of-squares error metric, where  $\gamma$  is a regularisation parameter controlling the bias-variance trade-off (Geman et al., 1992; Hoerl & Kennard, 1970). This corresponds to penalised maximum likelihood estimation of  $\mathbf{w}$ , assuming the targets have been corrupted by an independent and identically distributed (i.i.d.) sample drawn from a Gaussian noise process, with zero mean and fixed variance a  $\sigma^2$

$$t_i = \mathbf{w} \cdot \mathbf{x}_i + b + \varepsilon_i, \quad \varepsilon \sim \mathcal{N}(0, \sigma^2).$$

### 2.2. The least-squares support vector machine

A non-linear form of ridge regression (Saunders et al., 1998), known as the Least-Squares Support Vector Machine (LS-SVM) (Suykens et al., 2000, 2002), can be obtained via the so-called ‘kernel trick’, whereby a linear ridge regression model is constructed in a feature space,  $\mathcal{F}(\phi : \chi \rightarrow \mathcal{F})$ , induced by a Mercer kernel defining the inner product  $\mathcal{K}(\mathbf{x}, \mathbf{x}') = \phi(\mathbf{x}) \cdot \phi(\mathbf{x}')$ . The kernel function,  $\mathcal{K} : \chi \times \chi \rightarrow \mathbb{R}$  may be any positive definite kernel. For an introduction to kernel methods in machine learning see Cristianini and Shawe-Taylor (2000). The objective function minimised in constructing a least-squares support vector machine is given by

$$L(\mathbf{w}, b) = \frac{1}{2} \|\mathbf{w}\|^2 + \gamma \sum_{i=1}^{\ell} (t_i - \mathbf{w} \cdot \phi(\mathbf{x}_i) - b)^2.$$

The representer theorem (Kimeldorf & Wahba, 1971) indicates that the solution of an optimisation problem of this nature can be written in the form of an expansion over training patterns, i.e.  $\mathbf{w} = \sum_{i=1}^{\ell} \alpha_i \phi(\mathbf{x}_i)$ . The output of the least-squares support vector machine is then given by

the kernel expansion

$$f(\mathbf{x}) = \sum_{i=1}^{\ell} \alpha_i \mathcal{K}(\mathbf{x}_i, \mathbf{x}) + b.$$

It can easily be shown (Saunders et al., 1998; Suykens, Lukas, & Vandewalle, 2000) that the optimal coefficients of this expansion are given by the solution of a set of linear equations

$$\begin{bmatrix} \mathbf{K} + \gamma^{-1} \mathbf{I} & \mathbf{1} \\ \mathbf{1}^T & 0 \end{bmatrix} \begin{bmatrix} \boldsymbol{\alpha} \\ b \end{bmatrix} = \begin{bmatrix} \mathbf{t} \\ 0 \end{bmatrix},$$

where  $\mathbf{K} = [k_{ij} = \mathcal{K}(\mathbf{x}_i, \mathbf{x}_j)]_{i,j=1}^{\ell}$ ,  $\mathbf{t} = (t_1, t_2, \dots, t_{\ell})^T$ ,  $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \dots, \alpha_{\ell})^T$  and  $\mathbf{1} = (1, 1, \dots, 1)^T$ .

### 2.3. The sparse least-squares support vector machine

Unfortunately, unlike the original support vector machine (Boser, Guyon, & Vapnik, 1992; Cortes & Vapnik, 1995), the kernel expansions defining least-squares support vector machines are typically fully dense, i.e.  $\alpha_i \neq 0$ ,  $\forall i \in \{1, 2, \dots, \ell\}$ . This severely limits the utility of the LS-SVM in large scale applications as the complexity of the basic training algorithm is of order  $\mathcal{O}(\ell^3)$  (although more efficient methods have been proposed, e.g. Suykens, Lukas, Van Dooren, De Moor, & Vandewalle 1999). Fortunately a range of methods for obtaining sparse least-squares support vector machines are available, see e.g. Cawley & Talbot, 2002a,b; Williams & Seeger, 2001. The sparsity of the original support vector machine is perhaps best viewed as a convenient consequence of the use of the hinge loss, instead of the least-squares loss, in defining the primal optimisation problem, rather than as a deliberate design aim of the algorithm. As a result, the kernel expansions for support vector machine in general are not maximally sparse either and so several algorithms have been proposed for simplifying existing support vector expansions (Burges, 1996; Burges & Schölkopf, 1997; Downs, Gates, & Masters, 2001), or for training support vector machines using a reduced set of candidate support vectors (Lee & Mangasarian, 2001). The latter is similar in spirit to the sparse least-squares support vector machine used here.

In order to construct a sparse least-squares support vector machine, it is assumed that the weight vector,  $\mathbf{w}$ , can be represented as a weighted sum of a limited number of basis vectors  $\{\phi(\mathbf{z}_i)\}_{i=1}^n$ ,  $\mathbf{z} \in \chi \subset \mathbb{R}^d$ ,  $n \ll \ell$ , such that

$$\mathbf{w} = \sum_{i=1}^n \beta_i \phi(\mathbf{z}_i).$$

The objective function minimised by the sparse least-squares support vector machine then becomes

$$L(\beta, b) = \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \beta_i \beta_j k_{ij} + \gamma \sum_{i=1}^{\ell} \left( t_i - \sum_{j=1}^n \beta_j \hat{k}_{ij} - b \right)^2, \quad (2)$$

where  $k_{ij} = \mathcal{K}(\mathbf{z}_i, \mathbf{z}_j)$  and  $\hat{k}_{ij} = \mathcal{K}(\mathbf{x}_i, \mathbf{z}_j)$ . It is straightforward to show that this optimisation problem can again be expressed in the form of a system of linear equations (see Appendix A.3)

$$(\mathbf{R} + \mathbf{Z}^T \mathbf{Z}) \mathbf{p} = \mathbf{Z}^T \mathbf{t} \quad (3)$$

