Available online at www.sciencedirect.com
SCIENCE DIRECT®
Neurocomputing 64 (2005) 137–159
NEUROCOMPUTING
www.elsevier.com/locate/neucom

Building sparse representations and structure determination on LS-SVM substrates

K. Pelckmans*, J.A.K. Suykens, B. De Moor

K.U. Leuven, ESAT-SCD/SISTA, Kasteelpark Arenberg 10, B-3001 Leuven (Heverlee), Belgium

Available online 20 January 2005

Abstract

This paper proposes a new method to obtain sparseness and structure detection for a class of kernel machines related to least-squares support vector machines (LS-SVMs). The key method is to adopt an hierarchical modeling strategy. Here, the first level consists of an LS-SVM substrate which is based upon an LS-SVM formulation with additive regularization trade-off. This regularization trade-off is determined at higher levels such that sparse representations and/or structure detection are obtained. Using the necessary and sufficient conditions for optimality given by the Karush–Kuhn–Tucker conditions, one can guide the interaction between different levels via a well-defined set of hyper-parameters. From a computational point of view, all levels can be fused into a single convex optimization problem. Furthermore, the principle is applied in order to optimize the validation performance of the resulting kernel machine. Sparse representations as well as structure detection are obtained, respectively, by using an L_1 regularization scheme and a measure of maximal variation at the second level. A number of case studies indicate the usefulness of these approaches both with respect to interpretability of the final model as well as generalization performance.
© 2005 Elsevier B.V. All rights reserved.

Keywords: Least-squares support vector machines; Regularization; Structure detection; Model selection; Convex optimization; Sparseness

*Corresponding author. Tel.: +32 16 32 86 58; fax: +32 16 32 19 70.
E-mail addresses: kristiaan.pelckmans@esat.kuleuven.ac.be (K. Pelckmans), johan.suykens@esat.kuleuven.ac.be (J.A.K. Suykens).
URL: http://www.esat.kuleuven.ac.be/sista/lssvmlab/.

0925-2312/$ - see front matter © 2005 Elsevier B.V. All rights reserved.
doi:10.1016/j.neucom.2004.11.029

1. Introduction

The problem of inference of a model based on a finite set of observational data is often ill-posed [19]. To address this problem, typically a form of capacity control is introduced which is often expressed mathematically in the form of regularization [28]. Regularized cost functions have been applied successfully, e.g. in splines, multilayer perceptrons, regularization networks [18], support vector machines (SVM) and related methods (see, e.g. [11]). SVM [29] is a powerful methodology for solving problems in nonlinear classification, function estimation and density estimation which has also led to many other recent developments in kernel-based learning methods in general [22]. SVMs have been introduced within the context of statistical learning theory and structural risk minimization. In the methods, one solves convex optimization problems, typically quadratic programs. Least-Squares Support Vector Machines (LS-SVMs) [21,26] are reformulations to standard SVMs which lead to solving linear Karush–Kuhn–Tucker (KKT) systems for classification tasks as well as regression. Primal–dual LS-SVM formulations have also been given for KFDA, KPCA, KCCA, KPLS, recurrent networks and control [27].[1] Recently, LS-SVM methods were studied in combination with additive models [10] resulting in so-called componentwise LS-SVMs [15] which are suitable for component selection. So-called additive models consisting of a sum of lower dimensional nonlinearities per component (variable) have become one of the widely used non-parametric techniques as they offer a compromise between the somewhat conflicting requirements of flexibility, dimensionality and interpretability (see e.g. [11]).

The relative importance between the *smoothness* of the solution (as defined in different ways) and the norm of the residuals in the cost function involves a tuning parameter, usually called the regularization constant. The determination of regularization constants is important in order to achieve good generalization performance with the trained model and is an important problem in statistics and learning theory (see e.g. [11–13,22,25]). Several model selection criteria have been proposed in the literature to tune this constant. Special attention was given in the machine learning community to cross-validation and leave-one-out-based methods [24] and fast implementations were studied in the context of kernel machines (see e.g. [2]). In the following paper, the performance on an independent validation dataset is considered. The optimization of the regularization constant in LS-SVMs with respect to this criterion can be non-convex (and even non-smooth) in general. In order to overcome this difficulty, a re-parameterization of the regularization trade-off has been recently introduced in [16] referred to as *additive regularization trade-off*. When applied to the LS-SVM formulation, this leads to LS-SVM substrates. In [16], it was illustrated how to employ these LS-SVM substrates to obtain models which were optimal in training and validation or cross-validation sense.

[1] The Internet portal for LS-SVM related research and software can be found at http://www.esat.kuleuven.ac.be/sista/lssvmlab.

This paper investigates these methods towards hierarchical modeling based on optimization theory [1]. As in a Bayesian evidence framework [13], different hierarchical levels can be considered for the inference of model (parameters), regularization constants and model structure. The proposed hierarchical kernel machines are not formulated within a Bayesian context and lead to convex optimization problems (while application of the Bayesian framework often results in non-convex optimization problems with many local minima). Fig. 1 shows the hierarchical structure. The LS-SVM substrate constitutes the basis level of the hierarchy (*level* 1). This paper shows how one can consider additional levels resulting in sparse representations and structure detection by taking suitable cost-functions at level 2. By fusing all levels into one single problem, the resulting hierarchical kernel machine can be obtained by solving a convex optimization problem. Sparse representations are obtained through the use of an L_1-based regularization scheme at level 2 and are motivated by a sensitivity argument from optimization theory. This hierarchy is visualized in Fig. 2. Structure detection is expressed at the second level through a measure of maximal variation of a component, leading to the hierarchical kernel machine as represented in Fig. 3. The hierarchical kernel model is finalized by an additional third level which tunes the remaining hyper-parameters of the second level with a validation criterion.

This paper is organized as follows: Section 2 reviews the formulations of the LS-SVM regressor and its extensions towards componentwise LS-SVMs. Section 3 explains the structure of hierarchical kernel machines based on LS-SVM substrates

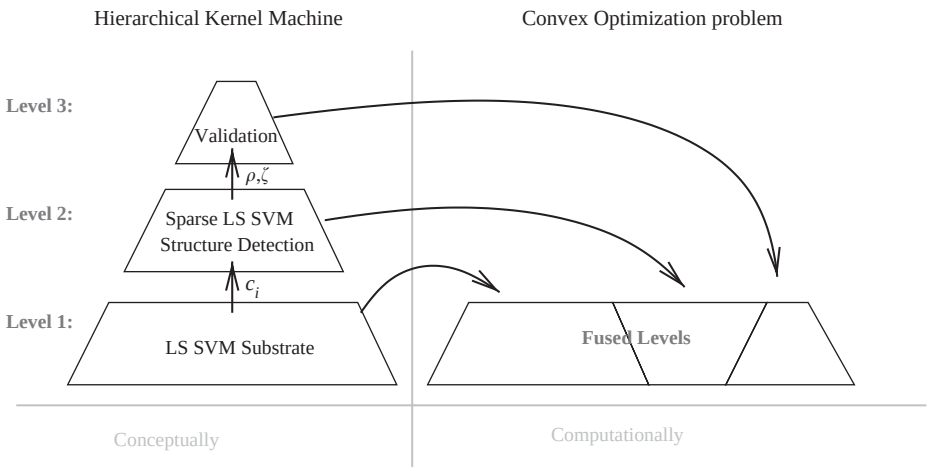

Fig. 1. Schematical representation of the design of hierarchical kernel machines. From a conceptual point of view, one builds upon the LS-SVM substrate which takes care of the classical ingredients of kernel machines. Hyper-parameters guide the interaction between different levels while KKT conditions enforce the individual levels to be defined properly. From a computational point of view, the whole can be fused into one convex optimization problem.

