

A Walk from 2-Norm SVM to 1-Norm SVM

Jussi Kujala and Timo Aho and Tapio Elomaa
Department of Software Systems
Tampere University of Technology
P. O. Box 553, FI-33101 Tampere, Finland
Email: *firstname.lastname@tut.fi*

Abstract—This paper studies how useful the standard 2-norm regularized SVM is in approximating the 1-norm SVM problem. To this end, we examine a general method that is based on iteratively re-weighting the features and solving a 2-norm optimization problem. The convergence rate of this method is unknown. Previous work indicates that it might require an excessive number of iterations. We study how well we can do with just a small number of iterations. In theory the convergence rate is fast, except for coordinates of the current solution that are close to zero. Our empirical experiments confirm this. In many problems with irrelevant features, already one iteration is often enough to produce accuracy as good as or better than that of the 1-norm SVM. Hence, it seems that in these problems we do not need to converge to the 1-norm SVM solution near zero values. The benefit of this approach is that we can build something similar to the 1-norm regularized solver based on any 2-norm regularized solver. This is quick to implement and the solution inherits the good qualities of the solver such as scalability and stability.

I. INTRODUCTION

Minimizing empirical error over some model class is a basic machine learning approach which is usually complemented with regularization to counterattack overfitting. The support vector machine (SVM) is an approach for building a linear separator, which performs well in tasks such as letter recognition and text categorization. In this paper the linear SVM is taken to solve the following minimization problem:

$$\min_w \underbrace{\frac{1}{2} \|w\|^2}_{\text{regularizer}} + C \underbrace{\sum_{i=1}^m \text{loss}(y_i, f_w(x_i))}_{\text{error}}. \quad (1)$$

Here $\{x_i, y_i\}_{i=1}^m \in \mathbb{R}^d \times \{-1, 1\}$ is a training set of examples x_i with binary labels y_i . The classifier $f_w(x)$ that we wish to learn is $w \cdot x$ (plus unregularized bias, if necessary), where $w \in \mathbb{R}^d$ is the normal of the separating hyperplane. The function $\text{loss}(y, f(x))$ is the hinge loss $\max(0, 1 - y f(x))$ (L1-SVM), or its square (L2-SVM).

The squared 2-norm $\|w\|^2$ is not always the best choice as the regularizer. In principle the 1-norm $\|w\|_1 = \sum_{i=1}^d |w_i|$ can handle a larger number of irrelevant features before overfitting [1]. In the context of least-squares regression Tibshirani [2] gives evidence that L1-regularization (lasso regression) is particularly well-suited when the problem domain has small to medium number of relevant features.

However, for a very small number of relevant features a subset selection method outperformed L1-regularization in these experiments and for a large number of relevant features the L2-regularization (ridge regression) was the best.

In a 1-norm SVM [3] the regularizer is the 1-norm $\|w\|_1$. The resulting classifier is a linear classifier without an embedding to an implicit high-dimensional space given by a non-linear kernel. Some problem domains do not require a non-linear kernel function. For example, this could be the case if the input dimension is already large [4]. Furthermore, we can try to map the features to an explicit high-dimensional space if the linear classifier on original features is not expressive enough.

In this paper we study a simple iterative scheme which approaches the 1-norm SVM by solving a series of standard 2-norm SVM problems. Each 2-norm SVM solves a problem where the features are weighted depending on the solution of the previous 2-norm SVM problem. Hence, we will refer to this algorithm as the *re-weighting algorithm* (RW). More generally, we can apply it to minimize any convex error function regularized with 1-norm.

In this scheme most of the complexity resides in the regular SVM solver. Hence, the desired features of any standard SVM solver are readily available. Such features include performance, scalability, stability, and minimization of different objective functions (like L1-SVM and L2-SVM). Several fast approximate solvers for linear SVMs have been proposed recently [5], [6]. They are sufficiently quick to justify solving several linear SVM problems for a single 1-norm SVM problem. For example, Pegasos [5] trains the Reuters data set with ca. 800,000 examples and 47,000 sparse features in five seconds (not counting the time to read the data to memory).

