

Training SVMs without Offset

Ingo Steinwart*, Don Hush, and Clint Scovel
MS B256
Information Sciences Group, CCS-3
Los Alamos National Laboratory
Los Alamos, NM 87545
{ingo,dhush,jcs}@lanl.gov

May 26, 2009

Abstract

We develop, analyze, and test a training algorithm for support vector machine classifiers without offset. Key features of this algorithm are a new stopping criterion and a set of inexpensive working set selection strategies that need almost as few iterations than the optimal working set selection strategy. For these working set strategies, we establish convergence rates that coincide with the best known rates for SVMs with offset. We further conduct various experiments that investigate both the run time behavior and the performed iterations of the new training algorithm. It turns out, that the new algorithm needs significantly less iterations and run-time than standard training algorithms for SVMs with offset.

Keywords

Support Vector Machines, Decomposition Algorithms

1 Introduction

Historically, support vector machines (SVMs) were motivated by a geometrical illustration of their linear decision surface in the feature space. This illustration motivated the use of an offset b that moves the decision surface from the origin. However, in recent years it has become increasingly evident that this geometrical interpretation has serious flaws (see, e.g., [19] for some illustrations) when considering kernels that have a large feature space such as the Gaussian RBF kernels. In addition, the current approach (see, e.g., [20]) for investigating the generalization performance of SVMs does not suggest that the offset offers any improvement for such kernels. On the other hand, the SVM optimization problem with offset imposes more restrictions on solvers than the optimization problem without offset does. For example, the offset leads to an additional equality constraint in the dual optimization problem, which in turn makes it necessary to update at least two dual variables at each iteration of commonly used solvers such as sequential minimal optimization (SMO). In addition, such solvers can only update certain pairs of dual variables, namely the ones whose update still satisfy the equality constraint. Moreover, the offset makes it relatively expensive to calculate the duality gap, see [3, p. 110], which may serve as a stopping criterion for these solvers, and hence one usually considers upper bounds of this gap such as the one from the maximal violating pair algorithm, see e.g. [14].

Despite these issues, research on algorithmic solutions has, with a few exceptions (see, e.g., [7]), so far focused on SVM formulations with offset, see e.g. [12, 11, 13, 9, 16, 5, 17, 2, 8, 6, 15, 18] and the references therein. The goal of this work is to fill this gap by developing algorithms for SVMs without offset. As it turns out, these algorithms not only achieve a classification accuracy that is comparable to the one for SVMs with offset, but also run significantly faster. This improvement is made possible by a couple of new algorithmic ideas that are not straightforward to implement for SVMs with offset. The first idea is a new stopping criterion that is inspired by recent results on the statistical performance of SVMs, see [20, Chapter 7.4], and that is, roughly speaking, a clipped duality gap. The second idea is a new working set selection strategy. As mentioned above, SMO type approaches for SVMs without offset can, in principle, update a single dual variable at each iteration. However, our experiments show that this approach does not lead to sufficiently fast training algorithms, and hence we will describe in detail, how an SMO type approach for two dual variables works. Of course, such an approach requires a good working set selection strategy. To identify one, we describe and test various strategies that try to find a pair of dual variables whose update approximately maximizes the gain in the dual objective function. Basically all these strategies first identify *one* dual variable whose update maximizes the gain in the dual objective and then search for a second variable that matches well to the first variable. Clearly, the first search is $\mathcal{O}(n)$, where n is the number of samples, while the order for the second search will be between $\mathcal{O}(1)$ and $\mathcal{O}(n)$ depending on the particular strategy. Interestingly, we will see that certain *combinations* of $\mathcal{O}(1)$ strategies for finding the second variable need almost as few iterations as an $\mathcal{O}(n^2)$ search over *all* pairs. In particular, these combinations need essentially the same amount of iterations than some natural $\mathcal{O}(n)$ strategies for the second dual variable. Since each iteration for the latter are obviously more expensive, these combinations enjoy significantly shorter run times as will be seen in the experiments.

We further establish theoretical guarantees on the number of performed iterations for solvers using the stopping criterion above and these working set strategies. It turns out that these guarantees coincide with the best known guarantees for solvers with offset, which can be obtained by combining the analysis of so-called rate certifying algorithms, see [17, 8, 18], with a recent analysis of the duality gap, see [15]. Unlike the rate certifying algorithms for SVMs with offset, however, our algorithms not only possess these guarantees, but also run faster than typically implemented training algorithms, as our experimental section shows.

We also consider the possibility to initialize the solver with (transformed) previous solutions when working on a grid of hyper-parameters. Here it first turns out that the missing equality constraint gives us more freedom to transform these solutions. We describe and test several such transformations ranging from relatively simple to quite complex procedures. In the experiments, we then see that SVMs without offset profit more from simple warm start initializations than SVMs with offset. In addition, the more complex warm start strategies, which cannot be directly implemented for SVMs with offset, lead to further improvements. In particular, for datasets containing a few thousand samples, SVMs without offset profit about twice as much from a good warm start strategy than SVMs with offset do. As a result, our SVMs without offset are approximately 8 times faster than SVMs with offset on these datasets, if the hyper-parameters are determined by a cross-validation approach.

This work is organized as follows: Section 2 introduces an SMO type algorithm for SVMs without offset that performs one dual variable update per iteration. We further describe the new stopping criterion based on a clipped duality gap as well as several warm start strategies. Section 3 then generalizes this algorithm to handle two variables at each iter-

ation. In particular, we describe how to exactly solve the corresponding two dimensional optimization problem. Furthermore, we present several working set selection strategies. Section 4 contains our theoretical analysis, while the experiments can be found in Section 5. Finally, concluding remarks can be found in Section 6.

2 The Basic Algorithm: Optimizing One Coordinate

Throughout this paper we write $[t]_a^b := \max\{a, \min\{b, t\}\}$, $t \in \mathbb{R}$, $b > a$, for the clipping operation that clips a real number t whenever it is outside the interval $[a, b]$. In order to introduce support vector machines without offset term let us consider a training set $T = ((x_1, y_1), \dots, (x_n, y_n)) \in (\mathbb{R}^d \times \{-1, 1\})^n$ and a function $f : X \rightarrow \mathbb{R}$. Then the empirical hinge risk of f is defined by

$$\mathcal{R}_{L,T}(f) := \frac{1}{n} \sum_{i=1}^n w_i L(y_i, f(x_i)),$$

where L denotes the hinge loss $L(y, t) := \max\{0, 1 - yt\}$, and $w_i > 0$ is a weight associated to the sample (x_i, y_i) . For example, in ordinary binary classification we have $w_i = 1$ for all $i = 1, \dots, n$, whereas in weighted binary classification we have two real numbers $w_{\text{pos}} > 0$ and $w_{\text{neg}} > 0$ such that $w_i = w_{\text{pos}}$ if $y_i = 1$ and $w_i = w_{\text{neg}}$ if $y_i = -1$. In the following we will exclusively consider the case of weighted binary classification, which, of course, includes the case of ordinary binary classification. Now the SVM without offset solves the problem

$$f_{T,\lambda} \in \operatorname{argmin}_{f \in H} \lambda \|f\|_H^2 + \mathcal{R}_{L,T}(f), \quad (1)$$

where H is the reproducing kernel Hilbert space (RKHS) of a kernel k . In the following, we always assume that k is strictly positive definite, that is, the Gram matrix $(k(x_i, x_j))_{i,j=1}^n$ is strictly positive definite whenever the data points x_1, \dots, x_n are mutually distinct. From this it is easy to conclude that $k(x, x) > 0$ for all $x \in X$, and hence we may additionally assume that k is *normalized*, i.e., $k(x, x) = 1$ for all $x \in X$. Note that polynomial kernels do not satisfy this assumption while the Gaussian RBF kernel $k(x, x') := \exp(-\sigma^2 \|x - x'\|_2^2)$ with width parameter $\sigma > 0$ does. Furthermore, note that for strictly positive definite and normalized kernels we have $|k(x, x')| = 1$ if and only if $x = x'$. For the Gaussian kernel, this characterization is, of course, trivial.

In order to derive an algorithm that produces an approximate solution of (1) we first multiply the objective function in (1) by $\frac{1}{2\lambda}$ and introduce slack variables.

$$\begin{aligned} \operatorname{argmin}_{(f,\xi)} \quad P_C(f, \xi) &:= \frac{1}{2} \|f\|_H^2 + \sum_{i=1}^n C_i \xi_i \\ \text{s.t.} \quad \xi_i &\geq 0, & i = 1, \dots, n, \\ \xi_i &\geq 1 - y_i f(x_i), & i = 1, \dots, n, \end{aligned} \quad (2)$$

where $C_i := \frac{w_{\text{pos}}}{2\lambda n}$ if $y_i = 1$ and $C_i := \frac{w_{\text{neg}}}{2\lambda n}$ otherwise. Analogously to the offset case (see, e.g., [3, p. 107f]) one can then show that the dual of this problem is

$$\begin{aligned} \max_{\alpha \in \mathbb{R}^n} \quad W(\alpha) &:= \langle e, \alpha \rangle - \frac{1}{2} \langle \alpha, K\alpha \rangle \\ \text{s.t.} \quad 0 &\leq \alpha_i \leq C_i, & i = 1, \dots, n, \end{aligned} \quad (3)$$

where $e := (1, \dots, 1) \in \mathbb{R}^n$ and K is the $n \times n$ matrix with entries $K_{i,j} := y_i y_j k(x_i, x_j)$, $i, j = 1, \dots, n$. In addition, the Karush-Kuhn-Tucker (KKT) conditions are

$$\begin{aligned} (y_i f(x_i) + \xi_i - 1) \alpha_i &= 0, & i = 1, \dots, n, \\ (C_i - \alpha_i) \xi_i &= 0, & i = 1, \dots, n, \end{aligned}$$

and a solution $\alpha^* \in [0, C] := [0, C_1] \times \dots \times [0, C_n]$ of (3) yields a solution (f^*, ξ^*) of (2) by setting

$$f^* := \sum_{i=1}^n y_i \alpha_i^* k(x_i, \cdot)$$

and $\xi_i^* := \max\{0, 1 - y_i f^*(x_i)\}$, $i = 1, \dots, n$. Obviously, (3) is identical to the standard dual SVM problem besides the missing equality constraint $\langle y, \alpha \rangle = 0$. Now recall that this equality constraint makes it necessary to update at least two coordinate values at a time, while in (3) we can update *single* coordinates. Some ideas for such a single direction update will be recalled in the following subsection to provide the background for working sets of size two considered in the following section.

2.1 Working sets of size one

In order to recall (see [3, p. 131ff]) the one-dimensional update step we define

$$\nabla W_i(\alpha) := \frac{\partial W}{\partial \alpha_i}(\alpha) = 1 - \sum_{j=1}^n \alpha_j K_{i,j}.$$

Moreover, for an $\alpha = (\alpha_1, \dots, \alpha_n) \in \mathbb{R}^n$ and an index $i \in \{1, \dots, n\}$ we write $\alpha^{\setminus i} := \alpha - \alpha_i e_i$, where e_i denotes the i -th vector of the standard basis of \mathbb{R}^n , i.e. $\alpha^{\setminus i}$ equals α in all coordinates other than i where it equals zero. Now basic calculus together with $K_{i,i} = 1$ for the normalized kernels shows that

$$\tilde{\alpha}_i \mapsto W(\alpha^{\setminus i} + \tilde{\alpha}_i e_i) = \langle \alpha^{\setminus i}, e \rangle + \tilde{\alpha}_i - \frac{1}{2} \langle \alpha^{\setminus i}, K \alpha^{\setminus i} \rangle - \tilde{\alpha}_i \langle e_i, K \alpha^{\setminus i} \rangle - \frac{1}{2} \tilde{\alpha}_i^2$$

attains its *global* minimum over \mathbb{R} at

$$\alpha_i^* = 1 - \langle e_i, K \alpha^{\setminus i} \rangle = 1 - \sum_{j \neq i} \alpha_j K_{i,j} = \nabla W_i(\alpha) + \alpha_i.$$

Obviously, if $\alpha_i^* \in [0, C_i]$ the function $\alpha_i \mapsto W(\alpha^{\setminus i} + \alpha_i e_i)$ also attains its maximum over $[0, C_i]$ at α_i^* . On the other hand, if, e.g., $\alpha_i^* > C_i$ then a simple concavity argument shows that the function attains its maximum over $[0, C_i]$ at C_i . By this and an analogous consideration in the case $\alpha_i^* < 0$ we hence see that the function $\alpha_i \mapsto W(\alpha^{\setminus i} + \alpha_i e_i)$ attains its maximum over $[0, C_i]$ at

$$\alpha_i^{new} := [\nabla W_i(\alpha) + \alpha_i]_0^{C_i}. \quad (4)$$

The next question is in which coordinate i should we perform the update. A simple and straightforward approach (see e.g. [3, p. 132]) is to update for each coordinate $i = 1, \dots, n$ iteratively, while a more advanced idea is to choose KKT violators (see [7, Chapter 3]) for the update, i.e., indices that, for a specified $\epsilon > 0$, satisfy

$$\begin{aligned} & \alpha_i < C_i \quad \text{and} \quad \nabla W_i(\alpha) > \epsilon \\ \text{or} & \alpha_i > 0 \quad \text{and} \quad \nabla W_i(\alpha) < -\epsilon. \end{aligned} \quad (5)$$

Obviously, the extreme case of this approach is to look for the indices

$$\begin{aligned} i_{\text{up}}^* &\in \arg \max \{ \nabla W_i(\alpha) : \alpha_i < C_i \} \\ i_{\text{down}}^* &\in \arg \min \{ \nabla W_i(\alpha) : \alpha_i > 0 \}, \end{aligned}$$

and to pick the index of these two candidates whose gradient has the larger absolute value. Another idea, which is motivated by [8, 9, 17], is to choose the coordinate i^* whose update achieves the largest improvement for the objective dual value $W(\alpha)$. In other words, it performs the update in the direction

$$i^* \in \arg \max_{i=1, \dots, n} W(\alpha + \delta_i e_i) - W(\alpha), \quad (6)$$

where $\delta_i := \alpha_i^{\text{new}} - \alpha_i$ denotes the difference between the new and the old value of α_i . Using the following trivial lemma, it is easy to see that Procedure 1 solves (6).

Lemma 2.1 *For $\delta \in \mathbb{R}$ and $i = 1, \dots, n$ we have*

$$W(\alpha + \delta e_i) - W(\alpha) = \delta \cdot (\nabla W_i(\alpha) - \delta/2).$$

Proof: By the symmetry of K we find

$$\langle \alpha, K\alpha \rangle - \langle \alpha + \delta e_i, K(\alpha + \delta e_i) \rangle = -2\delta \langle \alpha, K e_i \rangle - \delta^2.$$

Combining this with $\langle e, \alpha + \delta e_i \rangle - \langle e, \alpha \rangle = \delta$ yields the assertion. ■

Procedure 1 Calculate $i^* \in \arg \max_{i=1, \dots, n} \delta_i \cdot (\nabla W_i(\alpha) - \delta_i/2)$

bestgain $\leftarrow -1$

for $i = 1$ to n **do**

$\alpha_i^* \leftarrow [\nabla W_i(\alpha) + \alpha_i]_0^{C_i}$

$\delta \leftarrow \alpha_i^* - \alpha_i$

gain $\leftarrow \delta \cdot (\nabla W_i(\alpha) - \delta/2)$

if *gain* > *bestgain* **then**

$\text{bestgain} \leftarrow \text{gain}$

$i^* \leftarrow i$

end if

end for

2.2 Stopping Criteria

Several stopping criteria for the SVM *with* offset have been proposed and a straightforward approach is to adapt one of these. For example, one can stop if both $\nabla W_{i_{\text{up}}^*}(\alpha) \leq \epsilon$ and $\nabla W_{i_{\text{down}}^*}(\alpha) \geq -\epsilon$, i.e. if the KKT conditions are satisfied up to some predefined $\epsilon > 0$. Another simple idea is to use the duality gap as a stopping criterion, see e.g. [3, p. 109 & 128]. For SVMs *without* offset this duality gap is of the form

$$\text{gap}(\alpha) := \langle \alpha, K\alpha \rangle - \langle e, \alpha \rangle + \sum_{i=1}^n C_i [\nabla W_i(\alpha)]_0^\infty \leq \epsilon, \quad (7)$$

where $\epsilon > 0$ does not necessarily have the same value as above.

In this work, however, we consider a little more involved stopping criterion that is based on recent results from the statistical analysis of SVMs in [21]. Namely, it was shown in [21] that an $f^* \in H$ satisfying

$$\lambda \|f^*\|_H^2 + \mathcal{R}_{L,T}([f^*]_{-1}^1) \leq \min_{f \in H} \lambda \|f\|_H^2 + \mathcal{R}_{L,T}(f) + \epsilon \quad (8)$$

for yet another pre-defined $\epsilon > 0$ satisfies the same oracle inequality up to 4ϵ as the true solution $f_{T,\lambda}$. Moreover, a more careful analysis of [21] shows that the factor 4 can be essentially removed, so that for say $\epsilon := 0.001$ the learning guarantees for the approximate solution f^* are at most 0.1% worse than those for the true solution $f_{T,\lambda}$. In order to develop a stopping criterion from this observation we denote the minimum of the objective function P_C in (2) by P_C^* . Moreover, for a dual point $\alpha \in [0, C]$ we define, as usual, a corresponding primal function by

$$f := \sum_{i=1}^n \alpha_i y_i k(x_i, \cdot)$$

and its corresponding slack variables by $\xi_i := \max\{0, 1 - y_i f(x_i)\}$, $i = 1, \dots, n$. Using $1 - y_i f(x_i) = \nabla W_i(\alpha)$ and $\|f\|_H^2 = \langle \alpha, K\alpha \rangle$ as well as $P_C^* \geq W(\alpha) = \langle e, \alpha \rangle - \langle \alpha, K\alpha \rangle / 2$ and

$$\max\{0, 1 - y[t]_{-1}^1\} = 1 - y[t]_{-1}^1 = [1 - yt]_0^2$$

for all $y = \pm 1$, $t \in \mathbb{R}$ we hence see that (8) is satisfied if

$$S(\alpha) := \langle \alpha, K\alpha \rangle - \langle e, \alpha \rangle + \sum_{i=1}^n C_i [\nabla W_i(\alpha)]_0^2 \leq \frac{\epsilon}{2\lambda}. \quad (9)$$

Note that the statistical analysis of [21] also suggests that the right hand side of (7) can be replaced by $\frac{\epsilon}{2\lambda}$, where ϵ has the same value as in (9). Consequently, the difference between these two stopping criteria is the fact that (9) considers *clipped* slack variables which may be substantially smaller than the unclipped slack variables used in (7). Moreover, unlike the duality gap stopping criterion for SVMs *with* offset, see [3, p. 109f], both (7) and (9) are directly computable since they do not require the offset.

To efficiently compute $S(\alpha)$ we first observe that the first two terms of the updated $S(\alpha + \delta e_i)$ can be easily computed from the first two terms of $S(\alpha)$. Indeed, if we write

$$\begin{aligned} T(\alpha) &:= \langle \alpha, K\alpha \rangle - \langle e, \alpha \rangle \\ E(\alpha) &:= \sum_{i=1}^n C_i [\nabla W_i(\alpha)]_0^2 \end{aligned}$$

then we have $S(\alpha) = T(\alpha) + E(\alpha)$, and the calculations in the proof of Lemma 2.1 immediately show

$$T(\alpha + \delta e_i) = T(\alpha) - \delta(2\nabla W_i(\alpha) - 1 - \delta).$$

Now it is easy to derive an $\mathcal{O}(n)$ procedure that updates $\nabla W(\alpha)$ and calculates $S(\alpha)$. Procedure 2 provides pseudocode for this task.

Now the basic idea of the 1D-SVM described in Algorithm 1 is to repeatedly look for the best direction i^* and update in this direction until the stopping criterion (9) is satisfied. A closer look at this algorithm shows that it contains one piece that has not been discussed so far, namely the initialization of the solver. This initialization will be considered in the following subsection.

Procedure 2 Update $\nabla W(\alpha)$ in direction i by δ and calculate $S(\alpha)$

$$T(\alpha) \leftarrow T(\alpha) - \delta(2\nabla W_i(\alpha) - 1 - \delta)$$
$$E(\alpha) \leftarrow 0$$
for $j = 1$ to n **do**
$$\nabla W_j(\alpha) \leftarrow \nabla W_j(\alpha) - \delta K_{i,j}$$
$$E(\alpha) \leftarrow E(\alpha) + C_i \cdot [\nabla W_i(\alpha)]_0^2$$
end for
$$S(\alpha) \leftarrow T(\alpha) + E(\alpha)$$

Algorithm 1 1D-SVM solver

initialize α , $\nabla W(\alpha)$, $T(\alpha)$, and $S(\alpha)$ by one of the Procedures from Subsection 2.3

while $S(\alpha) > \frac{\epsilon}{2\lambda}$ **do**
$$i^* \leftarrow \arg \max_{i=1,\dots,n} W(\alpha + \delta_i e_i) - W(\alpha)$$
$$\delta \leftarrow [\nabla W_{i^*}(\alpha) + \alpha_{i^*}]_0^C - \alpha_{i^*}$$
$$\alpha_{i^*} \leftarrow [\nabla W_{i^*}(\alpha) + \alpha_{i^*}]_0^C$$
use Procedure 2 to update $\nabla W(\alpha)$ in direction i^* by δ and calculate $S(\alpha)$ **end while**

2.3 Initialization

We also have to decide how to initialize α . Of course, there exists various approaches for this task, and in the following, we describe a few methods we have tested in this work.

10 & W0. Cold start with zeros

Obviously, the most simple initialization is the *cold start* $\alpha \leftarrow 0$. Procedure 3 provides the pseudocode for this approach, which in the following we call 10 or W0.

Procedure 3 Initialize by $\alpha_i \leftarrow 0$ and compute $\nabla W(\alpha)$, $S(\alpha)$, and $T(\alpha)$.

$$T(\alpha) \leftarrow 0$$
$$S(\alpha) \leftarrow 0$$
for $i = 1$ to n **do**
$$\alpha_i \leftarrow 0$$
$$\nabla W_i(\alpha) \leftarrow 1$$
$$S(\alpha) \leftarrow S(\alpha) + C_i$$
end for

11 & W1. Cold start with kernel rule

Another simple cold start is to initialize with $\alpha_i \leftarrow C_i$ for all $i = 1, \dots, n$. Procedure 4 provides the pseudocode for this approach. In the following, we call this approach 11 or W1.

Obviously, Procedure 3 is of order $\mathcal{O}(n)$, whereas Procedure 4 is of order $\mathcal{O}(n^2)$, and hence the latter seems to be prohibitive. On the other hand, Procedure 4 basically initializes with a classical kernel rule, see [4, Chapter 10], and hence its initial training error may be smaller than that of Procedure 3. This in turn might lead to a smaller initial stopping criterion value $S(\alpha)$ and hence to less iterations of the solver. Of course, there is much room for speculation, and hence we will investigate the efficiency of both approaches in the

Procedure 4 Initialize by $\alpha_i \leftarrow C_i$ and compute $\nabla W(\alpha)$, $S(\alpha)$, and $T(\alpha)$.

```

 $T(\alpha) \leftarrow 0$ 
 $E(\alpha) \leftarrow 0$ 
for  $i = 1$  to  $n$  do
   $\alpha_i \leftarrow C_i$ 
   $\nabla W_i(\alpha) \leftarrow 1$ 
  for  $j = 1$  to  $n$  do
     $\nabla W_i(\alpha) \leftarrow \nabla W_i(\alpha) - C_j \cdot K_{i,j}$ 
  end for
   $T(\alpha) \leftarrow T(\alpha) - C_i \cdot \nabla W_i(\alpha)$ 
   $E(\alpha) \leftarrow E(\alpha) + C_i \cdot [\nabla W_i(\alpha)]_0^2$ 
end for
 $S(\alpha) \leftarrow T(\alpha) + E(\alpha)$ 

```

experiments. However, it is worth noting that unlike Procedure 3, Procedure 4 cannot be directly implemented for SVMs *with* offsets. In addition, Procedure 4 requires the entire kernel matrix to be computed, and hence it may actually be prohibitive if this matrix does not fit into memory.

W2. Warm start by recycling old solution

Besides the cold starts mentioned above, there are also a couple of simple *warm starts* possible. In order to explain these let us recall that often the hyper-parameter λ is chosen by a search over a grid $\Lambda = \{\lambda_1, \dots, \lambda_m\}$ of candidate values. Let us assume that these values are ordered in the form $\lambda_1 > \dots > \lambda_m$, and that we train the SVM in the order $\lambda_1, \dots, \lambda_m$. Then the resulting n -dimensional vectors $C^{(1)}, \dots, C^{(m)}$ defined by

$$C_i^{(j)} := \begin{cases} \frac{w_{\text{pos}}}{2\lambda_j n} & \text{if } y_i = 1 \\ \frac{w_{\text{neg}}}{2\lambda_j n} & \text{if } y_i = -1 \end{cases}$$

have the property $C_i^{(j)} < C_i^{(j+1)}$ for all $j = 1, \dots, m-1$ and $i = 1, \dots, n$. For $C^{(1)}$ we can then initialize with one of the above cold starts. Now observe that for $j \geq 2$ the approximate solution α^* obtained by training with $C^{\text{old}} := C^{(j-1)}$ is feasible for $C^{\text{new}} := C^{(j)}$, i.e. $\alpha^* \in [0, C^{\text{new}}]$. Consequently, for $j \geq 2$ we can either initialize with a cold start, or with the warm start $\alpha \leftarrow \alpha^*$. Obviously, in this case we can also recycle $\nabla W(\alpha)$ and $T(\alpha)$. In addition, the ratio

$$\frac{C_i^{\text{new}}}{C_i^{\text{old}}} = \frac{\lambda_{j-1}}{\lambda_j}$$

is *independent* of i and hence this warm start can be very easily implemented as Procedure 5 shows.

Procedure 5 Initialize by $\alpha_i \leftarrow \alpha_i^*$ and compute $\nabla W(\alpha)$, $S(\alpha)$, and $T(\alpha)$.

```

 $S(\alpha) \leftarrow T(\alpha^*) + \frac{C_1^{\text{new}}}{C_1^{\text{old}}} \cdot (S(\alpha^*) - T(\alpha^*))$ 

```

W4. Warm start by partially expanding and partially recycling old solution

Apart from the simple warm start above there is another conceptionally simple warm start for expanding box constraints. Namely, if α^* denotes an approximate solution to C^{old} and $C^{\text{old}} < C^{\text{new}}$ this warm start initializes by $\alpha_i \leftarrow \alpha_i^*$ if $\alpha_i^* < C_i^{\text{old}}$ and by $\alpha_i \leftarrow C_i^{\text{new}}$ if $\alpha_i^* = C_i^{\text{old}}$. The idea behind this warm start is that *bounded* support vectors, i.e., vectors in

$$bSV := \{j : \alpha_j^* = C_j^{\text{old}}\}$$

may have the tendency to become larger if the box constraint is loosened, while *unbounded* support vectors, i.e., vectors in

$$uSV := \{j : 0 < \alpha_j^* < C_j^{\text{old}}\}$$

may not have this tendency.