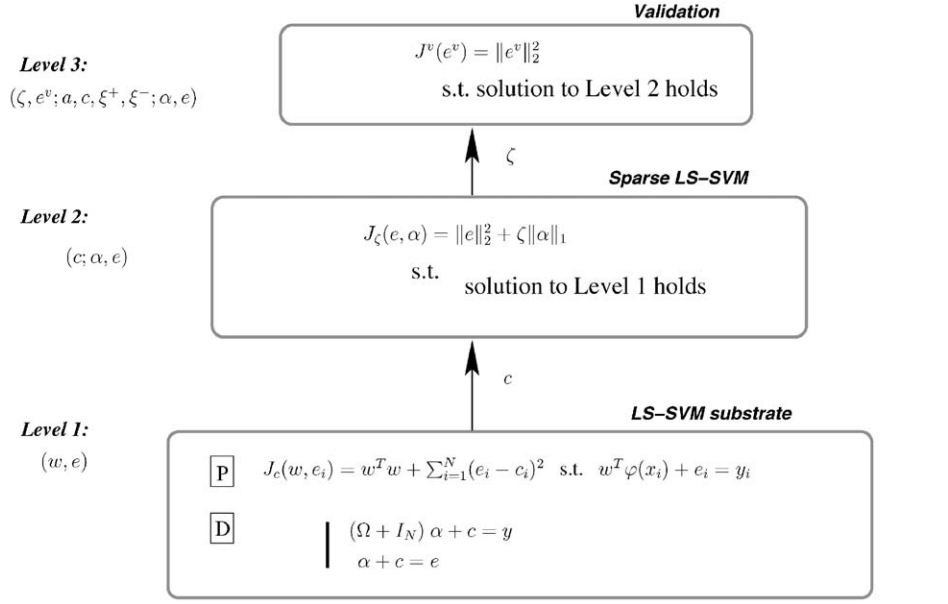

Computationally

(Convex Optimization Problem obtained after Fusion)

$$\min_{\zeta,\xi^+,\xi^-,a;\alpha} \|\Omega\alpha - y^v\|^2 + b^- \xi^{-T}(a - \alpha) + b^+ \xi^{+T}(a + \alpha)$$

s.t.

$$\Omega^{PT}\Omega^P\alpha - y^T\Omega^P = (\xi^- - \xi^+)$$
$$\xi^-, \xi^+ \geq 0$$
$$-a \leq \alpha \leq a$$
$$\zeta = 1_N^T(\xi^- + \xi^+)$$

Fig. 2. Schematical representation of the hierarchical kernel machine for sparse representations. From a conceptual point of view, inference is done at different levels and interaction is guided via a set of hyper-parameters. The first level constitutes of an LS-SVM substrate. On the second level, inference of the c_i is defined in terms of a cost function inducing sparseness, while ζ is optimized on a third level using a validation criterion.

and employs this idea for obtaining structure detection and LS-SVMs with sparse support values. Section 4 discusses how one can optimize the LS-SVM substrate or the sparse LS-SVM with respect to the performance on a validation set. Section 5 presents numerical results on both artificial and benchmark datasets. This paper is an extended version of the ESANN 2004 contribution [17], with a discussion of more efficient methods. A main improvement of this paper is the use of a three-level hierarchy instead of a two-level one as in the previous papers on the subject.

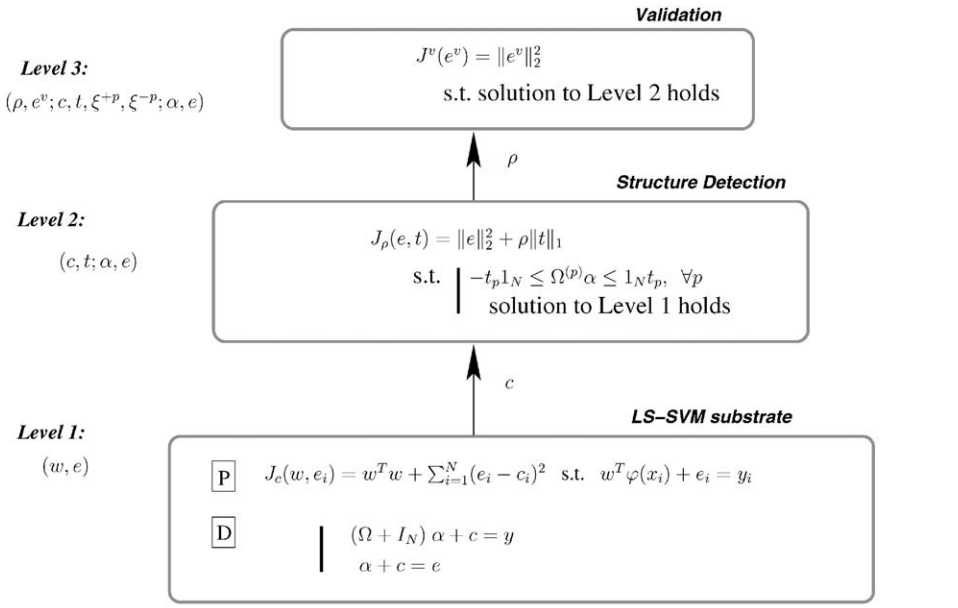

Computationally

(Convex Optimization Problem obtained after Fusion)

$$\min_{\rho,\xi^+,\xi^-,t,\alpha} \|\Omega^{P,v}\alpha - y^v\|^2 + \sum_{p=1}^{P}\left[b_p^- \xi^{-pT}\left(t_p 1_N - \Omega^{(p)}\alpha\right) + b_p^+ \xi^{+pT}\left(t_p 1_N + \Omega^{(p)}\alpha\right)\right]$$

$$\text{s.t.} \quad \begin{array}{l} \Omega^{PT}\Omega^P\alpha - y^T\Omega^P = \sum_{p=1}^{P}(\xi^{-p} - \xi^{+p}) \\ \xi^{-p}, \xi^{+p} \geq 0, \ \forall p \\ -t_p 1_N \leq \Omega^{(p)}\alpha \leq 1_N t_p, \ \forall p \\ \rho = 1_N^T(\xi^{-p} + \xi^{+p}), \ \forall p \end{array}$$

Fig. 3. Schematical representation of the hierarchical kernel machine for structure detection. On the second level, inference of the c_i is expressed in terms of a least-squares cost function with a minimal amount of maximal variation, while ρ is optimized on a third level using a validation criterion.

2. Model training

Let $\{(x_i, y_i)\}_{i=1}^N \subset \mathbb{R}^d \times \mathbb{R}$ be the training data with inputs x_i and outputs y_i. Consider the regression model $y_i = f(x_i) + e_i$, where $x_1, ..., x_N$ are deterministic points, $f : \mathbb{R}^d \to \mathbb{R}$ is an unknown real-valued smooth function and $e_1, ..., e_N$ are uncorrelated random errors with $E[e_i] = 0$, $E[e_i^2] = \sigma_e^2 < \infty$. The n data points of the validation set are denoted as $\{(x_j^v, y_j^v)\}_{j=1}^n$. In the case of classification, $y_i, y_j^v \in \{-1, 1\}$ for all $i = 1, ..., N$ and $j = 1, ..., n$. Let y denote $(y_1, ..., y_N)^{\mathrm{T}} \in \mathbb{R}^N$ and $y^v = (y_1^v, ..., y_n^v)^{\mathrm{T}} \in \mathbb{R}^n$.

2.1. Least-Squares Support Vector Machines

The LS-SVM model is given as $f(x) = w^{\mathrm{T}}\varphi(x) + b$ in the primal space, where $\varphi(\cdot) : \mathbb{R}^d \to \mathbb{R}^{n_h}$ denotes the potentially infinite ($n_h = \infty$) dimensional feature map. The regularized least-squares cost function is given as [21,27]

$$\min_{w,b,e_i} \mathscr{J}_\gamma(w,e) = \frac{1}{2}w^{\mathrm{T}}w + \frac{\gamma}{2}\sum_{i=1}^{N} e_i^2$$
$$\text{s.t. } w^{\mathrm{T}}\varphi(x_i) + b + e_i = y_i, \quad \forall i = 1,\ldots,N. \tag{1}$$

In the sequel, we refer to this scheme as standard LS-SVM regressors with Tikhonov regularization. Let $e = (e_1,\ldots,e_N)^{\mathrm{T}} \in \mathbb{R}^N$ denote a vector of residuals. The Lagrangian of the constrained optimization problem becomes $\mathscr{L}_\gamma(w,b,e_i;\alpha_i) = 0.5w^{\mathrm{T}}w + 0.5\gamma\sum_{i=1}^{N} e_i^2 - \sum_{i=1}^{N}\alpha_i(w^{\mathrm{T}}\varphi(x_i) + b + e_i - y_i)$. By taking the conditions for optimality $\partial\mathscr{L}_\gamma/\partial\alpha_i = 0$, $\partial\mathscr{L}_\gamma/\partial b = 0$, $\partial\mathscr{L}_\gamma/\partial w = 0$, $\partial\mathscr{L}_\gamma/\partial e_i = 0$ and application of the kernel trick $K(x_i,x_j) = \varphi(x_i)^{\mathrm{T}}\varphi(x_j)$ with a positive definite (Mercer) kernel K, one gets $e_i\gamma = \alpha_i$, $w = \sum_{i=1}^{N}\alpha_i\varphi(x_i)$, $\sum_{i=1}^{N}\alpha_i = 0$ and $w^{\mathrm{T}}\varphi(x_i) + b + e_i = y_i$. The dual problem is summarized as

$$\left[\begin{array}{c|c} 0 & 1_N^{\mathrm{T}} \\ \hline 1_N & \Omega + I_N/\gamma \end{array}\right]\left[\begin{array}{c} b \\ \hline \alpha \end{array}\right] = \left[\begin{array}{c} 0 \\ \hline y \end{array}\right], \tag{2}$$

where $\Omega \in \mathbb{R}^{N\times N}$ with $\Omega_{ij} = K(x_i,x_j)$, $I_N \in \mathbb{R}^{N\times N}$ denotes the identity matrix and $1_N \in \mathbb{R}^N$ is a vector containing all ones. The estimated function $\hat{f}$ can be evaluated at a new point x^* by $\hat{f}(x^*) = \sum_{i=1}^{N}\hat{\alpha}_i K(x_i,x^*) + \hat{b}$ where $\hat{\alpha}$ and $\hat{b}$ are the unique solution to (2).