Our contribution is three-fold. First, we provide theoretical results on the speed of the convergence. It is known that similar optimization methods are equivalent to the 1-norm solution [7]. Unfortunately, the convergence rate is unknown. Neither we are able to prove hard bounds on it, but we can, though, provide intuition on the behavior of the convergence. More precisely, we will lower bound the decrease of the 1-norm objective in one iteration. This lower bound is higher when the current solution is poor in terms of the 1-norm objective function. On the other hand, the

bound also depends on our current solution: it shows that the convergence is slow on coordinates that are already near zero.

As the second contribution we experimentally demonstrate the efficiency of the resulting algorithm in the specific applications of the SVMs. Because theoretical results do not guarantee the speed of convergence, we experiment on how many iterations one needs to run the algorithm. Previous work [8] has suggested that a similar algorithm needs up to 200 iterations to converge in the case of the multinomial logistic regression. However, the experiments are complicated by the fact that minimizing the objective is merely a proxy of the real target—generalization accuracy. When measuring accuracy, the 1-norm SVM solution does not necessarily give the best performance on problems with irrelevant features. Each iteration in the RW algorithm solves an optimization problem. Thus, the accuracy of the solution given by any iteration may be good even if it has not yet reached the 1-norm SVM optimum. In fact, previous work [9] argues that the best performance is often somewhere between the 1-norm and 2-norm optima.

Finally, we provide a patch to `liblinear` [6] which implements the RW algorithm.

The structure of this paper is as follows. Section II presents the algorithm and provides both intuition and theory on why it works. Section III concerns empirical behavior of the algorithm, including results on the speed of convergence and how the number of relevant features affects the performance. In Section IV we discuss previous work and performance. Finally, Section V concludes this work.

II. THE RE-WEIGHTING ALGORITHM

Algorithm 1 re-weights the features in each iteration, which makes it possible to use the SVM solver as a black-box. In short, during t -th iteration the RW algorithm multiplies the feature i with a weight $v_i^{(t)}$. Then it obtains a 2-norm SVM solution w from these weighted features. The new weights $v_i^{(t+1)}$ are set to $\sqrt{|w_i v_i^{(t)}|}$.

We could improve the performance of the algorithm by tailoring the SVM solver. Here we are interested in simplicity rather than performance optimizations that have unclear value. Note that the algorithm is in fact oblivious to the choice of the error function, so the SVM solver could be either L1-SVM or L2-SVM (or any other convex error for that matter).

A. Intuition

Figure 1 gives a graphical justification for the benefits of L1-regularization over L2-regularization. In it the contour of the red tilted square finds a sparser solution, because it is more spiky in directions where the weights are zero. The blue dashed ellipse shows what effect weighting of the features has. If weighted correctly, the optimization with the

Algorithm 1 1-norm SVM via 2-norm SVM

Input: a training set $\{(x^i, y^i)\}_{i=1}^m$ and number of iterations N .

Output: a weight vector w .

Initialize vector $v^{(1)}$ to all ones.

for $t = 1$ **to** N **do**

for each example x^i **do**

for each coordinate j **do**

 Set $x_j'^i := x_j^i v_j^{(t)}$.

end for

end for

$w^{(t)} :=$

 solution to SVM with examples $\{(x'^i, y^i)\}_{i=1}^m$.

for each coordinate i **do**

 Set $v_i^{(t+1)} := \sqrt{|w_i^{(t)} v_i^{(t)}|}$.

end for

end for

for each coordinate i **do**

 Set $w_i := w_i^{(N)} v_i^{(N)}$.

end for

Return w

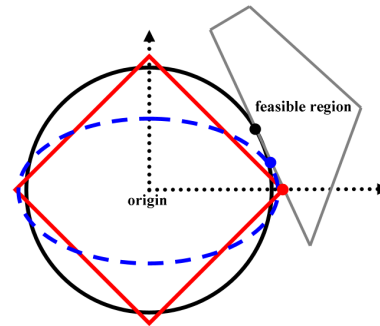


Figure 1. 2-norm contour as the black circle, 1-norm as the red tilted square, and scaled 2-norm with the blue dashed ellipse.

squeezed ellipse will approximate L1-regularized solution better than the L2-regularized one.