where  $\mathbf{p} = (\beta, b)^T$  is the vector of model parameters,  $\mathbf{Z} = [\hat{\mathbf{K}}, \mathbf{1}]$ ,  $\hat{\mathbf{K}} = [\hat{k}_{ij}]_{i=1, j=1}^{\ell, n}$ ,  $\hat{\mathbf{K}} = [k_{ij}]_{i,j=1}^{\ell, n}$  and  $\mathbf{1} = [1]_{i=1}^{\ell}$   $\mathbf{R}$  corresponds to the regularisation *Fast Leave-One-Out Cross-Validation of Sparse LS-SVMs* term in (2)

$$\mathbf{R} = \begin{bmatrix} \frac{1}{2} \gamma^{-1} \mathbf{K} & \mathbf{0} \\ \mathbf{0}^T & 0 \end{bmatrix}.$$

Note that the computational complexity of the resulting training algorithm is of order  $\mathcal{O}(\ell n^2)$ , being dominated by construction of the  $(n+1) \times (n+1)$  matrix  $(\mathbf{R} + \mathbf{Z}^T \mathbf{Z})$ , and therefore is far better suited to large scale application.

### 2.4. Selection of reference vectors

There are a number of viable approaches to the problem of selecting a set of reference vectors: Baudat and Anouar (2001) propose selection of a minimal set of training vectors forming a basis for the subspace of  $\mathcal{F}$  populated by the data (for alternative approaches, see Cawley & Talbot, 2002a, 2002b; Smola & Schölkopf, 2000; Williams & Seeger, 2001). Suykens et al. (2000, 2002) iteratively prune patterns having a small error from the kernel expansion (note however that this involves starting with a fully dense least-squares support vector machine, which is computationally demanding for large  $\ell$ ). Alternatively, one could simply select a random subsample of the training data, or even the entire training set. For this study, we adopt the method of Baudat and Anouar, as summarised in Section 2.4.1.

#### 2.4.1. Feature vector selection

The normalised Euclidean distance between the image of a datum,  $\mathbf{x}$ , in feature space,  $\phi(\mathbf{x})$ , and  $\hat{\phi}_S(\mathbf{x})$ , its optimal reconstruction using a sub-set of the training vectors  $\{\phi(\mathbf{x}_i)\}_{i \in S}$ , is given by

$$\delta_i(S) = \frac{\|\phi(\mathbf{x}_i) - \hat{\phi}_S(\mathbf{x}_i)\|^2}{\|\phi(\mathbf{x}_i)\|^2}. \quad (4)$$

This distance can be expressed in terms of inner products, and so via the ‘kernel trick’

$$\delta_i(S) = 1 - \frac{\mathbf{K}_{Si}^T \mathbf{K}_{SS}^{-1} \mathbf{K}_{Si}}{k_{ii}}.$$

where  $\mathbf{K}_{SS}$  is a square sub-matrix of the kernel matrix  $\mathbf{K} = [k_{ij} = \mathcal{K}(\mathbf{x}_i, \mathbf{x}_j)]_{i,j=1}^{\ell}$ , such that  $\mathbf{K}_{SS} = \{k_{ij}\}_{i,j \in S}$  and  $\mathbf{K}_{Si} = (k_{ji})_{j \in S}^T$  is a column vector of inner products in  $\mathcal{F}$ . To form a basis for the subspace in  $\mathcal{F}$  populated by the data, it is sufficient to minimise the mean reconstruction error  $\delta_i$  over

all patterns in the training set (Baudat & Anouar, 2001), i.e. maximize

$$J(\mathcal{S}) = \frac{1}{\ell} \sum_{i=1}^{\ell} \frac{\mathbf{K}_{\mathcal{S}i}^T \mathbf{K}_{\mathcal{S}\mathcal{S}}^{-1} \mathbf{K}_{\mathcal{S}i}}{k_{ii}}. \quad (5)$$

Starting with  $\mathcal{S} = \emptyset$ , a basis is constructed in a greedy manner, at each step adding the training vector maximising  $J(\mathcal{S})$ . Note that provided the feature vectors currently selected are linearly independent  $\mathbf{K}_{\mathcal{S}\mathcal{S}}$  will be invertible. This will always be the case; clearly one would not select a linearly dependent feature vector as this would not result in a reduction in the mean reconstruction error. The algorithm therefore terminates when  $\mathbf{K}_{\mathcal{S}\mathcal{S}}$  is no longer invertible, indicating that a complete basis has been identified, or when an upper bound on the number of basis vectors is reached, or when the mean relative reconstruction error falls below a given threshold. An efficient implementation of the feature vector selection algorithm is given in Cawley and Talbot (2002). The computational complexity of the standard feature vector selection is  $\mathcal{O}(\ell^2 |\mathcal{S}|^2)$ , where  $|\mathcal{S}|$  is the number of feature vectors extracted. The computational complexity of the improved algorithm is only  $\mathcal{O}(|\mathcal{S}|^3)$ , however this is achieved at the expense of introducing a stochastic element, such that there is no guarantee that the same set of feature vectors will be extracted each time the algorithm is executed. However, provided the feature vectors extracted (approximately) span the subspace in  $\mathcal{F}$  occupied by the training data, the mapping implemented by the sparse LS-SVM will be (approximately) the same as that of the full LS-SVM. The storage complexity of both algorithms is  $\mathcal{O}(|\mathcal{S}|^2)$ , being dominated by storage of  $\mathbf{K}_{\mathcal{S}\mathcal{S}}$ .