Optimization of the regularization constant γ with respect to a validation performance in the regression case can be written as

$$\min_\gamma \sum_{j=1}^{n}\left(y_j^v - \hat{f}_\gamma(x_j^v)\right)^2 = \sum_{j=1}^{n}\left(y_j^v - \left[\begin{array}{c} 1 \\ \hline \Omega^v \end{array}\right]^{\mathrm{T}}\left[\begin{array}{c|c} 0 & 1_N^{\mathrm{T}} \\ \hline 1_N & \Omega + I_N/\gamma \end{array}\right]^{-1}\left[\begin{array}{c} 0 \\ \hline y \end{array}\right]\right)^2, \tag{3}$$

where $\Omega^v \in \mathbb{R}^{N\times n}$ with $\Omega_{ij}^v = K(x_i,x_j^v)$ for all $i = 1,...,N$ and $j = 1,...,n$. This optimization problem in the hyper-parameter γ is usually non-convex. For the choice of the kernel $K(\cdot,\cdot)$, see e.g. [4,8]. Typical examples are the use of a linear kernel $K(x_i,x_j) = x_i^{\mathrm{T}}x_j$ or the RBF kernel $K(x_i,x_j) = \exp(-\|x_i - x_j\|_2^2/\sigma^2)$, where σ denotes the bandwidth of the kernel and $\|\cdot\|_2^2$ denotes the squared Euclidean norm.

A derivation of LS-SVMs was given originally for the classification task [26]. The LS-SVM classifier $\mathrm{sign}(w^{\mathrm{T}}\varphi(x)+b)$ is optimized with respect to

$$\min_{w,b,e_i} \mathscr{J}_\gamma(w,e) = \frac{1}{2}w^{\mathrm{T}}w + \frac{\gamma}{2}\sum_{i=1}^{N} e_i^2$$
$$\text{s.t. } y_i(w^{\mathrm{T}}\varphi(x_i)+b) = 1-e_i, \quad \forall i=1,\ldots,N, \tag{4}$$

where $e=(e_1,\ldots,e_N)^{\mathrm{T}}\in\mathbb{R}^N$ are slack-variables to tolerate misclassifications. Using a primal–dual optimization interpretation, the unknowns $\hat{\alpha}$ and $\hat{b}$ of the estimated classifier $\mathrm{sign}(\sum_{i=1}^{N}\hat{\alpha}_i y_i K(x_i,x)+\hat{b})$ are found by solving the dual set of linear equations

$$\left[\begin{array}{c|c} 0 & y^{\mathrm{T}} \\ \hline y & \Omega^y + I_N/\gamma \end{array}\right]\left[\begin{array}{c} b \\ \hline \alpha \end{array}\right] = \left[\begin{array}{c} 0 \\ \hline 1_N \end{array}\right], \tag{5}$$

where $\Omega^y\in\mathbb{R}^{N\times N}$ with $\Omega^y_{ij}=y_iy_jK(x_i,x_j)$ for all $i,j=1,\ldots,N$. The remainder focuses mainly on the regression case, although it is applicable also to the classification problem as illustrated in Section 5.

2.2. Componentwise LS-SVMs

It is often useful to consider less general classes of nonlinear models as the additive model class [10] in order to overcome the curse of dimensionality as they offer a compromise between the somewhat conflicting requirements of flexibility, dimensionality and interpretability (see e.g. [11]). This subsection reviews results obtained in [15]. Let a superscript (p) denote the pth component $x_i^{(p)}=(x_i^{p_1},\ldots,x_i^{p_{D_p}})^{\mathrm{T}}\in\mathbb{R}^{D_p}$ such that every datasample consists of the union of P different components $x_i=\bigcup\{x_i^{(1)},\ldots,x_i^{(P)}\}$. In the simplest case, $n_p=1$ such that $x_i^{(p)}=x_i^p$. Consider the following model:

$$f(x)=\sum_{p=1}^{P} w_p^{\mathrm{T}}\varphi_p(x^{(p)})+b, \tag{6}$$

where $\varphi_p:\mathbb{R}\to\mathbb{R}^{d_f}$ is a possibly infinite dimensional mapping of the pth component. One considers the regularized least-squares cost-function as introduced in [15],

$$\min_{w_p,b,e_i} \mathscr{J}_\gamma(w_p,e) = \frac{1}{2}\sum_{p=1}^{P} w_p^{\mathrm{T}}w_p + \frac{\gamma}{2}\sum_{i=1}^{N} e_i^2$$
$$\text{s.t.} \sum_{p=1}^{P} w_p^{\mathrm{T}}\varphi_p(x_i^{(p)})+b+e_i=y_i, \quad \forall i=1,\ldots,N. \tag{7}$$

Construction of the Lagrangian and taking the conditions for optimality as in the previous subsection results in the following linear system

$$\left[\begin{array}{c|c} 0 & 1_N^{\mathrm{T}} \\ \hline 1_N & \Omega^P + I_N/\gamma \end{array}\right] \left[\begin{array}{c} b \\ \hline \alpha \end{array}\right] = \left[\begin{array}{c} 0 \\ \hline y \end{array}\right], \tag{8}$$

$$\gamma e = \alpha,$$

$$w_p^{\mathrm{T}} \varphi_p\left(x_i^{(p)}\right) = \Omega_i^{(p)} \alpha, \quad \forall i = 1, \ldots, N, \quad \forall p = 1, \ldots, P,$$

where $\Omega^P = \sum_{p=1}^{P} \Omega^{(p)} \in \mathbb{R}^{N \times N}$ and $\Omega_{ij}^{(p)} = K_p(x_i^{(p)}, x_j^{(p)})$ is the kernel evaluated between the pth component of the ith and the jth data point. A qualitative comparison of the componentwise LS-SVM and the iterative back-fitting procedure [10], the kernel ANOVA decomposition [9,23] and the splines technique for additive models MARS [11,31] was given in [15]. Note that the difference between (2) and (8) can be stated entirely in terms of the used kernels, though the starting point was a different model and optimality principle.

3. Hierarchical kernel models for structure detection and sparse representations

3.1. Level 1: substrate LS-SVMs with additive regularization trade-off

An alternative way to parameterize the regularization trade-off associated with the model $f(x) = w^{\mathrm{T}}\varphi(x) + b$ is by means of a tuning vector $c \in \mathbb{R}^N$ [16] which penalizes the error variables in an additive way, i.e.

$$\min_{w,b,e_i} \mathcal{J}_c(w, e) = \frac{1}{2} w^{\mathrm{T}} w + \frac{1}{2} \sum_{i=1}^{N} (e_i - c_i)^2$$
$$\text{s.t. } w^{\mathrm{T}}\varphi(x_i) + b + e_i = y_i, \quad \forall i = 1, \ldots, N, \tag{9}$$

where the elements of the vector c serve as tuning parameters, called additive regularization constants. Note that at this level, the tuning parameters c_i are fixed. Motivations for this approach are seen in the corresponding problem of model selection, while additional interpretations of the additive scheme were given in [16] as the constants c_i are also directly influencing the distribution of the residuals.

After constructing the Lagrangian with multipliers α and taking the conditions for optimality w.r.t. w, b, e_i, α_i (being $e_i = c_i + \alpha_i$, $w = \sum_{i=1}^{N} \alpha_i \varphi(x_i)$, $\sum_{i=1}^{N} \alpha_i = 0$ and $w^{\mathrm{T}}\varphi(x_i) + b + e_i = y_i$), the following dual linear system is obtained:

$$\left[\begin{array}{c|c} 0 & 1_N^{\mathrm{T}} \\ \hline 1_N & \Omega + I_N \end{array}\right] \left[\begin{array}{c} b \\ \hline \alpha \end{array}\right] + \left[\begin{array}{c} 0 \\ \hline c \end{array}\right] = \left[\begin{array}{c} 0 \\ \hline y \end{array}\right], \tag{10}$$

and $\alpha + c = e$. Note that at this point, the value of c is not considered as an unknown to the optimization problem: once c is fixed, the solution α, b to (10) is uniquely determined. The determination of c is postponed to higher levels. Note at this stage that the extension of this derivation to the case of componentwise models follows straightforwardly.