The basic idea of an efficient implementation of this warm start method is to avoid calculating the gradient from scratch by recycling parts of the gradient from C^{old} . To be more precise, observe that, for fixed i , the sum $\sum_{j \in uSV} \alpha_j^* K_{i,j}$ remains unchanged by the described warm start, while $\sum_{j \in bSV} \alpha_j^* K_{i,j}$ is simply multiplied by $C_i^{\text{new}}/C_i^{\text{old}}$. Recall that the latter ratio is independent of i , and consequently we can update the gradients by either

Procedure 6 Initialize bounded SVs by $\alpha_i \leftarrow C_i^{\text{new}}$ while keeping the rest unchanged and compute $\nabla W(\alpha)$, $S(\alpha)$, and $T(\alpha)$.

```

 $T(\alpha) \leftarrow 0$ 
 $E(\alpha) \leftarrow 0$ 
for  $i = 1$  to  $n$  do
  if  $\alpha_i = C_i^{\text{old}}$  then
     $\alpha_i \leftarrow C_i^{\text{new}}$ 
  end if
end for
if  $2 \cdot \#uSV < \#bSV$  then
  for  $i = 1$  to  $n$  do
     $\nabla W_i(\alpha) \leftarrow \frac{C_1^{\text{new}}}{C_1^{\text{old}}} \cdot \nabla W_i(\alpha) + (1 - \frac{C_1^{\text{new}}}{C_1^{\text{old}}})(1 - \sum_{j \in uSV} \alpha_j K_{i,j})$ 
     $T(\alpha) \leftarrow T(\alpha) - \alpha_i \cdot \nabla W_i(\alpha)$ 
     $E(\alpha) \leftarrow E(\alpha) + C_i^{\text{new}} \cdot [\nabla W_i(\alpha)]_0^2$ 
  end for
else
  for  $i = 1$  to  $n$  do
     $\nabla W_i(\alpha) \leftarrow \nabla W_i(\alpha) + (C_i^{\text{old}} - C_i^{\text{new}}) \sum_{j \in bSV} K_{i,j}$ 
     $T(\alpha) \leftarrow T(\alpha) - \alpha_i \cdot \nabla W_i(\alpha)$ 
     $E(\alpha) \leftarrow E(\alpha) + C_i^{\text{new}} \cdot [\nabla W_i(\alpha)]_0^2$ 
  end for
end if
 $S(\alpha) \leftarrow T(\alpha) + E(\alpha)$ 

```

$$\nabla W_i(\alpha) \leftarrow 1 - \frac{C_1^{\text{new}}}{C_1^{\text{old}}} \left(1 - \nabla W_i(\alpha^*) - \sum_{j \in uSV} \alpha_j^* K_{i,j} \right) - \sum_{j \in uSV} \alpha_j^* K_{i,j}$$

for all $i = 1, \dots, n$, or

$$\nabla W_i(\alpha) \leftarrow \nabla W_i(\alpha) + (C_i^{\text{old}} - C_i^{\text{new}}) \sum_{j \in bSV} K_{i,j}, \quad i = 1, \dots, n,$$

where in the first formula we used

$$1 - \nabla W_i(\alpha^*) - \sum_{j \in uSV} \alpha_j^* K_{i,j} = \sum_{j \in bSV} \alpha_j^* K_{i,j}. \quad (10)$$

Note that the first method implicitly recycles $\sum_{j \in bSV} \alpha_j^* K_{i,j}$ by (10), while the second method implicitly recycles $\sum_{j \in uSV} \alpha_j^* K_{i,j}$. Obviously, depending on the number of bounded and unbounded support vectors either the first or the second method is more efficient, and hence should be chosen. We decided to pick the first or second method depending on whether $2 \cdot \#uSV < \#bSV$ or not. This decision was based on counts of the involved floating point operations and the fact that in all our experiments we stored the entire kernel matrix in the memory. However note that both methods require to access some rows of the kernel matrix, and hence there is most likely a more efficient cut-off if only parts of the kernel matrix are stored in memory by caching. Since in general, the costs of computing a row of the kernel matrix depends on data set specific features, such as its dimensionality when using Gaussian kernels, there does not seem to exist a simple rule of thumb in this case, though. Consequently, we decided not to analyze this case carefully. Procedure 6 displays the corresponding pseudocode for this warm start, which we call W4. It is not hard to see, that in the worst case Procedure 6 is of order $\mathcal{O}(n^2)$, while in the best case it is only of order $\mathcal{O}(n)$. Since the average case cannot be easily analyzed, we will experimentally evaluate whether this warm start is efficient or not.

W6. Warm start by partially shrinking and partially recycling old solution

Let us now assume that we run through the λ -grid in reverse order. Then we have $C^{\text{old}} > C^{\text{new}}$, and hence we cannot immediately recycle the old approximate solution α^* . Nonetheless, there is a certain analogue to Procedure 6 possible. Indeed, we can initialize by $\alpha_i \leftarrow \alpha_i^*$ if $\alpha_i^* \leq C^{\text{new}}$ and by $\alpha_i \leftarrow C^{\text{new}}$ if $\alpha_i^* > C^{\text{new}}$. Again, the corresponding warm start needs some work to find an efficient implementation that recycles suitable parts of the gradient. In order to explain such an implementation we split the set uSV into

$$\begin{aligned} nuSV &:= \{j : 0 < \alpha_j^* \leq C_j^{\text{new}}\} \\ nbSV &:= \{j : C_j^{\text{new}} < \alpha_j^* < C_j^{\text{old}}\}, \end{aligned}$$

where we note that we use a slight abuse of the letters u and b in this notation. Now note that the initialization above multiplies all $\alpha_j^* \in bSV$ by the factor $C_1^{\text{new}}/C_1^{\text{old}}$, while it keeps all $\alpha_j^* \in nuSV$ unchanged. Obviously, both update rules make it possible to recycle parts of the gradient. Unfortunately, however, for $\alpha_j^* \in nbSV$, the situation is more complicated and no simple recycling is possible. Procedure 7, which displays the corresponding pseudocode is thus a little more complicated than Procedure 6. Note that all remarks concerning the computational requirements of Procedure 6 also apply to Procedure 7, and the same holds true for the rule that decides which part of the gradient is recycled. In the following we call this approach displayed in Procedure 7, W6.

Procedure 7 Initialize directions that violate the new box constrained by $\alpha_i \leftarrow C_i^{\text{new}}$ while keeping the rest unchanged and compute $\nabla W(\alpha)$, $S(\alpha)$, and $T(\alpha)$.

```

for  $i = 1$  to  $n$  do
  if  $\alpha_i > C_i^{\text{new}}$  then
     $\alpha_i \leftarrow C_i^{\text{new}}$ 
  end if
end for
 $T(\alpha) \leftarrow 0$ 
 $E(\alpha) \leftarrow 0$ 
if  $\#nuSV < \#bSV$  then
  for  $i = 1$  to  $n$  do
     $\nabla W_i(\alpha) \leftarrow 1 - \frac{C_1^{\text{new}}}{C_1^{\text{old}}} \cdot (1 - \nabla W_i(\alpha) - \sum_{j \in nuSV} \alpha_j^* K_{i,j} - \sum_{j \in nbSV} \alpha_j^* K_{i,j})$ 
     $\nabla W_i(\alpha) \leftarrow \nabla W_i(\alpha) - \sum_{j \in nuSV} \alpha_j^* K_{i,j} - \sum_{j \in nbSV} C_j^{\text{new}} K_{i,j}$ 
     $T(\alpha) \leftarrow T(\alpha) - \alpha_i \cdot \nabla W_i(\alpha)$ 
     $E(\alpha) \leftarrow E(\alpha) + C_i^{\text{new}} \cdot [\nabla W_i(\alpha)]_0^2$ 
  end for
else
  for  $i = 1$  to  $n$  do
     $\nabla W_i(\alpha) \leftarrow \nabla W_i(\alpha) + \sum_{j \in bSV} (C_j^{\text{old}} - C_j^{\text{new}}) K_{i,j}$ 
     $\nabla W_i(\alpha) \leftarrow \nabla W_i(\alpha) + \sum_{j \in nbSV} (\alpha_j^* - C_j^{\text{new}}) K_{i,j}$ 
     $T(\alpha) \leftarrow T(\alpha) - \alpha_i \cdot \nabla W_i(\alpha)$ 
     $E(\alpha) \leftarrow E(\alpha) + C_i^{\text{new}} \cdot [\nabla W_i(\alpha)]_0^2$ 
  end for
end if
 $S(\alpha) \leftarrow T(\alpha) + E(\alpha)$ 

```

W3 & W5. Warm start by scaling old solution

Finally, there is an easy warm start option that works regardless of the direction we run through the λ -grid. Indeed, we can always initialize by $\alpha_i \leftarrow \alpha_i^* \cdot C_1^{\text{new}} / C_1^{\text{old}}$. The Procedure 8 shows the corresponding $\mathcal{O}(n)$ pseudocode. Depending on whether $C_1^{\text{old}} < C_1^{\text{new}}$ or $C_1^{\text{old}} > C_1^{\text{new}}$ we call this approach W3 or W5, respectively.

Procedure 8 Initialize by $\alpha_i \leftarrow \alpha_i^* \cdot C_1^{\text{new}} / C_1^{\text{old}}$ and compute $\nabla W(\alpha)$, $S(\alpha)$, and $T(\alpha)$.

```

 $T(\alpha) \leftarrow 0$ 
 $E(\alpha) \leftarrow 0$ 
for  $i = 1$  to  $n$  do
   $\alpha_i \leftarrow \frac{C_1^{\text{new}}}{C_1^{\text{old}}} \cdot \alpha_i^*$ 
   $\nabla W_i(\alpha) \leftarrow 1 - \frac{C_1^{\text{new}}}{C_1^{\text{old}}} \cdot (1 - \nabla W_i(\alpha))$ 
   $T(\alpha) \leftarrow T(\alpha) - \alpha_i \cdot \nabla W_i(\alpha)$ 
   $E(\alpha) \leftarrow E(\alpha) + C_i^{\text{new}} \cdot [\nabla W_i(\alpha)]_0^2$ 
end for
 $S(\alpha) \leftarrow T(\alpha) + E(\alpha)$ 

```

3 Working sets of size two

So far, our algorithm performs an update in one coordinate per iteration. Let us now consider an algorithm which performs an update in *two* coordinates per iteration. Let us first present a simple lemma that computes the gain of a 2-dimensional update.

Lemma 3.1 *For $\delta_i, \delta_j \in \mathbb{R}$ and $i, j = 1, \dots, n$ we have*

$$W(\alpha + \delta_i e_i + \delta_j e_j) - W(\alpha) = \delta_i \cdot (\nabla W_i(\alpha) - \delta_i/2) + \delta_j \cdot (\nabla W_j(\alpha) - \delta_j/2) - \delta_i \delta_j K_{i,j}.$$

Proof: Applying Lemma 2.1 twice we find the assertion by $\nabla W_j(\alpha + \delta_i e_i) = \nabla W_j(\alpha) - \delta_i K_{i,j}$. ■

3.1 Solving the two dimensional problem exactly

In order to describe an algorithm that updates two variables at each iteration we first have to investigate how the two-variable update looks like in detail. To this end, we fix two coordinates $i, j \in \{1, \dots, n\}$ with $i \neq j$ and consider the function

$$(\tilde{\alpha}_i, \tilde{\alpha}_j) \mapsto W_{i,j}(\tilde{\alpha}_i, \tilde{\alpha}_j) := W(\alpha^{\setminus i,j} + \tilde{\alpha}_i e_i + \tilde{\alpha}_j e_j),$$

where $\alpha^{\setminus i,j} := \alpha - \alpha_i e_i - \alpha_j e_j$ is a fixed vector whose i -th and j -th coordinates equal zero. Simple calculation then shows

$$\begin{aligned} W_{i,j}(\tilde{\alpha}_i, \tilde{\alpha}_j) &= \langle e, \alpha^{\setminus i,j} \rangle + \tilde{\alpha}_i + \tilde{\alpha}_j - \frac{1}{2} \langle \alpha^{\setminus i,j}, K \alpha^{\setminus i,j} \rangle - \tilde{\alpha}_i \langle e_i, K \alpha^{\setminus i,j} \rangle - \tilde{\alpha}_j \langle e_j, K \alpha^{\setminus i,j} \rangle \\ &\quad - \frac{1}{2} (\tilde{\alpha}_i^2 + 2\tilde{\alpha}_i \tilde{\alpha}_j K_{i,j} + \tilde{\alpha}_j^2), \end{aligned}$$

where we used $K_{i,i} = K_{j,j} = 1$. Consequently, the partial derivatives are given by

$$\begin{aligned} \frac{\partial W_{i,j}(\tilde{\alpha}_i, \tilde{\alpha}_j)}{\partial \tilde{\alpha}_i} &= 1 - \langle e_i, K \alpha^{\setminus i,j} \rangle - \tilde{\alpha}_i - \tilde{\alpha}_j K_{i,j} \\ \frac{\partial W_{i,j}(\tilde{\alpha}_i, \tilde{\alpha}_j)}{\partial \tilde{\alpha}_j} &= 1 - \langle e_j, K \alpha^{\setminus i,j} \rangle - \tilde{\alpha}_j - \tilde{\alpha}_i K_{i,j}. \end{aligned}$$

In order to derive the maximum of $W_{i,j}$ on $[0, C_i] \times [0, C_j]$ from these derivatives, we need to consider three different cases.

The case $K_{i,j} = 1$

By setting the above derivatives to zero, we obtain the following system of linear equations

$$\begin{aligned} \alpha_i^* + \alpha_j^* &= 1 - \langle e_i, K \alpha^{\setminus i,j} \rangle \\ \alpha_i^* + \alpha_j^* &= 1 - \langle e_j, K \alpha^{\setminus i,j} \rangle \end{aligned}$$

that have to be satisfied for all global maxima $(\alpha_i^*, \alpha_j^*) \in \mathbb{R}^2$ of $W_{i,j}$. Now recall that we assumed that the kernel k is strictly positive definite, and therefore we see that $K_{i,j} = 1$ implies $x_i = x_j$, and hence $y_i = y_j$. From this we conclude $K_{i,\ell} = K_{j,\ell}$ for all $\ell = 1, \dots, n$, and thus we obtain $1 - \langle e_i, K \alpha^{\setminus i,j} \rangle = 1 - \langle e_j, K \alpha^{\setminus i,j} \rangle$. Consequently, $W_{i,j}$ attains its global maximum at every point of the affine subspace

$$\{(\alpha_i^*, \alpha_j^*) : \alpha_i^* + \alpha_j^* = 1 - \langle e_i, K \alpha^{\setminus i,j} \rangle\}, \quad (11)$$

which is a translated version of the anti-diagonal subspace $\{(\alpha, -\alpha) : \alpha \in \mathbb{R}\}$.

Now recall that $y_i = y_j$ implies $C_i = C_j$, and hence we are actually interested in finding a pair $(\tilde{\alpha}_i, \tilde{\alpha}_j)$ that maximizes $W_{i,j}$ on the square $[0, C_i]^2$. If $1 - \langle e_i, K\alpha^{i,j} \rangle \in [0, 2C_i]$, it is easy to see that the subspace (11) intersects the square, and hence $W_{i,j}$ attains the desired maximum at every element in this intersection. In particular, (α_i^*, α_j^*) , where

$$\alpha_i^* := \frac{1 - \langle e_i, K\alpha^{i,j} \rangle}{2}$$

is such a pair. Let us now assume that $1 - \langle e_i, K\alpha^{i,j} \rangle > 2C_i$. Then the subspace (11) lies ‘‘above’’ the square $[0, C_i]^2$, and since $W_{i,j}$ is concave, $W_{i,j}$ then attains its maximum over $[0, C_i]^2$ at a point of the set of edges $\{C_i\} \times [0, C_i] \cup [0, C_i] \times \{C_i\}$. Let us fix a pair $(\tilde{\alpha}_i, \tilde{\alpha}_j) \in \{C_i\} \times [0, C_i]$. Then we have

$$\frac{\partial W_{i,j}(\tilde{\alpha}_i, \tilde{\alpha}_j)}{\partial \tilde{\alpha}_j} = 1 - \langle e_j, K\alpha^{i,j} \rangle - \tilde{\alpha}_j - \tilde{\alpha}_i K_{i,j} = 1 - \langle e_j, K\alpha^{i,j} \rangle - \tilde{\alpha}_j - C_i > 0,$$

and hence $W_{i,j}$ attains its maximum over $\{C_i\} \times [0, C_i]$ at the corner (C_i, C_i) . Interchanging the roles of i and j we can thus conclude that $W_{i,j}$ attains its maximum over $[0, C_i]^2$ at (C_i, C_i) . Since we can analogously show that, for $1 - \langle e_i, K\alpha^{i,j} \rangle < 0$, the function $W_{i,j}$ attains its maximum over $[0, C_i]^2$ at $(0, 0)$, we finally find the update rule

$$\alpha_i^{new} := \alpha_j^{new} := \left[\frac{1 - \langle e_i, K\alpha^{i,j} \rangle}{2} \right]_0^{C_i} = \left[\frac{\nabla W_i(\alpha) + \alpha_i + \alpha_j}{2} \right]_0^{C_i}.$$

The case $K_{i,j} = -1$

In this case, we have $x_i = x_j$, and hence $y_i = -y_j$. From this we conclude $K_{i,\ell} = -K_{j,\ell}$ for all $\ell = 1, \dots, n$, and thus we obtain $\langle e_i, K\alpha^{i,j} \rangle = -\langle e_j, K\alpha^{i,j} \rangle$. Consequently, the derivatives above reduce to

$$\begin{aligned} \frac{\partial W_{i,j}(\tilde{\alpha}_i, \tilde{\alpha}_j)}{\partial \tilde{\alpha}_i} &= 1 - \langle e_i, K\alpha^{i,j} \rangle - \tilde{\alpha}_i + \tilde{\alpha}_j \\ \frac{\partial W_{i,j}(\tilde{\alpha}_i, \tilde{\alpha}_j)}{\partial \tilde{\alpha}_j} &= 1 + \langle e_i, K\alpha^{i,j} \rangle - \tilde{\alpha}_j + \tilde{\alpha}_i, \end{aligned}$$

and from this it is easy to conclude that $W_{i,j}$ does not have a global maximum. However, a closer inspection of $W_{i,j}$ yields the formula

$$W_{i,j}(\tilde{\alpha}_i, \tilde{\alpha}_j) = \langle e, \alpha^{i,j} \rangle + \tilde{\alpha}_i + \tilde{\alpha}_j - \frac{1}{2} \langle \alpha^{i,j}, K\alpha^{i,j} \rangle - (\tilde{\alpha}_i - \tilde{\alpha}_j) \langle e_i, K\alpha^{i,j} \rangle - \frac{1}{2} (\tilde{\alpha}_i - \tilde{\alpha}_j)^2,$$

and hence we see that, for fixed $\beta \in \mathbb{R}$, we have

$$W_{i,j}(\tilde{\alpha}_i, \tilde{\alpha}_i + \beta) = \langle e, \alpha^{i,j} \rangle + 2\tilde{\alpha}_i + \beta - \frac{1}{2} \langle \alpha^{i,j}, K\alpha^{i,j} \rangle + \beta \langle e_i, K\alpha^{i,j} \rangle - \frac{1}{2} \beta^2.$$

In other words, $W_{i,j}$ is an affine linear function with positive slope on the affine subspaces

$$\{(\tilde{\alpha}_i, \tilde{\alpha}_i + \beta) : \tilde{\alpha}_i \in \mathbb{R}\}, \quad \beta \in \mathbb{R},$$

and therefore $W_{i,j}$ attains its maximum over $[0, C_i] \times [0, C_j]$ at a point from the set of edges $\{C_i\} \times [0, C_j] \cup [0, C_i] \times \{C_j\}$. Let us first consider a pair $(\tilde{\alpha}_i, \tilde{\alpha}_j) \in \{C_i\} \times [0, C_j]$. Then we have

$$\frac{\partial W_{i,j}(\tilde{\alpha}_i, \tilde{\alpha}_j)}{\partial \tilde{\alpha}_j} = 1 - \langle e_j, K\alpha^{i,j} \rangle - \tilde{\alpha}_j + C_i,$$

and hence $W_{i,j}$ attains its maximum over $\{C_i\} \times [0, C_j]$ at (C_i, α_j^*) , where

$$\alpha_j^* = [1 - \langle e_j, K\alpha^{\setminus i,j} \rangle + C_i]_0^{C_j} = [\nabla W_j(\alpha) + \alpha_j - \alpha_i + C_i]_0^{C_j}.$$

Moreover, for $\delta_i := C_i - \alpha_i$ and $\delta_j := \alpha_j^* - \alpha_j$ we obtain the gain of this update by Lemma 3.1. Analogously, we can show that $W_{i,j}$ attains its maximum over $[0, C_i] \times \{C_j\}$ at (α_i^*, C_j) , where

$$\alpha_i^* = [1 - \langle e_i, K\alpha^{\setminus i,j} \rangle + C_j]_0^{C_i} = [\nabla W_i(\alpha) + \alpha_i - \alpha_j + C_j]_0^{C_i}.$$

Again, the gain of the corresponding update can be computed by Lemma 3.1, and by comparing both gains we can then decide which two-dimensional update yields the larger gain. The corresponding update is chosen in the algorithm.

The case $K_{i,j} \neq \pm 1$

To solve the two dimensional problem in this case we fix an $\alpha \in \mathbb{R}^n$ and write

$$\begin{aligned} \gamma_i &:= 1 - \langle e_i, K\alpha^{\setminus i,j} \rangle = 1 - \sum_{\ell \neq i,j} \alpha_\ell K_{i,\ell} = \nabla W_i(\alpha) + \alpha_i + \alpha_j K_{i,j} \\ \gamma_j &:= 1 - \langle e_j, K\alpha^{\setminus i,j} \rangle = 1 - \sum_{\ell \neq i,j} \alpha_\ell K_{j,\ell} = \nabla W_j(\alpha) + \alpha_j + \alpha_i K_{i,j} \end{aligned}$$

Using the derivatives of $W_{i,j}$ it is then easy to see that $W_{i,j}$ attains its global maximum at each point (α_i^*, α_j^*) that satisfies $\gamma_i = \alpha_i^* + \alpha_j^* K_{i,j}$ and $\gamma_j = \alpha_j^* + \alpha_i^* K_{i,j}$. Simple algebraic transformations now show

$$\alpha_i^* = \frac{\gamma_i - \gamma_j K_{i,j}}{1 - K_{i,j}^2} \quad \text{and} \quad \alpha_j^* = \frac{\gamma_j - \gamma_i K_{i,j}}{1 - K_{i,j}^2},$$

and by re-substituting the definition of γ_i and γ_j we hence obtain

$$\begin{aligned} \alpha_i^* &= \alpha_i + \frac{\nabla W_i(\alpha) - \nabla W_j(\alpha) K_{i,j}}{1 - K_{i,j}^2} \\ \alpha_j^* &= \alpha_j + \frac{\nabla W_j(\alpha) - \nabla W_i(\alpha) K_{i,j}}{1 - K_{i,j}^2} \end{aligned} \tag{12}$$

for the uniquely determined point at which $W_{i,j}$ attains its global maximum. Now if $(\alpha_i^*, \alpha_j^*) \in [0, C_i] \times [0, C_j]$ we can simply update by $(\alpha_i^{new}, \alpha_j^{new}) := (\alpha_i^*, \alpha_j^*)$. However, if $(\alpha_i^*, \alpha_j^*) \notin [0, C_i] \times [0, C_j]$ we have to make further calculations. For example, for $\alpha_i^* > C_i$ and $\alpha_j^* \in [0, C_j]$, the function $W_{i,j}$ attains its maximum over $[0, C_i] \times [0, C_j]$ at a point of the line $\{C_i\} \times [0, C_j]$ by the concavity of $W_{i,j}$. Consequently, in this case the update is

$$(\alpha_i^{new}, \alpha_j^{new}) := (C_i, [\nabla W_j(\alpha) + (\alpha_i - C_i)K_{i,j} + \alpha_j]_0^{C_j}),$$

i.e., we first update the i coordinate, which leads to the temporary gradient

$$\nabla W_j(\alpha) + (\alpha_i - C_i)K_{i,j},$$

and then perform a one-dimensional optimization over the j coordinate. The remaining three cases where exactly one direction of (α_i^*, α_j^*) violates the box constraint can be handled analogously. Finally, let us consider the cases, where both coordinates violate the constraint, e.g., $\alpha_i^* > C_i$ and $\alpha_j^* > C_j$. In this case, the concavity of $W_{i,j}$ shows that $W_{i,j}$ attains its maximum over $[0, C_i] \times [0, C_j]$ at a point of the set $\{C_i\} \times [0, C_j] \cup [0, C_i] \times \{C_j\}$. Consequently, we have to temporarily perform the above one-dimensional optimization twice, namely one over the i coordinate and one over the j coordinate. By computing the resulting gain of W for both optimizations we then can decide which optimization we have to choose for the update. Again, the remaining three cases can be handled analogously.

3.2 Selecting a working set of size two

Obviously, the 2D-SVM-solver displayed in Algorithm 2 is conceptionally very similar to the 1D-SVM-solver presented in Algorithm 1. However, so far we have not addressed how to choose the directions i^* and j^* in which the 2D-SVM-solver performs an update. Several possibilities exists for this issue, and we discuss a few of them in the following.

Algorithm 2 2D-SVM solver

```

initialize  $(\alpha, \nabla W(\alpha), T(\alpha), S(\alpha))$ 
while  $S(\alpha) > \frac{\epsilon}{2\lambda}$  do
    select directions  $i^*$  and  $j^*$ 
    update  $\alpha$  in the directions  $i^*$  and  $j^*$ 
    update  $\nabla W(\alpha)$  in the directions  $i^*$  and  $j^*$  and calculate  $T(\alpha)$  and  $E(\alpha)$ 
     $S(\alpha) \leftarrow T(\alpha) + E(\alpha)$ 
end while

```

WSS 0. Choose the pair of directions with maximal gain

Given a pair of directions (i, j) , Lemma 3.1 can be used to compute the gain of W resulting from the exact two dimensional optimization described in Subsection 3.1. Now one could consider all pairs of directions and choose the one with the largest gain. Of course, in practice this approach is prohibitive since the search is an $\mathcal{O}(n^2)$ operation, which has to be performed in each iteration. Nonetheless, this approach intuitively yields the optimal pair of directions, and all subset selection strategies developed below can be interpreted as low cost approximations to this approach. Consequently, we tested it to get a baseline number of iterations, to which all other subset selection strategies are compared to.

WSS 1. Optimal 1D-direction and previously found 1D-direction

A careful analysis of the behavior of the 1D-SVM-solver shows that it often comes into a regime in which it picks alternating indices i^* and j^* for a while. In other words, it tries to approximately solve the 2D-problem in the directions i^* and j^* . In order to avoid this cost-intensive alternating we can look for the best 1D-direction i^* and then perform a 2D-update over i^* and the optimal 1D-direction i_{old}^* chosen in the previous iteration. Clearly, the advantage of this approach is that it preserves the low-cost search from the 1D-SVM-solver. On the downside, however, it may not reduce the number of iterations very effectively.

WSS 2. Two optimal 1D-directions from separate subsets

Another simple way to preserve the low cost search from the 1D-SVM-solver is to split the index set $\{1, \dots, n\}$ into two parts $\{1, \dots, n/2\}$ and $\{n/2 + 1, \dots, n\}$ and search for the optimal 1D-directions over these two parts. In other words, we can choose the directions i^* and j^* by

$$\begin{aligned}
 i^* &\in \arg \max_{i \leq n/2} W(\alpha + \delta_i e_i) - W(\alpha) \\
 j^* &\in \arg \max_{i > n/2} W(\alpha + \delta_i e_i) - W(\alpha),
 \end{aligned}$$

where δ_i is defined as in the 1D-SVM-solver. Clearly, this approach preserves the low cost search from the 1D-SVM-solver, but again it is not clear whether it reduces the number of iterations very effectively.