### 2.5. Fast exact leave-one-out cross-validation

At each step of the leave-one-out cross-validation procedure, a sparse least-squares support vector machine is constructed excluding a single training pattern from the data. The vector of model parameters,  $\mathbf{p}_{(i)}$  at the  $i$ th iteration is then given by the solution of the following system of linear equations

$$\mathbf{p}_{(i)} = \begin{bmatrix} \beta_{(i)} \\ b_{(i)} \end{bmatrix} = [\mathbf{R} + \mathbf{Z}_{(i)}^T \mathbf{Z}_{(i)}]^{-1} \mathbf{Z}_{(i)}^T \mathbf{t}$$

where  $\mathbf{Z}_{(i)}$  is the sub-matrix formed by omitting the  $i$ th row of  $\mathbf{Z}$ . Normally the most computationally expensive step is the inversion of the matrix  $\mathbf{C}_{(i)} = [\mathbf{R} + \mathbf{Z}_{(i)}^T \mathbf{Z}_{(i)}]$ , with a complexity of  $\mathcal{O}(n^3)$  operations. Fortunately  $\mathbf{C}_{(i)}$  can be written as a rank one modification of a matrix  $\mathbf{C}$

$$\mathbf{C}_{(i)} = [\mathbf{R} + \mathbf{Z}^T \mathbf{Z} - \mathbf{z}_i \mathbf{z}_i^T] = [\mathbf{C} - \mathbf{z}_i \mathbf{z}_i^T],$$

where  $\mathbf{z}_i$  is the  $i$ th row of  $\mathbf{Z}$  and  $\mathbf{C} = \mathbf{R} + \mathbf{Z}^T \mathbf{Z}$ . The following matrix inversion lemma, a corollary of the familiar Sherman–Woodbury–Morrison formula, can then be used

to find  $\mathbf{C}_{(i)}^{-1}$  in only  $\mathcal{O}(n^2)$  operations, given that  $\mathbf{C}^{-1}$  is already known.

**Lemma 2 (Bartlett, 1951).** *Given an invertible matrix  $\mathbf{A}$  and column vector  $\mathbf{v}$ , then assuming  $1 + \mathbf{v}^T \mathbf{A}^{-1} \mathbf{v} \neq 0$*

$$(\mathbf{A} + \mathbf{v} \mathbf{v}^T)^{-1} = \mathbf{A}^{-1} - \frac{\mathbf{A}^{-1} \mathbf{v} \mathbf{v}^T \mathbf{A}^{-1}}{1 + \mathbf{v}^T \mathbf{A}^{-1} \mathbf{v}}.$$

*This can be obtained straight-forwardly as a special case of the familiar Sherman–Woodbury–Morrison formula (see Appendix A.2).*

The solution to the matrix inversion problem can then be written as

$$\mathbf{C}_{(i)}^{-1} = [\mathbf{C} - \mathbf{z}_i \mathbf{z}_i^T]^{-1} = \mathbf{C}^{-1} + \frac{\mathbf{C}^{-1} \mathbf{z}_i \mathbf{z}_i^T \mathbf{C}^{-1}}{1 - \mathbf{z}_i^T \mathbf{C}^{-1} \mathbf{z}_i}.$$

The computational complexity of the leave-one-out cross-validation process is thus reduced to only  $\mathcal{O}(\ell n^2)$  operations, which is the same as that of the basic training algorithm for the sparse least-squares support vector machine.

### 2.6. A further refinement

For the purposes of model selection based on leave-one-out cross-validation, we are not generally interested in the values of the model parameters themselves, but only in the predicted output of the model for the training pattern excluded during each trial. For instance we might adopt Allen's PRESS (predicted residual sum of squares) statistic (Allen, 1974)

$$\text{PRESS} = \sum_{i=1}^{\ell} \{\mathbf{e}_{(i)}\}_i^2,$$

where  $\{\mathbf{e}_{(i)}\}_i = t_i - \{\mathbf{y}_{(i)}\}_i$  is the residual for the  $i$ th training pattern during the  $i$ th iteration of the leave-one-out cross-validation procedure. Fortunately, it is possible to compute these residuals without explicitly evaluating the model parameters in each trial. It is relatively straight-forward to show that

$$\{\mathbf{e}_{(i)}\}_i = \frac{e_i}{1 - h_{ii}},$$

(see Appendix A.4) where  $e_i = t_i - y_i$  is the residual for the  $i$ th training pattern for a sparse least-squares support vector machine trained on the entire dataset,  $\mathbf{H} = \mathbf{Z} \mathbf{C}^{-1} \mathbf{Z}^T$  is the *hat*, or *smoother*, matrix (Hastie, Tibshirani, & Friedman, 2001; Weisberg, 1985) of which  $h_{ii}$  the  $i$ th element of the leading diagonal (Cook & Weisberg, 1982). Allen's PRESS statistic can therefore be evaluated in closed form without explicit inversion of  $\mathbf{C}_{(i)}$ , again with a computational complexity of only  $\mathcal{O}(\ell n^2)$ .



### 3. Results

This section presents results obtained using conventional and fast leave-one-out cross-validation procedures for two well-studied non-linear regression bench-mark tasks: the Motorcycle and Boston Housing data sets. The Motorcycle data set, being a univariate regression task is easily visualised, whereas the Boston Housing data set is used to illustrate the scaling properties of the proposed procedure. All experiments were performed using the MATLAB programming environment on a personal computer with dual 1.7 GHz Xeon processors and 4 Gb of RDRAM (400 MHz bus) running under the Linux operating system.

