

The Planning-ahead SMO Algorithm

Tobias Glasmachers
tobias.glasachers@ini.rub.de
Optimization of Adaptive Systems Group
Institut für Neuroinformatik
Ruhr-Universität Bochum
Germany

August 1, 2013

Abstract

The sequential minimal optimization (SMO) algorithm and variants thereof are the de facto standard method for solving large quadratic programs for support vector machine (SVM) training. In this paper we propose a simple yet powerful modification. The main emphasis is on an algorithm improving the SMO step size by planning-ahead. The theoretical analysis ensures its convergence to the optimum. Experiments involving a large number of datasets were carried out to demonstrate the superiority of the new algorithm.

Keywords: Sequential Minimal Optimization, Quadratic Programming, Support Vector Machine Training

1 Introduction

Training a support vector machine (SVM) for binary classification is usually accomplished through solving a quadratic program. Assume we are given a training dataset $(x_1, y_1), \dots, (x_\ell, y_\ell)$ composed of inputs $x_i \in X$ and binary labels $y_i \in \{\pm 1\}$, a positive semi-definite Mercer kernel function $k : X \times X \rightarrow \mathbb{R}$ on the input space X and a regularization parameter value $C > 0$. Then the dual SVM training problem is given by

$$\text{maximize } f(\alpha) = y^T \alpha - \frac{1}{2} \alpha^T K \alpha \quad (1)$$

$$\text{s.t. } \sum_{i=1}^{\ell} \alpha_i = 0 \quad (\text{equality constraint})$$

$$\text{and } L_i \leq \alpha_i \leq U_i \quad \forall 1 \leq i \leq \ell \quad (\text{box constraint})$$

for $\alpha \in \mathbb{R}^\ell$. Here, the vector $y = (y_1, \dots, y_\ell)^T \in \mathbb{R}^\ell$ is composed of the labels, the positive semi-definite kernel Gram matrix $K \in \mathbb{R}^{\ell \times \ell}$ is given by $K_{ij} = k(x_i, x_j)$ and the lower and upper bounds are $L_i = \min\{0, y_i C\}$ and $U_i = \max\{0, y_i C\}$. We denote the feasible region by \mathcal{R} and the set of optimal points by \mathcal{R}^* . On \mathcal{R}^* the objective function f attains its maximum denoted by $f^* = \max\{f(\alpha) \mid \alpha \in \mathcal{R}\}$.

Solving this problem up to a sufficient accuracy seems to scale roughly quadratic in ℓ in practice [7]. This relatively bad scaling behavior is one of the major drawbacks of SVMs in general as the number ℓ of training examples may easily range to hundreds of thousands or even millions in today's pattern recognition problems.

2 State of the Art SMO Algorithm

The sequential minimal optimization (SMO) algorithm [13] is an iterative decomposition algorithm [12] using minimal working sets of size two. This size is minimal to keep the current solution feasible. The algorithm explicitly exploits the special structure of the constraints of problem (1) and shows very good performance in practice. For each feasible point $\alpha \in \mathcal{R}$ we define the index sets

$$I_{\text{up}}(\alpha) = \{i \in \{1, \dots, \ell\} \mid \alpha_i < U_i\}$$

$$I_{\text{down}}(\alpha) = \{i \in \{1, \dots, \ell\} \mid \alpha_i > L_i\} .$$

The canonical form of the SMO algorithm (using the common Karush-Kuhn-Tucker (KKT) violation stopping condition) can be stated as follows:

Algorithm 1: General SMO Algorithm	
Input: feasible initial point $\alpha^{(0)}$, accuracy $\varepsilon \geq 0$ compute the initial gradient $G^{(0)} \leftarrow \nabla f(\alpha^{(0)}) = y - K\alpha^{(0)}$ set $t \leftarrow 1$ do 1 select a working set $B^{(t)}$ 2 solve the sub-problem induced by $B^{(t)}$ and $\alpha^{(t-1)}$, resulting in $\alpha^{(t)}$ 3 compute the gradient $G^{(t)} \leftarrow \nabla f(\alpha^{(t)}) = G^{(t-1)} - K(\alpha^{(t)} - \alpha^{(t-1)})$ 4 stop if $\left(\max \{G_i^{(t)} \mid i \in I_{\text{up}}(\alpha^{(t)})\} - \min \{G_j^{(t)} \mid j \in I_{\text{down}}(\alpha^{(t)})\} \right) \leq \varepsilon$ set $t \leftarrow t + 1$ loop;	

If no additional information are available the initial solution is chosen to be $\alpha^{(0)} = (0, \dots, 0)^T$ resulting in the initial gradient $G^{(0)} = \nabla f(\alpha^{(0)}) = y$ which can be computed without any kernel evaluations.

It is widely agreed that the working set selection policy is crucial for the overall performance of the algorithm. This is because starting from the initial solution the SMO algorithm generates a sequence $(\alpha^{(t)})_{t \in \mathbb{N}}$ of solutions which is determined by the sequence of working sets $(B^{(t)})_{t \in \mathbb{N}}$. We will briefly discuss some concrete working set selection policies later on.

First we will fix our notation. In each iteration the algorithm selects a working set of size two. In this work we will consider (ordered) tuples instead of sets for a number of reasons. Of course we want our tuples to correspond to sets of cardinality two. Therefore a working set B is of the form (i, j) with $i \neq j$. Due to its wide spread we will stick to the term working *set*

instead of tuple as long as there is no ambiguity. Whenever we need to refer to the corresponding set, we will use the notation $\widehat{B} = \widehat{(i, j)} := \{i, j\}$. For a tuple $B = (i, j)$ we define the direction $v_B = e_i - e_j$ where e_n is the n -th unit vector of \mathbb{R}^ℓ . This direction has a positive component for α_i and a negative component for α_j . We will restrict the possible choices such that the current point α can be moved in the corresponding direction v_B without immediately leaving the feasible region. This is equivalent to restricting i to $I_{\text{up}}(\alpha)$ and j to $I_{\text{down}}(\alpha)$. We collect the allowed working sets in a point α in the set $\mathcal{B}(\alpha) = I_{\text{up}}(\alpha) \times I_{\text{down}}(\alpha) \setminus \{(n, n) \mid 1 \leq n \leq \ell\}$. With this notation a working set selection policy returns some $B^{(t)} \in \mathcal{B}(\alpha^{(t-1)})$.

The sub-problem induced by the working set $B^{(t)}$ solved in step 2 in iteration t is defined as

$$\begin{aligned} \text{maximize} \quad & f(\alpha^{(t)}) = y^T \alpha^{(t)} - \frac{1}{2}(\alpha^{(t)})^T K \alpha^{(t)} \\ \text{s.t.} \quad & \sum_{i=1}^{\ell} \alpha_i^{(t)} = 0 && \text{(equality constraint)} \\ & L_i \leq \alpha_i^{(t)} \leq U_i \text{ for } i \in \widehat{B}^{(t)} && \text{(box constraint)} \\ \text{and} \quad & \alpha_i^{(t)} = \alpha_i^{(t-1)} \text{ for } i \notin \widehat{B}^{(t)}. \end{aligned}$$

That is, we solve the quadratic program as good as possible while keeping all variables outside the current working set constant. We can incorporate the equality constraint into the parameterization $\alpha^{(t)} = \alpha^{(t-1)} + \mu^{(t)} v_{B^{(t)}}$ and arrive at the equivalent problem

$$\begin{aligned} \text{maximize} \quad & l_t \mu^{(t)} - \frac{1}{2} Q_{tt} (\mu^{(t)})^2 \\ \text{s.t.} \quad & \tilde{L}_t \leq \mu^{(t)} \leq \tilde{U}_t \end{aligned}$$

for $\mu^{(t)} \in \mathbb{R}$ with

$$\begin{aligned} Q_{tt} &= K_{ii} - 2K_{ij} + K_{jj} = v_{B^{(t)}}^T K v_{B^{(t)}} \\ l_t &= \frac{\partial f}{\partial \alpha_i}(\alpha^{(t-1)}) - \frac{\partial f}{\partial \alpha_j}(\alpha^{(t-1)}) = v_{B^{(t)}}^T \nabla f(\alpha^{(t-1)}) \\ \tilde{L}_t &= \max\{L_i - \alpha_i^{(t-1)}, \alpha_j^{(t-1)} - U_j\} \\ \tilde{U}_t &= \min\{U_i - \alpha_i^{(t-1)}, \alpha_j^{(t-1)} - L_j\} \end{aligned}$$

and the notation $B^{(t)} = (i, j)$. This problem is solved by clipping the Newton step $\mu^* = l_t/Q_{tt}$ to the bounds:

$$\mu^{(t)} = \max \left\{ \min \left\{ \frac{l_t}{Q_{tt}}, \tilde{U}_t \right\}, \tilde{L}_t \right\}. \quad (2)$$

For $\mu^{(t)} = l_t/Q_{tt}$ we call the SMO step free. In this case the SMO step coincides with the Newton step in direction $v_{B^{(t)}}$. Otherwise the step is said to hit the box boundary.