What is not apparent in Figure 1 is that if we know a non-optimal feasible point, then we can always choose the relative lengths of the axes of the ellipsis so that the L1-regularized objective value will decrease. This will be shown in the following theory section.

Remark: Figure 1 also suggests that the RW algorithm has difficulties to converge a coordinate to zero, because of the dull corners of the squeezed ellipse. The following theory section will also quantify this effect.

B. Theory

For vectors w and v we use $w \otimes v$ to denote an element-wise product (Hadamard product), where the i th coordinate $(w \otimes v)_i$ is $w_i v_i$. The absolute value $|w|$ of a vector w is

a vector containing the absolute values the components of the original vector: $|w|_i = |w_i|$. The error function $E(w)$ denotes the error given by a weight vector w . The modified error function $E_v(w)$ denotes the error, where features are weighted with $|v|$, i.e., $E_v(w) = E(w \otimes |v|)$. When the specific norm of $\|w\|$ is not indicated, it is always the 2-norm.

1) *1-norm objective decreases if 2-norm objective of the weighted problem decreases:* The following theorem gives a partial motivation for the minimization over weighted features. However, it does not guarantee that we find a better solution in each iteration. The theorem assumes that our current solution is a vector $v \otimes |v|$. We then minimize over the weighted problem where the weight on the i th feature is $|v_i|$. Now, one possible solution to the weighted problem is to set the solution w to v . Then the new solution to the original problem equals the previous solution $v \otimes |v|$. However, Theorem 1 tells that *if* the minimization finds a solution w to the weighted problem that has better 2-norm regularized objective function value, then $w \otimes |v|$ is a better solution to the unweighted 1-norm regularized problem.

Theorem 1. *If $u = v \otimes |v|$, $u^{new} = w \otimes |v|$, and*

$$\frac{1}{2} \|w\|^2 + E_v(w) < \frac{1}{2} \|v\|^2 + E_v(v), \quad (2)$$

then u^{new} is a better solution to the L1-regularized optimization than u :

$$\|u^{new}\|_1 + E(u^{new}) < \|u\|_1 + E(u).$$

Furthermore, the 1-norm objective decreases by at least as much as the weighted objective in (2).

The proof of Theorem 1 will appear in the full version of this paper.

2) *Convex error function guarantees that the 2-norm objective of the weighted problem decreases:* Theorem 1 does not state that we will find a vector w with smaller weighted objective value than v even if there is a solution u^* with a smaller value to the L1-regularized error. Theorem 2, though, will show that an iteration finds a solution with smaller weighted objective value. The theorem assumes that the error function is convex and that the current solution has no zero in the coordinates, where u^* has a non-zero value. Therefore, Theorems 1 and 2 together state that an iteration of the RW algorithm decreases the value of the L1-regularized objective function.

We will use a scaled 2-norm $\|w\|_u = \sum_{i=1}^d w_i^2 / |u_i|$. A substitution $w := w' \otimes |v|$ shows that regularization with $\|w\|_u / 2$ equals the weighted objective $\|w'\|^2 / 2 + E_v(w')$, if $u = v \otimes |v|$. Intuitively $\|w\|_u$ approximates $\|w\|_1$ in a neighborhood of u , which yields the following theorem.