#### 3.1. The motorcycle dataset

The Motorcycle benchmark consists of a sequence of accelerometer readings through time following a simulated motor-cycle crash during an experiment to determine the efficacy of crash-helmets (Silverman, 1985). Fig. 1 shows conventional and sparse support vector machine models of the motorcycle dataset, using a Gaussian radial basis function (RBF) kernel

$$\mathcal{K}(\mathbf{x}, \mathbf{x}') = \exp\{-\sigma^{-2}|\mathbf{x} - \mathbf{x}'|^2\}.$$

The hyper-parameters were selected to minimise a 10-fold cross-validation estimate of the test sum-of-squares error using a straight-forward Nelder–Mead simplex algorithm (Nelder & Mead, 1965) ( $\gamma = 1.2055 \times 10^{-3}$  and  $\sigma^{-2} = 13.1$ ). Note that the sparse model is functionally identical to the full least-squares support vector machine model with only 15 basis vectors comprising the sparse kernel expansion. The Gram matrix for an RBF kernel is of full-rank (Micchelli, 1986), and so it is not generally possible to form a complete sparse basis for the subspace of  $\mathcal{F}$  inhabited by the training data  $\mathcal{D}$ . In this case however, only 15 basis vectors can be




Fig. 1. Dense and sparse least-squares support vector machine models of the motorcycle data set; note the standard and sparse models are essentially identical.




Fig. 2. Leave-one-out cross-validation behaviour of a sparse least-squares support vector model of the Motorcycle benchmark dataset, using conventional and fast leave-one-out cross-validation algorithms.

extracted before the condition of  $\mathbf{K}_{S,S}$  deteriorates to the point at which it is no longer invertible, indicating that these vectors form an extremely close approximation to a complete basis.

Fig. 2 illustrates the leave-one-out cross-validation behaviour of the sparse least-squares support vector machine, i.e. the predicted value for each datapoint comprising the training data using a model trained only on the remaining datapoints. Note that the behaviour computed using the proposed fast algorithm (depicted by the diamonds) is essentially identical to that obtained by explicit leave-one-out cross-validation (depicted by crosses). The fast leave-one-out cross-validation algorithm is however more than two orders of magnitude faster than the conventional approach for this dataset, as shown in Table 1. Let the relative sum-of-squares error be defined as

$$E_r = \frac{\|\mathbf{e} - \hat{\mathbf{e}}\|}{\|\mathbf{e}\|}, \quad (6)$$

where  $\mathbf{e}$  is the vector of predicted (leave-one-out) residuals computed explicitly and  $\hat{\mathbf{e}}$  is the vector of predicted residuals computed via the proposed efficient method. The relative sum-of-squares error for the Motorcycle data set is negligible ( $E_r = 7.6 \times 10^{-6}$ ).

#### 3.2. The Boston housing dataset

The Boston housing dataset describes the relationship between the median value of owner occupied homes in

Table 1  
Mean run-time of standard and fast leave-one-out cross-validation algorithms for the Motorcycle benchmark dataset, over 100 trials

Algorithm	Mean run time (s)	Standard error
Explicit	0.396950	$1.212448 \times 10^{-3}$
Fast	0.003072	$1.479775 \times 10^{-5}$

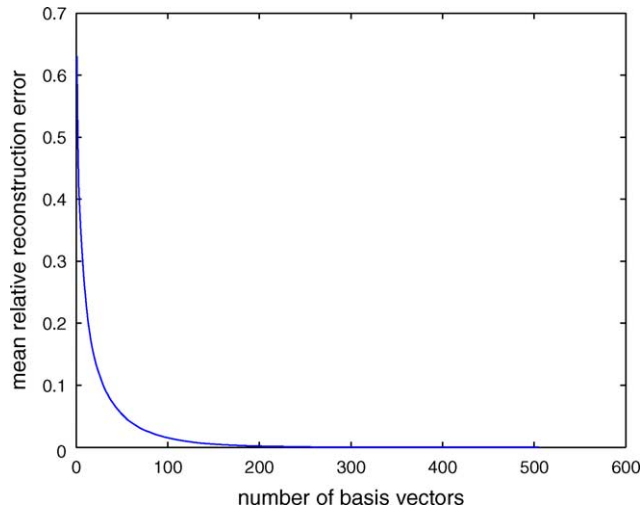


Fig. 3. Mean relative reconstruction error as a function of the number of basis vectors identified by the efficient feature vector selection algorithm for the Boston Housing data set.

the suburbs of Boston and 13 attributes representing environmental and social factors believed to be relevant (Harrison & Rubinfeld, 1978). The input patterns,  $\{\mathbf{x}_i\}_{i=1}^{\ell}$ , were first standardised, such that each element of the input vector has a zero mean and unit variance, allowing an isotropic Gaussian radial basis kernel (6) to be used. Again the hyper-parameters were chosen so as to minimise a 10-fold cross-validation estimate of the test sum-of-squares error ( $\gamma = 2.89 \times 10^4$  and  $\sigma^{-2} = 4.36$ ). Fig. 3 shows the mean relative reconstruction error (4) as a function of the number of basis vectors identified by the feature vector selection algorithm. The mean relative reconstruction error is negligible if more than  $\approx 200$  basis vectors are used.

Explicit and fast leave-one-out cross-validation procedures were then evaluated for sparse least-squares

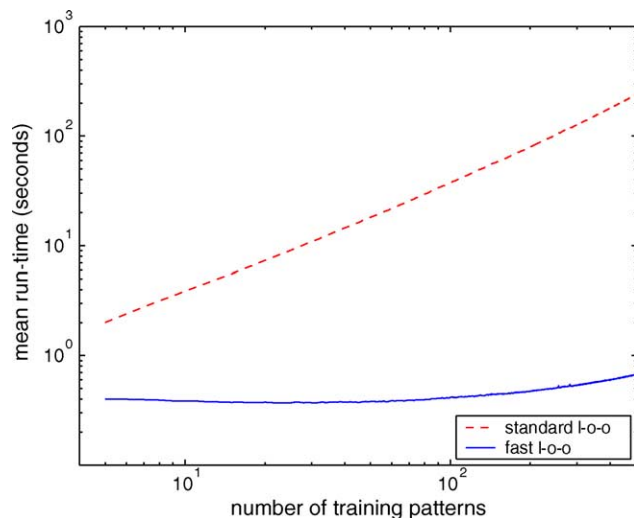


Fig. 4. Mean run-time for explicit and fast leave-one-out cross-validation procedures as a function of the number of training patterns (averaged over 20 trials).

regression networks based on the first 200 basis vectors identified by the feature vector selection algorithm. Fig. 4 shows the mean run-time in seconds for both algorithms as a function of the number of training patterns,  $\ell$ , over 20 trials. The fast leave-one-out cross-validation procedure is in all cases considerably faster than explicit evaluation, and has better scaling properties.