The estimated function $\hat{f}$ can be evaluated at a new point x^* by $\hat{f}(x^*) = \hat{w}^{\mathrm{T}}\varphi(x^*) + \hat{b} = \sum_{i=1}^{N} \hat{\alpha}_i K(x_i, x^*) + \hat{b}$ where $\hat{\alpha}$ and $\hat{b}$ are the unique solution to (10). The residuals $\hat{f}(x_j^v) - y_j^v$ from the evaluation on a validation dataset are denoted as e_j^v such that one can write

$$y_j^v = \hat{w}^{\mathrm{T}}\varphi(x_j^v) + \hat{b} + e_j^v = \sum_{i=1}^{N} \hat{\alpha}_i K(x_i, x_j^v) + \hat{b} + e_j^v, \quad \forall j = 1, \ldots, n. \tag{11}$$

By comparison of (2) and (10), LS-SVMs with Tikhonov regularization can be seen as a special case of LS-SVM substrates with the following additional constraint on the vectors α, c, γ

$$\gamma^{-1}\alpha = \alpha + c \text{ s.t. } \gamma > 0. \tag{12}$$

This means that the solutions to LS-SVM substrates are also solutions to LS-SVMs whenever the support values α_i are proportional to the residuals $e_i = \alpha_i + c_i$ for all $i = 1, \ldots, N$. This type of collinearity or rank constraints is known to be very hard to cast as convex optimization problems [1]. Finally, note that the additive regularization trade-off scheme does not replace the Tikhonov regularization scheme, but parameterizes the trade-off in a different way. This allows for inference of primal–dual kernel machines based on alternative criteria [14] as will be shown in the following sections. The aim of this paper is to make use of this mechanism in order to obtain sparse models and perform structure detection.

3.2. Level 2: sparse hierarchical kernel machines

Sparseness is often regarded as good practice in the machine learning community [29,30] as it gives an optimal solution with a minimal representation (from the viewpoint of VC theory and compression). The primal–dual framework also provides another line of thought based on sensitivity analysis as explained in [1]. Consider the derivation of the LS-SVM (2) or the LS-SVM substrate (10). The optimal Lagrange multipliers $\hat{\alpha}$ contain information of how much the (dual) optimal solution changes when the corresponding constraints are perturbed. This perturbation is proportional to c as can be seen from the constraints in (9). As such, one can write

$$\hat{\alpha}_i = -\frac{\partial p^*}{\partial c_i} \quad \text{and} \quad p^* = \inf_{w,b} \mathscr{J}_c \text{ s.t. (9) holds}, \tag{13}$$

as explained e.g. in [1]. In this respect, one can design a kernel machine that minimizes its own sensitivity to model misspecifications or atypical data observations by minimizing an appropriate norm on the Lagrange multipliers.

The 1-norm is considered

$$\min_{e,\alpha,b,c} \|e\|_2^2 + \zeta\|\alpha\|_1 \text{ s.t. (10) holds and } \alpha + c = e, \tag{14}$$

where $0<\zeta \in \mathbb{R}$ acts as a hyper-parameter. This criterion leads to sparseness [29] as already exploited under the name of v-SVM [3]. A similar principle is applied in the estimation of sparse parametric models known as basis pursuit [5] or LASSO [11], but these approaches lack the above interpretation in terms of sensitivity. The method in this paper is also treated at different hierarchical levels.

3.3. Level 2: hierarchical kernel machines for structure detection

The formulation of componentwise LS-SVMs suggests the use of a dedicated regularization scheme which is often very useful in practice. In the case where the nonlinear function consists of a sum of components, one may ask oneself which components have no contribution ($f_p(\cdot) = 0$) for prediction. Sparseness amongst the components is often referred to as structure detection. The described method is closely related to the kernel ANOVA decomposition [23] and the structure detection method of Gunn and Kandola [9], however, this paper starts from a clear optimality principle.

Previous results [15] were obtained using a 1-norm of the contribution (i.e. $\sum_{i=1}^N |f_p(x_i^{(p)})|$) which is somewhat similar to the measure of total variation [20]. This paper adopts a measure of maximal variation of components, defined as

$$t_p = \max_i |w_p^{\mathrm{T}} \varphi_p(x_i^{(p)})|, \quad \forall p = 1, \ldots, P. \tag{15}$$

Indeed, if t_p were zero, the corresponding component would not be able to contribute to the final model on the training data and $f_p(\cdot) = 0$. By using this measure as a regularization term, optimization problems are obtained which require much less (primal) variables and as such can handle much higher dimensions than the method proposed in e.g. [17]. The kernel machine for structure detection minimizes the following criterion for a given tuning constant $0<\rho \in \mathbb{R}$:

$$\begin{aligned} &\min_{e,t_p,\alpha,b,c} \|e\|_2^2 + \rho \sum_{p=1}^{P} t_p \text{ s.t. (10) holds} \\ &\quad \text{and } -t_p 1_N \leqslant \Omega^{(p)}\alpha \leqslant 1_N t_p, \quad \forall p = 1, \ldots, P, \end{aligned} \tag{16}$$

where the contribution of the individual components are described as $w_p^{\mathrm{T}} \varphi_p(x_i^{(p)}) = \Omega_i^{(p)}\alpha$. It is known that the use of 1-norms may lead to a sparse solution which is unnecessarily biased [7]. To overcome this drawback, one has proposed the use of norms such as the smoothly clipped absolute deviation (SCAD) penalty function as suggested by Fan [7] and which have been implemented in a kernel machine in [15]. This paper will not pursue this issue as it leads to non-convex optimization criteria in general. Instead, the use of the 1-norm is studied in order to detect structure, while the final predictions can be made based on a standard model using only the selected components (compare to basis pursuit, see e.g. [5]). In this paper, one works with a

standard componentwise LS-SVM as described in Section 2.2 based on the selected structure.

4. Fusion of validation with previous levels

The automatic tuning of hyper-parameters ζ or ρ with respect to the model selection criterion is highly desirable, especially in practice. The interplay between exact training and optimizing the regularization trade-off with respect to a validation criterion is studied. While fusion of a standard LS-SVM is often non-convex [16], the additive regularization trade-off scheme circumvents this problem. Fusion of the validation part and the LS-SVM substrate was considered in [16]

$$\min_{\alpha,b,c,e,e^v} \|e\|_2^2 + \|e^v\|_2^2 \text{ s.t. (10) holds and } e = \alpha + c. \tag{17}$$

This criterion is motivated by the assumption that the training as well as the validation data are independently sampled from the same distribution. In order to confine the effective degrees of freedom of the space of possible c values, in [17] the use of the following modified cost-function was proposed in order to obtain sparse solutions such as

$$\min_{e,e^v,\alpha,c,b} \|e\|_2^2 + \|e^v\|_2^2 + \xi\|\alpha\|_1 \text{ s.t. (10) holds and } e = \alpha + c. \tag{18}$$

While this problem is well defined, one may object that a new hyper-parameter $0 < \xi \in \mathbb{R}$ has popped up. Pelckmans et al. [17] considered the problem at two levels. In this paper, a three-level approach is introduced.

The following section extends these results by proposing a way to tune the hyper-parameter ρ in (16) and ζ in (14) of the second level with respect to a validation criterion using a third level of inference. This three-level architecture constitutes the hierarchical kernel machine. Fig. 1 gives a schematical representation of the conceptual idea of hierarchical kernel machines. The LS-SVM substrate constitutes the first level, while the sparse LS-SVM and the LS-SVM for structure detection make up the second level. The validation performance is used to tune the hyper-parameters ζ (or ρ) on a third level. For notational convenience, the bias term b is omitted from the derivations in the sequel.

4.1. Level 3: fusion of sparse LS-SVMs with validation

Consider the conceptual three-level hierarchical kernel machine for sparse LS-SVMs (see Fig. 1). Fig. 2 outlines the derivations of the hierarchical kernel machine for sparse representations and emphasizes the three-level hierarchical structure.

Consider the level 2 cost-function (14) where ζ acts as a hyper-parameter. Necessary and sufficient conditions for the global optimum of the cost-function (14)

are given by the KKT conditions. The Lagrangian becomes

$$\mathscr{L}_\zeta(\alpha, a; \xi^+, \xi^-) = \frac{1}{2}\left\|\Omega\alpha - y\right\|_2^2 + \rho\sum_{i=1}^{N} a_i + \xi^{-\mathrm{T}}(-a-\alpha) + \xi^{+\mathrm{T}}(-a+\alpha) \tag{19}$$

with positive multipliers $\xi^+, \xi^- \in \mathbb{R}^N$. The corresponding KKT conditions are necessary and sufficient as the primal problem is convex [1, p. 244], for the determination of the global optimum:

$$\begin{cases} \Omega^{\mathrm{T}}\Omega\alpha - y^{\mathrm{T}}\Omega = (\xi^- - \xi^+), & \text{(20a)}\\ \zeta = 1_N^{\mathrm{T}}(\xi^- + \xi^+), & \text{(20b)}\\ \xi^+, \xi^- \geqslant 0, & \text{(20c)}\\ -a \leqslant \alpha \leqslant a, & \text{(20d)}\\ \xi_i^-(a_i + \alpha_i) = 0, \quad \forall i = 1, \ldots, N, & \text{(20e)}\\ \xi_i^+(a_i - \alpha_i) = 0, \quad \forall i = 1, \ldots, N & \text{(20f)} \end{cases}$$

which are all linear (in)equalities except for the so-called complementary slackness conditions (20e) and (20f). Now consider the fusion of the validation criterion and these conditions with respect to the hyper-parameter ζ:

$$\min_{\zeta, a, \alpha, \xi^-, \xi^+} \mathscr{J}^v = \tfrac{1}{2}\|\Omega^v\alpha - y^v\|_2^2 \text{ s.t. (20) holds}, \tag{21}$$

Except for conditions (20e) and (20f), problem (21) can be solved as a quadratic programming (QP) problem. We propose the use of the modified QP below instead, which results in the same optimal solution as (21):

$$\min_{\zeta, a, \alpha, \xi^-, \xi^+} \mathscr{J}^v = \|\Omega^v\alpha - y^v\|_2^2 + b^-\xi^{-\mathrm{T}}(a+\alpha) + b^+\xi^{+\mathrm{T}}(a-\alpha) \text{ s.t. conditions (20a)}-\text{(20d) hold}, \tag{22}$$

by proper selection of the weighting terms $b^+ > 0$ and $b^- > 0$ such that the Hessian of the QP remains positive semi-definite and the complementary slackness conditions (20e) and (20f) are satisfied.