Theorem 2. *Let h be a vector which is normalized so that $\|h\|_u = 1$. Let vector $c^* h$, where c^* is a scalar, denote*

any direction in which the L1-regularized objective value decreases when starting from u . Hence,

$$\|u + c^* h\|_1 + E(u + c^* h) < \|u\|_1 + E(u).$$

Then the weighted objective function decreases to that same direction, i.e., there is a scalar $c > 0$ such that

$$\frac{1}{2} \|u + c h\|_u + E(u + c h) + \frac{1}{2} c^2 \leq \frac{1}{2} \|u\|_u + E(u).$$

More precisely, the step size c is at least the minimum of c^ and*

$$- \left(\sum_{i=1}^d \text{sign}(u_i) h_i \right) + \frac{1}{c^*} (E(u) - E(u + c^* h)).$$

The proof of Theorem 2 is also given in the full version of this paper.

Theorem 2 gives us some insight to the speed of convergence. It and Theorem 1 together show that the 1-norm objective function decreases by at least $c^2/2$ in one iteration. Let us now derive a more intuitive approximation to the expression of the step size c . If we assume that for all i the signs of h_i and u_i differ, then $-c^* \left(\sum_{i=1}^d \text{sign}(u_i) h_i \right) = \|u\| - \|u + c^* h\|_1$. This approximation is good, if the 1-norm of the solution is an important factor in the optimization. Hence,

$$c \gtrsim \min \left(\frac{\text{Obj}(u) - \text{Obj}(u^*)}{c^*}, c^* \right),$$

where $\text{Obj}(x)$ is the 1-norm objective function. Thus, L1-regularized error drops quickly if our current solution u is poor in comparison to optimal u^* . On the other hand, c has an inverse dependence on c^* . This implies that convergence is slow along a coordinate, in which our current solution is already close to zero. Therefore, it might be impossible to obtain hard limits to the convergence rate, because the RW algorithm has trouble in converging near-zero coordinates. The next section experimentally tests how well we can manage with only a small number of iterations.

III. EMPIRICAL BEHAVIOR OF RE-WEIGHTING

A. Speed of Convergence

Let us present empirical evidence on how fast the RW algorithm converges to the 1-norm SVM solution. The experiments were performed with 16 data sets selected from UCI machine learning repository and Broad Institute Cancer Program. The data sets from UCI are abalone, glass, segmentation, australian, ionosphere, sonar, bupa liver, iris, vehicle, wine, ecoli, wisconsin, german, and page. The two data sets from Broad Institute are leukemia and DLBCL.¹

For each data set we let the regularization parameter C obtain powers of ten from 10^{-3} to 10^3 . In each iteration we

¹Available from <http://www.aillab.si/orange/>.

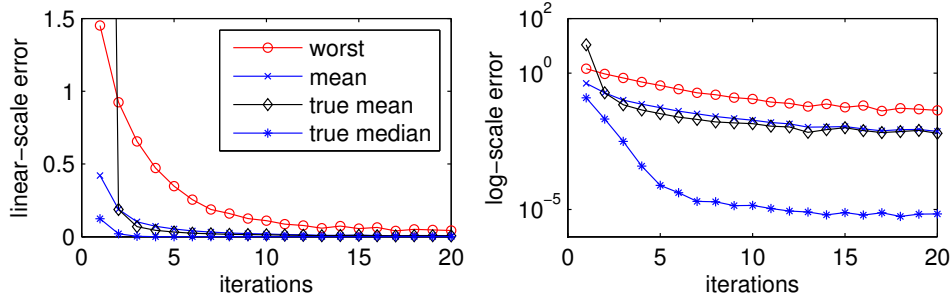


Figure 2. The behavior of the objective with different number of iterations. The first two curves are derived from worst-case performances over C on different data sets at the 20th iteration. The true mean and median are derived over all data sets and all values of the tradeoff parameter C .

recorded the objective value of the 1-norm regularized L1-SVM. For measuring the convergence we used `svmlight` [10] with default arguments for most runs (see below for a discussion on high values of C). We also measured the optimal value of the objective. For this we used a linear program and the `linprog` optimizer from Matlab (with default settings this is the `lipsol` solver).