#### 4. Discussion

The application of the Bartlett correction formula to the efficient evaluation of the leave-one-out cross-validation error of linear models minimising a least-squares criterion has been known for some time in the field of statistics (Cook & Weisberg, 1982). The ‘kernel trick’ allows this method to be extended to non-linear models formed via construction of a linear model in a feature space defined by a Mercer kernel. In principal it should be possible to apply a similar procedure to related kernel models, such as least-squares support vector classification networks (Suykens & Vandewalle, 1999), Kernel Fisher Discriminant (KFD) analysis (Mika, Rätsch, Weston, Schölkopf, & Müller, 1999; Xu et al., 2001) and 2-norm support vector machines (Boser et al., 1992; Cortes & Vapnik, 1995; Vapnik, 1998), where the Lagrange multipliers associated with the support vectors are given by the solution of a system of linear equations corresponding to the Karush–Kuhn–Tucker (KKT) conditions.

#### 5. Summary

In this paper we have shown that the Bartlett correction formula can be used to implement exact leave-one-out cross-validation of sparse least-squares support vector regression networks with a computational complexity of only  $\mathcal{O}(\ell n^2)$ , rather than the  $\mathcal{O}(\ell^2 n^2)$  of a naïve direct implementation. Since the computational complexity of the basic training algorithm is also  $\mathcal{O}(\ell n^2)$ , leave-one-out cross-validation becomes a practical criterion for model selection for this class of kernel machines, as demonstrated by experiments on the Motorcycle and Boston Housing benchmark data sets.

#### Acknowledgements

The authors would like to thank the anonymous reviewers, Danilo Mandic, Tony Bagnall and Rob Foxall for their helpful comments on previous drafts. This work was supported by Royal Society Research Grant RSRG-22270.

## Appendix A. Proofs and derivations

### A.1. Proof of Lemma 1 (Luntz & Brailovsky, 1969)

Following the approach of Vapnik (1998), the expectation of the leave-one-out estimator is simply the integral over all datasets,  $\mathcal{Z}$ , constituting an i.i.d sample from  $P(\mathbf{z})$  of size  $\ell$ :

$$E\{\mathcal{L}(\mathcal{H}, \mathcal{Z})\} = \int \frac{1}{\ell} \sum_{i=1}^{\ell} Q(\mathbf{z}_i, h_i | \mathcal{Z} \setminus \{\mathbf{z}_i\}) dP(\mathbf{z}_1) \cdots dP(\mathbf{z}_\ell).$$

Since the dataset is formed by an i.i.d. sample from  $P(\mathbf{z})$ , the  $i$ th integral may be moved inside the summation:

$$= \int \frac{1}{\ell} \sum_{i=1}^{\ell} (\int Q(\mathbf{z}_i, h_i | \mathcal{Z} \setminus \{\mathbf{z}_i\}) dP(\mathbf{z}_i)) dP(\mathbf{z}_1) \cdots dP(\mathbf{z}_{i-1}) dP(\mathbf{z}_{i+1}) \cdots dP(\mathbf{z}_\ell).$$

Let  $\mathcal{Z}'$  represent an i.i.d. sample from  $P(\mathbf{z})$  of size  $\ell-1$

$$= E\left\{\frac{1}{\ell} \sum_{i=1}^{\ell} R(\mathcal{H}, \mathcal{Z}')\right\},$$

and so

$$E\{\mathcal{L}(\mathcal{H}, \mathcal{Z})\} = E\{R(\mathcal{H}, \mathcal{Z}')\},$$

and the lemma is proven.

### A.2. Proof of Lemma 2 (Bartlett, 1951)

The inverse of a block matrix

$$\mathbf{M} = \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{bmatrix},$$

can be found through block Gauss–Jordan elimination; beginning by manipulating the matrix into upper echelon form, we get

$$\mathbf{M}^{-1} = \begin{bmatrix} \mathbf{A}^{-1} + \mathbf{A}^{-1}\mathbf{B}\mathbf{S}^{-1}\mathbf{C}\mathbf{A}^{-1} & -\mathbf{A}^{-1}\mathbf{B}\mathbf{S}^{-1} \\ -\mathbf{S}^{-1}\mathbf{C}\mathbf{A}^{-1} & \mathbf{S}^{-1} \end{bmatrix},$$

where  $\mathbf{S} = \mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B}$ . Equivalently, one could begin by manipulating the matrix into lower echelon form, giving

$$\mathbf{M}^{-1} = \begin{bmatrix} \mathbf{T}^{-1} & -\mathbf{T}^{-1}\mathbf{B}\mathbf{D}^{-1} \\ -\mathbf{D}^{-1}\mathbf{C}\mathbf{T}^{-1}\mathbf{D}^{-1} & +\mathbf{D}^{-1}\mathbf{C}\mathbf{T}^{-1}\mathbf{B}\mathbf{D}^{-1} \end{bmatrix},$$

where  $\mathbf{T} = \mathbf{A} - \mathbf{B}\mathbf{D}^{-1}\mathbf{C}$ . Equating the upper left element of the matrices on the right hand sides of the two preceding expressions gives

$$\mathbf{T}^{-1} = \mathbf{A}^{-1} + \mathbf{A}^{-1}\mathbf{B}\mathbf{S}^{-1}\mathbf{C}\mathbf{A}^{-1},$$

i.e.