4.2. Level 3: fusion of structure detection with validation

A third level is added to the LS-SVM for structure detection in order to tune the hyper-parameter ζ of the second level. Fig. 3 summarizes the derivation below and points out the hierarchical approach.

Let $t = (t_1, \ldots, t_P)^{\mathrm{T}} \in \mathbb{R}^P$ be a vector of bounds on the maximal variation per component. Consider optimization problem (16) where ρ acts as a hyper-parameter.

One can eliminate e and c from this optimization problem leading to

$$\min_{t,\alpha} \mathscr{J}_\rho(\alpha,t) = \frac{1}{2}\|\Omega^P\alpha - y\|_2^2 + \rho\sum_{p=1}^{P} t_p$$
$$\text{s.t. } -t_p 1_N \leqslant \Omega^{(p)}\alpha \leqslant t_p 1_N, \quad \forall p = 1,\ldots,P. \tag{23}$$

The Lagrangian becomes

$$\mathscr{L}_\rho(\alpha,t;\xi^{+p},\xi^{-p}) = \frac{1}{2}\|\Omega^P\alpha - y\|_2^2 + \rho\sum_{p=1}^{P} t_p$$
$$+\sum_{p=1}^{P}\xi^{-p\mathrm{T}}(-t_p 1_N - \Omega^{(p)}\alpha) + \sum_{p=1}^{P}\xi^{+p\mathrm{T}}(-t_p 1_N + \Omega^{(p)}\alpha) \tag{24}$$

with multipliers ξ^{+p} and $\xi^{-p} \in \mathbb{R}^{+,N}$ for all $p = 1,\ldots,P$. The corresponding KKT conditions are necessary and sufficient as the primal problem is convex [1, p. 244], for the determination of the global optimum:

$$\begin{cases} \Omega^{P\mathrm{T}}\Omega^P\alpha - y^{\mathrm{T}}\Omega^P = \sum_{p=1}^{r}(\xi^{-p} - \xi^{+p}), & \text{(25a)} \\ \rho = 1_N^{\mathrm{T}}(\xi^{-p} + \xi^{+p}), \quad \forall p = 1,\ldots,P, & \text{(25b)} \\ \xi^{+p},\xi^{-p} \geqslant 0, \quad \forall p = 1,\ldots,P, & \text{(25c)} \\ -t_p 1_N \leqslant \Omega^{(p)}\alpha \leqslant t_p 1_N, \quad \forall p = 1,\ldots,P, & \text{(25d)} \\ \xi_i^{-p}(t_p + \Omega_i^{(p)}\alpha) = 0, \quad \forall i = 1,...,N, \quad \forall p = 1,\ldots,P, & \text{(25e)} \\ \xi_i^{+p}(t_p - \Omega_i^{(p)}\alpha) = 0, \quad \forall i = 1,\ldots,N, \quad \forall p = 1,\ldots,P & \text{(25f)} \end{cases}$$

which are all linear (in)equalities except for the so-called complementary slackness conditions (25e) and (25f). Now consider the fusion of the validation criterion and these conditions with respect to the hyper-parameter ρ:

$$\min_{\rho,t,\alpha,\xi^-,\xi^+} \mathscr{J}^v = \tfrac{1}{2}\|\Omega^{P,v}\alpha - y^v\|_2^2 \text{ s.t. (25) holds}, \tag{26}$$

where $\Omega^{P,v} \in \mathbb{R}^{n\times N} = \sum_{p=1}^{P}\Omega^{(p),v}$ and $\Omega_{ij}^{(p),v} = K^p(x_i^{(p)}, x_j^{(p),v})$ for all $i = 1,...,n$ and $j = 1,...,N$. Except for conditions (25e) and (25f), the problem (26) can be solved as a QP problem. We propose the use of the modified QP below instead, which results in the same optimal solution as (26):

$$\min_{\rho,t,\alpha,\xi^-,\xi^+} \mathscr{J}^v = \|\Omega^{P,v}\alpha - y^v\|_2^2 + \sum_{p=1}^{P} b_p^-\,\xi^{-p\mathrm{T}}(t_p 1_N + \Omega_i^{(p)}\alpha) + \sum_{p=1}^{P} b_p^+\,\xi^{+p\mathrm{T}}$$
$$\times (t_p 1_N - \Omega_i^{(p)}\alpha) \quad \text{s.t. conditions (25a)}-\text{(25d) hold.} \tag{27}$$

Crucial in this formulation is the observation that the complementary slackness terms $\xi^{-p\mathrm{T}}(t_p 1_N + \Omega_i^{(p)}\alpha)$ and $\xi^{+p\mathrm{T}}(t_p 1_N - \Omega_i^{(p)}\alpha)$ are bounded from below by zero as all cross-product terms are positive. Indeed, if the minimum value 0 of the complementary slackness is attained in the global optimum, the solution to (27) would equal the solution to (26). For a proper choice of the weighting terms b_i, this will be attained. In order to keep the problem convex, the values of b_i should be chosen such that the Hessian of the quadratical programming problem remains positive definite. One can tune these weighting terms by checking the complementary slackness conditions (25e) and (25f) on the resulting solution. Similar as is the case for componentwise LS-SVM substrates, the model can be evaluated at new data points $x_* \in \mathbb{R}^d$ as

$$\hat{f}(x^*) = \hat{w}^{\mathrm{T}}\varphi(x^*) = \sum_{i=1}^{N} \hat{\alpha}_i \sum_{t_p \neq 0} K^p(x_i^{(p)}, x_*^{(p)}), \tag{28}$$

where $\hat{\alpha}$ is the solution to (27).

5. Experiments

5.1. Sparseness

The performance of the proposed sparse LS-SVM substrate was measured on a number of regression and classification datasets, respectively, an artificial dataset sinc (generated as $Y = \mathrm{sinc}(X) + e$ with $e \sim \mathcal{N}(0, 0.1)$ and $N = 100$, $d = 1$) and the motorcycle dataset [6] ($N = 100$, $d = 1$) for regression (see Fig. 4), the artificial Ripley dataset ($N = 250$, $d = 2$) (see Fig. 5) and the PIMA dataset ($N = 468$, $d = 8$) from UCI at classification problems. The models resulting from sparse LS-SVM substrates were tested against the standard SVMs and LS-SVMs where the kernel parameters and the other tuning-parameters (respectively C, ε for the SVM, γ for the LS-SVM and ξ for sparse LS-SVM substrates) were obtained from 10-fold cross-validation (see Table 1).

5.2. Structure detection

An artificial example is taken from [29] and the Boston housing dataset from the UCI benchmark repository was used for analyzing the practical relevance of the structure detection mechanism. This subsection considers the formulation from Section 3.3, where sparseness amongst the components is obtained by use of the sum of maximal variation. The performance on a validation set was used to tune the parameter ρ both via a naive line-search as well as using the method which is described in Section 4.2.