WSS 4. Optimal 1D-direction and a direction of a nearby sample

Yet another approach to preserve the low cost search from the 1D-SVM-solver is to first look for the optimal 1D-direction i^* , and then, in a second step, to pick a direction j^* such that x_{j^*} is close to x_{i^*} with respect to the metric induced by the kernel. The rationale behind this idea is that if one uses the Gaussian RBF kernel and optimizes in direction i^* , the gradients of samples close to x_{i^*} are most affected by the update in direction i^* because of the bell shape of the Gaussian. If these gradients were close to zero *before* the update, they are thus most likely no longer close to zero *after* the update. Consequently, the corresponding directions will have a good chance of being chosen in a subsequent iteration. In our experiments, we considered the k -nearest neighbors of x_{i^*} , where $k = 10$, and picked the neighbor x_{j^*} for which the 2D-update in the directions (i^*, j^*) yielded the largest gain. Note that, as soon as the direction i^* is found, it is clear that one needs to access the i^* -th kernel row for updating the gradient subsequently, and hence finding the k -nearest neighbors is relatively inexpensive. Moreover, computing the 2D-gain over k candidates is also relatively inexpensive, if k remains small. Nonetheless, initial experiments suggested that searching over the k -nearest neighbors only makes sense if the solver mainly updates *inner* support vectors, i.e., directions i with $0 < \alpha_i < C_i$. Consequently, we implemented a Boolean flag that was recomputed every 10 iterations. In this re-computation, the flag was set to true, if and only if in at least 5 of the previous 10 iterations the picked directions i^* and j^* both were inner support vectors. We then considered the k -nearest neighbors only if this Boolean flag was set, while in the other case we applied the working set selection strategy described in WSS 1.

WSS x. Combinations of optimal 1D-direction-based approaches

It is easy to see that one can combine the three methods that are based on finding the optimal 1D-direction. For example, in each iteration one can combine WSS 1 and WSS 2 by computing the 2D-gain of both methods and pick the one with the larger gain. Obviously, this still preserves the low cost search from the 1D-SVM-solver and only adds little cost for computing the 2D-gain for the two candidate pairs. Similarly, all three methods can be combined. Combinations of these methods are called WSS x , where x is the sum of the combined methods. For example, by combining WSS 1, WSS 2, and WSS 4 we obtain WSS 7, and by combining WSS 1, WSS 2, WSS 4 with WSS 512 below, we obtain WSS 519. In the following, we keep this binary numbering system which makes it possible to easily describe arbitrary combinations of basic working set selection strategies.

WSS 8. Optimal 1D-direction and optimal one-step-ahead 1D-direction

Another way to modify the 1D-SVM subset selection strategy to two directions is to look for the optimal 1D-direction i^* first, and then look for the optimal 1D-direction j^* that would be found after having updated in direction i^* . Obviously, this strategy, which we call WSS 8, is closely related to WSS 1 in that the update and search routines are partially permuted. However, it has a higher cost for the search part per iteration, while intuitively it should reduce the number of iterations.

WSS 16. Maximal violating pair

A completely different subset selection strategy is based on the maximal violating pair idea. For the SVM without offset, this means that the pair (i^*, j^*) is chosen that violates (5) most. In other words, for both index sets $\{i : \alpha_i < C_i\}$ and $\{i : \alpha_i > 0\}$ the two indices with the largest, respectively smallest, gradients are picked, and the final pair (i^*, j^*) consists of the indices that have the gradient with the largest absolute value among the four candidate directions. In order to implement this working set selection strategy efficiently, the sets $\{i : \alpha_i < C_i\}$ and $\{i : \alpha_i > 0\}$ should be kept in memory and updated in every iteration. This may add some cost per iteration compared to the previous working set selection strategies, while it is unclear how the number of iterations behave compared to these strategies.

WSS 32. Optimal 1D-direction and corresponding optimal 2D-direction

None of the methods introduced so far try to seriously approximate the 2D subset selection strategy WSS 0, which intuitively picks the best possible pair of indices. The first method that seriously tries such an approximation is WSS 32, which first picks the optimal 1D-direction i^* , and then searches for the $j^* \in \{1, \dots, n\}$ such that (i^*, j^*) maximizes the corresponding 2D-gain. Obviously, the cost for this search method is significantly higher than those of WSS 1 to WSS 7, but it is still $\mathcal{O}(n)$. On the other hand, the better choice of (i^*, j^*) may substantially reduce the number of iterations of the 2D-SVM-solver, and hence it is not a-priori clear how WSS 32 performs compared to the earlier methods.

WSS 64. Optimal 1D-direction and random optimal 2D-direction

Instead of considering all pairs (i^*, j) , $j = 1, \dots, n$, as WSS 32 does, it may suffice to reduce the search over the pairs (i^*, j) , $j \in J$, where $J \subset \{1, \dots, n\}$ is a random subset. In our experiments we considered the case $\#J = n/5$.

WSS 128. Optimal 1D-direction and approximately optimal 2D-direction

One of the disadvantages of WSS 32 is that computing the 2D-gain is quite expensive due to the relatively large number of branches and floating point operations. One way to address this issue is to compute the 2D-gain in WSS 32 only approximately. WSS 128 uses the following approximation: for indices i and j with $K_{i,j} = \pm 1$ it computes the exact gain, while for the other pairs it first computes α_i^* and α_j^* by (12), and then applies the simple clipping operation

$$\begin{aligned}\alpha_i^{new} &:= [\alpha_i^*]_0^{C_i} \\ \alpha_j^{new} &:= [\alpha_j^*]_0^{C_j}.\end{aligned}$$

For these new α 's, WSS 128 finally computes the gain by Lemma 3.1. Clearly, this gain is in general less than the exact gain, but it still may be a good approximation. In particular, if both α_i^* and α_j^* satisfy the box constraints, then the approximation is actually exact. On the other hand, the approximation is clearly less expensive, but we expect more iterations compared to WSS 32.

WSS 256. Optimal random 2D-directions

Another way to approximate WSS 0 is to consider k random pairs (i, j) , and pick the pair (i^*, j) that yields the largest exact 2D-gain among them. In WSS 256 we followed this idea for $k := n$.

WSS 512. Optimal 1D-direction and optimal 2D-direction over inner SVs

Although the approximate computation of the 2D-gain in WSS 128 is cheaper than the exact computation in WSS 32, it may still be too expensive. One way to further decrease these costs is based on the observation that the 2D-gain is given by

$$\frac{1}{2} \cdot \frac{|\nabla W_i(\alpha)|^2 + |\nabla W_j(\alpha)|^2 - 2\nabla W_i(\alpha)\nabla W_j(\alpha)K_{i,j}}{1 - K_{i,j}^2}$$

if $K_{i,j} \neq \pm 1$ and α_i^* and α_j^* computed by (12) satisfy the box constraints. WSS 512 uses this simplified formula in the following way. Again, it first searches for the optimal 1D-direction i^* . If α_{i^*} is an inner support vector, see WSS 4 for a definition, and the Boolean flag of WSS 4 is set, WSS 512 searches for the direction

$$j^* \in \{j : 0 < \alpha_j < C_j \text{ and } K_{i^*,j} \neq \pm 1\}$$

that optimizes the above formula of the 2D-gain for fixed $i := i^*$. Since in some iterations WSS 512 reduces to the 1D-SVM-solver we further considered some combinations with WSS 3, and WSS 7 in our experiments. Following the naming convention of combinations mentioned earlier, these strategies are called WSS 515 and WSS 519.

WSS 1024. Optimal 1D-direction and random 2D-direction over inner SVs

The next subset selection strategy, WSS 1024, is quite similar to WSS 512, except that it does not consider all inner support vectors in the search for j^* , but only k random inner support vectors. We set the routine up so that k equaled 20% of the current number of inner support vectors. In addition, we initiated the search whenever α_{i^*} was an inner support vector, i.e., the search was initiated *independently* of the Boolean flag of WSS 4. Again, in some iterations WSS 1024 reduces to the 1D-SVM-solver, and hence we further considered some combinations with WSS 1, WSS 2, and WSS 4, where again the naming convention above was used.

WSS 2048. Add random optimal 2D-directions over inner SVs

The final subset selection strategy, WSS 2048, is actually not a subset selection strategy of its own, but only a strategy that works in combination with others. Once one of the previous subset selection strategies has picked a pair (i^*, j^*) and α_{i^*} has turned out to be an inner support vector, WSS 2048 considers k random pairs of inner support vectors, and picks the pair (i^{**}, j^{**}) that has largest approximate gain, where the approximation was computed as in WSS 512. Then the exact gain of (i^*, j^*) and (i^{**}, j^{**}) is computed and the pair with the larger exact gain was chosen. We considered this method in combination with WSS 1, WSS 2, and WSS 4, where again the naming convention above was used.

4 Convergence Analysis

In this section we establish an upper bound on the number of iterations for both the 1D-SVM and the 2D-SVM. Our approach is heavily based on ideas developed for the analysis of rate-certifying decomposition algorithms, see e.g., [9, 17, 8, 18]. In particular, we need the σ -functional, which, for $\alpha \in [0, C] = [0, C_1] \times \cdots \times [0, C_n]$ and $I \subset \{1, \dots, n\}$, is defined by

$$\sigma(\alpha|I) := \sup_{\substack{\tilde{\alpha} \in [0, C] \\ \tilde{\alpha}_i = \alpha_i \forall i \notin I}} \langle \nabla W(\alpha), \tilde{\alpha} - \alpha \rangle.$$

Since our algorithms are based on gain optimization rather than rate certification, we further need the γ -functional

$$\gamma(\alpha|I) := \sup_{\substack{\tilde{\alpha} \in [0, C] \\ \tilde{\alpha}_i = \alpha_i \forall i \notin I}} W(\tilde{\alpha}) - W(\alpha),$$

which expresses the gain in the dual objective function resulting from an optimization over the directions contained in I . To simplify notations, we write $\sigma(\alpha|i) := \sigma(\alpha|\{i\})$ and $\gamma(\alpha|i) := \gamma(\alpha|\{i\})$. Note that we have

$$\sigma(\alpha|i) = \sup_{\tilde{\alpha}_i \in [0, C_i]} (\tilde{\alpha}_i - \alpha_i) \nabla W_i(\alpha),$$

while $\gamma(\alpha|i)$ expresses the gain

$$W(\alpha + (\alpha_i^{new} - \alpha_i)e_i) - W(\alpha)$$

of the 1D-update in direction i , where α_i^{new} is defined by (4). In addition, $\gamma(\alpha|\{i, j\})$ is the gain obtained by the update discussed in Subsection 3.1. Moreover, for $I = \{1, \dots, n\}$ we write $\sigma(\alpha) := \sigma(\alpha|I)$ and $\gamma(\alpha) := \gamma(\alpha|I)$, respectively. Note that both σ and γ are monotonic in I , i.e., for $I \subset J$ we have $\sigma(\alpha|I) \leq \sigma(\alpha|J)$ and $\gamma(\alpha|I) \leq \gamma(\alpha|J)$. Finally, we need the obvious relation

$$\gamma(\alpha) = W(\alpha^*) - W(\alpha),$$

where we recall from Section 2 that $\alpha^* \in [0, C]$ denotes a solution of the dual problem (3). In other words, $\gamma(\alpha)$ expresses the dual sub-optimality of α .

Let us now begin our analysis by presenting two lemmata that establish relationships between these quantities.

Lemma 4.1 *For all $\alpha \in [0, C]$ we have*

$$\sum_{i=1}^n \sigma(\alpha|i) = \sigma(\alpha) = \text{gap}(\alpha),$$

where $\text{gap}(\alpha)$ denotes the duality gap defined in (7). In particular, there exists an index $i^* \in \{1, \dots, n\}$ such that

$$\sigma(\alpha|i^*) \geq n^{-1}\sigma(\alpha).$$

This lemma can be easily derived from results in [15] and [18]. However, in the case of SVMs without offset, its proof is very elementary and hence we present it here.

Proof: For $i \in \{1, \dots, n\}$ it is easy to see that the supremum used to define $\sigma(\alpha|i)$ is attained at

$$\bar{\alpha}_i := \begin{cases} C_i & \text{if } \nabla W_i(\alpha) \geq 0 \\ 0 & \text{if } \nabla W_i(\alpha) < 0. \end{cases} \quad (13)$$

Moreover, the vector $\bar{\alpha} := (\bar{\alpha}_1, \dots, \bar{\alpha}_n) \in [0, C]$ realizes the supremum defining $\sigma(\alpha)$, and hence we obtain

$$\sum_{i=1}^n \sigma(\alpha|i) = \sum_{i=1}^n \langle \nabla W(\alpha), (\bar{\alpha}_i - \alpha_i)e_i \rangle = \langle \nabla W(\alpha), \bar{\alpha} - \alpha \rangle = \sigma(\alpha).$$

Furthermore, we have

$$\begin{aligned} \sigma(\alpha) = \langle \nabla W(\alpha), \bar{\alpha} - \alpha \rangle &= \langle \alpha, K\alpha \rangle - \langle e, \alpha \rangle + \sum_{i=1}^n \bar{\alpha}_i \cdot \nabla W_i(\alpha) \\ &= \langle \alpha, K\alpha \rangle - \langle e, \alpha \rangle + \sum_{i=1}^n C_i [\nabla W_i(\alpha)]_0^\infty, \end{aligned}$$

and hence we have shown $\sigma(\alpha) = \text{gap}(\alpha)$. The last assertion is a trivial consequence of the first assertion. \blacksquare

The second lemma relates $\sigma(\alpha|I)$ to the gain $\gamma(\alpha|I)$. For its formulation we need the quantity $B_{\max} := \max_{i=1, \dots, n} C_i$.

Lemma 4.2 *For all $\alpha \in [0, C]$ and $I \subset \{1, \dots, n\}$ we have*

$$\sigma(\alpha|I) \geq \gamma(\alpha|I) \geq \frac{\sigma(\alpha|I)}{2} \min \left\{ 1, \frac{\sigma(\alpha|I)}{|I|^2 B_{\max}^2} \right\},$$

where $|I|$ denotes the cardinality of I .

In a slightly different form, this lemma has been established in, e.g., [8], and it was somewhat implicitly used in [18]. Again, we present its proof for the sake of completeness.

Proof: Let $\bar{\alpha}_i$ be defined by (13) and $d := \sum_{i \in I} (\bar{\alpha}_i - \alpha_i)e_i$. For $\lambda \in [0, 1]$, we then have $\alpha + \lambda d \in [0, C]$, and a calculation analogous to the one in the proof of Lemma 2.1 yields

$$\gamma(\alpha|I) \geq W(\alpha + \lambda d) - W(\alpha) = \lambda \langle \nabla W(\alpha), d \rangle - \frac{\lambda^2}{2} \langle d, Kd \rangle \geq \lambda \sigma(\alpha|I) - \frac{\lambda^2 |I|^2 B_{\max}^2}{2}.$$

Now the right hand side is maximized at

$$\lambda^* := \begin{cases} 1 & \text{if } \sigma(\alpha|I) > |I|^2 B_{\max}^2 \\ |I|^{-2} B_{\max}^{-2} \sigma(\alpha|I) & \text{if } \sigma(\alpha|I) \leq |I|^2 B_{\max}^2. \end{cases}$$

In the case $\sigma(\alpha|I) > |I|^2 B_{\max}^2$ we hence find

$$\gamma(\alpha|I) \geq \sigma(\alpha|I) - \frac{|I|^2 B_{\max}^2}{2} > \frac{\sigma(\alpha|I)}{2},$$

while in the other case $\sigma(\alpha|I) \leq |I|^2 B_{\max}^2$ we obtain

$$\gamma(\alpha|I) \geq \frac{\sigma^2(\alpha|I)}{2|I|^2 B_{\max}^2}.$$

Combining all estimates we then obtain the inequality on the right hand side.

To show the inequality on the left hand side we fix an $\tilde{\alpha} \in [0, C]$ such that $\tilde{\alpha}_i = \alpha_i$ for all $i \notin I$. Then we have

$$W(\tilde{\alpha}) - W(\alpha) = \langle \nabla W(\alpha), \tilde{\alpha} - \alpha \rangle - \frac{1}{2} \langle \tilde{\alpha} - \alpha, K(\tilde{\alpha} - \alpha) \rangle \leq \langle \nabla W(\alpha), \tilde{\alpha} - \alpha \rangle \leq \sigma(\alpha|I),$$

and by maximizing the left hand side of this inequality over $\tilde{\alpha}$ we find $\gamma(\alpha|I) \leq \sigma(\alpha|I)$. ■

With these preparations we can now present a preliminary result on iterative algorithms that have a certain control of their gain.

Proposition 4.3 *Let $\alpha^{(0)}, \alpha^{(1)}, \dots \in [0, C]$ be a sequence of feasible vectors that satisfies*

$$W(\alpha^{(\ell+1)}) - W(\alpha^{(\ell)}) \geq \gamma(\alpha^{(\ell)}|i_\ell^*), \quad \ell \geq 0, \quad (14)$$

where for each ℓ the index $i_\ell^* \in \{1, \dots, n\}$ is the one described in Lemma 4.1, that is, it satisfies $\sigma(\alpha^{(\ell)}|i_\ell^*) \geq n^{-1}\sigma(\alpha^{(\ell)})$. Then for all $\ell \geq 1$ we have

$$\gamma(\alpha^{(\ell+1)}) \leq \gamma(\alpha^{(\ell)}) \left(1 - \frac{1}{2n} \min \left\{ 1, \frac{\gamma(\alpha^{(\ell)})}{nB_{\max}^2} \right\} \right).$$

Moreover, for all $\varepsilon > 0$ and all $\ell \geq \ell_\varepsilon$ we have $\gamma(\alpha^{(\ell)}) \leq \varepsilon$, where

$$\ell_\varepsilon := \left\lceil \frac{2n^2 B_{\max}^2}{\varepsilon} \right\rceil + \max \left\{ 0, \left\lceil 2n \ln \frac{W(\alpha^*) - W(\alpha^{(0)})}{\varepsilon} \right\rceil \right\}.$$

Proof: By Lemmas 4.2 and 4.1 we find

$$\begin{aligned} \gamma(\alpha^{(\ell)}) - \gamma(\alpha^{(\ell+1)}) = W(\alpha^{(\ell+1)}) - W(\alpha^{(\ell)}) &\geq \gamma(\alpha^{(\ell)}|i_\ell^*) \\ &\geq \frac{\sigma(\alpha^{(\ell)}|i_\ell^*)}{2} \min \left\{ 1, \frac{\sigma(\alpha^{(\ell)}|i_\ell^*)}{B_{\max}^2} \right\} \\ &\geq \frac{\sigma(\alpha^{(\ell)})}{2n} \min \left\{ 1, \frac{\sigma(\alpha^{(\ell)})}{nB_{\max}^2} \right\} \\ &\geq \frac{\gamma(\alpha^{(\ell)})}{2n} \min \left\{ 1, \frac{\gamma(\alpha^{(\ell)})}{nB_{\max}^2} \right\}. \end{aligned}$$

From this we easily obtain the first assertion.

The second assertion has already been shown in the second part of the proof of the first assertion of [18, Theorem 4], which can be found on the pages 312 and 313 of [18]. ■

Note that $1/n$ -rate certifying algorithms considered in [18] clearly satisfy assumption (14). Moreover, Proposition 4.3 can also be applied to the 1D-SVM and 2D-SVM:

Theorem 4.4 *Consider the 1D-SVM described in Algorithm 1 or a 2D-SVM in the sense of Algorithm 2 that uses a working set selection strategy whose gain at each iteration is not less than that of the 1D-SVM. Furthermore, assume that $\max\{w_{\text{neg}}, w_{\text{pos}}\} \leq 1$. Then for all $\varepsilon > 0$, $n \geq 1$, and all $\lambda > 0$ these algorithms terminate after at most*

$$\left\lceil \frac{2}{\lambda \varepsilon \min\{1, 2\lambda \varepsilon\}} \right\rceil + \max \left\{ 0, \left\lceil 2n \ln \frac{4\lambda(W(\alpha^*) - W(\alpha^{(0)}))}{\varepsilon \min\{1, 2\lambda \varepsilon\}} \right\rceil \right\}$$

iterations. In particular, in the most likely scenario $2\lambda\epsilon \leq 1$ these algorithms do not need more iterations than

$$\left\lceil \frac{1}{\lambda^2 \epsilon^2} \right\rceil + \max \left\{ 0, \left\lceil 2n \ln \frac{2(W(\alpha^*) - W(\alpha^{(0)}))}{\epsilon^2} \right\rceil \right\}.$$

Proof: The 1D-SVM chooses at each iteration ℓ a direction i_ℓ^* that maximizes the 1D-gain $\gamma(\alpha^{(\ell)}|i)$. Consequently, we have

$$W(\alpha^{(\ell+1)}) - W(\alpha^{(\ell)}) = \gamma(\alpha^{(\ell)}|i_\ell^*) \geq \gamma(\alpha^{(\ell)}|i_\ell^*),$$

where i_ℓ^* is the direction described in Lemma 4.1. In other words, (14) is satisfied for this algorithm, and from this it is not hard to see that the considered 2D-SVM's also satisfy assumption (14). Let us now define

$$h(\sigma) := \frac{\sigma}{2} \min \left\{ 1, \frac{\sigma}{n^2 B_{\max}^2} \right\}, \quad \sigma > 0.$$

For $\epsilon := h(\frac{\epsilon}{2\lambda})$ Proposition 4.3 together with Lemma 4.2 then shows that

$$h\left(\frac{\epsilon}{2\lambda}\right) = \epsilon \geq \gamma(\alpha^{(\ell)}) \geq h(\sigma(\alpha^{(\ell)}))$$

for all $\ell \geq \ell_\epsilon$ and hence we obtain $S(\alpha^{(\ell)}) \leq \text{gap}(\alpha^{(\ell)}) = \sigma(\alpha^{(\ell)}) \leq \frac{\epsilon}{2\lambda}$ by the monotonicity of the function h . Using $B_{\max} \leq \frac{1}{2\lambda n}$ we then obtain the assertion by simple algebraic transformations. \blacksquare

Note that the working set selection strategies WSS 1, WSS 2, WSS 4, WSS 8, WSS 32, WSS 64, WSS 128, WSS 512, and WSS 1024, satisfy the assumptions of Theorem 4.4. Moreover, the same is true for all combinations of working set selection strategies that include at least one of the strategies listed. Finally, note that the upper bound established in Proposition 4.3 coincide (modulo constants that come from different working set sizes) with the bounds for rate certifying algorithms presented in [17, 8, 18]. Moreover, the step from dual ϵ -optimality to primal ϵ -optimality considered in the proof of Theorem 4.4 coincides with the analysis [15] for SVMs with offset. Consequently, the bound presented in Theorem 4.4 equals the best known guarantees for solvers for SVMs with offset.

5 Experiments

The described 1D-SVM-solver and 2D-SVM-solver enjoy nice theoretical properties with respect to both generalization performance and required training time. However, it is unclear how tight these bounds are, so it remains unclear whether the proposed SVMs also perform well in practice. Therefore, we performed several experiments that address the following questions:

- i) Which subset selection strategies lead to the smallest number of iterations or the shortest runtime? How many more iterations than WSS 0 do these strategies perform?
- ii) How many less iterations needs the stopping criterion (9) compared to standard duality gap (7) and is there also an advantage in terms of runtime?

- iii) How much more efficient is the 2D-SVM-solver than the technically much easier 1D-SVM-solver?
- iv) How well does the 2D-SVM-solver work compared to standard software packages such as LIBSVM [1]?
- v) What is the advantage of warm start initializations when the parameter search is performed over a grid?

To answer these questions we implemented the 1D-SVM and the 2D-SVM solver in C++, and downloaded LIBSVM [1] version 2.82. The algorithms were compiled by LINUX's gcc version 4.3 with various software and hardware optimizations enabled. All experiments were conducted on a computer with INTEL XEON X5355 (2.66 GHz) quad core processor and 8GB RAM under a 64bit version of RedHat Linux Enterprise 4. During all experiments that incorporated measurements of runtime, one core was used solely for the experiments, and the number of other processes running on the system was minimized. The runtime itself was measured by the C function `clock()` from the library `time.h`. The resulting resolution was 0.01 seconds.

In some preliminary experiments we made a couple of observations that changed the described implementation strategy slightly: First, it turned out that the auto-vectorization of gcc only gave mediocre and sometimes even contradicting results even if the implementation guidelines of gcc 4.3's auto-vectorization were strictly followed. Therefore, we decided to manually code SSE2 vectorized versions of the most important routines, namely: computing kernel values, searching for the optimal 1D-direction, updating the gradient, and computing the weighted sum $E(\alpha)$ of clipped slack variables. To this end, we used the library `emmintrin.h` together with properly aligned arrays of doubles. Some of our preliminary experiments not reported here indicated that this specialized hardware instruction set yielded a runtime improvement by a factor between 1.3 and 1.8 depending on the working set selection strategy and the data set. In addition, the initial experiments suggested substantial numerical instabilities on a few datasets when using single floats, so we decided to use double floats throughout the experiments. Second, we were rather disappointed by the runtime behavior of LIBSVM, even when we enabled its shrinking heuristic.¹ After some investigations we found that the main reason for the disappointing runtime performance was the fact that LIBSVM copies kernel rows into the kernel cache if one uses pre-computed kernel matrices, which, as discussed below, we did throughout the experiments. This copying mechanism results in a small number of iterations per second whenever the LIBSVM solver is started on a new parameter point, while with the kernel cache being filled up during the optimization, the solver is actually able to process more iterations per second. To ensure a fair comparison, we thus decided to implement our own version of LIBSVM's solver (without shrinking strategy). As a side effect, this new implementation also benefited from the SSE2 instructions for upgrading the gradient. Unlike the subset selection strategy of the 1D-SVM-solver, however, LIBSVM's subset selection strategy, though implementable, does not benefit from vectorization since not all indices are considered, and hence the relatively slow RAM access of the CPU outweighs the speed improvement of the SSE2 instructions.

We downloaded all datasets for binary classification from LIBSVM's homepage whose number of features did not exceed 1000. We made this cut because having data sets with a huge

¹In fact, it turned out that neither the number of iterations nor the runtime was significantly affected by the shrinking heuristic.

	training size	test size	dimension	Test error 2D-SVM	Test error LIBSVM
SONAR	146	62	60	12.80 \pm 4.04	12.68 \pm 4.27
HEART	188	82	13	17.42 \pm 4.39	17.58 \pm 3.80
LIVER-DISORDERS	248	97	6	29.76 \pm 4.31	29.31 \pm 4.00
IONOSPHERE	248	103	34	8.59 \pm 2.85	5.43 \pm 2.16
AUSTRALIAN	484	206	14	14.54 \pm 2.09	14.76 \pm 2.21
BREAST-CANCER	493	190	10	3.15 \pm 1.07	3.30 \pm 1.06
DIABETES	544	334	8	23.68 \pm 2.49	23.43 \pm 2.38
FOURCLASS	623	239	2	0.04 \pm 0.14	0.09 \pm 0.18
GERMAN.NUMER	718	282	24	24.84 \pm 2.29	24.95 \pm 2.27
SVMGUIDE3	892	392	21	16.60 \pm 1.77	16.48 \pm 1.70
COVTYPE-2000	1392	616	54	23.92 \pm 1.69	24.06 \pm 1.60
IJCNN1-2000	1424	584	33	4.38 \pm 0.93	4.38 \pm 0.91
A1A	1605	30956	123	15.89 \pm 0.21	15.78 \pm 0.17
SPLICE	2176	999	60	8.93 \pm 0.88	8.68 \pm 0.87
A2A	3365	30296	123	15.74 \pm 0.27	15.76 \pm 0.30
W1A	2477	47332	300	2.18 \pm 0.06	2.20 \pm 0.07
A3A	3185	29336	123	15.82 \pm 0.21	15.57 \pm 0.08
W2A	3470	46339	300	1.95 \pm 0.06	1.94 \pm 0.09
COVTYPE-5000	3472	1536	54	20.73 \pm 0.83	20.77 \pm 0.88
IJCNN1-5000	3486	1514	33	2.70 \pm 0.45	2.73 \pm 0.42
A4A	4781	33780	123	15.80 \pm 0.30	15.52 \pm 0.07
W3A	4912	44833	300	1.75 \pm 0.05	1.75 \pm 0.05
SVMGUIDE1	4959	2130	4	3.01 \pm 0.33	2.97 \pm 0.32
MUSHROOMS	5773	2351	112	0.00 \pm 0.00	0.00 \pm 0.01

Table 1: Size and dimensionality of the considered datasets together with the test errors (\pm standard deviations) on 100 random runs for the 2D-SVM with WSS 7 working set selection and l1-W4 initialization and LIBSVM. The hyper-parameters were selected by 10-fold cross-validation on the 10 by 10 grid described in the text. Besides IONOSPHERE, both algorithms performed almost indistinguishable. Moreover, the training and test set sizes refer to the splits used in the experiments on the run time behavior of the SVM solvers.

number of features would have required substantial extra effort for implementing our algorithms, and this effort was clearly out of the scope of this paper. In all cases we used the scaled versions of these datasets, and if they were not available, we scaled the unscaled datasets with the help of LIBSVM’s scaling tool. For datasets that were not split into a training and test set we generated a random split that contained approximately 70% training and 30% test samples. Moreover, for the already split datasets SPLICE, SVMGUIDE1, SVMGUIDE3, we decided to first merge the corresponding training and test set, and then generate the random split above. For the large datasets COVTYPE and IJCNN1, we generated random subsets of the two datasets of sizes $n = 2000, 5000$, and then applied the random split above. Finally, we ignored some versions with larger training set of the AXA and WXA families, namely A5A – A9A, and W4A – W8A because of time or memory constraints. Moreover, for these two families of data sets we kept the split between training and test sets. Table 1 shows the corresponding characteristics of the considered datasets together with classification errors of the fastest version of the 2D-SVM and LIBSVM, respectively.