Recently it has been observed that the SMO step itself can be used for working set selection resulting in so called second order algorithms [2, 5].

We can formally define the gain of a SMO step as the function $g_B(\alpha)$ which computes the difference $f(\alpha') - f(\alpha)$ of the objective function before and after a SMO step with starting point α on the working set B , resulting in α' . For each working set B this function is continuous and piecewise quadratic (see [5]). Then these algorithms greedily choose a working set $B^{(t)}$ promising the largest functional gain $g_{B^{(t)}}(\alpha^{(t-1)}) = f(\alpha^{(t)}) - f(\alpha^{(t-1)})$ by heuristically evaluating a subset of size $\mathcal{O}(\ell)$ of the possible working sets $\mathcal{B}(\alpha^{(t-1)})$.

Fan et al. [2] propose to choose the working set according to

$$\begin{aligned} i &= \arg \max \left\{ \frac{\partial f}{\partial \alpha_n}(\alpha) \mid n \in I_{\text{up}}(\alpha) \right\} \\ j &= \arg \max \left\{ \tilde{g}_{(i,n)}(\alpha) \mid n \in I_{\text{down}}(\alpha) \setminus \{i\} \right\} \\ \text{with } \tilde{g}_B(\alpha) &= \frac{1}{2} \frac{(v_B^T \nabla f(\alpha))^2}{v_B^T K v_B} \in \mathbb{R}^{\geq 0} \cup \{\infty\} \end{aligned} \quad (3)$$

where $\tilde{g}_B(\alpha)$ is an upper bound on the gain which is exact if and only if the step starting from α with working set B is not constrained by the box.¹ Note that in this case the Newton step $\mu^* = (v_B^T \nabla f(\alpha)) / (v_B^T K v_B)$ in direction v_B is finite and we get the alternative formulation

$$\tilde{g}_B(\alpha) = \frac{1}{2} (v_B^T K v_B) (\mu^*)^2 . \quad (4)$$

This formula can be used to explicitly compute the exact SMO gain $g_B(\alpha)$ by plugging in the clipped step size (2) instead of the Newton step μ^* .

The stopping condition in step 4 checks if the Karush-Kuhn-Tucker (KKT) conditions of problem (1) are fulfilled with the predefined accuracy ε . List et al. [9] have shown that this is a meaningful stopping criterion. The accuracy ε is usually set to 0.001 in practice.

SMO is a specialized version of the more general decomposition algorithm which imposes the weaker condition $|B^{(t)}| \leq q \ll \ell$ on the working set size. The main motivation for decomposition is that in each step only the rows of the kernel matrix K which correspond to the working set indices are needed. Therefore the algorithm works well even if the whole matrix K does not fit into the available working memory. The SMO algorithm has the advantage over decomposition with larger working sets that the sub-problems in step 2 can be solved very easily. Because of its minimal working set size the algorithm makes less progress in a single iteration compared to larger working sets. On the other hand single iterations are faster. Thus, there is a trade-off between the time per iteration and the number of iterations needed to come close enough to the optimum. The decisive advantage of SMO in this context is that it can take its decisions which working set B (corresponding to the optimization direction v_B) to choose more frequently between its very fast iterations. This strategy has proven beneficial in practice.

In elaborate implementations the algorithm is accompanied by a kernel cache and a shrinking heuristic [7]. The caching technique exploits the

¹The software LIBSVM [2] sets the denominator of $\tilde{g}_B(\alpha)$ to $\tau = 10^{-12} > 0$ whenever it vanishes. This way the infinite value is avoided. However, this trick was originally designed to tackle indefinite problems.

fact that the SMO algorithm needs the rows of the kernel matrix which correspond to the indices in the current working set $B^{(t)}$. The kernel cache uses a predefined amount of working memory to store rows of the kernel matrix which have already been computed. Therefore the algorithm needs to recompute only those rows from the training data evaluating the possibly costly kernel function which have not been used recently. The shrinking heuristic removes examples from the problem that are likely to end up at the box boundaries in the final solution. These techniques perfectly cooperate and result in an enormous speed up of the training process. We will later use the fact that the most recently used rows of the kernel matrix K are available from the cache.

The steps 1, 3, and 4 of the SMO optimization loop take $\mathcal{O}(\ell)$ operations, while the update 2 of the current solution is done in constant time.

There has not been any work on the improvement of step 2 of Algorithm 1. Of course, it is not possible to considerably speed up a computation taking $\mathcal{O}(1)$ operations, but we will see in the following how we may replace the optimal (greedy) truncated Newton step with other approaches.

3 Behavior of the SMO Algorithm

We want to make several empirical and theoretical statements about the overall behavior of the SMO algorithm. This includes some motivation for the algorithm presented later. We start with theoretical results.

It is well known that the algorithm converges to an optimum for a number of working set selection strategies. Besides convergence proofs for important special cases [8, 15, 5, 3] proof techniques for general classes of selection policies have been investigated [6, 10, 1].

Chen et al. [1] have shown that under some technical conditions on problem (1) there exists t_0 such that no SMO step ends up at the box bounds for iterations $t > t_0$. For these iterations the authors derive a linear convergence rate. However, the prerequisites exclude the relevant case that the optimum is not isolated. Upper bounds for t_0 are not known, and in experiments the algorithm rarely seems to enter this stage.

From an empirical point of view we can describe the qualitative behavior of SMO roughly as follows. In the first iterations, starting from the initial solution $\alpha^{(0)} = (0, \dots, 0)^T$, many steps move variables α_i to the lower or upper bounds L_i and U_i . After a while, these steps become rare and most iterations are spent on a relatively small number of variables performing free steps. In this phase the shrinking heuristics removes most bounded variables from the problem. Then working set selection, gradient update and stopping condition need to be computed only on the relatively small set of active variables, leading to extremely fast iterations.

Many common benchmark problems and real world applications are simple in the sense that the algorithm performs only a number of iterations comparable to the number of examples. In this case only very few variables (if any) are changed many times. This indicates that there are only very few free support vectors or that the dependencies between variables are weak, making the optimization easy. On the other hand, for harder problems the

algorithm spends most of its iterations on free SMO steps to resolve complicated dependencies between the free variables. In fact, in some cases we can observe large blocks of iterations spent on a small number of variables. Due to its finite number of optimization directions the SMO algorithm is prone to oscillate while compensating the second order cross terms of the objective function. This oscillatory behavior observed in case of difficult problems is the main motivation for the consideration presented in the next section.

4 Planning Ahead

Without loss of generality we consider the iteration $t = 1$ in this section.