Figs. 6 and 7 shows results obtained on an artificial dataset consisting of 100 samples and dimension 25, uniformly sampled from the interval $[0, 1]^{25}$. The

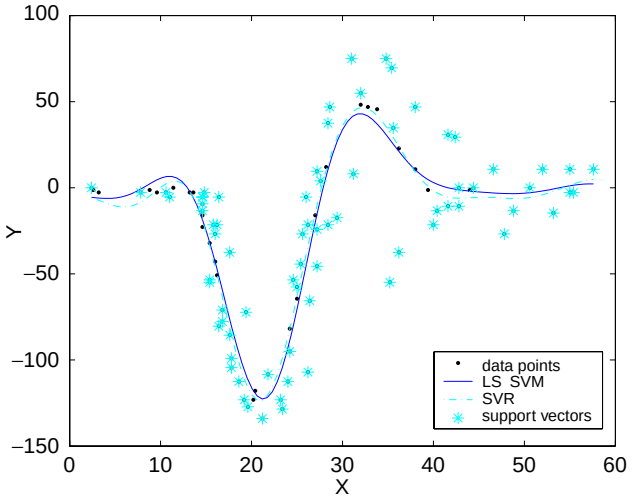

(a) Motorcycle:SVM

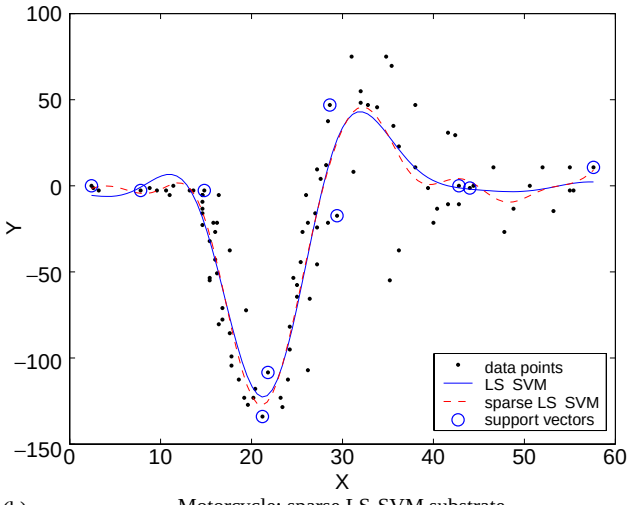

(b) Motorcycle: sparse LS-SVM substrate

Fig. 4. Comparison of the SVM, LS-SVM and sparse LS-SVM substrate of Section 3.2 on the Motorcycle regression dataset. One sees the difference in selected support vectors of (a) a standard SVM and (b) a sparse hierarchical kernel machine.

underlying function takes the following form:

$$f(x) = 10\sin(X^1) + 20(X^2 - 0.5)^2 + 10X^3 + 5X^4, \tag{29}$$

such that $y_i = f(x_i) + e_i$ with $e_i \sim \mathcal{N}(0,1)$ for all $i = 1, \ldots, 100$. Fig. 7 gives the nontrivial components ($t_p > 0$) associated with the LS-SVM substrate with ρ optimized in validation sense. Fig. 6 presents the evolution of values of t when ρ

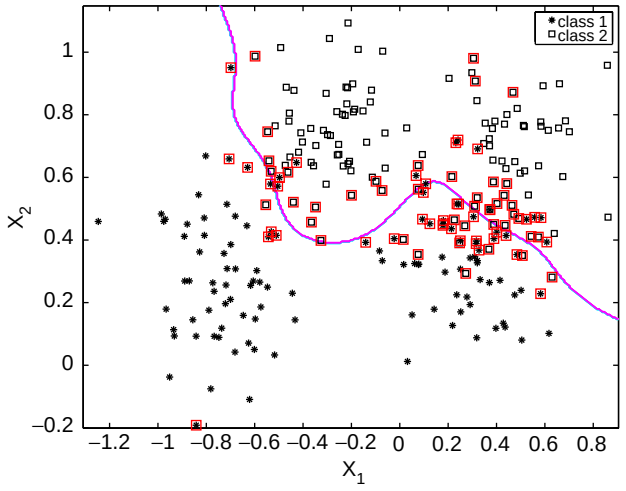

(a) Ripley dataset: SVM

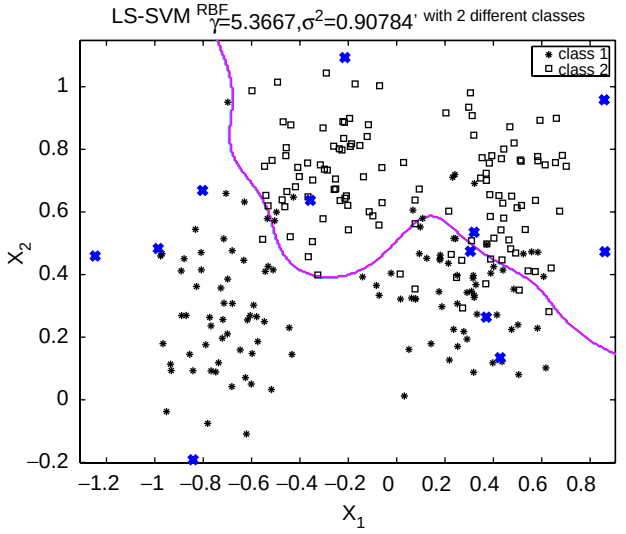

(b) Ripley dataset: sparse LS-SVM substrate

Fig. 5. Comparison of the SVM, LS-SVM and sparse LS-SVM substrate of Section 3.2 on the Ripley classification dataset. One can see the difference in selected support vectors of (a) a standard SVM and (b) a sparse hierarchical kernel machine. The support vectors of the former concentrate around the margin while the sparse hierarchical kernel machine will provide a more global support.

is increased from 1 to 1000 in a maximal variation evolution diagram (similarly as used for LASSO [11]).

The Boston housing dataset was taken from the UCI benchmark repository. This dataset concerns the housing values in suburbs of Boston. The dependent continuous

Table 1
Performances of SVMs, LS-SVMs and the sparse LS-SVM substrates of Section 3.2 expressed in mean squared error (MSE) on a test set in the case of regression or percentage correctly classified (PCC) in the case of classification

	SVM		LS-SVM	Sparse LS-SVM substr.	
	MSE	Sparse	MSE	MSE	Sparse
Sinc	0.0052	68%	0.0045	0.0034	9%
Motorcycle	516.41	83%	444.64	469.93	11%
	PCC	Sparse	PCC	PCC	Sparse
Ripley	90.10%	33.60%	90.40%	90.50%	4.80%
Pima	73.33%	43%	72.33%	74%	9%

Sparseness is expressed in percentage of support vectors w.r.t. number of training data. The kernel machines were tuned for the kernel parameter and the respective hyper-parameters C, ε; γ and ζ with 10-fold cross-validation. These results indicate that sparse LS-SVM substrates are at least comparable in generalization performance with existing methods, but are often more effective in achieving sparseness.

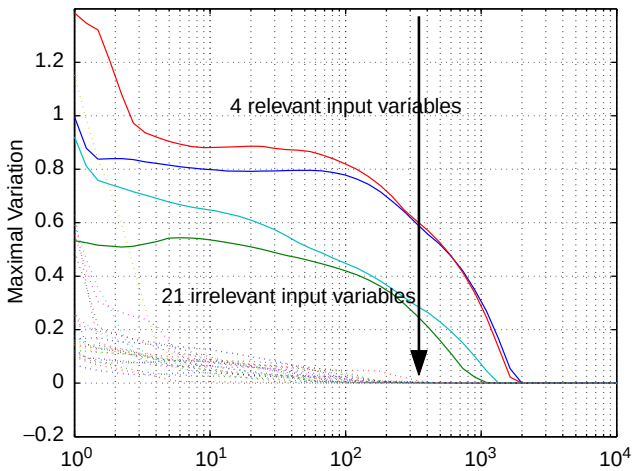

Fig. 6. Results of structure detection on an artificial dataset as used in [29], consisting of 100 data samples generated by four componentwise non-zero functions of the first 4 inputs and 21 irrelevant inputs and perturbed by i.i.d. unit variance Gaussian noise. This diagram shows the evolution of the maximal variations per component when increasing the hyper-parameter ρ from 1 to 10,000. The black arrow indicates a value ρ corresponding with a minimal cross-validation performance. Note that for the corresponding value of ρ, the underlying structure is indeed detected successfully.

variable expresses the median value of owner-occupied homes. From 13 given inputs, an additive model was built using the mechanism of maximal variation for detection of which input variables have a non-trivial contribution. Two hundred and fifty data

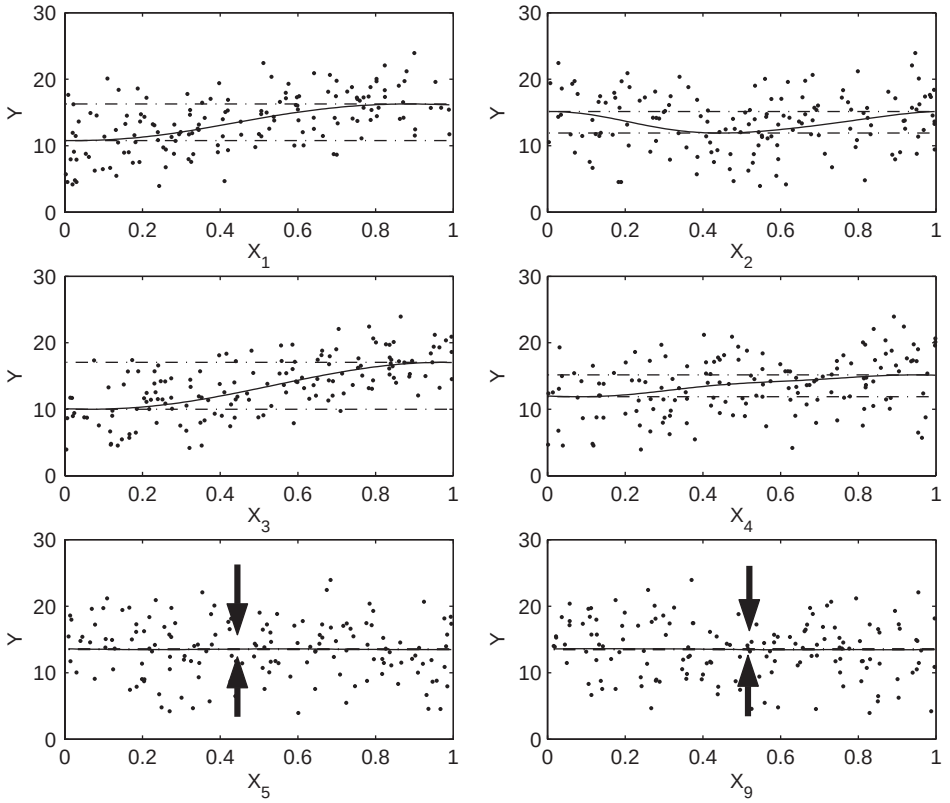

Fig. 7. Results of structure detection on an artificial dataset as used in [29]. The estimated componentwise LS-SVM (7) with $\rho = 300$ as tuned by cross-validation. All inputs except the four first are irrelevant to the problem. The structure detection algorithm detects the irrelevant inputs by measuring a maximal variation of zero (as indicated by the black arrows).