In all our experiments we considered k -fold cross validation with randomly generated folds performed on the training set. In our choice for the hyper-parameter grid we were guided by recent theoretical results from [21], which show that *asymptotically* good values of λ and σ are contained in the intervals $[c_1 n^{-2}, 1]$ and $[c_2, c_3 n^{1/d}]$, respectively, where n is the number of training samples, d is the input dimension and c_1, c_2 , and c_3 are arbitrarily specifiable constants. Based on this result, we considered a geometrically spaced 10 by 10 grid in $[10n^{-2}, 1] \times [0.1, 2n^{1/d}]$, i.e., the ratio of consecutive grid points was constant. Moreover, it is worth mentioning that during the k -fold cross validation λ was internally converted to C

by the formula $C := \frac{k}{2(k-1)\lambda n}$ to accommodate the fact that the *actual* training set size for k-fold cross validation is approximately $(k-1)n/k$.

In all experiments with the 1D-SVM and the 2D-SVM we used $\epsilon := 0.001$ for the stopping criterion (9), while for our version of LIBSVM’s solver we used, like the original LIBSVM, the classical MVP stopping criterion with value $\epsilon = 0.001$. Here we note that this was necessary since LIBSVM’s solver deals with SVMs *with* offset b , and hence the stopping criterion (9) is no longer applicable. In addition, an appropriately modified stopping criterion seems to be computationally inefficient, while by [15, Lemma 8] the MVP stopping criterion with value $\epsilon = 0.001$ also ensures (8) for $\epsilon := 0.001$ and f^* instead of $[f^*]_{-1}^1$. In other words, LIBSVM’s default value, which we picked throughout our experiments, actually has a good interpretation in terms of learning. Of course, the different stopping criteria used raise the question whether the results reported below are due to differences in the working set selection strategy, the different nature of the optimization problem, *or* the stopping criteria. In this regard, we note that in the experiments with LIBSVM our goal is to compare the *entire* 2D-SVM-solver with a state-of-the-art solver, rather than to, e.g., compare different working set selection strategies. Obviously, for this purpose it is irrelevant whether the working set selection strategy, the nature of the optimization problem, or the different stopping criteria are more responsible for differences in the runtime. Nonetheless, it remains an interesting question for future work whether solver’s for SVMs with offset can also benefit from some of the ideas of the working set selection strategies introduced for the 2D-SVM-solver.

In all experiments we pre-computed the kernel matrix in order to avoid that these solver independent but dataset dependent computations are contained in the reported training time. Obviously, this approach gives us a clearer view of the performance of the core solver, on the downside however, this may be an unrealistic setting for large datasets whose kernel matrices do not fit into the computers memory. On the other hand, for all considered datasets the matrices *did* fit into memory, and in addition, it turned out that for all datasets there were parameter regions of the grid where all or basically all vectors were support vectors. Clearly, the corresponding kernel rows would have been computed if we had not precomputed the kernel matrices, and consequently, training over the grid would have required the solver to compute the kernel matrix anyway. In other words, our experiments suggests that training over a grid with medium-sized datasets whose kernel matrix still fits into memory, there is no need to implement a caching strategy. In fact, we strongly conjecture that without pre-computing the kernel matrices, our experiments would have rendered computationally infeasible with one computer only. It is, of course, needless to say, that the situation may change, if other parameter selection strategies are used, or the datasets are too large.

5.1 Comparing classification performance

Comparing the standard SVM optimization problem with the version in (1), which does not have an offset, the first question is probably, whether the absence of the offset term has an influence on the classification performance. To answer this question we performed on each data set 100 runs for both a version of the 2D-SVM-solver and our implementation of LIBSVM’s solver. We performed these experiments, though we report them first, actually at the very end of our investigations. This way, we could use for each solver the fastest version. For the 2D-SVM-solver it turned out, as we will see below, that this is the WSS 7 strategy together with l1-W4 initialization, while for the LIBSVM’s solver we used, depending on the dataset, either l1-W2 or l1-W5 initialization. Besides for the datasets of the AXA and

WXA families, we generated for each dataset 100 random splits, where each training set contained, modulo randomness, 70% of the samples. Moreover, on each of these training sets the hyper-parameter selection was performed by 10 fold cross-validation over the parameter grid described above. The test error was then computed on the test set part of the split, which, modulo randomness, contained 30% of the samples. The resulting average classification errors are reported in Table 1. As one quickly observes, LIBSVM yielded the better classification performance on the dataset IONOSPHERE and the two larger versions of the AXA family. On all other datasets, however, both algorithms performed almost indistinguishable. Therefore, it seems fair to conclude that the classification performance is not significantly influenced by the absence of the offset.

5.2 Comparisons to the optimal 2D subset selection strategy

In our first set of experiments on 2D subset selection strategies, we investigated the number of iterations needed for the different strategies of selecting working sets. Our baselines were the 1D-SVM-solver and the optimal 2D-SVM strategy WSS 0. Since the latter is computationally very expensive we decided to use only half of the samples of each training set for actual training. Besides that, we followed the approach of k -fold cross validation with $k = 2$ outlined earlier. Finally, in all experiments of this subsection, we initialized by $\alpha \leftarrow 0$.

Let us now have a closer look at the results that are displayed in Figures 1 to 4. Figure 1 compares the 1D-SVM, WSS 0 and the simple 2D-modifications of the 1D-SVM. Not

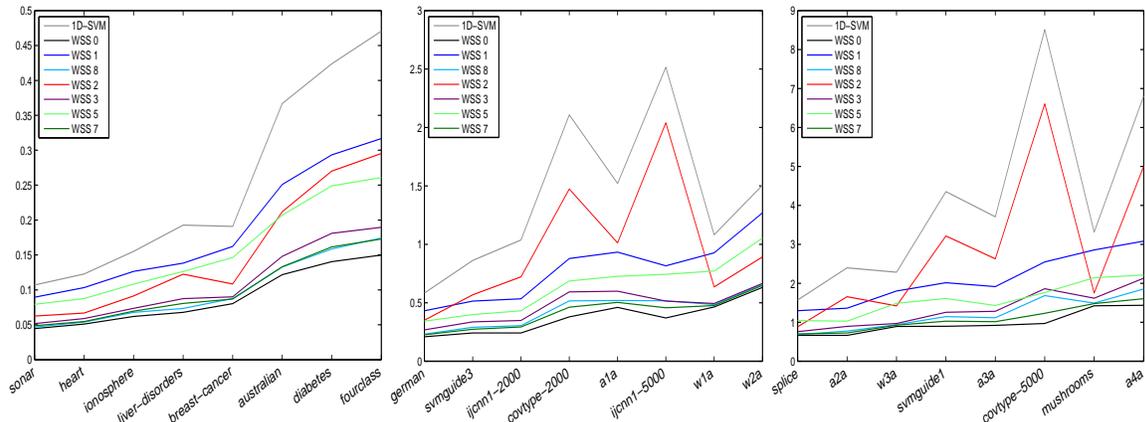


Figure 1: Performance of methods based on simple extensions of the 1D-search strategy for small (left), mid-sized (middle), and relatively large datasets (right). The graphic displays the average number of iterations in thousands for the different methods over the entire 10 by 10 parameter grid. All 2D-methods perform better than the 1D-SVM (gray), but the amount of improvement differs significantly. WSS 2 (red) performs sometimes better and sometimes substantially worse than WSS 1 (blue), but combining both methods into WSS 3 (dark magenta) leads to uniform improvements. The same holds for WSS 5 (light green), though with less improvements. The combination WSS 7 (dark green) uniformly yields the lowest number of iterations.

surprisingly, WSS 0 needs substantially less iterations than its one-dimensional equivalent 1D-SVM, while all of the simple 2D-modifications perform somewhere in between. More precisely, WSS 1 yields some significant improvement over the 1D-SVM. For WSS 2 the message is mixed; while on some datasets, WSS 2 performs significantly better, the difference is more marginal on other datasets. However, combining WSS 1 and WSS 2 into WSS 3 yields a clear overall improvement over both methods and the 1D-SVM. Another combination, WSS

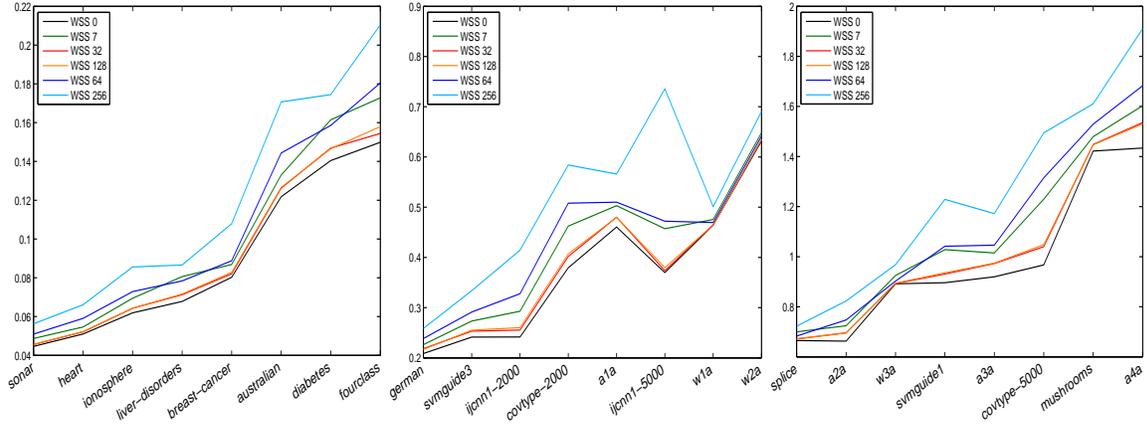


Figure 2: Performance of methods based on approximations of the optimal 2D-search strategy. The graphic displays the average number of iterations in thousands for the different methods over the entire 10 by 10 parameter grid. WSS 0 (black) performs uniformly best, but both deterministic strategies WSS 32 (red) and WSS 128 (orange), which are basically indistinguishable, closely follow the performance of WSS 0. WSS 7 (dark green) and the hybrid WSS 64 (dark blue) still captures most of the behavior of the previous methods with small advances for WSS 7, while the complete randomization (light blue) performs uniformly worst.

5 that combines WSS 1 with a search over 10 nearest neighbors, also needs substantially less iterations than WSS 1 and the 1D-SVM, but the improvements are less than those of WSS 3. However, the combination of all, WSS 7, does not only perform uniformly better than all participating methods, but also needs in most cases only a few more iterations than the optimal WSS 0. Finally, WSS 8, which is a variant of WSS 1, also reduces the number of iterations substantially, yet it fails to perform as well as WSS 7. Let us now have a closer

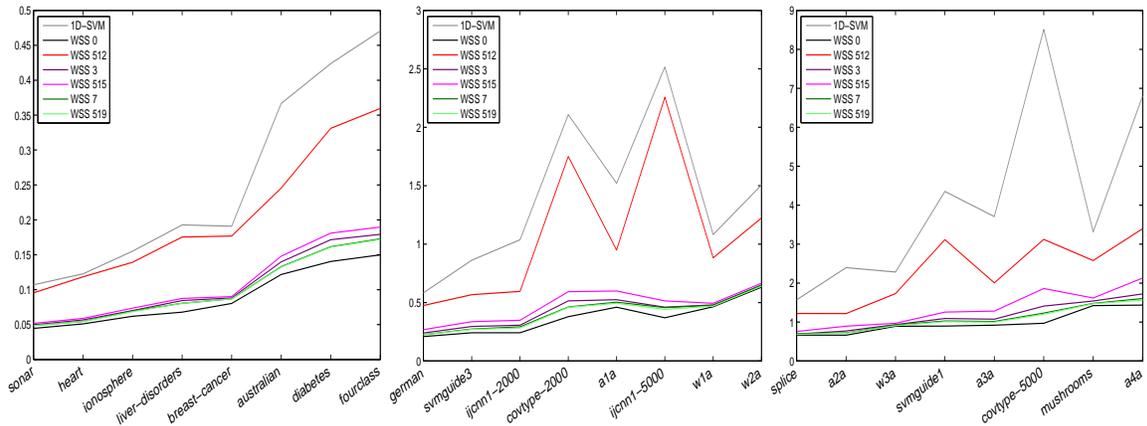


Figure 3: Combining methods based on simple 1D-extensions with the approximate gain on inner SVs. The graphic displays the average number of iterations in thousands for the different methods over the entire 10 by 10 parameter grid. Without combining WSS 512 (red) with other methods, it performs quite poorly, while combining WSS 512 with WSS 3 (dark magenta) to WSS 515 (light magenta) yields an improvement over both methods. In contrast, combining WSS 512 and WSS 7 to WSS 519 (light green) does not give an improvement over WSS 7 (dark green) as the almost indistinguishable two green lines show.

look at Figure 2 that shows how the methods based on an approximation of the optimal WSS 0 perform. Here it turns that WSS 32, which uses the exact computation of the 2D-

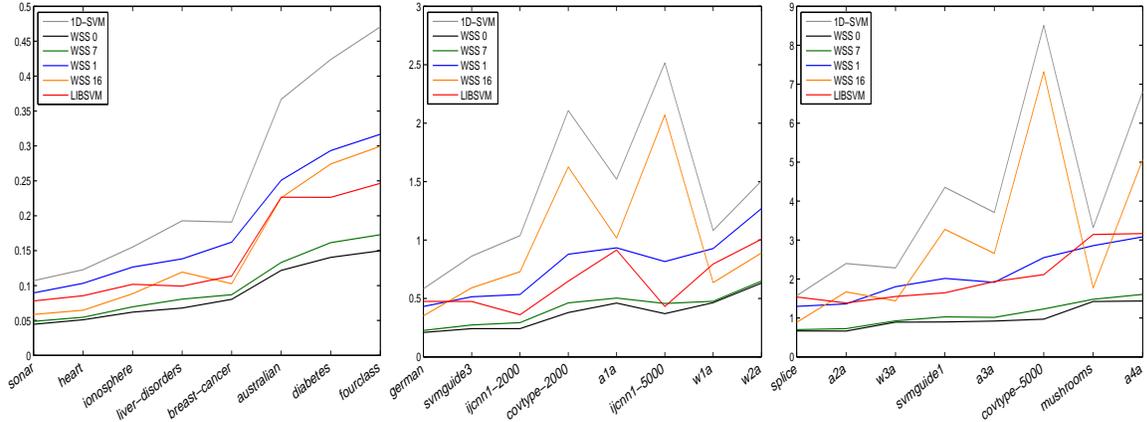


Figure 4: LIBSVM and MVP compared to some other approaches. The graphic displays the average number of iterations in thousands for the different methods over the entire 10 by 10 parameter grid. On all datasets considered, the 2D-MVP strategy WSS 16 (orange) has some advantage over the simple 1D-SVM (gray), while LIBSVM (red) often needs substantially less iterations and performs comparably to the WSS 1 (blue). However, neither of the methods approach the close-to-optimal performance of WSS 7 (dark green) or even the optimal performance of WSS 0 (black).

gain, and WSS 128, which uses an approximation of the 2D-gain, perform indistinguishably. In addition, they only need a few more iterations than WSS 0, and constantly outperform WSS 7, yet the latter improvement is in most cases only marginal. Finally, the random approaches WSS 64 and WSS 256 do not need less iterations than WSS 7, and the complete random approach of WSS 256 performs worse than the hybrid strategy of WSS 64. However, by comparing with Figure 1 we see that WSS 256 still needs significantly less iterations than the 1D-SVM.

Another way to approximately compute the 2D-gain is implemented in WSS 512. Figure 3 compares the number of iterations of this method to the 1D-SVM, WSS 0, and some combinations of WSS 512 with simple 2D-extensions of the 1D-SVM approach. A closer look at this figure shows that WSS 512 alone is not a very good alternative to the 1D-SVM, while combinations do yield significant improvement. However, these improvements are not significantly better than WSS 7.

Finally, let us compare the 1D-SVM and the optimal 2D strategy WSS 2 with the MVP approach of WSS 16 and LIBSVM. Figure 4 shows that the 2D-MVP approach of WSS 16 performs only slightly better than the 1D-SVM. In contrast, LIBSVM needs, not surprisingly, substantially less iterations than the 1D-SVM, but it fails to perform as well as the simple WSS 3, and the more complicated WSS 7.

5.3 Comparisons of different 2D subset selection strategies

The experiments of the previous subsection identified some working set selection strategies that perform close to the optimal WSS 0 in terms of iterations. All these strategies were of order $\mathcal{O}(n)$, yet it seems obvious, that their computational requirements in terms of runtime may substantially differ. Therefore, the goal of the experiments in this sections is to evaluate the selection strategies in terms of their runtime. To this end, we performed 10-fold cross validation training on our datasets, where the details of the cross validation procedure were

already outlined above. In the following, we not only report the runtime of the different strategies, but also their iterations. This way, it becomes easier to judge whether a strategy suffers from its large number of iterations or only from its computational requirements for selecting the working set. We always report the average requirements per grid point, where the average is either taken with respect to all 10 folds and the grid, or just with respect to the 10 folds and the grid points whose validation error is close to the minimal validation error. The latter averages are of particular interest, if one does not use grid search for the hyper-parameter selection, but some other methods, such as [10], which may find a good hyper-parameter pair faster. In addition, the latter averages are also interesting for grid search since after such a search one usually retrains the SVM on the entire training set with the hyper-parameters that performed best in terms of validation error.

Let us now have a closer look at the results. The first observation from Figures 5 and 6 is that WSS 2, which needs less iterations than the 1D-SVM, does not run substantially faster. However, this behavior can be relatively easily explained by the fact that in each iteration the 1D-SVM updates the gradients for one direction only, whereas WSS 2, due to its 2D-nature, performs two such updates per iteration. Similarly, WSS 8 cannot translate its advantage over WSS 1 in terms of iterations into a substantial advantage in terms of runtime. In this case, a closer look reveals that, compared to WSS 1, WSS 8 performs an additional, implicit gradient upgrade when looking for the second direction j . The other results displayed in Figures 5 and 6 confirm our results from Figure 1. In particular, WSS 7 not only need the fewest number of iterations, but also runs fastest on almost all datasets. Finally, Figure 7 reveals, where some combined methods achieve their speed-up compared to WSS 1. In particular, for large values of λ , WSS 3 and WSS 7 need only half of the iterations of WSS 1 and WSS 5, which indicates that in this regime, WSS 2 is the dominating strategy in the former two combinations. On the other hand, for small values of λ , the nearest neighbor strategy WSS 4 seems to be the dominating working set selection strategy of WSS 5 and WSS 7 since both methods need substantially less iterations than the methods WSS 1 and WSS 3, which do not include the nearest neighbor strategy. Finally, these advantages in terms of iterations do translate into almost the same advantage in terms of runtime, since the additional costs of the nearest neighbor strategy only depend on the number k of considered nearest neighbors, which, in general, is quite small compared to the sample size. Nonetheless it is worth mentioning that for a few hyper-parameter pairs, it is faster not to use the nearest neighbor strategy.

Let us now turn to Figures 8 and 9 that consider the methods that try to approximate the working set strategy of the optimal WSS 0. Here it turns out, that WSS 32 and WSS 128, whose required number of iterations were closest to WSS 0, have a significant higher runtime than WSS 7. Since the number of iterations of these three methods behave quite similarly, the only explanation for this different runtime behavior is the additional cost per iteration for computing all (approximate) 2D-gains. This explanation is further confirmed by the fact that WSS 128, which involves the cheaper approximate 2D-gain has a better runtime behavior than WSS 32, which uses the exact computation of the 2D-gain. Furthermore, WSS 64, which computes only a fifth of the 2D-gains WSS 32 computes, runs substantially faster than WSS 32, despite the fact the the former needs more iterations. In this direction we finally note that WSS 256 runs overproportionally slowly compared to, e.g., WSS 128. Most likely, this behavior can be explained by less effective hardware caching for the random pair selection of WSS 256. To get a better impression, on how effective WSS 7 chooses its working sets, let us now have a closer look at the number of iterations of the different working set selection strategies. The bottom graphics of Figure 8 show that over the entire grid, WSS

7 only needs 5% to 20% more iterations than the best performing WSS 32. However, if one considers only the grid points with small validation error, this good behavior becomes worse. Indeed, the bottom graphics of Figure 9 show that for such hyper-parameters, WSS 7 typically needs more than 20% more iterations than WSS 32, and in some cases even more than 50% more. Finally, Figure 10 reveals that in particular for small values of λ and flat kernels, WSS 7 requires substantially more iterations than WSS 32. However, at least on the dataset SVMGUIDE1 this worse behavior takes place at grid points that do not need a lot of iterations anyway, and hence the advantage of WSS 32 is marginal. The next question, which naturally arises from the observations above, is whether the number of iterations used in WSS 7 can be reduced by combining WSS 7 with some methods that mimic WSS 32 on the inner support vectors. Here, Figure 8 shows that the number of iterations can be reduced by such combinations in a few cases, but this never pays off in terms of runtime if one considers the entire grid. On the grid points with small validation error, however, the situation is slightly more involved. Clearly, the combination with WSS 2048 performs worst, yet combining WSS 7 with WSS 512 or WSS 1024 sometimes yield a shorter runtime. Finally, Figure 13 shows that, at least for the dataset SVMGUIDE1, the improvements achieved by these combinations are mainly at grid points that do not require a lot of iterations. On the other hand, it also illustrates that the computational overhead of these combinations is significant.

The last figures of this subsection, Figures 14 to 16, compare LIBSVM with some subset selection strategies such as the MVP approach of WSS 16 and the overall best performing WSS 7. Here the most interesting observation is that although WSS 1 and LIBSVM have comparable behavior in terms of iterations, they substantially differ in runtime. Because we used our own implementation of LIBSVM’s solver, which employed the same SSE2 optimizations as the 2D-SVM methods, the only way to explain this behavior is that the subset selection strategy of LIBSVM is significantly more expensive than the simple WSS 1. To understand the latter, recall that LIBSVM’s strategy is based on computing an approximate 2D-gain, which is quite expensive as we have seen in Figures 8 and 9 for the 2D-SVM methods WSS 32, WSS 64, WSS 128, and WSS 256. In addition, LIBSVM’s strategy cannot be efficiently vectorized, which is another disadvantage compared to WSS 1. Finally, it is interesting to note that WSS 7 is between 2 and 4 times faster than LIBSVM, when the average over all grid points is considered. Moreover, on the grid points with small cross validation error the improvement is rarely less than by a factor of 4, and as Figure 16 illustrates, this is most likely not an artefact caused by different optimal grid points. Indeed, on some grid points LIBSVM needs more than 10 times the run time WSS 7 requires.

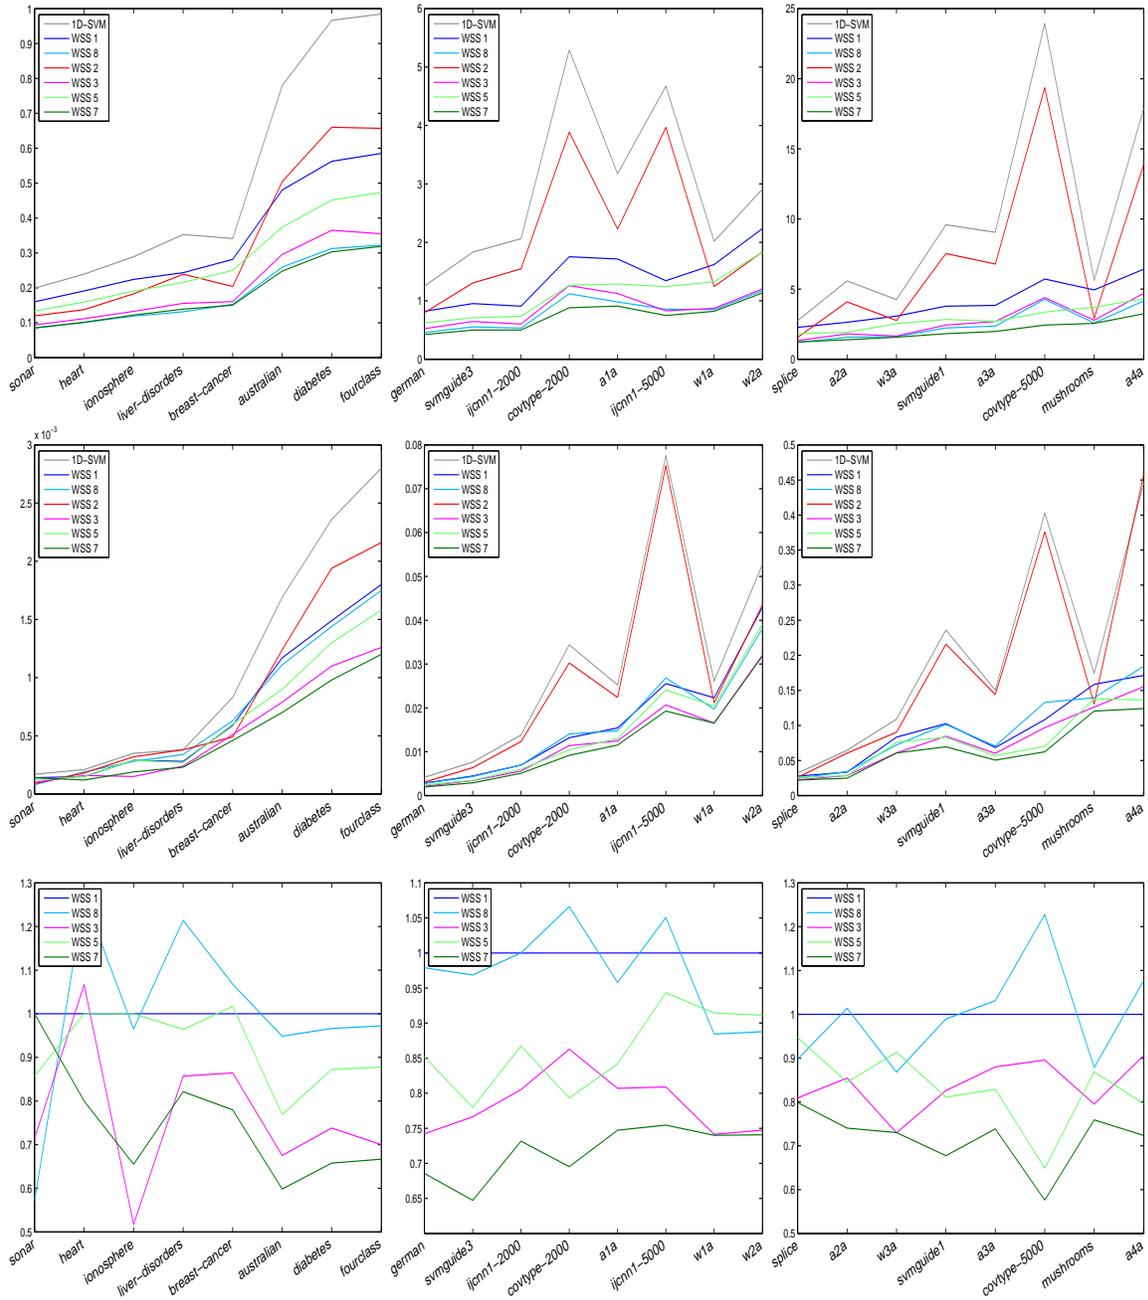


Figure 5: Average computational requirements per grid point of methods based on simple extensions of the 1D-search strategy for small (left), mid-sized (middle), and relatively large datasets (right) over the entire 10 by 10 grid. The graphics display the number of iterations in thousands (top), the runtime in seconds (middle), and the ratios WSS_x/WSS_1 of the runtimes (bottom). WSS 7 (dark green) performs almost uniformly the best in both metrics, followed by WSS 3 (dark magenta) and WSS 5 (light green) in terms of runtime.