Assume we are given the current working set $B^{(1)} = (i^{(1)}, j^{(1)})$ and for some reason we already know the working set $B^{(2)} = (i^{(2)}, j^{(2)})$ to be selected in the next iteration. In addition, we presume that the solutions of both sub-problems involved are not at the bounds. That is, we can simply ignore the box constraints. From equation (4) we know with $\mu^{(1)} = l_1/Q_{11}$ and $\mu^{(2)} = l_2/Q_{22}$ that both free steps together result in the gain

$$g^{2\text{-step}} := f(\alpha^{(2)}) - f(\alpha^{(0)}) = \frac{1}{2}Q_{11}(\mu^{(1)})^2 + \frac{1}{2}Q_{22}(\mu^{(2)})^2 . \quad (5)$$

Under the assumption that we already know the working set $B^{(2)}$ we can of course precompute the second step. To stress this point of view we introduce the quantities

$$w_t = \frac{\partial f}{\partial \alpha_{i^{(t)}}}(\alpha^{(0)}) - \frac{\partial f}{\partial \alpha_{j^{(t)}}}(\alpha^{(0)}) = v_{B^{(t)}} \nabla f(\alpha^{(0)}) \quad \text{for } t \in \{1, 2\}$$

which only depend on $\alpha^{(0)}$ and are thus known in iteration $t = 1$. We rewrite

$$\begin{aligned} l_1 &= w_1 \\ l_2 &= w_2 - Q_{12}\mu^{(1)} \end{aligned}$$

with $Q_{12} = Q_{21} = K_{i^{(1)}i^{(2)}} - K_{i^{(1)}j^{(2)}} - K_{j^{(1)}i^{(2)}} + K_{j^{(1)}j^{(2)}} = v_{B^{(1)}}^T K v_{B^{(2)}}$.

Then we can express the step size

$$\mu^{(2)} = l_2/Q_{22} = w_2/Q_{22} - Q_{12}/Q_{22}\mu^{(1)} \quad (6)$$

in these terms. The above notation suggests the introduction of the 2×2 matrix

$$Q = \begin{pmatrix} Q_{11} & Q_{21} \\ Q_{12} & Q_{22} \end{pmatrix}$$

which is symmetric and positive semi-definite. If we drop the assumption that both steps involved are Newton steps the computation of the gain is more complicated:

$$g^{2\text{-step}}(\mu^{(1)}, \mu^{(2)}) := f(\alpha^{(2)}) - f(\alpha^{(0)}) = \begin{pmatrix} w_1 \\ w_2 \end{pmatrix}^T \begin{pmatrix} \mu^{(1)} \\ \mu^{(2)} \end{pmatrix} - \frac{1}{2} \begin{pmatrix} \mu^{(1)} \\ \mu^{(2)} \end{pmatrix}^T Q \begin{pmatrix} \mu^{(1)} \\ \mu^{(2)} \end{pmatrix}$$

Plugging everything in, and in particular substituting $\mu^{(2)}$ according to eq. (6) we express the gain as a function of the single variable $\mu^{(1)}$ resulting in

$$g^{2\text{-step}}(\mu^{(1)}) = -\frac{1}{2} \cdot \frac{\det(Q)}{Q_{22}} (\mu^{(1)})^2 + \frac{Q_{22}w_1 - Q_{12}w_2}{Q_{22}} \mu^{(1)} + \frac{1}{2} \cdot \frac{w_2^2}{Q_{22}}. \quad (7)$$

For $\mu^{(1)} = w_1/Q_{11}$ we obtain the gain computed in (5), but of course the maximizer of the quadratic function (7) will in general differ from this value. Thus, under the above assumption that we already know the next working set $B^{(2)}$ we can achieve a better functional gain by computing the optimal step size

$$\mu^{(1)} = \frac{Q_{22}w_1 - Q_{12}w_2}{\det(Q)} \quad (8)$$

where we again assume that we do not hit the box constraints which are dropped. It is easy to incorporate the constraints into the computation, but this has two drawbacks in our situation: First, it leads to a large number of different cases, and second it complicates the convergence proof. Further, dropping the constraints will turn out to be no restriction, as the algorithms resulting from these considerations will handle the box constrained case separately. Figure 1 illustrates the resulting step. We call the step $\mu^{(1)} \cdot v_{B^{(1)}}$ the planning-ahead step, because we need to simulate the current and the next step in order to determine the step size $\mu^{(1)}$. Analogously we refer to $\mu^{(1)}$ as the planning-ahead step size.

Just like the usual SMO update this step can be computed in constant time, that is, independent of the problem dimension ℓ . However, the kernel values of an up to 4×4 principal minor of the kernel Gram matrix K are needed for the computation, in contrast to a 2×2 minor for the standard SMO update step.

Note the asymmetry of the functional gain as well as the optimal step size w.r.t. the iteration indices 1 and 2. We control the length of the first step, which of course influences the length of the second step. The asymmetry results from the fact that the second step is greedy in contrast to the first one. The first step is optimal given the next working set $B^{(2)}$ and planning one step ahead, while the second step is optimal in the usual sense of doing a single greedy step without any planning-ahead.

Another interesting property is that for $\mu^{(1)} \notin [0, 2l_1/Q_{11}]$ the first step actually results in a decay of the dual objective, that is, $f(\alpha^{(1)}) < f(\alpha^{(0)})$, see Figure 2. Nevertheless, such steps can be extremely beneficial in some situations, see Figure 1. Of course, by construction both planned-ahead steps together result in an increase of the objective function, which is even maximized for the given working sets.

5 Algorithms

In this section we will turn the above consideration into algorithms. We will present a first simple version and a refinement which focuses on the




Figure 1: Optimization path with (dark gray) and without (dotted) planning ahead in a minimal scenario composed of only two possible working sets $B^{(1)}$ and $B^{(2)}$. The light gray ellipses indicate niveau sets of the objective function f . From these it is obvious that the first step of the dark gray path results in a decrease of the objective function. However, while the usual SMO procedure oscillates inside a cone bounded by the hyperplanes $\{\alpha \mid v_{B^{(t)}}^T \nabla f(\alpha) = 0\}$, planning ahead one step results in the optimal step size to solve this low dimensional problem.

convergence of the overall algorithm to an optimal solution. These modifications are all based on the SMO Algorithm 1. Thus, we will only state replacements for the working set selection step 1 and the update step 2.

In the previous section it is left open how we can know the choice of the working set in the forthcoming iteration. If we try to compute this working set given a step size μ , it turns out that we need to run the working set selection algorithm. That is, the next working set depends on the current step size and it takes linear time to determine the working set for a given step size. This makes a search for the optimal combination of step size and working set impractical. We propose a very simple heuristic instead. For two reasons we suggest to reuse the previous working set: First, the chance that the corresponding kernel evaluations are cached is highest for this working set. Second, as already stated in section 3, the SMO algorithm sometimes tends to oscillate within a small number of variables. Figure 1 gives a low-dimensional example. Now, if we are in a phase of oscillation, the previous working set is a very good candidate for planning-ahead.

These considerations result in a new algorithm. It differs from the SMO algorithm only in step 2. The basic idea is to use the previous working set for planning-ahead. However, we revert to the standard SMO step if the previous step was used for planning-ahead or the planned steps are not free. This proceeding is formalized in Algorithm 2, which is a replacement for




Figure 2: Gain of the step with size μ compared to the Newton step size μ^* . As long as the quantity μ/μ^* is in the open interval $(0, 2)$ the step makes some progress. For the interval $[1 - \eta, 1 + \eta]$ the gain is even strictly lower bounded by a fraction of $1 - \eta^2$ of the Newton step gain \tilde{g}^* . Note that if the working set direction v_B is in the kernel of the matrix K the parabola is degenerated to a linear function. Then we distinguish two cases: If the linear term of \tilde{g} vanishes we set $\mu^* = 0$ and get $\tilde{g}^* = 0$. Otherwise the graph grows linearly corresponding to $\mu^* = \pm\infty$ and $\tilde{g}^* = \infty$.

step 2 of Algorithm 1.