points were used for training purposes and 100 were randomly selected for validation. The analysis works with standardized data (zero mean and unit variance), while results are expressed in the original scale. The structure detection algorithm as proposed in Section 3.3 was used to construct the maximal variation evolution diagram (see Fig. 8). Fig. 9 displays the contributions of the individual components. The performance on the validation dataset was used to tune the kernel parameter and ρ. The latter was determined both manually (by a line-search) as automatically by fusion as described in Section 4.2. For the optimal parameter ρ, the following inputs have a maximal variation of zero:

1 CRIM: per capita crime rate by town,
2 ZN: proportion of residential land zoned for much over 25,000 sq.ft.,
4 CHAS: Charles River dummy variable (= 1 if tract bounds river; 0 otherwise),

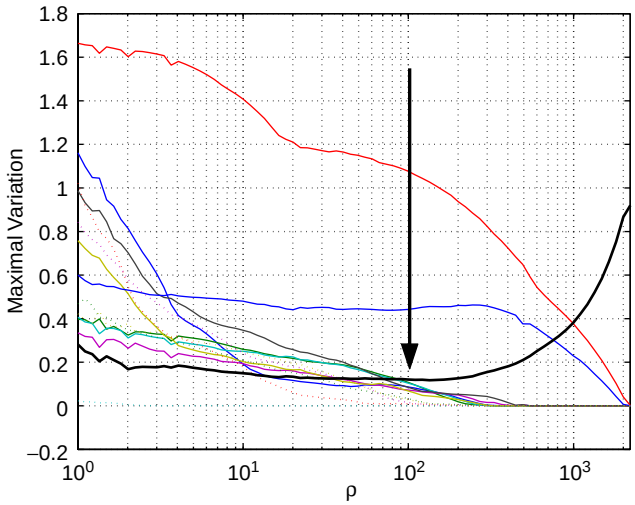

Fig. 8. Results of structure detection on the Boston housing dataset consisting of 250 training, 100 validation and 156 randomly selected testing samples. The evolution of the maximal variation of the variables when increasing ρ. The arrow indicates the choice of ρ made by the fusion argument minimizing the validation performance (solid line).

10 TAX: full-value property-tax rate per 10,000,
12 B: $1000(Bk - 0.63)^2$ where Bk is the proportion of blacks.

Testing was done by retraining a componentwise LS-SVM based on only the selected inputs. The resulting additive model increases in performance expressed in MSE on an independent test set with 22%. The improvement is even more significant (32%) with respect to a standard nonlinear LS-SVM model with an RBF-kernel.

6. Conclusions

This paper discussed an hierarchical method to build kernel machines on LS-SVM substrates resulting in sparseness and/or structure detection. The hierarchical modeling strategy is enabled by the use of additive regularization and exploiting the necessary and sufficient KKT conditions, while interactions between the levels are guided by a proper set of hyper-parameters. Higher levels are based on the one hand, on L_1 regularization and a measure of maximal variation, and on the other hand, on maximization of the validation performance. While the resulting hierarchical kernel machine has separated conceptual levels, the machine can be fused into a single convex optimization problem resulting in the training solution and the hyper-

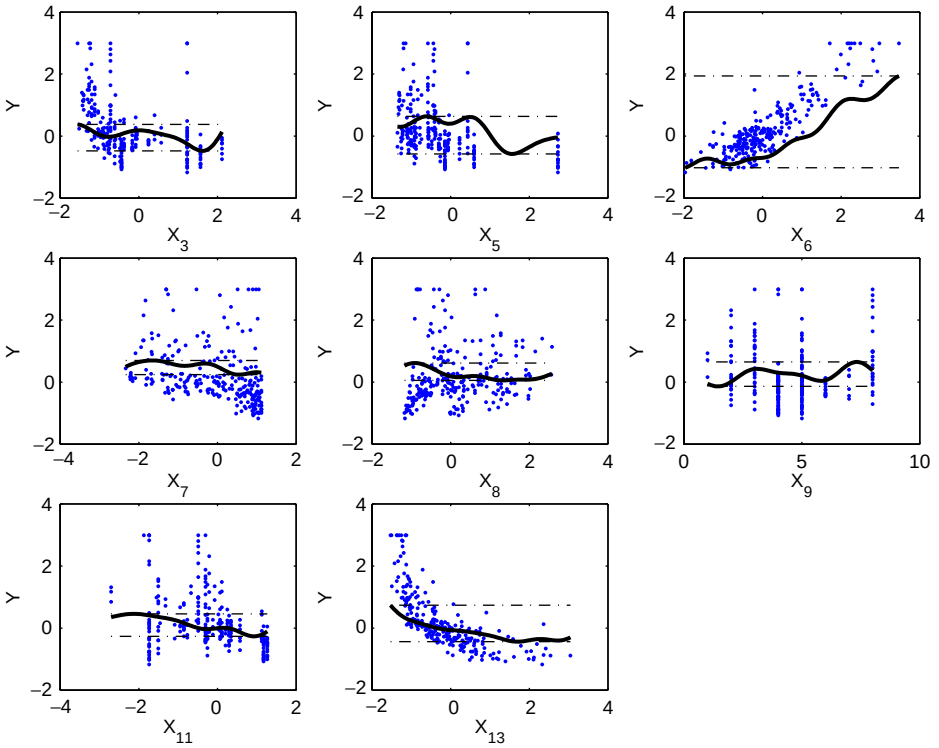

Fig. 9. Results of structure detection on the Boston housing dataset consisting of 250 training, 100 validation and 156 randomly selected testing samples. The contributions of the variables which have a non-zero maximal variation are shown. The fusion argument as described in Section 4.2 was used to tune the parameter ρ.

parameters at once. A number of experiments illustrate the use of the proposed method both with respect to interpretability as well as generalization performance.

Acknowledgements

Research is supported by Research Council KUL: GOA-Mefisto 666, GOA AMBioRICS, several Ph.D./postdoc. & fellow Grants; Flemish Government: FWO: Ph.D./postdoc. Grants, projects, G.0240.99 (multilinear algebra), G.0197.02 (power islands), G.0141.03 (identification and cryptography), G.0491.03 (control for intensive care glycemia), G.0120.03 (QIT), G.0452.04 (new quantum algorithms), G.0407.02 (support vector machines), G.0499.04 (robust statistics), G.0211.05 (nonlinear identification), G.0080.01 (collective behavior) AWI: Bil. Int. Collaboration Hungary/Poland; research communities (ICCoS, ANMMM, MLDM); IWT: Ph.D. Grants, GBOU (McKnow); Belgian Federal Science Policy Office:

IUAP P5/22 ('Dynamical Systems and Control: Computation, Identification and Modelling', 2002–2006), IUAP V; PODO-II (CP/40: TMS and Sustainability); EU: FP5-Quprodis; ERNSI; Eureka 2063-IMPACT; Eureka 2419-FliTE; Contract Research/agreements: ISMC/IPCOS, Data4s, TML, Elia, LMS, Mastercard.