Figure 2 gives a summary of the findings. The plotted objective value is (attained objective value – optimal value)/(optimal value). For each data set, we selected the *worst* convergence over C at the 20th iteration. From these 16 curves we formed two curves, the worst-case and the mean. Additionally, we show the mean and median over all problem domains and values of C (true mean and true median).

The plots show that in absolute terms the convergence is fast on an average problem. We can see this from the behavior of the true median. However, the worst-case curve never decreases below 0.1. The worst convergence was obtained for the gene expression data sets, which have thousands of features out of which many are zero at the optimum. Thus, the small non-zero errors over many features lead to slow convergence in the objective.

We had a problem with `svmlight` for high values of the trade-off parameter C . The objective we measure is the primal objective. However, `svmlight` actually optimizes the dual objective to a given error [10]. If zero hinge loss is attainable, then the primal objective is unstable for high values of the parameter C . This is because the difference between hinge loss of 0 and 0.01 is large after being multiplied with $C = 1,000$. However, this does not appear in the error of the dual. Therefore, we used more strict error parameters in two of the problem domains (leukemia and DLBCL) for value 100 of the parameter C .

B. Accuracy of Classifiers

Let us now turn our attention to our true objective: building accurate classifiers. In this section we chart out how quickly the accuracy changes during iterations. The experiments include both problem domains with many relevant

Table I
THE NUMBER OF EXAMPLES AND FEATURES IN EACH DATA SET.
NUMBER OF EXAMPLES IN FORMAT X/Y DENOTES X EXAMPLES IN THE GIVEN TRAINING SET AND Y EXAMPLES IN THE GIVEN TEST SET.

DATA SET	EXAMPLES	FEATURES
REUTERS	23,149/199,328	47,236
GISETTE	6,000/1,000	5,000
DLBCL	77	7,070
LEUKEMIA	72	5,147
LUNG	203	12,600
PROSTATA	102	12,533
SRBCT	83	2,308
SONAR	208	59

features and those that have many irrelevant features. This selection criterion should make sure that there is a difference between accuracies of 1-norm and 2-norm SVM.

1) *Datasets*: `Reuters`² is a text categorization data set and `Reuters-sampled` is a synthetic problem domain obtained from it; each experiment samples 250 examples from the original training set. We form a binary classification task by training the CCAT category versus all other categories. `Gisette`³ is a digit classification task, which contains many irrelevant features, because 50% of its features are synthetic noise features. We also experiment with several gene expression data sets⁴ which should have many irrelevant features. Recall that the gene expression data sets had a slow convergence in the experiments of the previous section. The data sets are DLBCL, Leukemia, Lung, Prostate, and SRBCT. Additionally we experiment on classical Sonar data set from UCI. Table I summarizes the properties of these data sets.

2) *Experimental setup*: For `Reuters` and `Gisette` we use the given split into training set and a test or validation set. For the other domains we perform 30 experiments, in which we randomly split the data set half and half into

²Available from <http://jmlr.csail.mit.edu/papers/volume5/lewis04a/>

³Available from <http://www.nipsfsc.ecs.soton.ac.uk/datasets/>

⁴Available from <http://www.ailab.si/orange/>

Table II

CLASSIFICATION ACCURACIES OF THE RW ALGORITHM AND 1-NORM SVM RUNS WITH DIFFERENT NUMBERS OF ITERATIONS. THE EMPIRICAL STANDARD DEVIATION OF THE ACCURACY IS GIVEN. THE t -TH ITERATION OF THE RW ALGORITHM IS DENOTED BY RW(t). FOR GISETTE AND REUTERS OUR 1-NORM SVM SOLVER RUNS OUT OF MEMORY.