Algorithm 2: Modification of step 2 of the SMO Algorithm	
if	<i>previous iteration performed a SMO step (eq. (2))</i> then
	Compute the planning-ahead step size $\mu = \frac{Q_{22}w_1 - Q_{12}w_2}{\det(Q)}$ (eq. (8)) assuming $B^{(t-1)}$ as the next working set
	if <i>the current or the planned step ends at the box boundary</i> then
	perform a SMO step (eq. (2))
	else
	perform the step of size μ as planned
	end
else	
	perform a SMO step (eq. (2))
end	

As already indicated in section 4 the algorithm uses planning-ahead only if both steps involved do not hit the box boundaries. This means that we need to check the box constraints while planning ahead, and we turn to the standard SMO algorithm whenever the precomputed steps become infeasible. Thus, there is no need to incorporate the box constraints into the planning-ahead step size given in equation (8).

The algorithm works well in experiments. However, it is hard to prove

its convergence to an optimal solution for a number of reasons. The main difficulty involved is that we can not prove the strict increase of the objective function for the planning-ahead step, even if we additionally consider the subsequent iteration. Therefore we additionally replace the working set selection step 1 of Algorithm 1 with Algorithm 3.

Algorithm 3: Modification of step 1 of the SMO Algorithm

```

Input:  $\eta \in (0, 1)$ 
Input:  $\mu^{(t-1)}$ : step size of the previous iteration  $t - 1$ 
Input:  $\mu^*$ : Newton step size of the previous iteration  $t - 1$ 
if previous step resulted from planning-ahead then
  // standard selection, see equation (3)
   $i^{(t)} \leftarrow \arg \max \{ \frac{\partial f}{\partial \alpha_n}(\alpha^{(t-1)}) \mid n \in I_{\text{up}}(\alpha^{(t-1)}) \}$ 
   $j^{(t)} \leftarrow \arg \max \{ \tilde{g}_{(i^{(t)}, n)}(\alpha^{(t-1)}) \mid n \in I_{\text{down}}(\alpha^{(t-1)}) \setminus \{i\} \}$ 
   $B^{(t)} \leftarrow (i^{(t)}, j^{(t)})$ 
else
  if  $1 - \eta \leq \mu^{(t-1)} / \mu^* \leq 1 + \eta$  then
    // selection with additional candidate  $B^{(t-2)}$ 
     $i^{(t)} \leftarrow \arg \max \{ \frac{\partial f}{\partial \alpha_n}(\alpha^{(t-1)}) \mid n \in I_{\text{up}}(\alpha^{(t-1)}) \}$ 
     $j^{(t)} \leftarrow \arg \max \{ \tilde{g}_{(i^{(t)}, n)}(\alpha^{(t-1)}) \mid n \in I_{\text{down}}(\alpha^{(t-1)}) \setminus \{i\} \}$ 
     $B^{(t)} \leftarrow (i^{(t)}, j^{(t)})$ 
    if  $\tilde{g}_{B^{(t-2)}}(\alpha^{(t-1)}) > \tilde{g}_{B^{(t)}}(\alpha^{(t-1)})$  then
       $B^{(t)} \leftarrow B^{(t-2)}$ 
    end
  else
    // selection with additional candidate  $B^{(t-2)}$ 
    // based on  $g$  instead of  $\tilde{g}$ 
     $i^{(t)} \leftarrow \arg \max \{ \frac{\partial f}{\partial \alpha_n}(\alpha^{(t-1)}) \mid n \in I_{\text{up}}(\alpha^{(t-1)}) \}$ 
     $j^{(t)} \leftarrow \arg \max \{ g_{(i^{(t)}, n)}(\alpha^{(t-1)}) \mid n \in I_{\text{down}}(\alpha^{(t-1)}) \setminus \{i\} \}$ 
     $B^{(t)} \leftarrow (i^{(t)}, j^{(t)})$ 
    if  $g_{B^{(t-2)}}(\alpha^{(t-1)}) > g_{B^{(t)}}(\alpha^{(t-1)})$  then
       $B^{(t)} \leftarrow B^{(t-2)}$ 
    end
  end
end

```

At a first glance this algorithm looks more complicated than it is. The selection basically ensures that the planning-ahead step and the next SMO step together have a positive gain: Recall that for $\mu^{(t-1)} / \mu^* \in [1 - \eta, 1 + \eta]$ the planning step itself makes some progress, see Figure 2. The following SMO step has always positive gain. Now consider the case that the planning-step does not make a guaranteed progress, that is, $\mu^{(t-1)} / \mu^* \notin [1 - \eta, 1 + \eta]$. The planned double-step gain (7) is by construction lower bounded by the Newton step gain. Thus, if the previous working set is reused in the following iteration the total gain is positive. Now the usage of the SMO gain g instead of the Newton step gain \tilde{g} for working set selection ensures that this gain can only increase if another working set is actually selected in the step following planning-ahead. Thus, both steps together have positive gain in any case. In the following we will arbitrarily fix $\eta = 0.9$. Thus, we will not consider η

as a free hyper-parameter of Algorithm 3.

It obviously makes sense to provide the working set which was used for planning-ahead as a candidate to the working set selection algorithm. As explained above, this property together with the usage of the SMO gain function g instead of the approximation \tilde{g} ensures positive gain of the double-step. Of course, positive gain is not sufficient to show the convergence to an optimal point. The following section is devoted to the convergence proof.

Although planning-ahead is done in constant time, it takes considerably longer than the computation of the Newton step. For simple problems where planning-ahead does not play a role because most steps end up at the box bounds the unsuccessful planning steps can unnecessarily slow down the algorithm. As discussed in section 3 this is mainly the case at the beginning of the optimization. We introduce the following simple heuristic: If the previous iteration was a free SMO step, then we perform planning ahead, otherwise we perform another SMO step. Thus, we use the previous SMO step as a predictor for the current one. Algorithm 4 captures this idea.

Algorithm 4: Modification of step 2 of the SMO Algorithm. The only difference compared to Algorithm 2 is the first condition that the SMO step must be *free*.

```

if previous iteration performed a free SMO step then
    Compute the planning-ahead step size  $\mu = \frac{Q_{22}w_1 - Q_{12}w_2}{\det(Q)}$  (eq. (8))
        assuming  $B^{(t-1)}$  as the next working set
    if the current or the planned step ends at the box boundary then
        | perform a SMO step (eq. (2))
    else
        | perform the step of size  $\mu$  as planned
    end
else
    | perform a SMO step (eq. (2))
end

```

The SMO algorithm with modified steps 1 and 2 as defined in Algorithms 4 and 3, respectively, will be referred to as the planning-ahead SMO (PA-SMO) algorithm in the following. For completeness, we state the complete PA-SMO algorithm at the end of the paper.

It is not really clear whether the consideration of the working set $B^{(t-1)}$ for planning ahead is a good choice. In fact, it would be good to know whether this choice has a critical impact on the performance. To evaluate this impact we need to introduce a variant of PA-SMO. Because the planning-step takes only constant time we could afford to perform $N > 1$ planning steps with the working sets $B^{(t-n)}$ for $1 \leq n \leq N$ and choose the step size with the largest double-step gain. In this case we should also provide these sets to the working set selection algorithm as additional candidates. We call this variant the multiple planning-ahead algorithm using the $N > 1$ most recent working sets.

6 Convergence of the Method

In this section we will show that the PA-SMO algorithm converges to the optimum f^* of problem (1). First we introduce some notation and make some definitions.

Definition 1. *We say that a function $s : \mathcal{R} \setminus \mathcal{R}^* \rightarrow \mathbb{R}$ has property (*) if it is positive and lower semi-continuous. We extend this definition to functions $s : \mathcal{R} \rightarrow \mathbb{R}$ which are positive and lower semi-continuous on $\mathcal{R} \setminus \mathcal{R}^*$. We say that a function h has property (**) if there exists a function s with property (*) which is a lower bound for h on $\mathcal{R} \setminus \mathcal{R}^*$.*

Recall two important properties of a lower semi-continuous function s : First, if $s(\alpha) > 0$ then there exists an open neighborhood U of α such that $s(\alpha) > 0$ for all $\alpha \in U$. Second, a lower semi-continuous function attains its minimum on a (non-empty) compact set.