$$(\mathbf{A} - \mathbf{B}\mathbf{D}^{-1}\mathbf{C})^{-1} = \mathbf{A}^{-1} + \mathbf{A}^{-1}\mathbf{B}(\mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B})^{-1}\mathbf{C}\mathbf{A}^{-1}. \quad (\text{A.1})$$

Eq. (A.1) is known as the Sherman–Woodbury–Morrison formula (Golub & Van Loan, 1996; Hager, 1989; Sherman & Morrison, 1949, 1950; Woodbury, 1950). An important special case results from the assumption that  $\mathbf{D}$  is the identity matrix,  $\mathbf{I}$ , and  $\mathbf{B}$  and  $\mathbf{C}$  are column vectors,  $\mathbf{u}$  and  $\mathbf{v}$ , respectively (Bartlett, 1951):

$$(\mathbf{A} - \mathbf{u}\mathbf{v}^T)^{-1} = \mathbf{A}^{-1} + \frac{\mathbf{A}^{-1}\mathbf{u}\mathbf{v}^T\mathbf{A}^{-1}}{1 - \mathbf{v}^T\mathbf{A}^{-1}\mathbf{u}}.$$

### A.3. Derivation of training procedure for the sparse LS-SVM

The objective function minimised by the sparse least-squares support vector machine is given by

$$L(\beta, b) = \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n \beta_i \beta_j k_{ij} + \gamma \sum_{i=1}^{\ell} \left( t_i - \sum_{j=1}^n \beta_j \hat{k}_{ij} - b \right)^2,$$

where  $k_{ij} = \mathcal{K}(\mathbf{z}_i, \mathbf{z}_j)$  and  $\hat{k}_{ij} = \mathcal{K}(\mathbf{x}_i, \mathbf{z}_j)$ . Setting the partial derivatives of the objective function with respect to  $\beta$  and  $b$  to zero, and dividing through by  $2\gamma$ , yields

$$\sum_{i=1}^n \beta_i \sum_{j=1}^{\ell} \hat{k}_{ji} + b = \sum_{j=1}^{\ell} t_j$$

and  $\forall r \in \{1, 2, \dots, n\}$

$$\sum_{i=1}^n \beta_i \left( \frac{1}{2} \gamma^{-1} k_{ir} + \sum_{j=1}^{\ell} \hat{k}_{jr} \hat{k}_{ji} \right) + b \sum_{i=1}^{\ell} \hat{k}_{ir} = \sum_{i=1}^{\ell} t_i \hat{k}_{ir},$$

These equations can be expressed as a system of  $n+1$  linear equations in  $n+1$  unknowns

$$\begin{bmatrix} \mathbf{\Omega} & \mathbf{\Phi} \\ \mathbf{\Phi}^T & 1 \end{bmatrix} \begin{bmatrix} \beta \\ b \end{bmatrix} = \begin{bmatrix} \mathbf{c} \\ \sum_{k=1}^{\ell} t_k \end{bmatrix}, \quad (\text{A.2})$$

where the  $(i, j)$ th element of the sub-matrix  $\mathbf{\Omega}$  is

$$\omega_{ij} = \frac{1}{2} \gamma^{-1} k_{ij} + \sum_{l=1}^{\ell} \hat{k}_{lj} \hat{k}_{li}, \quad 1 \leq i, j \leq n,$$

the  $j$ th element of the vector  $\mathbf{\Phi}$  is

$$\phi_j = \sum_{i=1}^{\ell} \hat{k}_{ij}, \quad 1 \leq j \leq n,$$

and the  $j$ th element of the vector  $\mathbf{c}$  is:

$$c_j = \sum_{i=1}^{\ell} t_i \hat{k}_{ij}, \quad 1 \leq j \leq n.$$

This matrix equation can be simplified by referring to the vector consisting of all the model parameters as  $\mathbf{p} = (\beta, b)^T$ ,

and by defining the matrices  $\mathbf{R}$  and  $\mathbf{Z}$  as

$$\mathbf{R} = \begin{bmatrix} \frac{1}{2\gamma}\mathbf{K} & \mathbf{0} \\ \mathbf{0}^T & \mathbf{0} \end{bmatrix}, \quad \mathbf{Z} = [\hat{\mathbf{K}} \quad \mathbf{1}]$$

where the  $(i, j)$ th element of  $\mathbf{K}$  is  $k_{ij}$ , and the  $(i, j)$ th element of  $\hat{\mathbf{K}}$  is  $\hat{k}_{ij}$ . Eq. (A.2) can then be written as:

$$(\mathbf{R} + \mathbf{Z}^T\mathbf{Z})\mathbf{p} = \mathbf{Z}^T\mathbf{t}. \quad (\text{A.3})$$

#### A.4. Derivation of closed form expression for predicted residuals

From Eq. (A.3) we know that the vector of model parameters  $\mathbf{p} = (\boldsymbol{\beta}, b)$  is given by

$$\mathbf{p} = (\mathbf{R} + \mathbf{Z}^T\mathbf{Z})^{-1}\mathbf{Z}^T\mathbf{t}$$

where  $\mathbf{Z} = [\hat{\mathbf{K}} \quad \mathbf{1}]$ . For convenience, let  $\mathbf{C} = \mathbf{R} + \mathbf{Z}^T\mathbf{Z}$  and  $\mathbf{d} = \mathbf{Z}^T\mathbf{t}$ , such that  $\mathbf{p} = \mathbf{C}^{-1}\mathbf{d}$ . Furthermore, let  $\mathbf{Z}_{(i)}$  and  $\mathbf{t}_{(i)}$  represent the data with the  $i$ th observation deleted, then

$$\mathbf{C}_{(i)} = \mathbf{C} - \mathbf{z}_i\mathbf{z}_i^T,$$

and

$$\mathbf{d}_{(i)} = \mathbf{d} - t_i\mathbf{z}_i.$$

The Bartlett matrix inversion formula then gives

$$\mathbf{C}_{(i)}^{-1} = \mathbf{C} + \frac{\mathbf{C}^{-1}\mathbf{z}_i\mathbf{z}_i^T\mathbf{C}^{-1}}{1 - \mathbf{z}_i^T\mathbf{C}^{-1}\mathbf{z}_i},$$

such that the vector of model parameters during the  $i$ th iteration of the leave-one-out cross-validation procedure becomes

$$\mathbf{p}_{(i)} = \left( \mathbf{C} + \frac{\mathbf{C}^{-1}\mathbf{z}_i\mathbf{z}_i^T\mathbf{C}^{-1}}{1 - \mathbf{z}_i^T\mathbf{C}^{-1}\mathbf{z}_i} \right) (\mathbf{d} - t_i\mathbf{z}_i).$$