References

[1] S. Boyd, L. Vandenberghe, Convex Optimization, Cambridge University Press, Cambridge, 2004.
[2] G.C. Cawley, N.L.C. Talbot, Efficient leave-one-out cross-validation of kernel Fisher discriminant classifiers, Pattern Recognition 36 (11) (2003) 2585–2592.
[3] C.C. Chang, C.J. Lin, Training nu-support vector regression: theory and algorithms, Neural Comput. 14 (8) (2002) 1959–1977.
[4] O. Chapelle, V. Vapnik, O. Bousquet, S. Mukherjee, Choosing multiple parameters for support vector machines, Mach. Learn. 46 (1–3) (2002) 131–159.
[5] S.S. Chen, D.L. Donoho, M.A. Saunders, Atomic decomposition by basis pursuit, SIAM Rev. 43 (1) (2001) 129–159.
[6] R.L. Eubank, Nonparametric Regression and Spline Smoothing, vol. 157, Marcel Dekker, New York, 1999.
[7] J. Fan, Comments on wavelets in statistics: a review, J. Italian Statist. Assoc. (6) (1997) 131–138.
[8] M.G. Genton, Classes of kernels for machine learning: a statistics perspective, J. Mach. Learn. Res. 2 (2001) 299–312.
[9] S.R. Gunn, J.S. Kandola, Structural modelling with sparse kernels, Mach. Learn. 48 (1) (2002) 137–163.
[10] T. Hastie, R. Tibshirani, Generalized Additive Models, Chapman & Hall, London, 1990.
[11] T. Hastie, R. Tibshirani, J. Friedman, The Elements of Statistical Learning, Springer, Heidelberg, 2001.
[12] A.E. Hoerl, R.W. Kennard, K.F. Baldwin, Ridge regression: some simulations, Comm. Statist. Part A-Theory Methods 4 (1975) 105–123.
[13] D.J.C. MacKay, The evidence framework applied to classification networks, Neural Comput. 4 (1992) 698–714.
[14] K. Pelckmans, M. Espinoza, J. De Brabanter, J.A.K. Suykens, B. De Moor, Primal–dual monotone kernel machines, Internal Report 04-108, ESAT-SISTA, K.U.Leuven, Leuven, Belgium, 2004, submitted for publication.
[15] K. Pelckmans, I. Goethals, J. De Brabanter, J.A.K. Suykens, B. De Moor, Componentwise least squares support vector machines, in: L. Wang (Ed.), Support Vector Machines: Theory and Applications, Springer, Berlin, 2004.
[16] K. Pelckmans, J.A.K. Suykens, B. De Moor, Additive regularization trade-off: fusion of training and validation levels in kernel methods, Internal Report 03-184, ESAT-SISTA, K.U.Leuven, Leuven, Belgium, 2003, submitted for publication.
[17] K. Pelckmans, J.A.K. Suykens, B. De Moor, Sparse LS-SVMs using additive regularization with a penalized validation criterion, in: Proceedings of the 12th European Symposium on Artificial Neural Networks, vol. 12, 2004, pp. 435–440.
[18] T. Poggio, F. Girosi, Networks for approximation and learning, Proc. IEEE 78 (9) (1990) 1481–1497.
[19] T. Poggio, V. Torre, C. Koch, Computational vision and regularization theory, Nature 317 (6035) (1985) 314–319.
[20] L. Rudin, S.J. Osher, E. Fatemi, Nonlinear total variation based noise removal algorithms, Physica D 60 (1992) 259–268.
[21] C. Saunders, A. Gammerman, V. Vovk, Ridge regression learning algorithm in dual variables, in: Proceedings of the 15th International Conference on Machine Learning (ICML'98), Morgan Kaufmann, Los Altos, CA, 1998, pp. 515–521.
[22] B. Schölkopf, A. Smola, Learning with Kernels, MIT Press, Cambridge, MA, 2002.

[23] M. Stitson, A. Gammerman, V. Vapnik, V. Vovk, C. Watkins, J. Weston, Advanced in Kernel Methods: Support Vector Learning, Chap. Support Vector Regression with ANOVA Decomposition Kernels, MIT Press, Cambridge, MA, 1999.
[24] M. Stone, Cross-validatory choice and assessment of statistical predictions, J. Roy. Statist. Soc. Ser. B (36) (1974) 111–147.
[25] J.A.K. Suykens, G. Horvath, S. Basu, C. Micchelli, J. Vandewalle (Eds.), Advances in Learning Theory: Methods, Models and Applications, vol. 190 of NATO Science Series III: Computer & Systems Sciences, IOS Press, Amsterdam, 2003.
[26] J.A.K. Suykens, J. Vandewalle, Least squares support vector machine classifiers, Neural Process. Lett. 9 (3) (1999) 293–300.
[27] J.A.K. Suykens, T. Van Gestel, J. De Brabanter, B. De Moor, J. Vandewalle, Least Squares Support Vector Machines, World Scientific, Singapore, 2002.
[28] A.N. Tikhonov, V.Y. Arsenin, Solution of Ill-Posed Problems, Winston, Washington, DC, 1977.
[29] V.N. Vapnik, Statistical Learning Theory, Wiley, New York, 1998.
[30] U. von Luxburg, O. Bousquet, B. Schölkopf, A compression approach to support vector model selection, J. Mach. Learn. Res. (5) (2004) 293–323.
[31] G. Wahba, Spline Models for Observational Data, SIAM, Philadelphia, PA, 1990.

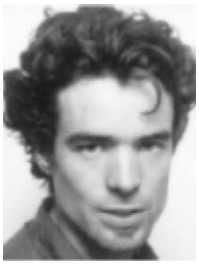

Kristiaan Pelckmans was born at 3 november 1978 in Merksplas, Belgium. He received a M.Sc. degree in Computer Science in 2000 from the Katholieke Universiteit Leuven. After a projectwork for an implementation of kernel machines and LS-SVMs (LS-SVMlab), he currently pursues a PhD at the KULeuven in the faculty of Applied Sciences, department of Electrical Engineering in the SCD/ SISTA laboratory. His research mainly focusses on machine learning and statistical inference using primal-dual kernel machines.



Johan A.K. Suykens was born in Willebroek Belgium, May 18 1966. He received the degree in Electro-Mechanical Engineering and the Ph.D. degree in Applied Sciences from the Katholieke Universiteit Leuven, in 1989 and 1995, respectively. In 1996 he has been a Visiting Postdoctoral Researcher at the University of California, Berkeley. He has been a Postdoctoral Researcher with the Fund for Scientific Research FWO Flanders and is currently an Associate Professor with K.U.Leuven. His research interests are mainly in the areas of the theory and application of neural networks and nonlinear systems. He is author of the books "Artificial Neural Networks for Modelling and Control of Non-linear Systems" (Kluwer Academic Publishers) and "Least Squares Support Vector Machines" (World Scientific) and editor of the books "Nonlinear Modeling: Advanced Black-Box Techniques" (Kluwer Academic Publishers) and "Advances in Learning Theory: Methods, Models and Applications" (IOS Press). In 1998 he organized an International Workshop on Nonlinear Modelling with Time-series Prediction Competition. He is a Senior IEEE member and has served as associate editor for the IEEE Transactions on Circuits and Systems- Part I (1997-1999) and Part II (since 2004) and since 1998 he is serving as associate editor for the IEEE Transactions on Neural Networks. He received an IEEE Signal Processing Society 1999 Best Paper (Senior) Award and several Best Paper Awards at International Conferences. He is a recipient of the International Neural Networks Society INNS 2000 Young Investigator Award for significant contributions in the field of neural networks. He has served as Director and Organizer of a NATO Advanced Study Institute on Learning Theory and Practice (Leuven 2002) and as a program co-chair for the International Joint Conference on Neural Networks IJCNN 2004.



Bart De Moor was born Tuesday July 12, 1960 in Halle, Belgium. He is married and has three children. In 1983, he obtained his Master (Engineering) Degree in Electrical Engineering at the Katholieke Universiteit Leuven, Belgium, and a PhD in Engineering at the same university in 1988. He spent 2 years as a Visiting Research Associate at Stanford University (1988–1990) at the departments of EE (ISL, Prof. Kailath) and CS (Prof.Golub). Currently, he is a full professor at the Department of Electrical Engineering (http://www.esat.kuleuven.ac.be) of the K.U.Leuven. His research interests are in numerical linear algebra and optimization, system theory and identification, quantum information theory, control theory, data-mining, information retrieval and bio-informatics, areas in which he has (co) authored several books and hundreds of research papers (consult the publication search engine at http://www.esat.kuleuven.ac.be/sista-cosic-docarch/template.php). Currently, he is leading a research group of 39 PhD students and 8 postdocs and in the recent past, 16 Ph.Ds were obtained under his guidance. He has been teaching at and been a member of PhD jury's in several universities in Europe and the US. He is also a member of several scientific and professional organizations. His work has won him several scientific awards (Leybold-Heraeus Prize (1986), Leslie Fox Prize (1989), Guillemin-Cauer best paper Award of the IEEE Transaction on Circuits and Systems (1990), Laureate of the Belgian Royal Academy of Sciences (1992), bi-annual Siemens Award (1994), best paper award of Automatica (IFAC, 1996), IEEE Signal Processing Society Best Paper Award (1999). Since 2004 he is a fellow of the IEEE (www.ieee.org). He is an associate editor of several scientific journals. From 1991–1999 he was the chief advisor on Science and Technology of several ministers of the Belgian Federal Government (Demeester, Martens) and the Flanders Regional Governments(Demeester, Van den Brande). He was and/or is in the board of 3 spin-off companies (www.ipcos.be, www.data4s.com, www.tml.be), of the Flemish Interuniversity Institute for Biotechnology (www.vib.be), the Study Center for Nuclear Energy (www.sck.be) and several other scientific and cultural organizations. He was a member of the Academic Council of the Katholieke Universiteit Leuven, and still is a member of its Research Policy Council.