The gap function

$$\psi(\alpha) = \max \left\{ v_B^T \nabla f(\alpha) \mid B \in \mathcal{B}(\alpha) \right\} \in \mathbb{R}$$

will play a central role in the following. Note that this function is used in the stopping condition in step 4 of Algorithm 1 because it is positive on $\mathcal{R} \setminus \mathcal{R}^*$ and zero or negative on \mathcal{R}^* .

Lemma 1. *The function ψ has property (*).*

Proof. On \mathcal{R} we introduce the equivalence relation $\alpha_1 \sim \alpha_2 \Leftrightarrow \mathcal{B}(\alpha_1) = \mathcal{B}(\alpha_2)$ and split the feasible region into equivalence classes, denoted by $[\alpha]$. Obviously, ψ is continuous on each equivalence class $[\alpha]$. Now, the topological boundary $\partial[\alpha]$ of a class $[\alpha]$ is the union of those classes $[\alpha']$ with $\mathcal{B}([\alpha']) \subset \mathcal{B}([\alpha])$. Because the argument of the maximum operation in the definition of ψ is a subset of $\mathcal{B}([\alpha])$ on the boundary the maximum can only drop down. Thus, ψ is lower semi-continuous. Further, it is well known that ψ is positive for non-optimal points. \square

We collect all possible working sets in

$$\mathcal{B} = \left\{ (i, j) \mid i, j \in \{1, \dots, \ell\} \text{ and } i \neq j \right\}$$

and write the working set selection (3) as a map $W : \mathcal{R} \rightarrow \mathcal{B}$. With this fixed working set selection we consider the Newton step gain

$$\tilde{g}_W : \mathcal{R} \rightarrow \mathbb{R}^{\geq 0} \cup \{\infty\}, \quad \tilde{g}_W(\alpha) = \tilde{g}_{W(\alpha)}(\alpha)$$

as a function of α , in contrast to the family of functions with variable working set defined in eq. (4).

Lemma 2. *There exists $\sigma > 0$ such that the function $\varphi(\alpha) = \sigma (\psi(\alpha))^2$ is a lower bound for \tilde{g}_W on $\mathcal{R} \setminus \mathcal{R}^*$. Thus, the Newton step gain \tilde{g}_W has property (**).*

Proof. Of course, φ inherits property $(*)$ from ψ . We split $\mathcal{R} \setminus \mathcal{R}^* = M \cup N$ into disjoint subsets

$$\begin{aligned} M &= \left\{ \alpha \in \mathcal{R} \setminus \mathcal{R}^* \mid v_{W(\alpha)} \notin \ker(K) \right\} \\ N &= \left\{ \alpha \in \mathcal{R} \setminus \mathcal{R}^* \mid v_{W(\alpha)} \in \ker(K) \right\} \end{aligned}$$

and introduce the constants

$$\begin{aligned} \sigma_1 &= \max \left\{ v_B^T K v_B \mid B \in \mathcal{B} \right\} \\ \sigma_2 &= \min \left\{ v_B^T K v_B \mid B \in \mathcal{B} \text{ with } v_B^T K v_B > 0 \right\} . \end{aligned}$$

On M we have $v_{W(\alpha)}^T K v_{W(\alpha)} > 0$. We proceed in two steps. First we define the gap of the working set $W(\alpha)$

$$\psi_W : \mathcal{R} \rightarrow \mathbb{R}, \quad \alpha \mapsto v_{W(\alpha)}^T \nabla f(\alpha) .$$

Then we make use of a result from [2], where the bound $\psi_W(\alpha) \geq \sqrt{\sigma_2/\sigma_1} \psi(\alpha)$ is derived in section 3. From the definition of $\tilde{g}_B(\alpha)$ applied to the working set $B(\alpha)$ we get the inequality

$$\tilde{g}_W(\alpha) = \frac{1}{2} \frac{(v_{W(\alpha)}^T \nabla f(\alpha))^2}{v_{W(\alpha)}^T K v_{W(\alpha)}} \geq \frac{1}{2\sigma_1} (\psi_W(\alpha))^2$$

resulting in the desired lower bound with

$$\sigma = \frac{1}{2\sigma_1} \left(\sqrt{\frac{\sigma_2}{\sigma_1}} \right)^2 = \frac{\sigma_2}{2\sigma_1^2} > 0 .$$

On N the situation is much simpler. For $v_{W(\alpha)}^T \nabla f(\alpha) = 0$ we can not make any progress on the working set $W(\alpha)$ which contradicts $\alpha \notin \mathcal{R}^*$. Thus we have $v_{W(\alpha)}^T \nabla f(\alpha) \neq 0$ which implies $\tilde{g}_W(\alpha) = \infty$. This is because the quadratic term of the objective function in direction $v_{W(\alpha)}$ vanishes and the function increases linearly. Of course we then have $\tilde{g}_W(\alpha) \geq \sigma(\psi(\alpha))^2$ for $\alpha \in N$. \square

In contrast to [2] there is no need to use an artificial lower bound $\tau > 0$ for vanishing quadratic terms in this proof. With the properties of φ at hand it is straight forward to prove the following theorem:

Theorem 1. *Consider the sequence $(\alpha^{(t)})_{t \in \mathbb{N}}$ in \mathcal{R} with $(f(\alpha^{(t)}))_{t \in \mathbb{N}}$ monotonically increasing. Let there exist a constant $c > 0$ and an infinite set $T_c \subset \mathbb{N}$ such that the steps from $\alpha^{(t-1)}$ to $\alpha^{(t)}$ have the property*

$$f(\alpha^{(t)}) - f(\alpha^{(t-1)}) \geq c \cdot \tilde{g}_W(\alpha^{(t-1)})$$

for all $t \in T_c$. Then we have $\lim_{t \rightarrow \infty} f(\alpha^{(t)}) = f^*$.

Proof. Because of the compactness of \mathcal{R} there exists a convergent subsequence $(\alpha^{(t-1)})_{t \in \tilde{T}}$ for some $\tilde{T} \subset T_c$. We will denote its limit point by $\alpha^{(\infty)}$. Assume the limit point is not optimal. Then $\varphi(\alpha^{(\infty)}) > 0$ by property (*). The lower semi-continuity of φ implies $\varphi(\alpha) > 0$ for all α in an open neighborhood U' of $\alpha^{(\infty)}$. We choose a smaller open neighborhood U of $\alpha^{(\infty)}$ such that its closure \bar{U} is contained in U' . Again by lower semi-continuity φ attains its minimum $m > 0$ on \bar{U} . There is t_0 such that $\alpha^{(t)} \in U$ for all $t \in \tilde{T}$ with $t > t_0$. Then we have

$$\begin{aligned} f(\alpha^{(\infty)}) &\geq f(\alpha^{(t_0)}) + \sum_{t \in \tilde{T}, t > t_0} f(\alpha^{(t)}) - f(\alpha^{(t-1)}) \\ &\geq f(\alpha^{(t_0)}) + \sum_{t \in \tilde{T}, t > t_0} c \cdot \tilde{g}(\alpha^{(t-1)}) \\ &\geq f(\alpha^{(t_0)}) + \sum_{t \in \tilde{T}, t > t_0} c \cdot \varphi(\alpha^{(t-1)}) \\ &\geq f(\alpha^{(t_0)}) + \sum_{t \in \tilde{T}, t > t_0} c \cdot m = \infty > f^* \end{aligned}$$

which is a contradiction. Thus, $\alpha^{(\infty)}$ is optimal. \square

With the additional assumption that infinitely many SMO steps end up free we can use Theorem 1 to show the convergence of Algorithm 1 to an optimal solution. This was already proven in [2] without this assumption.