Let  $\mathbf{H} = \mathbf{Z}\mathbf{C}^{-1}\mathbf{Z}^T$  represent the *hat* matrix; note that the  $i$ th element of the leading diagonal can be written  $h_{ii} = \mathbf{z}_i^T\mathbf{C}^{-1}\mathbf{z}_i$ , so expanding the brackets we have

$$\begin{aligned} \mathbf{p}_{(i)} &= \mathbf{C}^{-1}\mathbf{d} - \mathbf{C}^{-1}t_i\mathbf{z}_i + \frac{\mathbf{C}^{-1}\mathbf{z}_i\mathbf{z}_i^T\mathbf{C}^{-1}}{1 - \mathbf{z}_i^T\mathbf{C}^{-1}\mathbf{z}_i}\mathbf{d} - \frac{\mathbf{C}^{-1}\mathbf{z}_i\mathbf{z}_i^T\mathbf{C}^{-1}}{1 - \mathbf{z}_i^T\mathbf{C}^{-1}\mathbf{z}_i}t_i\mathbf{z}_i \\ &= \mathbf{p} + \left( \frac{\mathbf{z}_i^T\mathbf{p} - t_i}{1 - h_{ii}} \right) \mathbf{C}^{-1}\mathbf{z}_i. \end{aligned}$$

The residual error for the  $i$ th training pattern is  $e_i = t_i - \mathbf{z}_i^T\mathbf{p}$  and so

$$\mathbf{p}_{(i)} = \mathbf{p} - \frac{e_i}{1 - h_{ii}}\mathbf{C}^{-1}\mathbf{z}_i.$$

Noting that  $\mathbf{y} = \mathbf{Z}\mathbf{p}$ , the output of the model during the  $i$ th iteration of the leave-one-out cross-validation procedure can be written as

$$\mathbf{y}_{(i)} = \mathbf{Z}\mathbf{p}_{(i)} = \mathbf{y} - \frac{e_i}{1 - h_{ii}}\mathbf{h}_i$$

where  $\mathbf{h}_i$  is the  $i$ th column of  $\mathbf{H}$ . The vector of residuals for the training patterns during the leave-one-out cross-validation procedure can then be written in closed form as

$$\mathbf{e}_{(i)} = \mathbf{t} - \mathbf{y}_{(i)} = \mathbf{e} + e_i \frac{1}{1 - h_{ii}}\mathbf{h}_i.$$

The  $i$ th element of  $\mathbf{e}_{(i)}$  is given by

$$\{\mathbf{e}_{(i)}\}_i = \frac{1}{1 - h_{ii}}e_i.$$

## References

- Allen, D. M. (1974). The relationship between variable selection and prediction. *Technometrics*, 16, 125–127.
- Bartlett, M. S. (1951). An inverse matrix adjustment arising in discriminant analysis. *Annals of Mathematical Statistics*, 22(1), 107–111.
- Baudat, G., & Anouar, F. (2001). Kernel-based methods and function approximation. In *Proceedings of IJCNN, Washington, DC, July 2001* (pp. 1244–1249).
- Boser, B., Guyon, I., & Vapnik, V. N. (1992). A training algorithm for optimal margin classifiers. In *Proceedings of the fifth annual workshop on computational learning theory, Pittsburgh*. ACM (pp. 144–152).
- Burges, C. J. C. (1996). Simplified support vector decision rules. In *Proceedings of the 13th international conference on machine learning, Bari, Italy* (pp. 71–77).
- Burges, C. J. C., & Schölkopf, B. (1997). Improving speed and accuracy of support vector learning machines. In *Advances in neural information processing systems* (Vol. 9) (pp. 375–381). Cambridge, MA: MIT Press.
- Cawley, G. C., & Talbot, N. L. C. (2002a). A greedy training algorithm for sparse least-squares support vector machines. In *Proceedings of the international conference on artificial neural networks (ICANN-2002), Madrid, Spain, August 27–30 2002. Vol. 2415. Lecture Notes in Computer Science (LNCS)* (pp. 681–686). Berlin: Springer.
- Cawley, G. C., & Talbot, N. L. C. (2002b). Improved sparse least-squares support vector machines. *Neurocomputing*, 48, 1025–1031.
- Cawley, G. C., & Talbot, N. L. C. (2002c). Efficient formation of a basis in a kernel induced feature space. In *Proceedings of the European symposium on artificial neural networks (ESANN-2002), Bruges, Belgium, April 24–26 2002* (pp. 1–6).
- Chapelle, O., Vapnik, V., Bousquet, O., & Mukherjee, S. (2002). Choosing multiple parameters for support vector machines. *Machine Learning*, 46(1), 131–159.
- Cook, R. D., & Weisberg, S. (1982). *Residuals and influence in regression. Monographs on statistics and applied probability*. New York: Chapman & Hall.
- Cortes, C., & Vapnik, V. (1995). Support vector networks. *Machine Learning*, 20, 1–25.
- Cristianini, N., & Shawe-Taylor, J. (2000). *An introduction to support vector machines (and other kernel-based learning methods)*. Cambridge: Cambridge University Press.
- Downs, T., Gates, K. E., & Masters, A. (2001). Exact simplification of support vector solutions. *Journal of Machine Learning Research*, 2, 293–297.
- Geman, S., Bienenstock, E., & Doursat, R. (1992). Neural networks and the bias/variance dilemma. *Neural Computation*, 4, 1–58.
- Golub, G. H., & Van Loan, C. F. (1996). *Matrix computations* (3rd ed.). Baltimore, MD: Johns Hopkins University Press.
- Hager, W. W. (1989). Updating the inverse of a matrix. *SIAM Review*, 31(2), 221–239.
- Harrison, D., & Rubinfeld, D. L. (1978). Hedonic prices and the demand for clean air. *Journal Environmental Economics and Management*, 5, 81–102.