DATA SET	2-NORM SVM	RW(2)	RW(3)	RW(5)	RW(10)	1-NORM SVM
REUTERS	93.4	93.5	93.5	93.4	93.3	NA
REUTERS-SAMPLED	85.9 ± 0.2	84.9 ± 0.3	83.7 ± 0.3	82.2 ± 0.3	80.6 ± 0.3	72.8 ± 0.8
GISETTE	97.8	98.3	98.5	98.2	98.1	NA
DLBCL	68.6 ± 1.0	91.4 ± 0.7	93.2 ± 0.7	93.5 ± 0.8	91.7 ± 1.1	85.3 ± 1.3
LEUKEMIA	82.7 ± 1.0	95.2 ± 0.6	96.0 ± 0.5	95.8 ± 0.6	94.7 ± 0.8	90.1 ± 1.5
LUNG	92.9 ± 0.4	95.1 ± 0.3	95.0 ± 0.4	94.5 ± 0.5	93.8 ± 0.6	92.6 ± 0.5
PROSTATA	90.8 ± 0.6	91.4 ± 0.5	92.0 ± 0.5	92.5 ± 0.5	91.6 ± 0.6	88.5 ± 0.8
SRBCT	88.8 ± 1.1	98.1 ± 0.4	98.7 ± 0.3	98.0 ± 0.5	96.8 ± 0.6	95.7 ± 0.5
SONAR	73.3 ± 0.7	73.0 ± 0.6	73.1 ± 0.7	72.8 ± 0.8	72.6 ± 0.7	72.2 ± 0.8

a training set and a test set. We select the parameter C with 5-fold cross-validation. Different iterations of the RW algorithm may use different C . The best value for C is the rounded-down median of those values that attain the best cross-validated error. The range of C is the powers of ten in $[10^{-9}, 10^2]$ for the gene expression data sets and in $[10^{-5}, 10^5]$ for the other data sets (the best accuracy is on these intervals for all algorithms). The 2-norm SVM solver is `liblinear` [6]. We use default settings, except that the solver is set to L1-SVM dual optimizer. The default settings include an additional bias feature that has a constant value of 1. The 1-norm SVM is solved with Matlab, as in the previous section. We train the most frequent label versus the remaining labels, if the data set contains more than two labels. Each example in the data set `Gisette` is normalized to unit norm and each feature in the gene expression data sets is normalized to zero mean and unit variance.

Table II presents the results. Let us discuss a few observations. First, the difference in accuracy during iterations is small in `Reuters`, but large in `Reuters-sampled`. This suggests that if a problem domain has a large number of examples, then the regularization has only a small effect. Second, the best accuracy on the gene expression data sets is in between the solutions of 2-norm SVM and 1-norm SVM. In some of these data sets the difference between 1-norm SVM and the RW algorithm is surprisingly large. Friedman and Popescu [9] make a similar observation in their experiments with linear regression on both synthetic and proteomics data.

Of course, these experiments still leave open the question on how to determine the right number of iterations. In our experiments, a good number of iterations was easy to find with cross-validation. We already perform a cross-validation over the trade-off parameter C . Hence, we have an access to a table that gives the cross-validated error for each value of C and for each iteration.

IV. RELATED WORK AND DISCUSSION

A. Related Methods

Zhu et al. [3] put forward the 1-norm SVM and solved the optimization problem with a linear program. A linear programming solution is only available for the L1-SVM objective function, because the L2-SVM objective function is non-linear. Mangasarian [11] describes a specialized solver for the 1-norm linear programming problem. In it the problem is transformed to an unconstrained quadratic problem for which there are efficient solvers.

Breiman [12] is to the best of our knowledge the first one to suggest a RW algorithm similar to the one studied in this paper. His non-negative garotte solves a linear regression problem by first solving an ordinary least squares problem, which gives a solution w . This solution is then used to weight another regression problem, where the i -th feature is weighted with w_i . The solution to this new regression problem is limited to a positive weight vector, and the sum of these weights is constrained.