Corollary 1. *Consider the sequence $(\alpha^{(t)})_{t \in \mathbb{N}}$ in \mathcal{R} with $(f(\alpha^{(t)}))_{t \in \mathbb{N}}$ monotonically increasing. If there are infinitely many t such that the step from $\alpha^{(t-1)}$ to $\alpha^{(t)}$ is a free SMO step with working set (3) then we have $\lim_{t \rightarrow \infty} f(\alpha^{(t)}) = f^*$.*

Proof. For a free SMO step we have $f(\alpha^{(t)}) - f(\alpha^{(t-1)}) = \tilde{g}(\alpha^{(t-1)})$. Thus we can simply apply the above theorem with $c = 1$. \square

The allurement of this approach is that we do not need any assumption on the steps which differ from free SMO steps as long as the objective function does not decrease. This is an ideal prerequisite to tackle the convergence of hybrid algorithms which need to distinguish qualitatively different branches, like for example the PA-SMO algorithm. Consequently, the following lemma will be helpful when applying the above results to PA-SMO.

Lemma 3. *Consider two iterations t and $t + 1$ of the PA-SMO algorithm where planning-ahead is active in iteration t . The double-step gain $g^{2\text{-step}} = f(\alpha^{(t+1)}) - f(\alpha^{(t-1)})$ is then lower bounded by $(1 - \eta^2) \cdot \tilde{g}_W(\alpha^{(t-1)})$.*

Proof. Let μ^* denote the Newton step size in iteration t and let $\tilde{g}^* = \tilde{g}_W(\alpha^{(t-1)})$ be the gain achieved by this (possibly infeasible) step. Just like in Algorithm 3 we distinguish two cases:

1. The step size $\mu^{(t)}$ satisfies $1 - \eta \leq \mu^{(t)}/\mu^* \leq 1 + \eta$:
We write the gain in iteration t in the form $(2\mu^{(t)}/\mu^* - (\mu^{(t)}/\mu^*)^2) \cdot \tilde{g}^*$, see Figure 2. Together with the strict increase of the objective function in iteration $t + 1$ we get $g^{2\text{-step}} \geq (1 - \eta^2) \cdot \tilde{g}^*$.

2. The step size $\mu^{(t)}$ satisfies $\mu^{(t)}/\mu^* \notin [1 - \eta, 1 + \eta]$:

By construction the planned ahead gain (7) is lower bounded by \tilde{g}^* (see Section 4). The planning-step assumes that the working set $B^{(t-1)}$ is selected in iteration $t + 1$. However, another working set may actually be chosen. Because Algorithm 3 uses the SMO gain $g_B(\alpha^{(t)})$ for working set selection the gain may only improve due to the choice of $B^{(t+1)} \neq B^{(t-1)}$. Therefore $g^{2\text{-step}}$ is even lower bounded by \tilde{g}^* . With $1 - \eta^2 \leq 1$ the desired bound follows. □

The first case seems to complicate things unnecessarily. Further, it reduces the progress by a factor of $1 - \eta^2$. We could simply skip the second if-condition in Algorithm 3 and in all cases turn to the else-part. From a purely mathematical point of view this is clearly true. However, the usage of the exact gain $g_B(\alpha)$ instead of $\tilde{g}_B(\alpha)$ is an unfavorable choice for working set selection in practice. For performance reasons we want to allow the algorithm to use the working set selection objective $\tilde{g}_B(\alpha)$ as often as possible. Thus we have to cover case 1 in Lemma 3, too.

Theorem 2. *Let $\alpha^{(t)}$ denote the sequence of feasible points produced by the PA-SMO algorithm starting from $\alpha^{(0)}$ and working at perfect accuracy $\varepsilon = 0$. Then the algorithm either stops in finite time with an optimal solution or produces an infinite sequence with $\lim_{t \rightarrow \infty} f(\alpha^{(t)}) = f^*$.*

Proof. Because the algorithm checks the exact KKT conditions the finite stopping case is trivial. For the infinite case we distinguish two cases. If the sequence contains only finitely many steps which are planning ahead then there exists $t_0 > 0$ such that in all iterations $t > t_0$ the algorithm coincides with Algorithm 1 and the convergence proof given in [2] holds. Otherwise there exists an infinite sequence $(t_n)_{n \in \mathbb{N}}$ of planning steps. Now there are at least two possibilities to apply the above results. An easy one is as follows: From Lemma 3 we obtain a constant $c = 1 - \eta^2$ such that Theorem 1 implies the desired property. Alternatively we can argue that the double-step gain is non-negative by Lemma 3. Algorithm 4 ensures that the SMO steps in iterations $t_n - 1$, $n \in \mathbb{N}$ just before the planning-ahead steps are free. Then we can apply Corollary 1. However, the second variant of the proof does not hold if we replace Algorithm 4 by Algorithm 2. □

As already noted above Theorem 1 and Corollary 1 resolve the separate handling of different cases by the algorithm in a general manner. In the case of an infinite sequence of planning-ahead steps the proof does not consider the other iterations at all. This technique is similar to the convergence proof for the Hybrid Maximum-Gain second order algorithm presented in [5] which needs to cover different cases to ensure convergence, too.

7 Experiments

The main emphasis of the experiments is to compare the PA-SMO algorithm with the standard (greedy) SMO algorithm. The most recent LIBSVM

version 2.84 implements Algorithm 1. For comparison, we implemented the modifications described in Algorithm 3 and Algorithm 4 directly into LIBSVM.

Note that in the first iteration starting from $\alpha^{(0)} = (0, \dots, 0)^T$ the components $y_i = \pm 1$ of the gradient $\nabla f(\alpha^{(0)}) = y$ take only two possible values. The absolute values of these components are equal and they all point into the box. Therefore the working set selection algorithm could select any $i^{(1)} \in I_{\text{up}}(\alpha^{(0)})$ as the first index, because the gradient components of all indices are maximal. Thus, there is a freedom of choice for the first iteration. LIBSVM arbitrarily chooses $i^{(1)} = \max(I_{\text{up}}(\alpha^{(0)}))$. Of course, this choice influences the path taken by the optimization. Experiments indicate that this choice can have a significant impact on the number of iterations and the runtime of the algorithm. Now, on a fixed dataset, an algorithm may appear to be superior to another one just because it is lucky to profit more from the asymmetry than the competitor. To reduce random effects, we created 100 random permutations of each dataset. All measurements reported are mean values over these 100 permutations. Because the permutations were drawn i.i.d. we can apply standard significance tests to our measurements.

We collected a set of 22 datasets for the performance comparison. For the 13 benchmark datasets from [14] we merged training and test sets. The artificial chess-board problem [4] was considered because it corresponds to quadratic programs which are very difficult to solve for SMO-type decomposition algorithms. Because this problem is described by a known distribution, we are in the position to sample datasets of any size from it. We arbitrarily fixed three datasets consisting of 1,000, 10,000, and 100,000 examples. Six more datasets were taken from the UCI benchmark repository [11]: The datasets `connect-4`, `king-rook-vs-king`, and `tic-tac-toe` are extracted from games, while `ionosphere`, `spambase`, and `internet-ads` stem from real world applications.

In all experiments we use the Gaussian kernel

$$k(x_i, x_j) = \exp(-\gamma \|x_i - x_j\|^2)$$

with the single kernel parameter $\gamma > 0$. The complexity control parameter C and the kernel parameter γ were selected with grid search on the cross-validation error to ensure that the parameters are in a regime where the resulting classifiers generalize reasonably well, see Table 1. All experiments were carried out on a Xeon 3 GHz CPU running Fedora Linux.

7.1 Results

We performed 100 runs (corresponding to the 100 permutations) per dataset for both algorithms and measured the runtime and the number of iterations. The results are summarized in Table 2.

There is a clear trend in these results. For some datasets the PA-SMO algorithm significantly outperforms the SMO algorithm, while for other datasets there is no significant difference. Most important, PA-SMO performs *in no case* worse than standard SMO.