- Hastie, T., Tibshirani, R., & Friedman, J. (2001). *The elements of statistical learning: data mining, inference, and prediction*. Springer texts in statistics. Berlin: Springer.
- Hoerl, A. E., & Kennard, R. W. (1970). Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1), 55–67.
- Keerthi, S. S. (2002). Efficient tuning of SVM hyperparameters using radius/margin bound and iterative algorithms. *IEEE Transactions on Neural Networks*, 13(5), 1225–1229.
- Kimeldorf, G. S., & Wahba, G. (1971). Some results on Tchebycheffian spline functions. *Journal of Mathematical and Analytical Applications*, 33, 82–95.
- Kohavi, R. (1995). A study of cross-validation and bootstrap for accuracy estimation and model selection. In *Proceedings of the 14th international conference on artificial intelligence (IJCAI)* (pp. 1137–1143). San Mateo, CA: Morgan Kaufmann.
- Lee, Y.-J., & Mangasarian, O. L. (2001). RSVM: Reduced support vector machines. In *Proceedings of the first SIAM international conference on data mining, Chicago, IL, USA, April 5–7 2001*.
- Luntz, A., & Brailovsky, V. (1969). On estimation of characters obtained in statistical procedure of recognition. *Tekhnicheskaya Kibernetika*, 3 (in Russian).
- Micchelli, C. A. (1986). Interpolation of scattered data: Distance matrices and conditionally positive definite functions. *Constructive Approximation*, 2, 11–22.
- Mika, S., Rätsch, G., Weston, J., Schölkopf, B., & Müller, K.-R. (1999). Fisher discriminant analysis with kernels, Vol. IX. In *Neural networks for signal processing* (pp. 41–48). New York: IEEE Press.
- Nelder, J. A., & Mead, R. (1965). A simplex method for function minimization. *Computer Journal*, 7, 308–313.
- Poggio, T., & Girosi, F. (1990). Networks for approximation and learning. *Proceedings of the IEEE*, 78(9).
- Saunders, C., Gammernan, A., & Vovk, V. (1998). Ridge regression learning algorithm in dual variables. In *Proceedings of the 15th International Conference on Machine Learning, Madison, WI, July 24–27 1998* (pp. 24–27).
- Sherman, J., & Morrison, W. J. (1949). Adjustment of an inverse matrix corresponding to changes in the elements of a given column or a given row of the original matrix. *Annals of Mathematical Statistics*, 20(4), 621.
- Sherman, J., & Morrison, W. J. (1950). Adjustment of an inverse matrix corresponding to a change in one element of a given matrix. *Annals of Mathematical Statistics*, 21(1), 124–127.
- Silverman, B. W. (1985). Some aspects of the spline smoothing approach to non-parametric regression curve fitting. *Journal of the Royal Statistical Society, B*, 47(1), 1–52.
- Smola, A. J., & Schölkopf, B. (2000). Sparse greedy matrix approximation for machine learning. In *Proceedings of the 17th international conference on machine learning, June 29–July 2 2000* (pp. 911–918). San Francisco, CA: Morgan Kaufmann.
- Suykens, J., Lukas, L., & Vandewalle, J. (2000). Sparse approximation using least-squares support vector machines. In *Proceedings of the IEEE International Symposium on Circuits and Systems, Geneva, Switzerland, May 2000* (pp. 11757–11760).
- Suykens, J., & Vandewalle, J. (1999). Least squares support vector machine classifiers. *Neural Processing Letters*, 9(3), 293–300.
- Suykens, J. A. K., De Brabanter, J., Lukas, L., & Vandewalle, J. (2002). Weighted least squares support vector machines: robustness and sparse approximation. *Neurocomputing*, 48(1–4), 85–105.
- Suykens, J. A. K., Lukas, L., Van Dooren, P., De Moor, B., & Vandewalle, J. (1999). Least squares support vector machine classifiers: a large scale algorithm. In *Proceedings of the European Conference on Circuit Theory and Design (ECCTD-99), Stresa, Italy, September 1999* (pp. 839–842).
- Tikhonov, A. N., & Arsenin, V. Y. (1977). *Solutions of ill-posed problems*. New York: Wiley.
- Vapnik, V. N. (1998). *Statistical learning theory Wiley series on adaptive and learning systems for signal processing, communications and control*. New York: Wiley.
- Vapnik, V., & Chapelle, O. (2000). Bounds on error expectation for support vector machines. *Neural Computation*, 12(9), 2013–2036.
- Weisberg, S. (1985). *Applied linear regression*. New York: Wiley.
- Williams, C. K. I., & Seeger, M. (2001). Using the Nyström method to speed up kernel machines. In T. K. Leen, T. G. Dietterich, & V. Tresp (Eds.), *Advances in Neural Information Processing Systems 13* (pp. 682–688). Cambridge, MA: MIT Press.
- Woodbury, M. (1950). *Inverting modified matrices. Memorandum report 42*. Princeton: Princeton University.
- Xu, J., Zhang, X., & Li, Y. (2001). Kernel MSE algorithm: A unified framework for KFD, LS-SVM and KRR. In *Proceedings of IJCNN, Washington, DC, July 2001* (pp. 1486–1491).