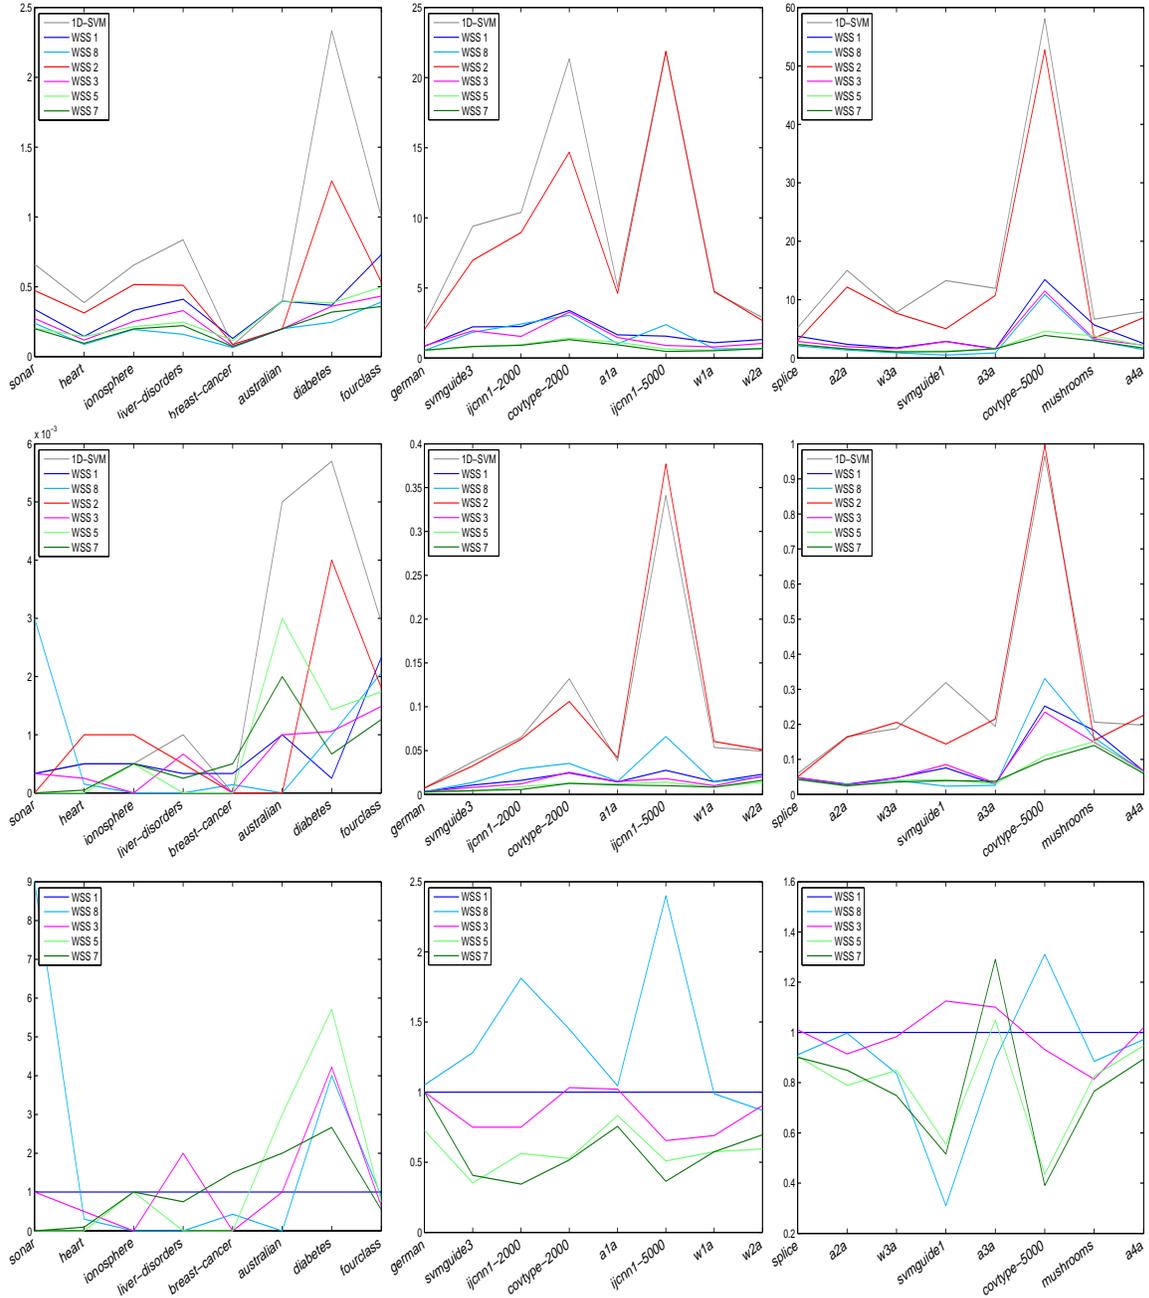


Figure 6: Computational requirements of methods based on simple extensions of the 1D-search strategy on the grid points whose cross validation error is not larger than 1.05 the minimal cross validation error. The graphics display the average number of iterations in thousands (top), the runtime in seconds (middle), and the ratios WSS_x / WSS_1 of the runtimes (bottom). For the small datasets, the runtime measurements are not very reliable. In addition, the set of considered grid points may slightly vary for the different methods, which in turn may influence the computational requirements and hence the graphic at the bottom has little informative value. It seems fair to say that overall, WSS 7 performs best in both metrics, but is closely followed by WSS 5 in terms of runtime. Moreover, on some datasets, such as SVMGUIDE1, WSS 7 performs worse than WSS 8. However, comparing this with Figure 7 shows that this behavior is due to different grid points with small cross validation error.

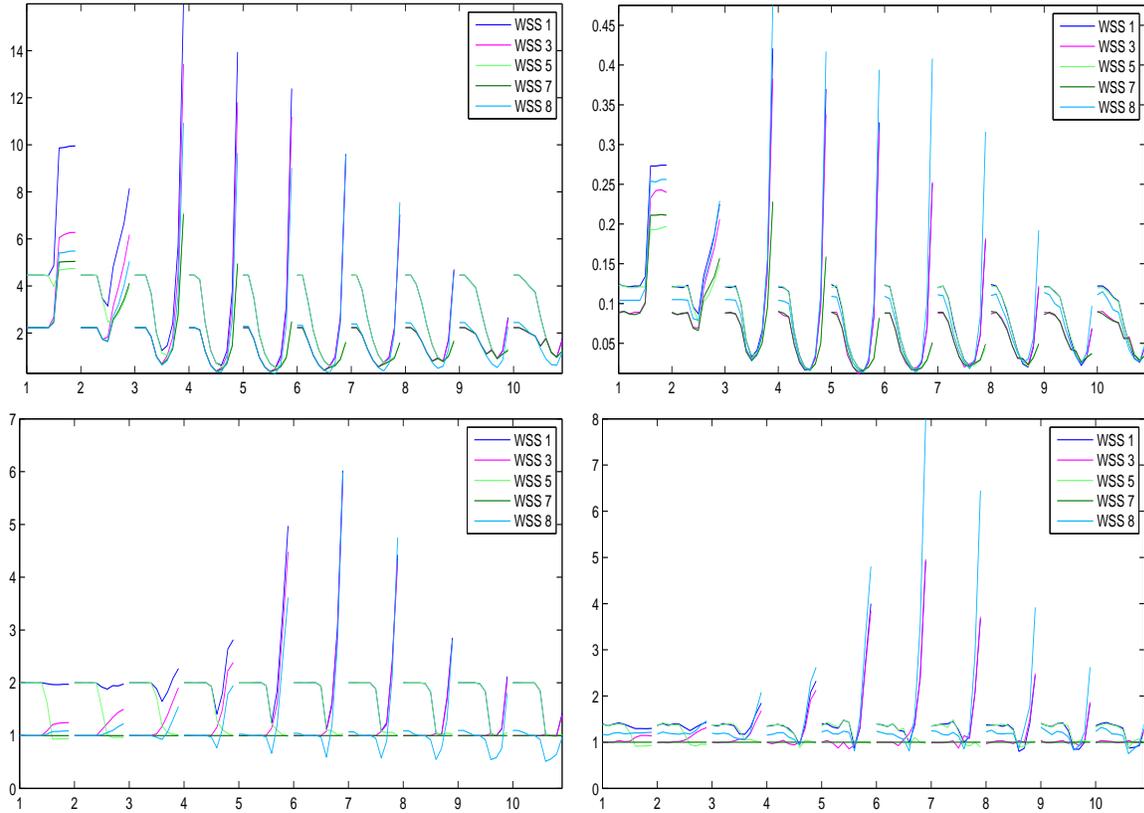


Figure 7: Computational requirements per single grid point of methods based on simple extensions of the 1D-search strategy for the svmGUIDE1 dataset. Each horizontal cell numbered by 1 to 10 corresponds to a single kernel parameter σ and an ordered run through the 10 λ -values, where the left of each cell corresponds to the largest λ -value, and the right to the smallest. Analogously, cell 1 corresponds to the largest σ -value, and cell 10 on the right corresponds to the smallest σ -value. The graphics at the top display the number of iterations in thousands (left) and the runtime in seconds (right), both averaged over the 10 folds, for WSS 1 (blue), WSS 3 (dark magenta), WSS 5 (green), WSS 7 (dark green), and WSS 8 (light blue). WSS 7 performs almost uniformly the best in both metrics. However, for large λ , WSS 3 performs comparable, while for small λ , WSS 7 is closely followed by WSS 5. The graphics at the bottom show the ratios $\text{WSS } x / \text{WSS } 7$, $x = 1, 3, 5, 7$, for the number of iterations (left) and the runtime (right) to illustrate the performance gain of WSS 7.

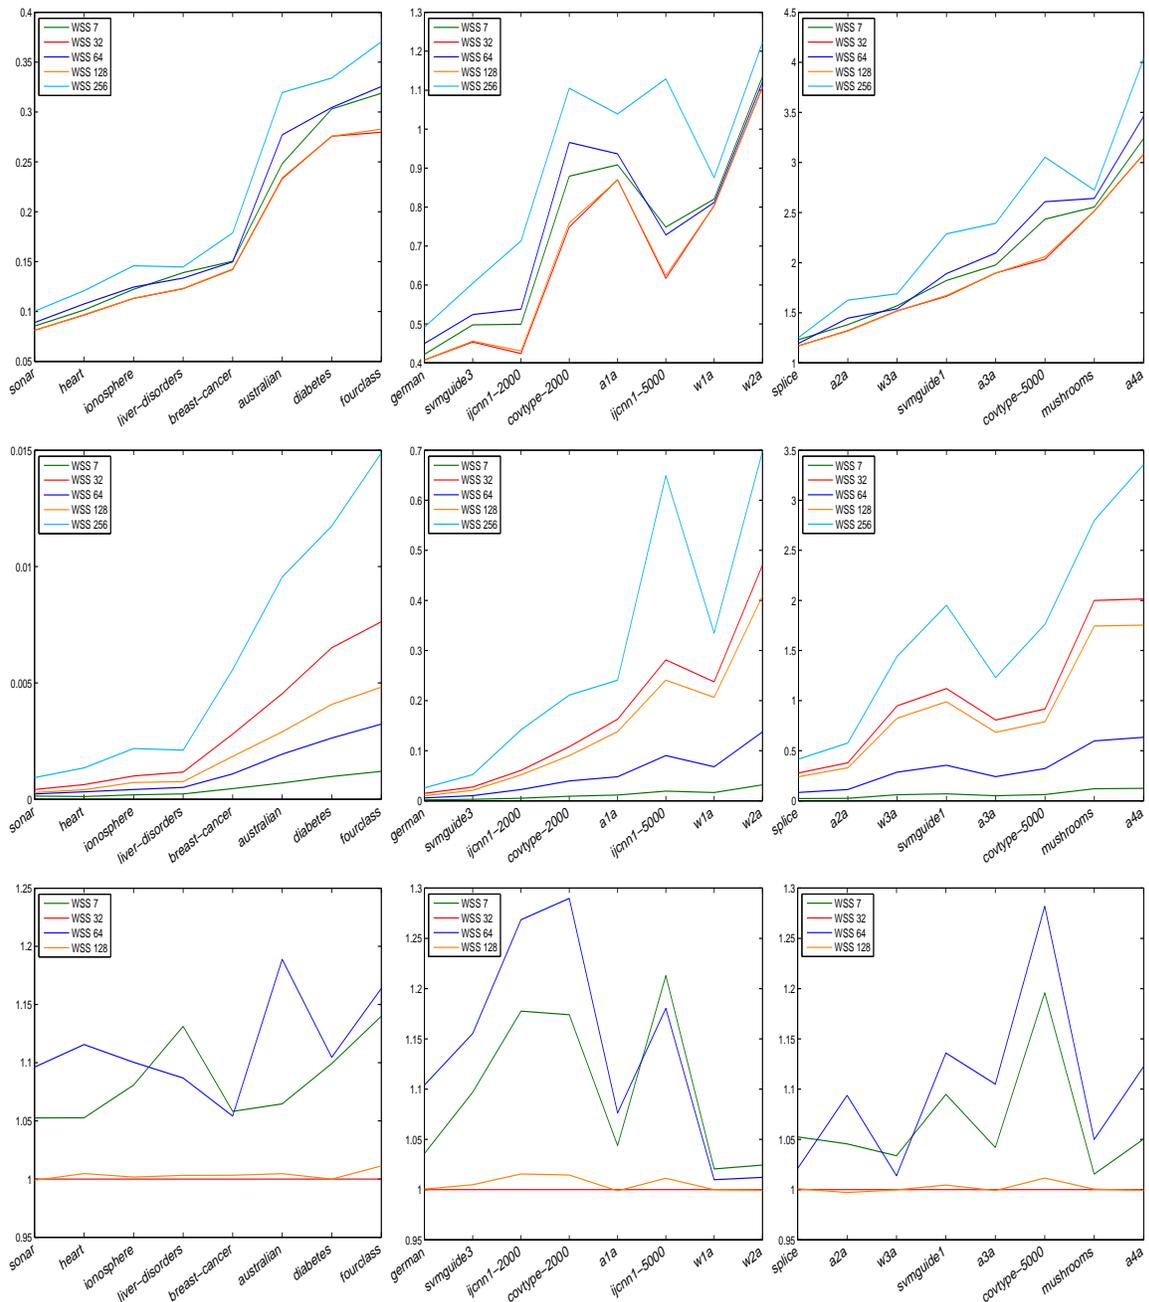


Figure 8: Average computational requirements per grid point of methods based on approximations of the optimal 2D-strategy. The graphics display the number of iterations in thousands (top), the runtime in seconds (middle), and the ratios $WSS \times / WSS 32$ of the number of iterations (bottom). Although WSS 7 (dark green) and the semi-random WSS 64 (blue) need slightly more iterations than WSS 32 (red) and WSS 128 (orange), their costs per iteration is substantially less, which results in a significantly shorter runtime. The completely random WSS 256 (light blue) needs over-proportionally more runtime, possibly because of the less effective hardware cache.

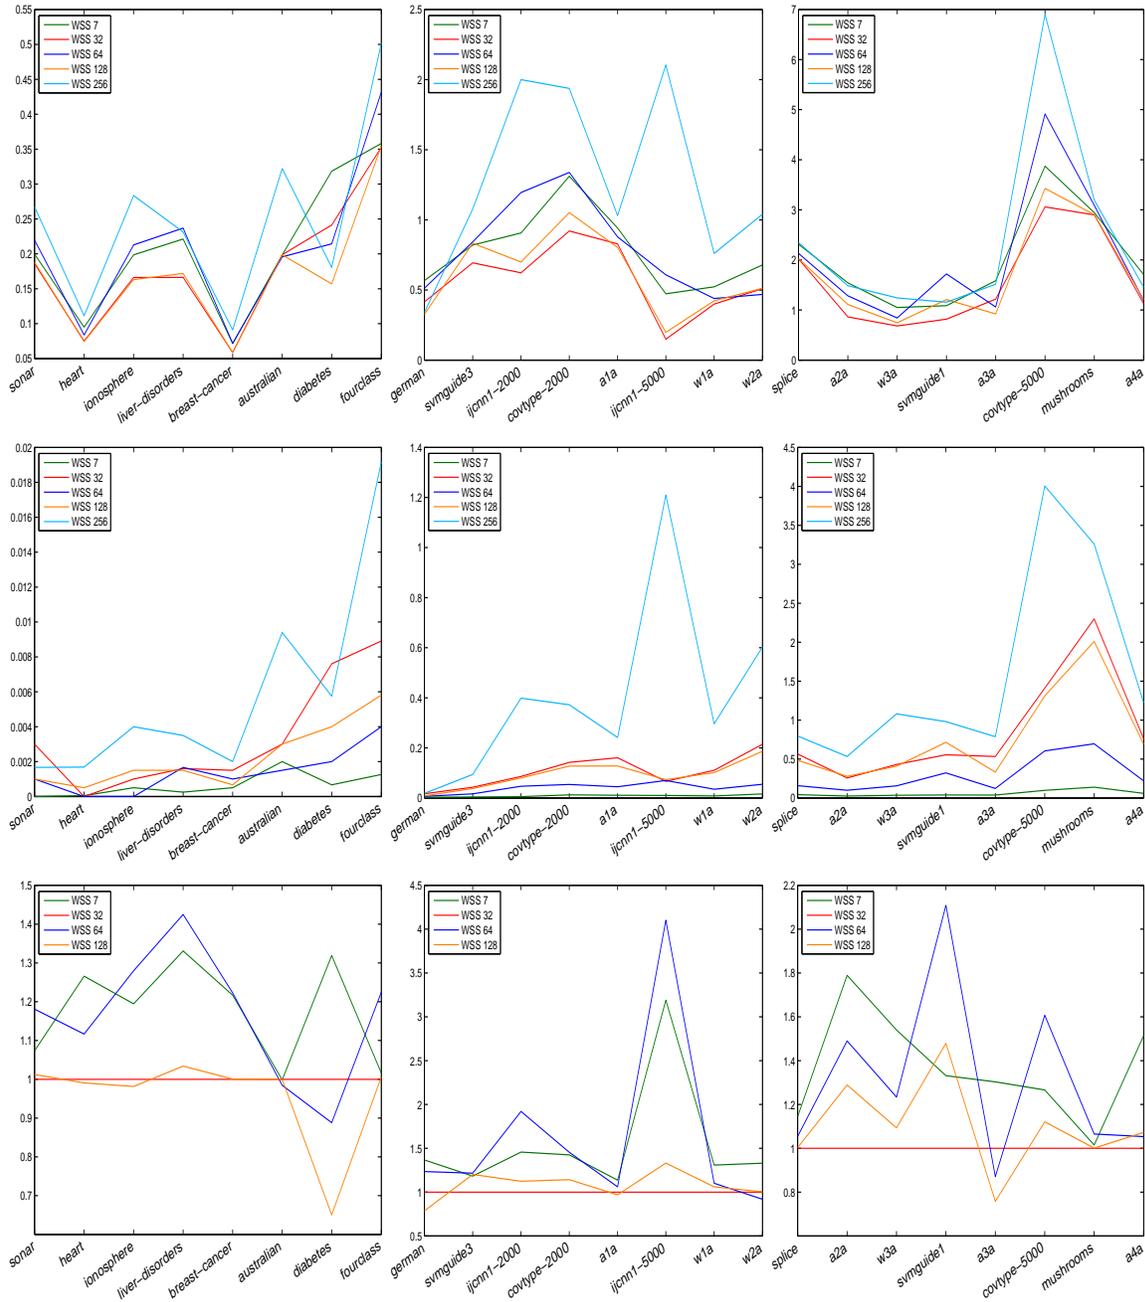


Figure 9: Computational requirements of methods based on approximations of the optimal 2D-strategy on the grid points whose cross validation error is not larger than 1.05 the minimal cross validation error. The graphics display the average number of iterations in thousands (top), the runtime in seconds (middle), and the ratios WSS_x/WSS_{32} of the number of iterations (bottom). For the small datasets, the runtime measurements are not very reliable. In addition, the set of considered grid points may slightly vary for the different methods, which in turn may influence the computational requirements.

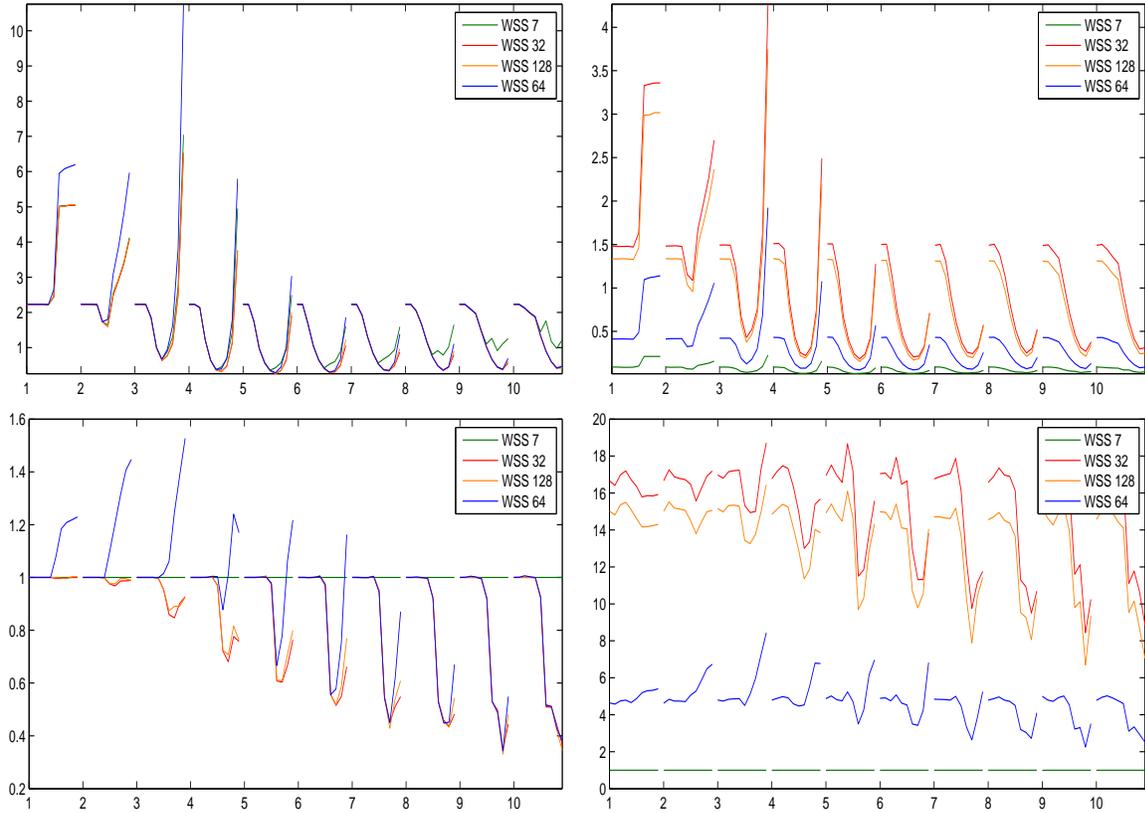


Figure 10: Computational requirements per single grid point of methods based on approximations of the optimal 2D-strategy for the SVMGUIDE1 dataset. The four graphics have the same format as the ones in Figure 7. The graphics at the top display the number of iterations in thousands (left) and the runtime in seconds (right), both averaged over the 10 folds, while the graphics at the bottom display the corresponding ratios $WSS\ x / WSS\ 7$. For some grid points, WSS 7 and WSS 32 need approximately the same number of iterations, while for some other grid points, WSS 7 needs significantly more. Nonetheless, the runtime behavior of WSS 32 is substantially worse than that of WSS 7.

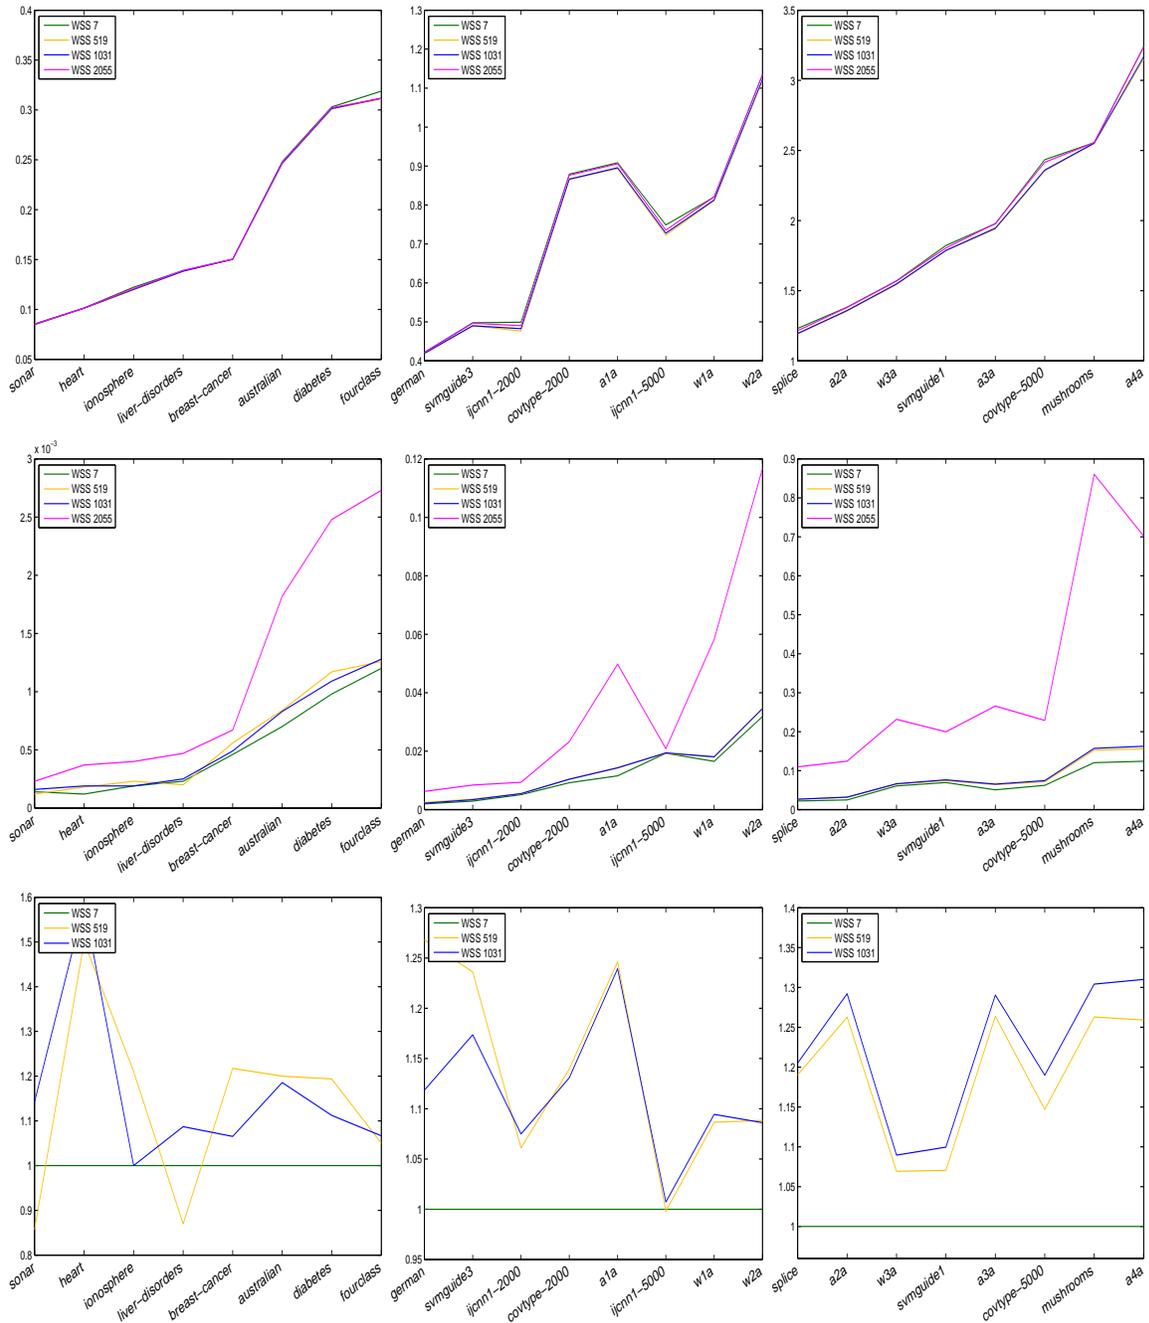


Figure 11: Average computational requirements per grid point of combining WSS 7 (dark green) with some methods that use the formula for the approximate gain on inner SVs over the entire 10 by 10 grid. The graphics display the number of iterations in thousands (top), the runtime in seconds (middle), and the ratios $WSS \times / WSS 7$ of the runtimes (bottom). Although the combinations need a slightly smaller number of iterations, their additional overhead per iteration leads to worse runtime behavior.