The number of iterations is significantly reduced in nearly all cases. This result is not surprising. It basically means that the algorithm works as ex-

dataset	ℓ	C	γ	SV	BSV
banana	5,300	100	0.25	1,223	1,199
breast-cancer	277	0.6	0.1	178	131
diabetis	768	0.5	0.05	445	414
flare-solar	1,066	1.5	0.1	744	709
german	1,000	1	0.05	620	426
heart	270	1	0.005	158	149
image	2,310	100	0.1	301	84
ringnorm	7,400	2	0.1	625	86
splice	3,175	10	0.01	1,426	7
thyroid	215	500	0.05	17	3
titanic	2,201	1,000	0.1	934	915
twonorm	7,400	0.5	0.02	734	662
waveform	5,000	1	0.05	1,262	980
chess-board-1000	1,000	1,000,000	0.5	41	3
chess-board-10000	10,000	1,000,000	0.5	129	84
chess-board-100000	100,000	1,000,000	0.5	556	504
connect-4	61,108	4.5	0.2	13,485	5,994
king-rook-vs-king	28,056	10	0.5	5,815	206
tic-tac-toe	958	200	0.02	104	0
internet-ads	2,358	10	0.03	1,350	6
ionosphere	351	3	0.4	190	8
spam-database	4,601	10	0.005	1,982	583

Table 1: Datasets used for the comparison. The dataset size, the regularization parameter C and the kernel parameter γ are given. The last two columns list the resulting total number of support vectors and the number of bounded support vectors. Due to the finite accuracy of the solutions these mean values are not always integers. For clarity we provide rounded values.

dataset	time		iterations		
	SMO	PA-SMO	SMO	>	PA-SMO
banana	2.07	2.08	23295	>	19721
breast-cancer	0.02	0.02	313	>	292
diabetis	0.08	0.10	361	>	358
flare-solar	0.18	0.19	792	>	744
german	0.20	> 0.19	908	>	879
heart	0.02	0.02	113		112
image	0.45	0.46	6553	>	6359
ringnorm	2.41	2.27	1569	>	1537
splice	4.04	> 3.92	6643	>	5854
thyroid	0.02	0.01	744	>	667
titanic	0.54	> 0.47	3375	>	1653
twonorm	2.67	2.65	641		642
waveform	3.03	2.99	1610	>	1539
chess-board-1000	3.86	> 2.98	1883310	>	1186963
chess-board-10000	76.72	75.36	32130476	>	24997371
chess-board-100000	475.37	> 428.18	145364030	>	105199379
connect4-0.2	1268.04	1243.56	82076	>	77690
king-rook-vs-king	272.80	273.06	69410	>	64067
tic-tac-toe	0.10	0.10	8321	>	7786
internet-ads	2.38	2.31	2785	>	2750
ionosphere	0.03	0.04	411	>	408
spambase	8.36	8.36	9641	>	9171

Table 2: Comparison of standard SMO (Algorithm 1) and planning-ahead SMO (Algorithm 5). Mean time in seconds and number of iterations are listed. The “>” sign indicates that the left value is statistically significantly larger than the right value (paired Wilcoxon rank sum test, $p = 0.05$ over 100 permutations of the datasets). The left value is in no case significantly smaller than the right one.

pected. However, early iterations working on the whole problem take much longer than late iterations after shrinking has more or less identified the interesting variables. Therefore it is natural that the number of iterations is only a weak indicator for the runtime. The runtime of the PA-SMO algorithm is usually slightly reduced in the mean. This difference is significant in 5 cases. However, the striking argument for the algorithm is that it never performs worse than the standard SMO algorithm.

Although both algorithms use the same stopping condition the dual objective values achieved slightly varies. A careful check of these values reveals that the PA-SMO algorithm consistently achieves better solutions (paired Wilcoxon rank sum test, $p = 0.05$) for all datasets but `chess-board-100000`. Thus, the speed up is not a trivial effect of reduced solution quality. The tests reveal that the contrary is the case, that is, the new algorithm outputs better solutions in less time.

7.2 Influence of Planning-Ahead vs. Working Set Selection

It is interesting to look a little bit behind the scenes. Recall that we changed two parts of the SMO algorithm. The truncated Newton step was replaced by the planning-ahead Algorithm 4 and the working set selection was modified accordingly by Algorithm 3. It is possible to use the second modification without the first one, but hardly vice versa. Therefore, we ran the SMO algorithm with the modified working set selection but without planning ahead to get a grip on the influence of these changes on the overall performance. That is, we made sure that the algorithm selects the working set used two iterations ago if it is a feasible direction and maximizes the Newton step gain \tilde{g} . While the results of the comparison to standard SMO were completely ambiguous, the PA-SMO algorithm turned out to be clearly superior. Thus, the reason for the speed up of PA-SMO is not the changed working set selection, but planning-ahead.

7.3 Planning-Ahead Step Sizes

To understand how planning-ahead is really used by the algorithm we measured the quantity $\mu/\mu^* - 1$, that is, the size of the planning-ahead step relative to the Newton step. For free SMO steps this quantity is always 0, for larger steps it is positive, for smaller steps negative, and for steps in the opposite direction it is even smaller than -1 . We present some representative histograms in Figure 3. These histograms reveal that most planning-steps are only slightly increased compared to the Newton step size, but there are cases where the algorithm chooses a step which is enlarged by a factor of several thousands. However, very few steps are reduced or even reversed, if any.

Obviously the step size histograms are far from symmetric. Therefore it is natural to ask whether a very simple increase of the Newton step size can be a good strategy. By heretically using

$$\mu^{(t)} = \max \left\{ \min \left\{ 1.1 \cdot \frac{l_t}{Q_{tt}}, \tilde{U}_t \right\}, \tilde{L}_t \right\}$$




Figure 3: Histograms (number of iterations) of the planning-step size μ divided by the Newton step size μ^* , minus 1. On both axes a logarithmic scale is used to increase the resolution for small values (to achieve this effect for the x -axis we used the parameterization $t \mapsto \text{sign}(t) \cdot (10^{t^2/2} - 1)$ which is symmetric around the Newton step size corresponding to the origin of $t = \mu/\mu^* - 1$, with a high resolution around this point). The rightmost bin counts all steps which exceeded the scale, which is actually the case for the `chess-board-1000` dataset.

instead of equation (2) we still achieve $1 - 0.1^2 = 99\%$ of the SMO gain in each iteration (see Figure 2) and avoid the drawback of the more complicated computations involved when planning-ahead. Further, this strategy can be implemented into an existing SMO solver in just a few seconds. Experiments indicate that it is surprisingly successful, no matter if the original working set selection or Algorithm 3 is used. For most simple problems it performs as good as the much more refined PA-SMO strategy. However, for the extremely difficult `chess-board` problem this strategy performs significantly worse.

7.4 Multiple Planning-Ahead

We now turn to the variant of the PA-SMO algorithm which uses more than one recent working set for planning-ahead. This variant, as explained at the

end of section 5, plans ahead with multiple candidate working sets. Further, these working sets are additional candidates for the working set selection. We can expect that the number of iterations decreases the more working sets are used this way. However, the computations per iteration of course increase, such that too many working sets will slow the entire algorithm down. Thus, there is a trade-off between the number of iterations and the time needed per iterations. Now the interesting question is whether there is a uniform best number of working sets for all problems.




Figure 4: The curves show the performance (mean runtime) of the PA-SMO algorithm with 1, 2, 3, 5, 10, and 20 most recent working sets. All results are normalized with the runtime of the standard variant (only the most recent working set is used for planning ahead). Only datasets with runtimes above 100 ms are plotted, because otherwise the deviation due to the measurement accuracy shadows the effect.

We performed experiments with the 2, 3, 5, 10, and 20 most recent working sets. It turned out that the strategies considering the most recent two or three working sets perform comparable to standard PA-SMO, and even slightly better. For 5, and more drastically, for 10 or 20 working set evaluations the performance drops, see Figure 4. This result makes clear that we do not lose much when completely ignoring the multiple working set selection strategy, and at the same time we stay at the safe side. Therefore it seems reasonable to stick to the standard form of the PA-SMO algorithm. On the other hand we can get another small improvement if the two or three most recent working sets are considered.