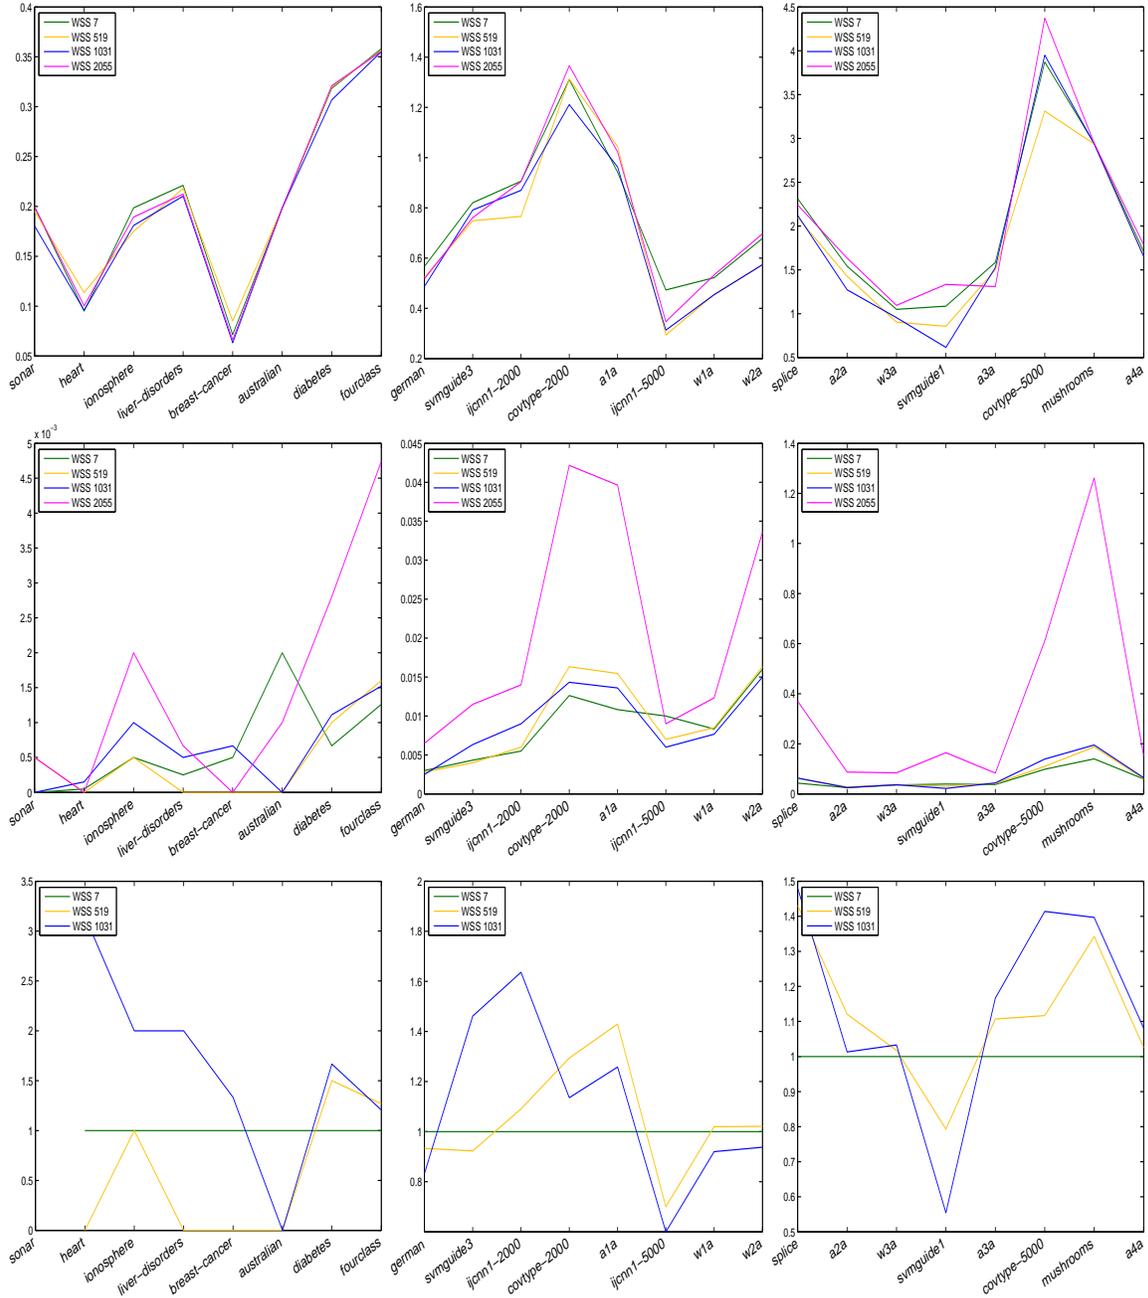


Figure 12: Computational requirements of combining WSS 7 with some methods that use the formula for the approximate gain on inner SVs on the grid points whose cross validation error is not larger than 1.05 the minimal cross validation error. The graphics display the number of iterations in thousands (top), the runtime in seconds (middle), and the ratios $WSS \times / WSS 7$ of the runtimes (bottom). For the small datasets, the runtime measurements are not very reliable. In addition, the set of considered grid points may vary slightly for the different methods, which in turn may influence the computational requirements. As in Figure 11 some of the combinations need a slightly smaller number of iterations, but their additional overhead per iteration leads to worse runtime behavior on most of the medium-sized and large datasets.

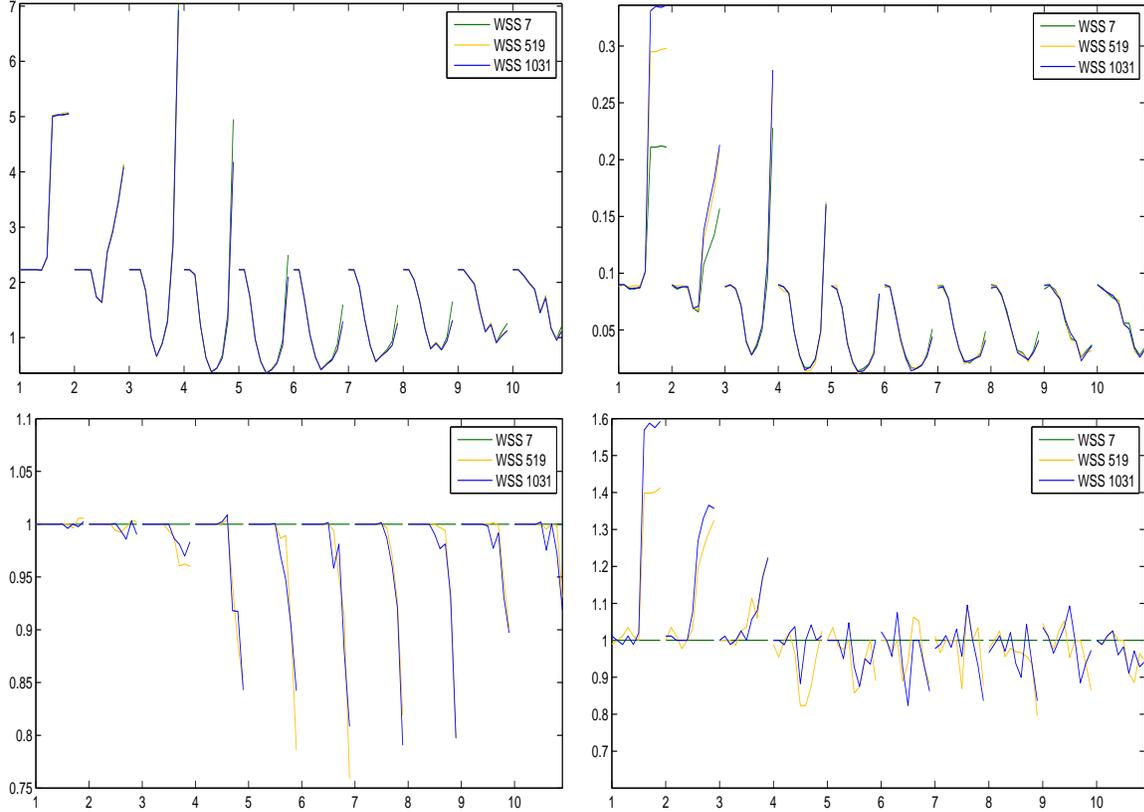


Figure 13: Computational requirements per single grid point of methods based on simple extensions of the 1D-search strategy for the SVMGUIDE1 dataset. The four graphics have the same format as the ones in Figure 7. The graphics at the top display the number of iterations in thousands (left) and the runtime in seconds (right), both averaged over the 10 folds, while the graphics at the bottom display the corresponding ratios WSS_x/WSS_7 . Note that for large λ the Boolean flag of WSS 4 is typically not set to true during the optimization, and hence all methods reduce to WSS 3. Analogously, for large λ and σ , the graphics nicely display the additional costs of WSS 512 and WSS 1024. Finally, the differences in the run time occur on a very low and hard to measure level, which explains the fluctuations in the bottom right graphics.

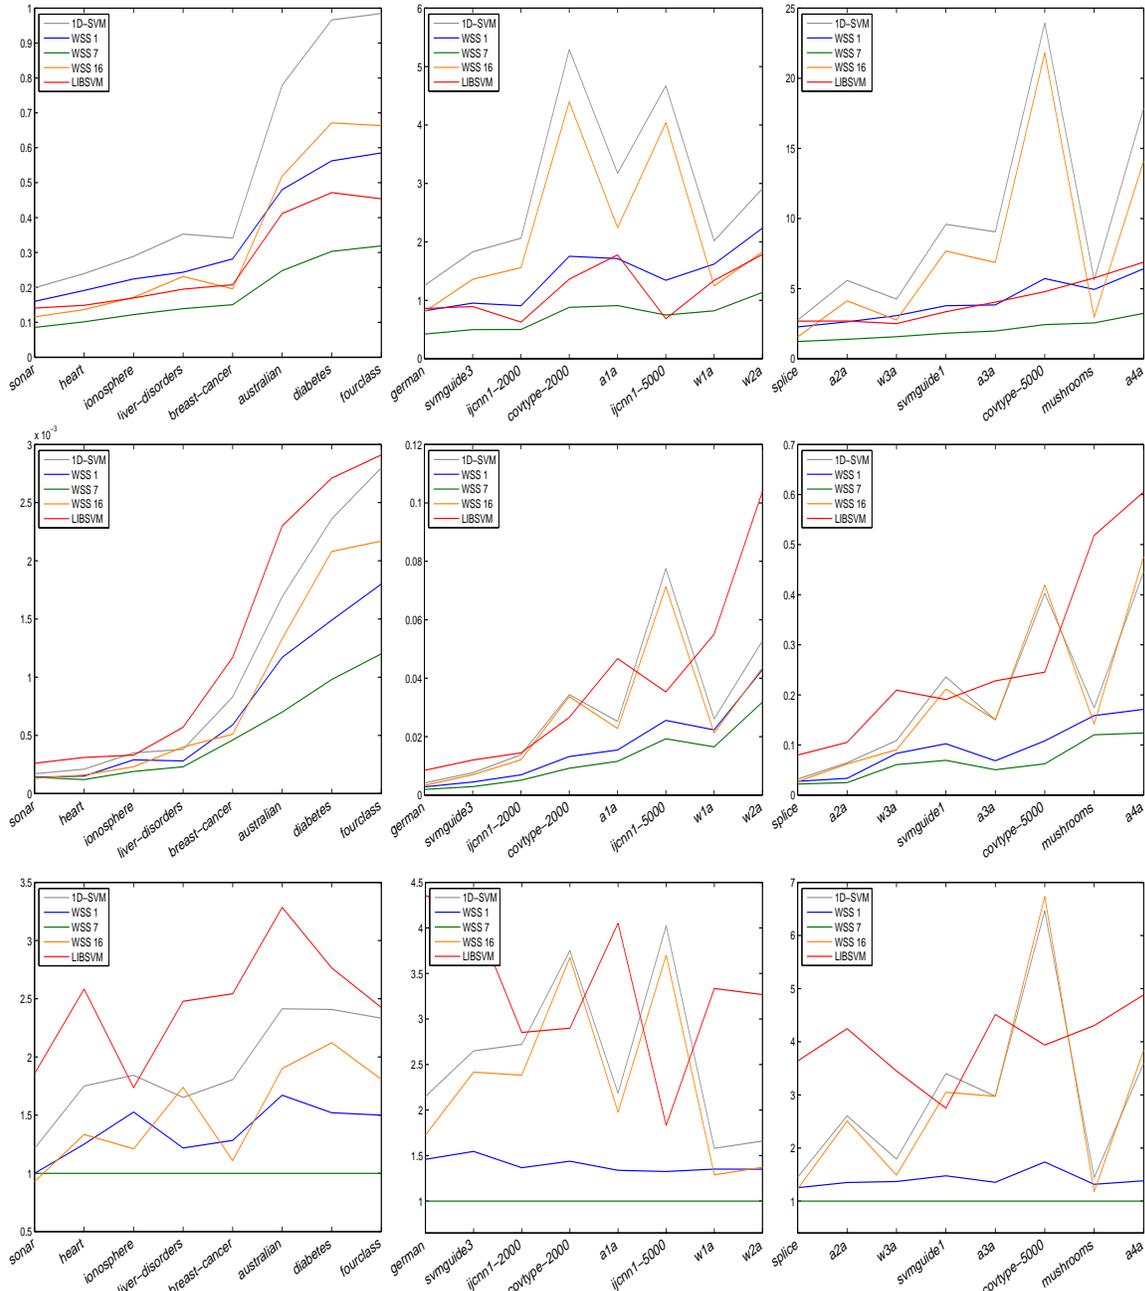


Figure 14: Average computational requirements per grid point of LIBSVM and MVP compared to some other approaches over the entire 10 by 10 grid. The graphics display the number of iterations in thousands (top), the runtime in seconds (middle), and the ratios \times /WSS 7 of the runtimes (bottom). The 2D-MVP approach of WSS 16 (orange) is not a good alternative to the 1D-SVM (gray) or even the two-dimensional WSS 7 (dark green). Moreover, although WSS 1 (blue) and LIBSVM (red) perform approximately the same number of iterations, their runtime is significantly different due to the more expensive working set strategy of LIBSVM (red).

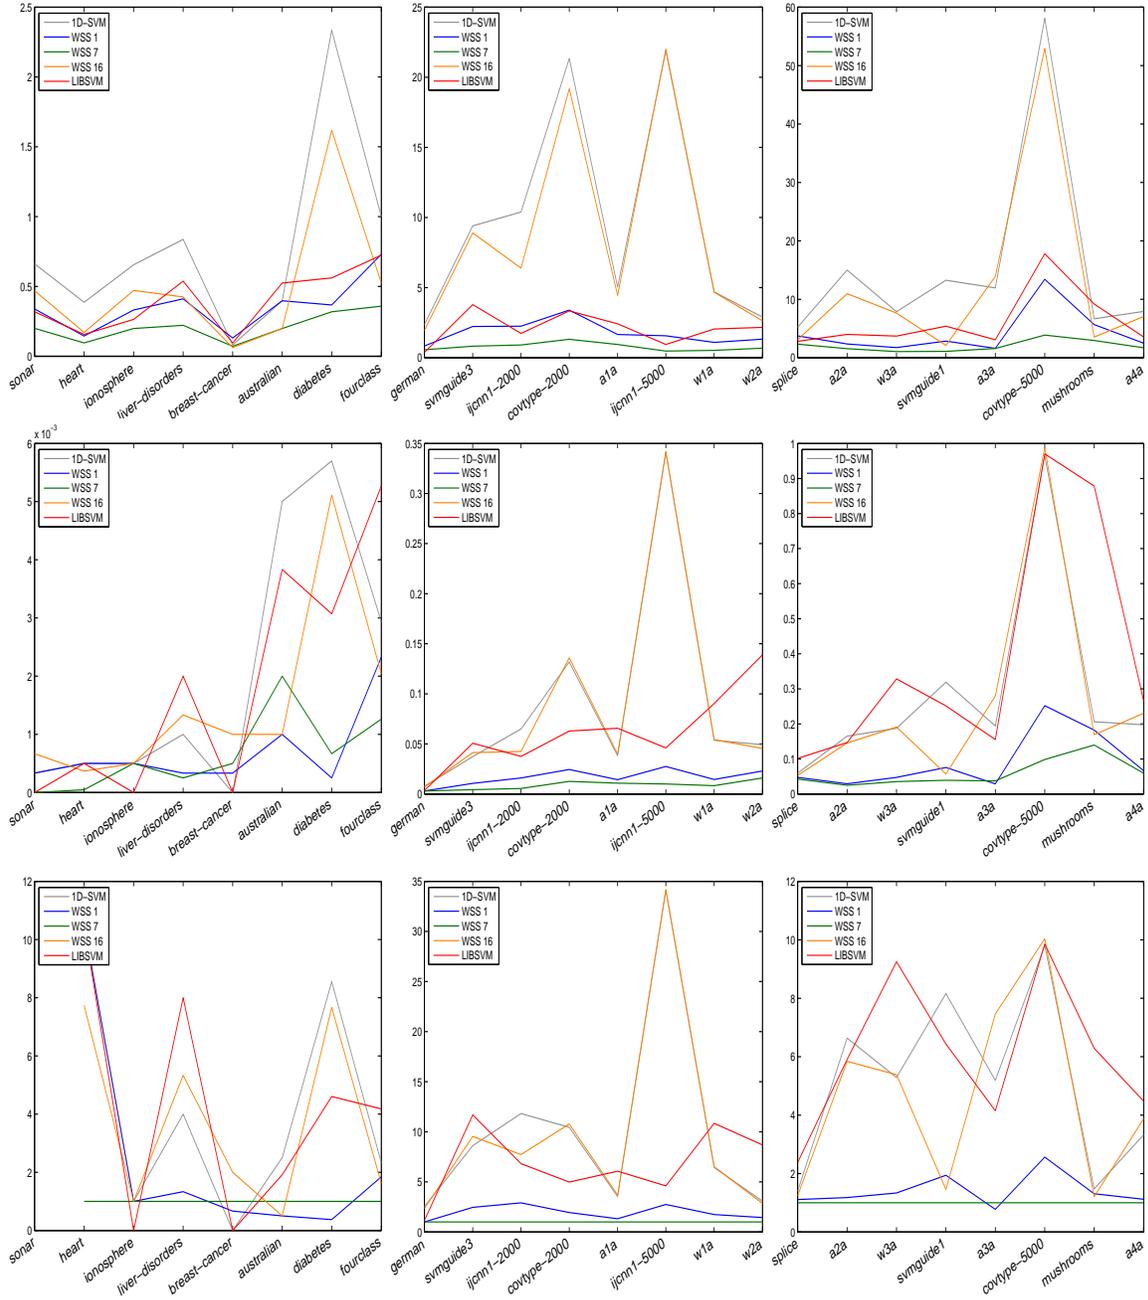


Figure 15: Computational requirements of LIBSVM and MVP compared to some other approaches on the grid points whose cross validation error is not larger than 1.05 the minimal cross validation error. The graphics display the average number of iterations in thousands (top), the runtime in seconds (middle), and the ratios \times /WSS 7 of the runtimes (bottom). Again, for the small datasets, the runtime measurements are not very reliable. In particular, for the SONAR dataset, the average *measured* runtime for WSS 7 was 0.00 seconds, and hence the corresponding ratios could not be plotted. Besides that the conclusions of Figure 14 are confirmed.

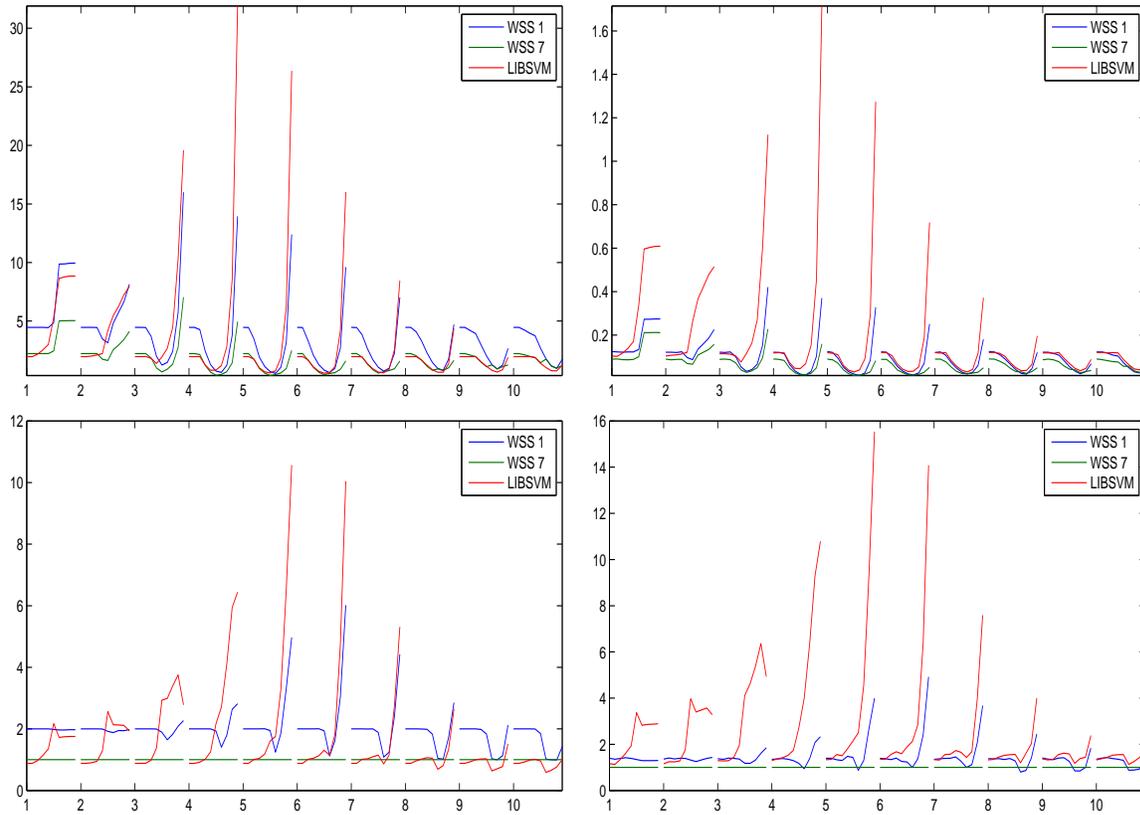


Figure 16: Computational requirements per single grid point of methods based on simple extensions of the 1D-search strategy and LIBSVM on the SVMGUIDE1 dataset. The four graphics have the same format as the ones in Figure 7. For flatter kernels, LIBSVM needs less iterations than WSS 7, possibly because it solves a different optimization problem, however the improvement is small in terms of absolute numbers. On the other hand, both WSS 1 and WSS 7 are less sensitive to small λ values in regions with high computational demand.

5.4 Influence of the stopping criterion

In this subsection, we investigate the influence of the stopping criterion (9) on the computational requirements. To this end, we considered the 10-fold cross validation procedure described earlier. Moreover, in order to save time, we only considered the best performing working set selection strategy, namely WSS 7. For this method we considered our stopping criterion (9) and the classical duality gap stopping criterion (7), where we set the right hand side of both stopping criteria to be $\epsilon/(2\lambda)$ with $\epsilon := 0.001$. Note this this is exactly the same set up as in our previous experiments, and it is not hard to show that for the duality gap (7), this choice again leads to the same theoretical bounds on the generalization performance.

The results of our experiments are summarized in Figures 17 to 19. A quick look shows that, not surprisingly, the stopping criterion (9) never leads to more iterations, but the improvements depend very much on the dataset. Moreover, these smaller number of iterations also pay off in terms of runtime, though the effect is less pronounced when we consider the entire grid. We believe this is due to the fact that computing (9) is a little more expensive than computing (2) since it involves two rather than just one clipping operations. In this regard, it is interesting to note that the SSE2 instruction set in `emmintrin.h` makes it possible to avoid expensive branches for the computation of the clipping by providing `min()` and `max()` operations. Moreover, when we only consider the grid points with small validation error, the positive effect of the clipped duality gap is amplified as Figure 18 shows. The reason for this behavior is illustrated in Figure 19 for the `SVMGUIDE1` dataset. Indeed, this figure shows that for small values of λ , the stopping criterion (9) leads to both substantially less iterations and shorter runtimes, whereas for larger λ , the computational requirements for both stopping criteria are essentially the same. In other words, although uniformly superior the effect of (9) is highly inhomogenously distributed over the parameter grid.

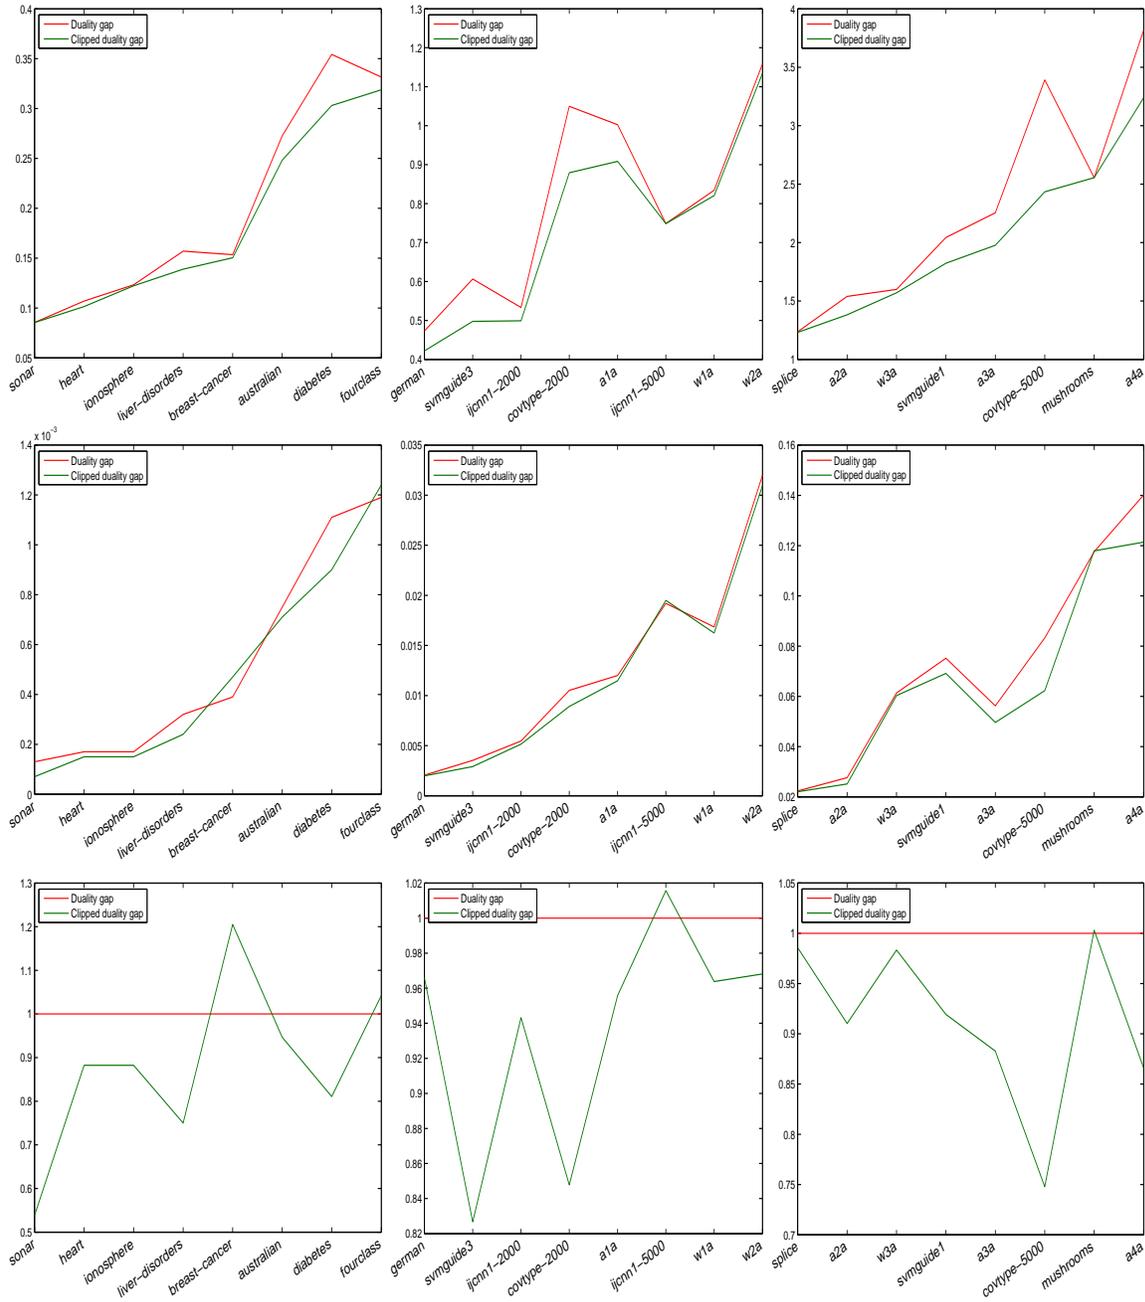


Figure 17: Average computational requirements per grid point of WSS 7 with different stopping criteria for small (left), mid-sized (middle), and relatively large datasets (right). The graphics at the top display the number of iterations in thousands for the different stopping criteria applied to the 2D-SVM with WSS 7, while the graphics in the middle show the corresponding runtime in seconds. The graphics at the bottom display the ratio of runtimes. Not surprisingly, the clipped duality gap (9) always leads to less iterations than the classical duality gap stopping criterion (7). Moreover, in most cases, this also results in a better runtime behavior.

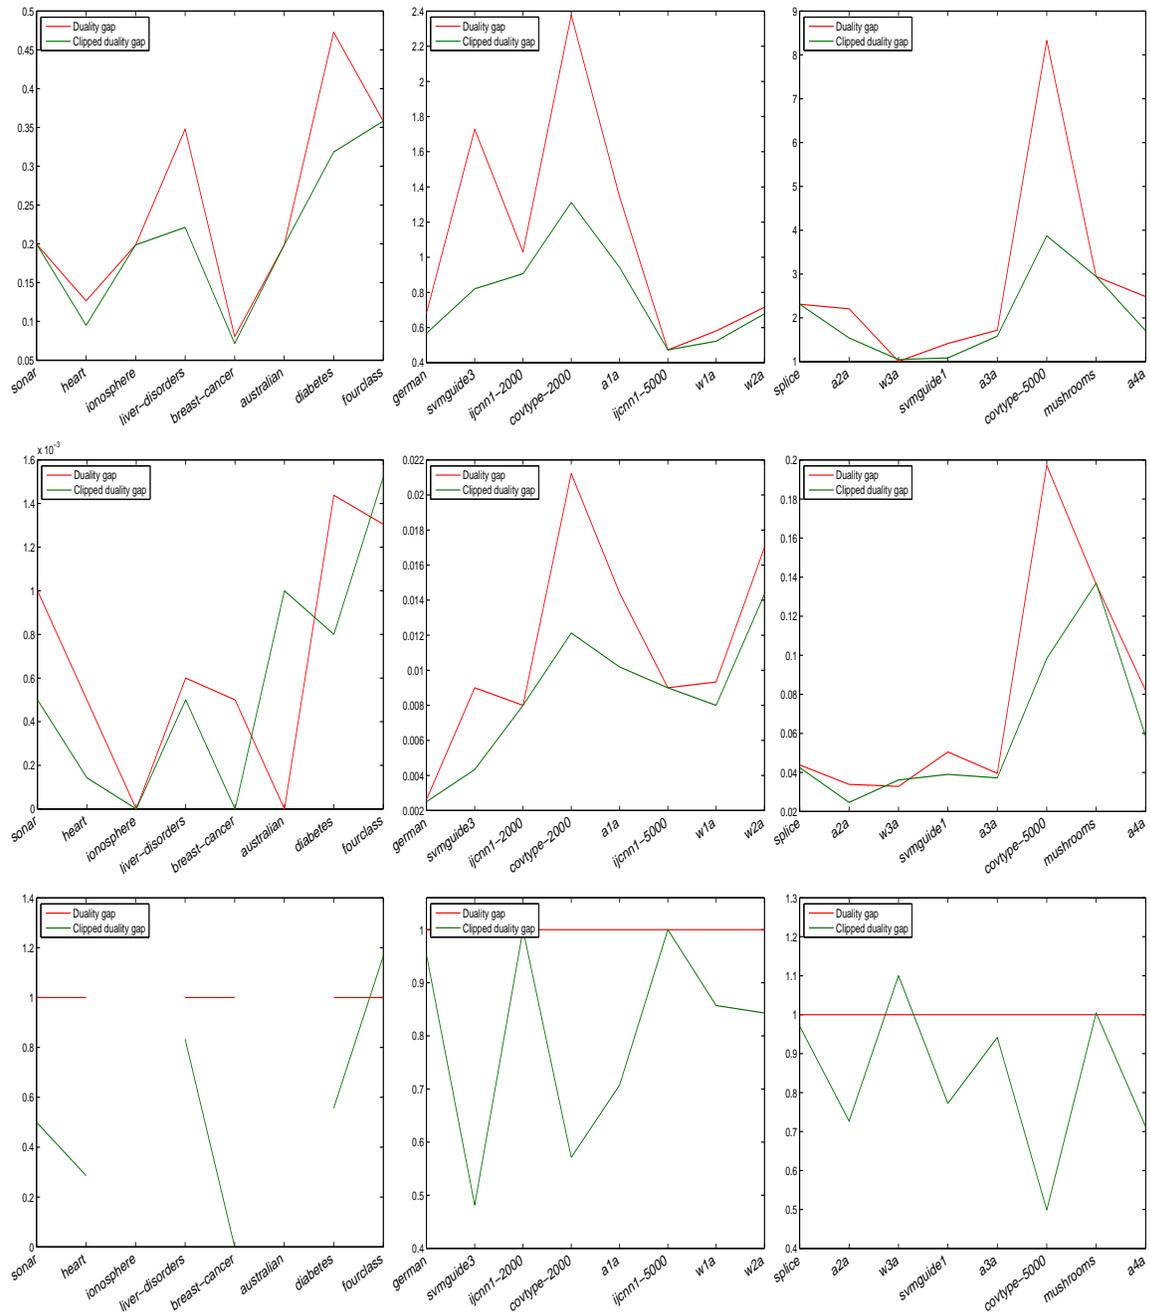


Figure 18: Computational requirements of WSS 7 with different stopping criteria on the grid points whose cross validation error is not larger than 1.05 the minimal cross validation error. Again, the graphics at the top display the number of iterations in thousands for the different stopping criteria applied to the 2D-SVM with WSS 7, while the graphics in the middle show the corresponding runtime in seconds. The graphics at the bottom display the ratio of runtimes, where we note that for some datasets in the bottom left graphic the ratio could not be computed since the *measured* run time was zero.

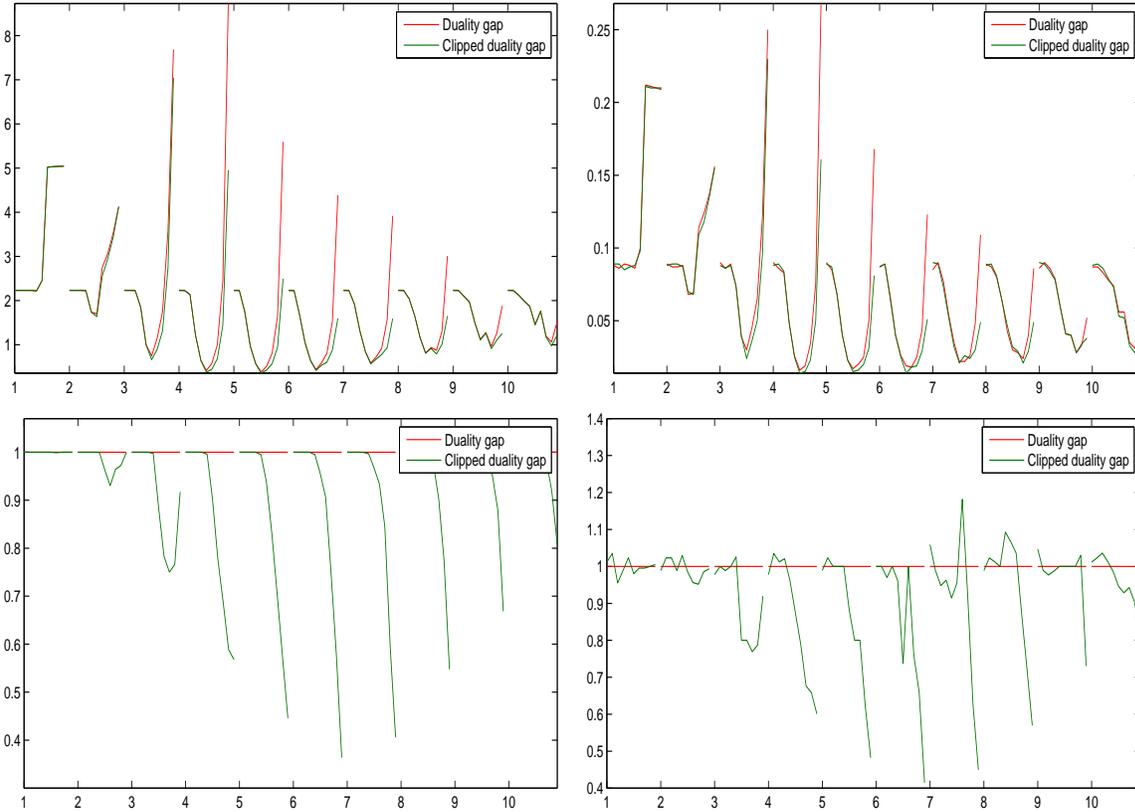


Figure 19: Computational requirements per single grid point of the two stopping criteria for the SVMGUIDE1 dataset. The four graphics have the same format as the ones in Figure 7. The graphics at the top display the number of iterations in thousands (left) and the runtime in seconds (right), both averaged over the 10 folds, while the graphics at the bottom display the corresponding ratios. The clipped stopping criteria (9) helps for small values of λ , whereas for larger values the behavior is basically identical. Again, some of the roughness in the bottom right graphic can be explained by the resolution of the time measurements. However, the general trend in this graphic is confirmed by the ratio of iterations displayed in the bottom left graphic.

5.5 Comparing different numbers of nearest neighbors

So far we have considered WSS 7 for 10 nearest neighbors only. Of course, this was a relatively arbitrary choice, and hence it is interesting to investigate how the computational requirements change with the number of nearest neighbors. This is the goal of this subsection.

To this end, we again used the 10 fold cross validation approach described earlier. We further considered the behavior of WSS 7 for N -nearest neighbors, where $N = 5, 10, 15, 20, 25, 30$. Our first observation was that for $N = 25$ and $N = 30$ there was rarely an improvement in terms of iterations, but the required runtime tended to slightly increase compared to smaller N . In order to keep the figures clean, we hence plot the results for $N = 5, 10, 15, 20$, only. Figures 20 and 21 show that WSS 7 behaves slightly worse for $N = 5$, while for larger N the behavior over the *entire* grid is essentially indistinguishable. The latter observation mildly changes if one only considers the hyper-parameters with small cross validation error, yet it is unclear to which extend this effect is caused by different hyper-parameters picked by the different methods. In addition, a detailed look at Figure 22 does not really clarify the situation since many of the run times measured are close to the finest resolution time.h could provide. Consequently, it seems safe to say that at least in the range $N = 10 \dots 20$ the performance of WSS 7 is essentially independent of N .

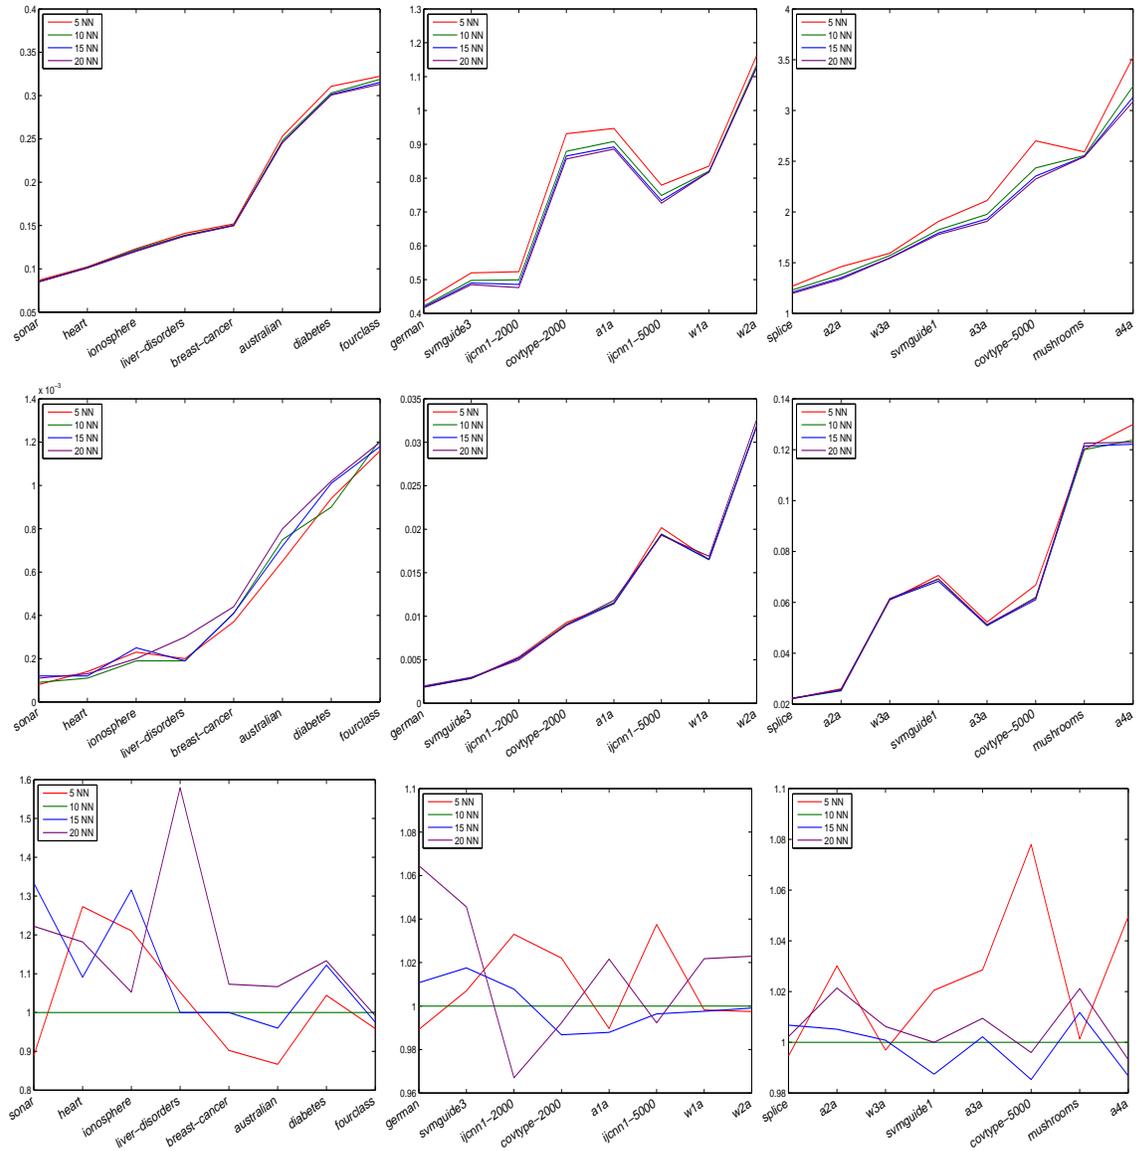


Figure 20: Average computational requirements per grid point of WSS 7 with different numbers N of nearest neighbors for small (left), mid-sized (middle), and relatively large datasets (right). The graphics display the number of iterations in thousands (top), the runtime in seconds (middle), and the corresponding ratios $xNN/10NN$ of the runtimes (bottom). Besides, for $N = 5$, and the small datasets on which the time measurements are not very reliable, the performance is basically identical.

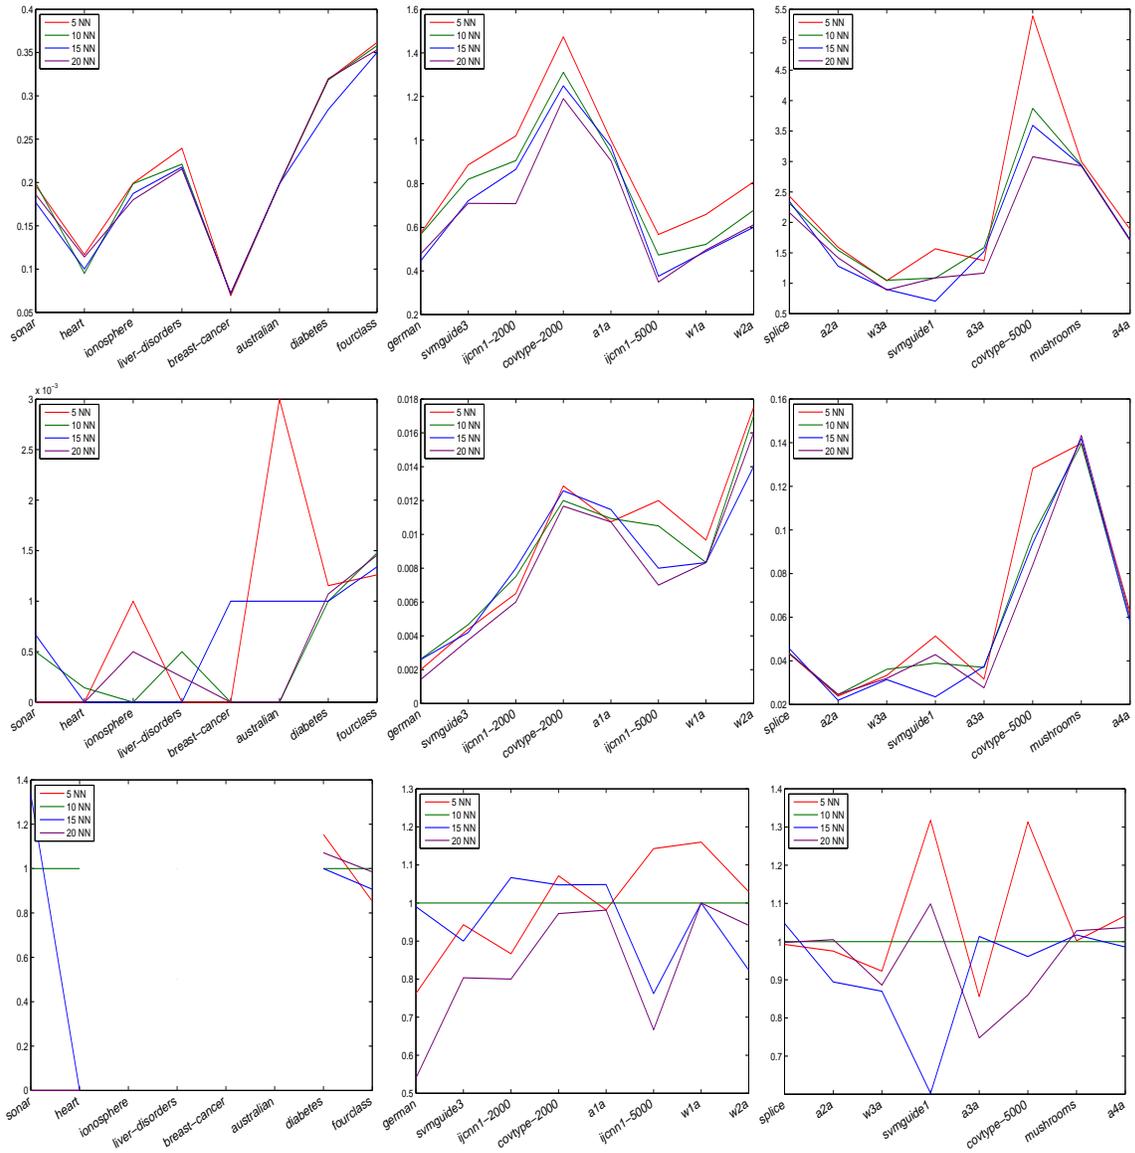


Figure 21: Computational requirements for WSS 7 with different numbers N of nearest neighbors on the grid points whose cross validation error is not larger than 1.05 the minimal cross validation error. The graphics display the average number of iterations in thousands (top), the runtime in seconds (middle), and the corresponding ratios $xNN/10NN$ of the runtimes (bottom). The plots suggest that for grid points with good validation error the number of nearest neighbors has a stronger influence than for the average grid point, yet it is unclear to which extent this effect is caused by different hyper-parameters picked by the different methods.

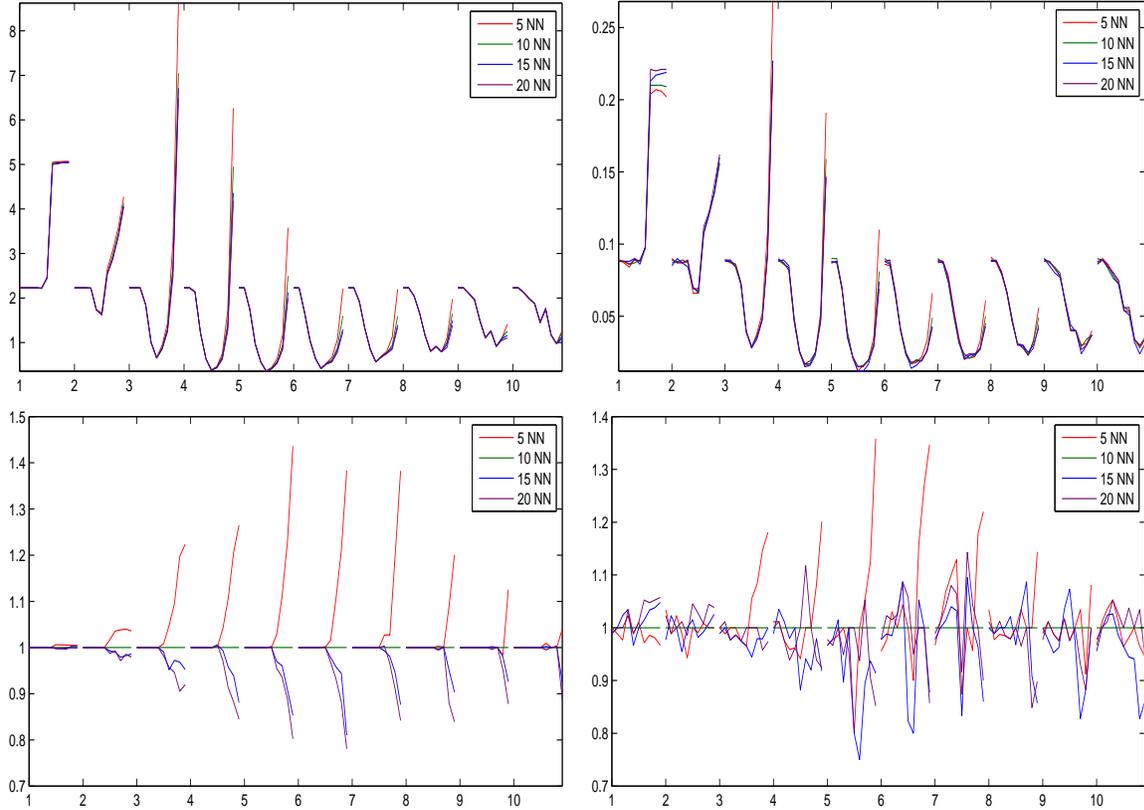


Figure 22: Average computational requirements per grid point of WSS 7 with different numbers N of nearest neighbors for the SVMGUIDE1 dataset. The four graphics have the same format as the ones in Figure 7. The graphics at the top display the number of iterations in thousands (left) and the runtime in seconds (right), both averaged over the 10 folds, while the graphics at the bottom display the corresponding ratios $xNN/10NN$. Using 5 nearest neighbors clearly results in a worse performance compared to using 10 nearest neighbors. Moreover, compared to $N = 10$ the number of iterations can be further reduced by using more nearest neighbors, but due to unreliable measurements of the run time, it remains somewhat unclear, if this results in a significantly shorter run time.

5.6 Comparing different initializations

Let us now investigate the influence of different initialization strategies on the computational requirements. To this end, we trained 2D-SVM with WSS 7 and with the different combinations of cold start and warm start options on all the datasets summarized in Table 1.

Figure 23 shows that initializing with zeros always leads to less iterations than initializing with a kernel rule. Surprisingly, however, the required runtime for both initialization strategies is substantial less different. A closer inspection revealed² that this is caused by the fact that the solver initialized with the kernel rule method I1-W1 spends most of its iterations during initialization, that is, most of the iterations counted are from the outer loop of Procedure 4. Since these iterations do not involve a working set selection step, they are relatively cheap compared to the iterations of the actual solver described in Algorithm 2. Moreover, Figure 23 shows that the simple warm start strategies W2, W3, and W5 do reduce the computational requirements significantly, where in almost all cases the scaling approach of W3 and W5 performs superior.

Interestingly, the computational requirements can often be further reduced by one of the more complicated initialization strategies W4 and W6 as Figure 24 illustrates. In particular, the combinations I0-W4, I1-W4, and I0-W6 run in most cases faster than the simple combination I0-W3, and overall it seems fair to say that I1-W4 performs best. However note that this approach requires access to the entire kernel matrix, and hence the combinations I0-W4 and I0-W6 may be the better choice, if storing this matrix is not an option.