8 Conclusion

We presented the planning-ahead SMO algorithm (PA-SMO), which is a simple yet powerful improvement of SMO. At a first glance it is surprising that the truncated Newton step used in all existing variants of the SMO algorithm can be outperformed. This becomes clear from the greedy character of decomposition iterations. The experimental evaluation clearly shows the benefits of the new algorithm. As we never observed a decrease in performance, we recommend PA-SMO as the default algorithm for SVM training. PA-SMO is easy to implement based on existing SMO solvers. Due to the guaranteed convergence of the algorithm to an optimal solution the method is widely applicable. Further, the convergence proof introduces a general technique to address the convergence of hybrid algorithms.

References

- [1] P.-H. Chen, R.-E. Fan, and C.-J. Lin. A Study on SMO-type Decomposition Methods for Support Vector Machines. *IEEE Transactions on Neural Networks*, 17:893–908, 2006.
- [2] R.-E. Fan, P.-H. Chen, and C.-J. Lin. Working Set Selection using the Second Order Information for Training Support Vector Machines. *Journal of Machine Learning Research*, 6:1889–1918, 2005.
- [3] T. Glasmachers. On related violating pairs for working set selection in SMO algorithms. Submitted to European Symposium on Artificial Neural Networks (ESANN), 2008.
- [4] T. Glasmachers and C. Igel. Gradient-based Adaptation of General Gaussian Kernels. *Neural Computation*, 17(10):2099–2105, 2005.
- [5] T. Glasmachers and C. Igel. Maximum Gain Working Set Selection for SVMs. *Journal of Machine Learning Research*, 7:1437–1466, 2006.
- [6] D. Hush and C. Scovel. Polynomial-time Decomposition Algorithms for Support Vector Machines. *Machine Learning*, 51:51–71, 2003.
- [7] T. Joachims. Making Large-Scale SVM Learning Practical. In B. Schölkopf, C. Burges, and A. Smola, editors, *Advances in Kernel Methods – Support Vector Learning*, chapter 11, pages 169–184. MIT Press, 1999.
- [8] S. S. Keerthi and E. G. Gilbert. Convergence of a Generalized SMO Algorithm for SVM Classifier Design. *Machine Learning*, 46:351–360, 2002.
- [9] N. List, D. Hush, C. Scovel, and I. Steinwart. Gaps in Support Vector Optimization. In *Learning Theory, 20th Annual Conference on Learning Theory, COLT 2007. Lecture Notes in Computer Science*, volume 4539, pages 336–348, 2007.

Algorithm 5: The complete PA-SMO Algorithm

Input: feasible initial point $\alpha^{(0)}$, accuracy $\varepsilon \geq 0$, $\eta \in (0, 1)$
 $G^{(0)} \leftarrow \nabla f(\alpha^{(0)}) = y - K\alpha^{(0)}$, $p \leftarrow \text{true}$, $t \leftarrow 1$

do

if $p = \text{true}$ **then**

$i^{(t)} \leftarrow \arg \max \{G_n^{(t-1)} \mid n \in I_{\text{up}}(\alpha^{(t-1)})\}$
 $j^{(t)} \leftarrow \arg \max \{\tilde{g}_{(i^{(t)}, n)}(\alpha^{(t-1)}) \mid n \in I_{\text{down}}(\alpha^{(t-1)}) \setminus \{i\}\}$
 $B^{(t)} \leftarrow (i^{(t)}, j^{(t)})$

else

if $1 - \eta \leq \mu^{(t-1)}/\mu^* \leq 1 + \eta$ **then**

$i^{(t)} \leftarrow \arg \max \{G_n^{(t-1)} \mid n \in I_{\text{up}}(\alpha^{(t-1)})\}$
 $j^{(t)} \leftarrow \arg \max \{\tilde{g}_{(i^{(t)}, n)}(\alpha^{(t-1)}) \mid n \in I_{\text{down}}(\alpha^{(t-1)}) \setminus \{i\}\}$
 $B^{(t)} \leftarrow (i^{(t)}, j^{(t)})$
 if $\tilde{g}_{B^{(t-2)}}(\alpha^{(t-1)}) > \tilde{g}_{B^{(t)}}(\alpha^{(t-1)})$ **then** $B^{(t)} \leftarrow B^{(t-2)}$

else

$i^{(t)} \leftarrow \arg \max \{G_n^{(t-1)} \mid n \in I_{\text{up}}(\alpha^{(t-1)})\}$
 $j^{(t)} \leftarrow \arg \max \{g_{(i^{(t)}, n)}(\alpha^{(t-1)}) \mid n \in I_{\text{down}}(\alpha^{(t-1)}) \setminus \{i\}\}$
 $B^{(t)} \leftarrow (i^{(t)}, j^{(t)})$
 if $g_{B^{(t-2)}}(\alpha^{(t-1)}) > g_{B^{(t)}}(\alpha^{(t-1)})$ **then** $B^{(t)} \leftarrow B^{(t-2)}$

end

end

$\mu^* \leftarrow l_t/Q_{tt} = (v_{B^{(t)}}^T G^{(t-1)}) / (v_{B^{(t)}}^T K v_{B^{(t)}})$

if *previous iteration performed a free SMO step* **then**

 Compute the planning-ahead step size $\mu = \frac{Q_{22}w_1 - Q_{12}w_2}{\det(Q)}$
 (eq. (8))
 assuming $B^{(t-1)}$ as the next working set

if *the current or the planned step ends at the box boundary* **then**

then

$\mu^{(t)} \leftarrow \max \left\{ \min \left\{ \mu^*, \tilde{U}_t \right\}, \tilde{L}_t \right\}$ (eq. (2))
 $p \leftarrow \text{false}$

else

$\mu^{(t)} \leftarrow \mu$
 $p \leftarrow \text{true}$

end

else

$\mu^{(t)} \leftarrow \max \left\{ \min \left\{ \mu^*, \tilde{U}_t \right\}, \tilde{L}_t \right\}$ (eq. (2))
 $p \leftarrow \text{false}$

end

$\alpha^{(t)} \leftarrow \alpha^{(t-1)} + \mu^{(t)} \cdot v_{B^{(t)}}$
 $G^{(t)} \leftarrow G^{(t-1)} - \mu^{(t)} K v_{B^{(t)}}$

stop if
 $\left(\max \{G_i^{(t)} \mid i \in I_{\text{up}}(\alpha^{(t)})\} - \min \{G_j^{(t)} \mid j \in I_{\text{down}}(\alpha^{(t)})\} \right) \leq \varepsilon$
 $t \leftarrow t + 1$

loop;

- [10] N. List and H. U. Simon. A General Convergence Theorem for the Decomposition Method. In J. Shawe-Taylor and Y. Singer, editors, *Proceedings of the 17th Annual Conference on Learning Theory, COLT 2004*, volume 3120 of *LNCS*, pages 363–377. Springer-Verlag, 2004.
- [11] D. J. Newman, S. Hettich, C. L. Blake, and C. J. Merz. UCI Repository of machine learning databases, 1998. <http://www.ics.uci.edu/mlearn/MLRepository.html>.
- [12] E. Osuna, R. Freund, and F. Girosi. An Improved Training Algorithm for Support Vector Machines. In J. Principe, L. Giles, N. Morgan, and E. Wilson, editors, *Neural Networks for Signal Processing VII*, pages 276–285. IEEE Press, 1997.
- [13] J. Platt. Fast Training of Support Vector Machines using Sequential Minimal Optimization. In B. Schölkopf, C. J. C. Burges, and A. J. Smola, editors, *Advances in Kernel Methods - Support Vector Learning*, chapter 12, pages 185–208. MIT Press, 1999.
- [14] G. Rätsch, T. Onoda, and K.-R. Müller. Soft Margins for AdaBoost. *Machine Learning*, 42(3):287–320, 2001.
- [15] N. Takahashi and T. Nishi. Rigorous Proof of Termination of SMO Algorithm for Support Vector Machines. *IEEE Transaction on Neural Networks*, 16(3):774–776, 2005.