Finally, we conducted a control experiment in which the warm start options available for SVMs with offset are compared. Figure 25, which displays the corresponding results, shows that in most cases scaling by W3 and W5 is better than keeping the solution, i.e., W2. This is similar to our results for SVMs without offset, but a closer look reveals, that the runtime gain for SVMs with offset is both less pronounced and less consistent. In particular for the larger datasets, the gain by using a warm start for SVMs with offset is about 20%, while for SVMs without offset it is about 45% even if only the simple warm start option W5 is used. Moreover, the more complex strategies for SVMs without offset can reduce the runtime by about 60% on these datasets. Consequently, it seems fair to say that SVMs without offset benefit substantially more from warm start strategies than SVMs with offset do.

Let us finally have a detailed look at the performance of some of the initialization strategies. Here Figure 26 reveals that the warm start options perform almost uniformly better than the cold start option I0-W0. Moreover, the complex warm start strategy W4 achieves its largest gains at small values of λ , which is not surprising since it starts with large values of λ . Moreover, the grid points that need the most computational requirements all have relatively small values of λ , which explains why the strategy W4 is successful. On the other hand, the strategies W5 and W6 start with small values of λ , and hence they do not achieve any improvement over I0-W0 for such λ . However, they achieve a significant improvement for large and medium values of λ , which in turn explains their success. By combining these observations and the fact that the cold start I0 requires a relatively small number of iterations on medium values for λ , it seems promising to use a hybrid strategy that starts with such a medium value for λ , and then performs W4 for smaller λ and W6 for larger values. However, investigating such a strategy is out of the scope of this paper.

²For brevity's sake we omitted the display of the corresponding plots.

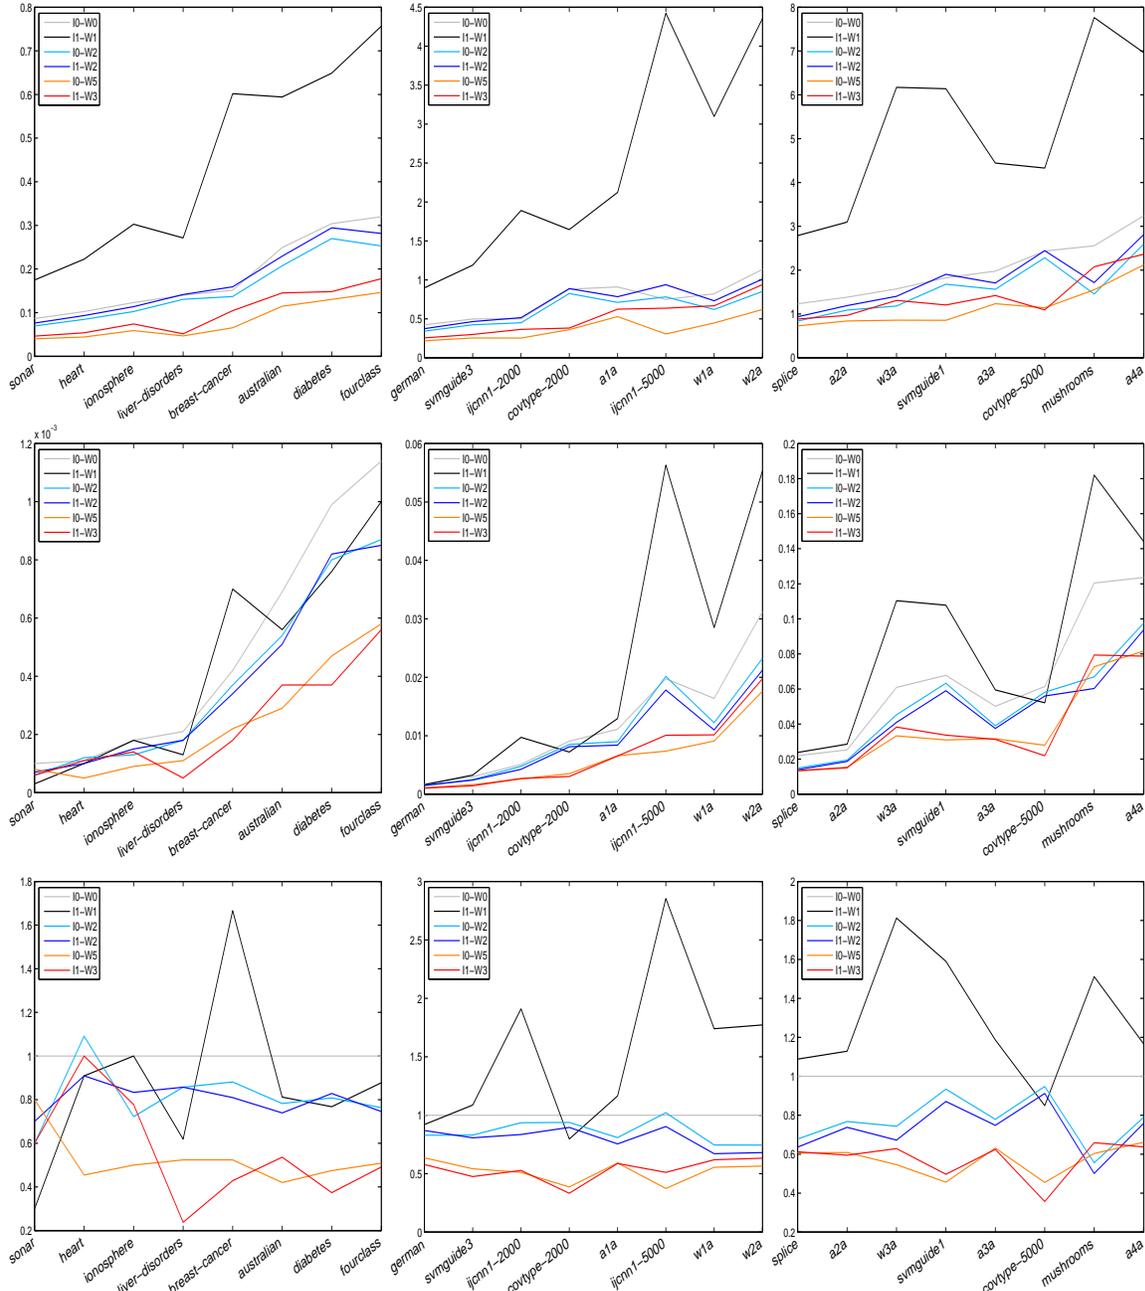


Figure 23: Average computational requirements per grid point of simple initialization strategies for the 2D-SVM with WSS 7 for small (left), mid-sized (middle), and relatively large datasets (right). The graphics display the number of iterations in thousands (top), the runtime in seconds (middle), and the ratios t_{x-Wy}/t_{0-W0} of the runtimes (bottom). The cold start initializations with zeros (I0-plots) almost always need less iterations but in some cases more runtime. The reason for this seemingly paradoxical behavior is the fact the most of the I1-iterations are performed during the initialization phase and hence do not require the costly search for the optimal directions to optimize over during the solver phase. On the other hand, I1-initializations require computing the entire kernel matrix which may be prohibitive for larger datasets. Finally, scaling the previous solution, i.e., W3 and W5, is in almost all cases faster than keeping the solution as done in W2.

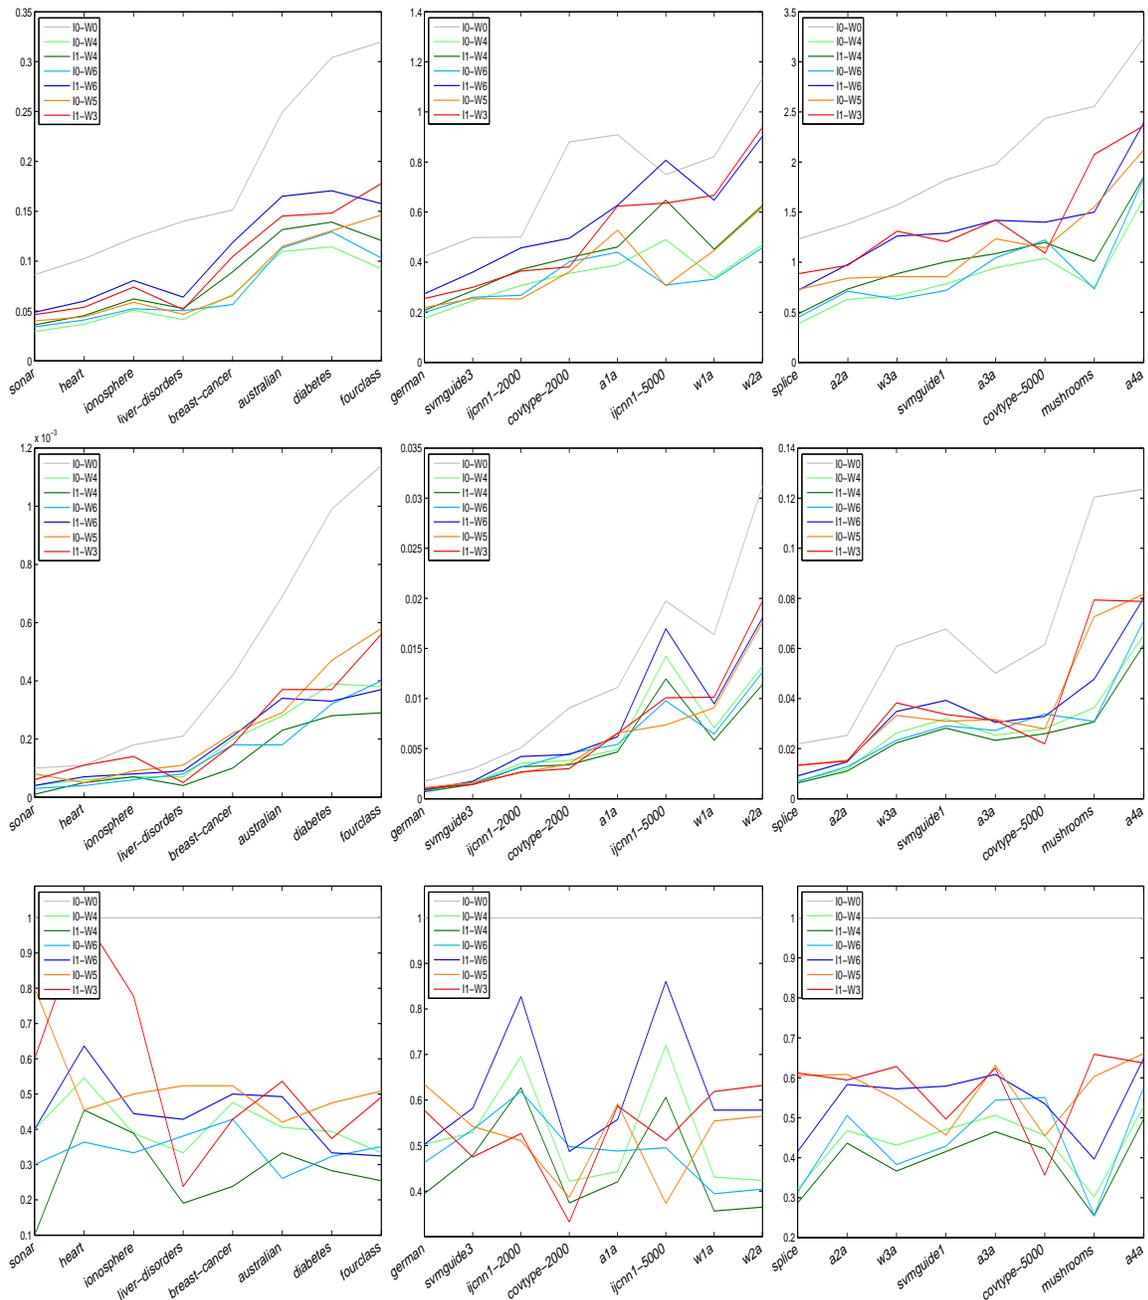


Figure 24: Average computational requirements per grid point of more complex initialization strategies for the 2D-SVM with WSS 7 for small (left), mid-sized (middle), and relatively large datasets (right). The graphics display the number of iterations in thousands (top), the runtime in seconds (middle), and the ratios t_{x-Wy}/t_{I0-W0} of the runtimes (bottom). Note that, again, the cold start initializations with zeros (I0- plots) need less iterations but in most cases more runtime. In almost all cases, the more complicated initialization strategies perform better than the simple warm start approaches. Overall, I0-W4, I1-W4, and I0-W6 are the most efficient methods in terms of runtime.

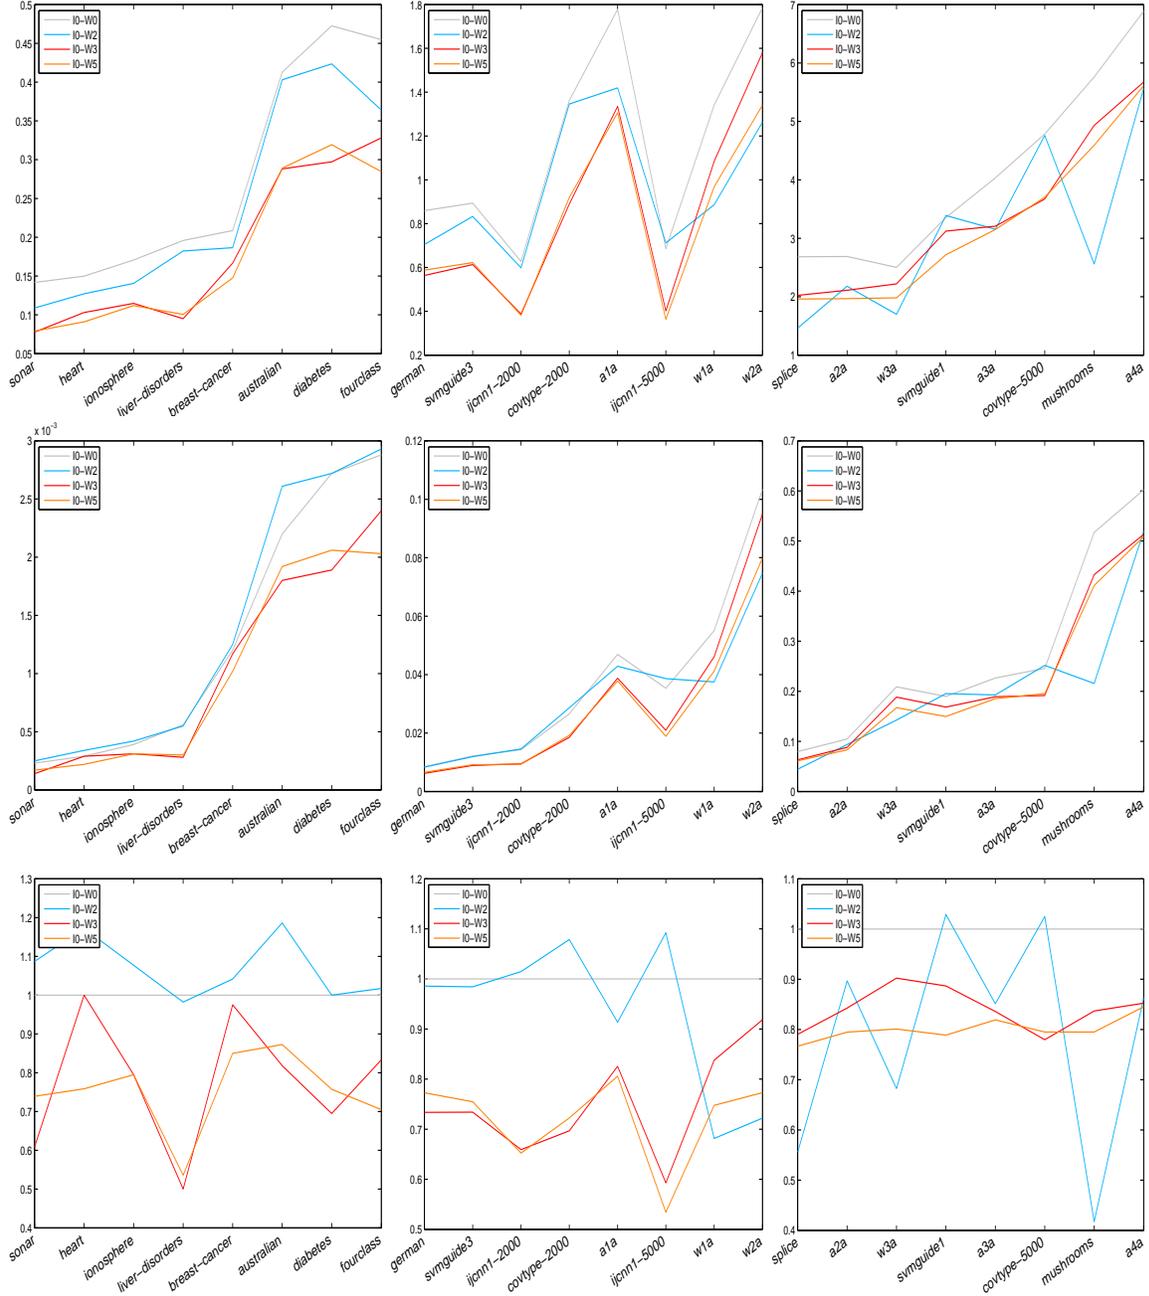


Figure 25: Average computational requirements per grid point of more complex initialization strategies for the LIBSVM for small (left), mid-sized (middle), and relatively large datasets (right). The graphics display the number of iterations in thousands (top), the runtime in seconds (middle), and the ratios $I_x-W_y/IO-W_0$ of the runtimes (bottom). Like for SVMs without offset, using a warm start pays off for this SVM with offset, but the gain is less pronounced.

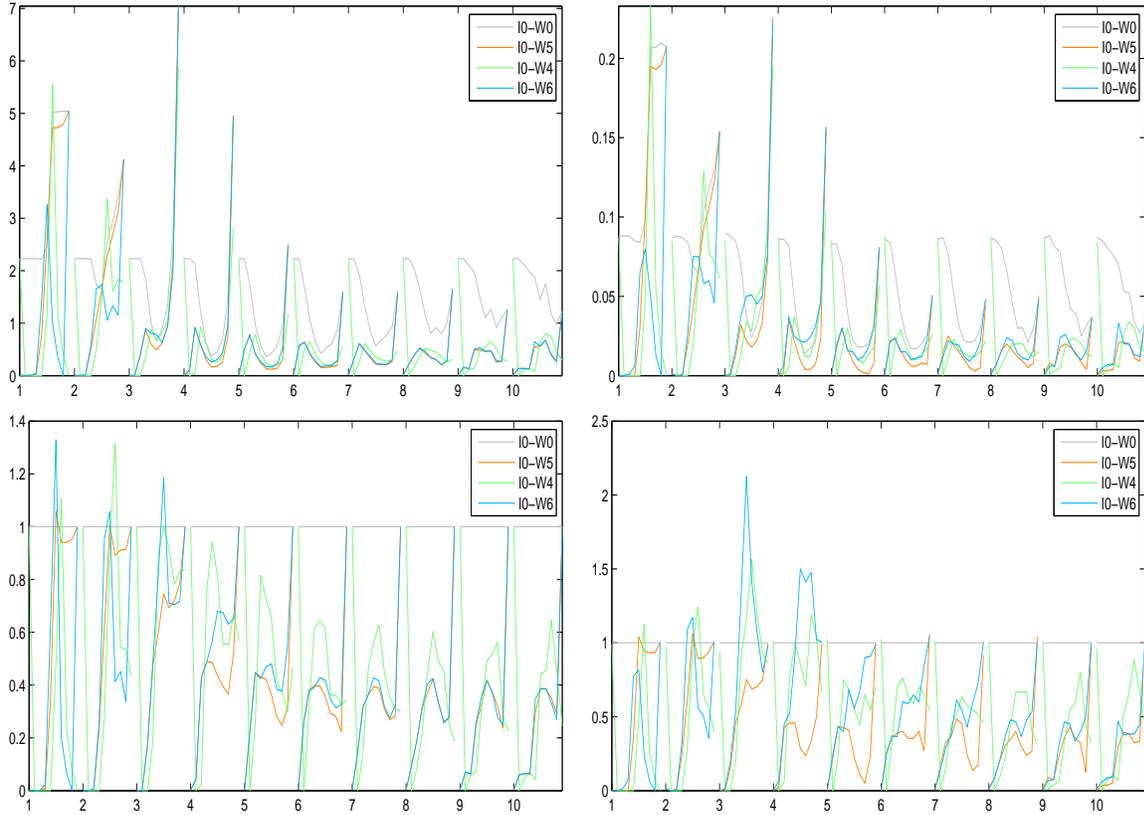


Figure 26: Computational requirements per single grid point of some initialization strategies for the SVMGUIDE1 dataset. The four graphics have the same format as the ones in Figure 7. The graphics at the top display the number of iterations in thousands (left) and the runtime in seconds (right), both averaged over the 10 folds, while the graphics at the bottom display the corresponding ratios $10\text{-}W_x/10\text{-}W_0$. All warm start strategies perform almost uniformly better than the cold start option $10\text{-}W_0$. Moreover, note that the strategies $10\text{-}W_5$ and $10\text{-}W_6$ start with the smallest λ , i.e. at the right hand side of each cell, whereas $10\text{-}W_4$ starts with the largest λ , i.e. on the left hand side of each cell.

6 Conclusions

We have thoroughly investigated SVMs without offset term b that use the hinge loss and Gaussian kernels. It turned out that these SVMs have convergence rates and classification performance that are comparable to SVMs with offset, while the absence of the offset gives more freedom in the algorithm design. In particular, we identified three areas, where this additional freedom results in faster algorithms:

- **Working set selection.** In principle, an SMO-type solver for SVMs without offset can update one variable at each iteration. However, we saw that this approach does not lead to run times that were shorter than those of an SMO-type solver for SVMs with offset. We then identified some selection methods for working sets of size two, that modified the search for working sets of size one only very slightly. It further turned out that these modifications decreased the number of iterations substantially, and since updating the gradient and computing the stopping criterion for two variables did not change the costs of an iteration dramatically, these modification also resulted in a significantly shorter run time. It is further worth mentioning that the most successful selection strategies for workings sets of size two were actually combinations of a few such simple modifications. The reason for the latter was that some strategies worked particularly well for large values of the regularization parameter λ , while others worked better for small values of λ . The good combinations then contained both types of strategies and identified the better one at each iteration automatically.
- **Stopping criterion.** Another improvement of the run time behavior of our algorithm came from a new stopping criterion that has a clear justification from recent statistical analysis of SVMs. This stopping criterion, which is essentially a relaxed duality gap, never leads to more iterations than the classical duality gap stopping criterion, but it often decreased the number of iterations. Moreover, its computational costs were almost identical to those of the classical duality gap, and hence it often resulted in shorter run times.
- **Warm start initializations.** SVMs without offset also allow more freedom in the design of warm start initializations when the hyper-parameters are determined over a grid of hyper-parameters. We investigated a couple of such initialization methods and saw that some of them led to a substantially shorter run time. Moreover, by comparing to some warm start initializations for SVMs with offset, we observed that SVMs without offset benefit significantly more from such strategies.

In our experiments we only considered datasets for which the kernel matrix fit completely in the RAM of a desktop computer. With present configurations of, say up to 8GB RAM, this limits the data set size somewhere between 25,000 and 30,000 samples. While such sizes are typically not considered to be extremely large, they already constitute a decent challenge for existing off-the-shelf SVM software, if the training time is an issue. Moreover, even for smaller data sets a fully automated hyper-parameter selection run for SVMs with offset is, for some applications, too time intensive. Our new SVM solver yields a significant reduction in time for medium-sized datasets, thus opening the applicability of SVMs to such problem domains. However, it seems fair to say that although many data sets actually fall in this range of size, some other applications demand processing substantially larger data sets. So far, it remains unclear, how well our new solver performs for such data sets,

and since our experimental study was already quite extensive, we postpone this question to future research.

Some other directions of future research include the following questions: *a)* Are there cheap modifications of our 2D-working set selection strategy that identify working sets of larger size for which the number of iterations and the run time is further reduced? *b)* Can some of our ideas be used or modified for other SVMs, that, e.g., use different kernels and/or loss functions? *c)* Can the run time of the solver be further improved by not only using vectorization via SSE instructions but by also distributing tasks between different cores of a modern processor?

References

- [1] Chih-Chung Chang and Chih-Jen Lin. LIBSVM: a library for support vector machines. <http://www.csie.ntu.edu.tw/~cjlin/papers/libsvm.ps.gz>, 2004 – 2009.
- [2] P.-H. Chen, R.-E. Fan, and C.-J. Lin. A study on SMO-type decomposition methods for support vector machines. *IEEE Trans. Neural Networks*, 17:893–908, 2006.
- [3] N. Cristianini and J. Shawe-Taylor. *An Introduction to Support Vector Machines*. Cambridge University Press, Cambridge, 2000.
- [4] L. Devroye, L. Györfi, and G. Lugosi. *A Probabilistic Theory of Pattern Recognition*. Springer, New York, 1996.
- [5] R.-E. Fan, P.-H. Chen, and C.-J. Lin. Working set selection using second order information for training support vector machines. *J. Mach. Learn. Res.*, 6:1889–1918, 2005.
- [6] T. Glasmachers and C. Igel. Maximum-gain working set selection for SVMs. *J. Mach. Learn. Res.*, 7:1437–1466, 2006.
- [7] T.-M. Huang, V. Kecman, and I. Kopriva. *Kernel Based Algorithms for Mining Huge Data Sets: Supervised, Semi-supervised, and Unsupervised Learning*. Springer, Berlin, 2006.
- [8] D. Hush, P. Kelly, C. Scovel, and I. Steinwart. QP algorithms with guaranteed accuracy and run time for support vector machines. *J. Mach. Learn. Res.*, 7:733–769, 2006.
- [9] D. Hush and C. Scovel. Polynomial-time decomposition algorithms for support vector machines. *Mach. Learn.*, 51:51–71, 2003.
- [10] S. Keerthi, V. Sindhwani, and O. Chapelle. An efficient method for gradient-based adaptation of hyperparameters in SVM models. In *Advances in Neural Information Processing Systems 19*, pages 673–680. MIT Press, Cambridge, MA, 2007.
- [11] S. S. Keerthi, S. K. Shevade, C. Battacharyya, and K. R. K. Murthy. Improvements to Platt’s SMO algorithm for SVM classifier design. *Neural Comput.*, 13:637–649, 2001.
- [12] C. J. Lin. On the convergence of the decomposition method for support vector machines. *IEEE Trans. Neural Networks*, 12:1288–1298, 2001.

- [13] C. J. Lin. Asymptotic convergence of an SMO algorithm without any assumptions. *IEEE Trans. Neural Networks*, 13:248–250, 2002.
- [14] C. J. Lin. A formal analysis of stopping criteria of decomposition methods for support vector machines. *IEEE Trans. Neural Networks*, 13:248–250, 2002.
- [15] N. List, D. Hush, C. Scovel, and I. Steinwart. Gaps in support vector optimization. In N. Bshouty and C. Gentile, editors, *Proceedings of the 20th Conference on Learning Theory*, pages 336–348. Springer, New York, 2007.
- [16] N. List and H.-U. Simon. A general convergence theorem for the decomposition method. In *Proceedings of the 17th Annual Conference on Learning Theory*, pages 363–377. Springer, Heidelberg, 2004.
- [17] N. List and H. U. Simon. General polynomial time decomposition algorithms. In S. Ben-David, J. Case, and A. Maruko, editors, *Proceedings of the 18th Annual Conference on Learning Theory, COLT 2005*, pages 308–322. Springer, Heidelberg, 2005.
- [18] N. List and H. U. Simon. General polynomial time decomposition algorithms. *J. Mach. Learn. Res.*, 8:303–321, 2007.
- [19] I. Steinwart. Sparseness of support vector machines. *J. Mach. Learn. Res.*, 4:1071–1105, 2003.
- [20] I. Steinwart and A. Christmann. *Support Vector Machines*. Springer, New York, 2008.
- [21] I. Steinwart, D. Hush, and C. Scovel. An oracle inequality for clipped regularized risk minimizers. In B. Schölkopf, J. Platt, and T. Hoffman, editors, *Advances in Neural Information Processing Systems 19*, pages 1321–1328. MIT Press, Cambridge, MA, 2007.