Several papers study optimizing the 1-norm regularized error with a 2-norm regularized solver. To the best of our knowledge Grandvalet and Canu [7], [13] were the first to suggest a connection between 1-norm regularization and 2-norm regularization. They showed that 1-norm regularized least squares regression equals 2-norm regularization of a certain error function. However, their work does not give the same updates as the ones in this paper. The more recent work of Argyriou et al. [14], though, implies the updates that we use.

B. Discussion on Performance

The iterative RW algorithm relies on a 2-norm SVM solver. Hence, the performance of this solver is important. Recently several fast approximate solvers for the linear 2-norm SVM have been developed, such as `Svm-perf` [15], `OCAS` [16], online stochastic gradient descent algorithms [5], [17], and `liblinear` [6]. These algorithms scale easily to large data sets containing hundreds

of thousands training examples. Typically the training time is bounded by the time needed to read the input.

In this paper we did not perform comprehensive experiments on the run time of the RW algorithm. Instead we gave evidence on how many iterations the algorithm needs. Thus, we can approximate the run time in units of “2-norm SVM problem”. This is more informative than measuring run times which are influenced by factors such as the termination criteria of the optimization and whether we use a subsample of the training set.

The experiments in Section III were performed by repeatedly calling the same implementation of either `svmlight` or `liblinear` with differently weighted inputs. However, we also integrated the RW algorithm directly to `liblinear` to assure ourself that our intuitions on performance are correct. As an example, ca. 200,000 examples from the Reuters data set took 33 seconds to train with a 2-norm SVM and 38 seconds to train with one additional weighted iteration (we set C to one). The fact that reading the data from a file took 27 seconds explains the small difference between these run-times. The computer on which we ran the experiments was a 2,8 Ghz Pentium 4 with 1 GB of main memory.

V. CONCLUSION

We have studied how a simple iterative re-weighting algorithm performs in problem domains with irrelevant features. In theory the re-weighting algorithm converges to a value that is close to the 1-norm SVM solution. The experimental results indicated that a small number of iterations is enough to attain the best accuracy. In fact, in many problem domains the re-weighting algorithm outperformed the 1-norm SVM and the standard 2-norm SVM. However, a close convergence to 1-norm SVM might require many more iterations. This work suggested that we can use popular 2-norm SVM solvers to derive a solver that is more resilient to irrelevant features. Hence, the good properties of these solvers are available for problem domains that contain such.

REFERENCES

- [1] A. Y. Ng, “Feature selection, L_1 vs. L_2 regularization, and rotational invariance,” in *Proceedings of the 21st International Conference on Machine Learning*. New York, NY: ACM, 2004, pp. 78–85.
- [2] R. Tibshirani, “Regression shrinkage and selection via the lasso,” *Journal of the Royal Statistical Society, Series B*, vol. 58, pp. 267–288, 1996.
- [3] J. Zhu, S. Rosset, T. Hastie, and R. Tibshirani, “1-norm support vector machines,” in *Advances in Neural Information Processing Systems*, S. Thrun, L. Saul, and B. Schölkopf, Eds. Cambridge, MA: MIT Press, 2004, vol. 16.
- [4] C. W. Hsu, C. C. Chang, and C. J. Lin, “A practical guide to support vector classification,” Department of Computer Science and Information Engineering, National Taiwan University, Taipei, Tech. Rep., 2003.
- [5] S. Shalev-Shwartz, Y. Singer, and N. Srebro, “Pegasos: Primal Estimated sub-GrAdient SOLver for SVM,” in *Proceedings of the 24th International Conference on Machine Learning*. New York, NY: ACM, 2007, pp. 807–814.
- [6] C.-J. Hsieh, K.-W. Chang, C.-J. Lin, S. S. Keerthi, and S. Sundararajan, “A dual coordinate descent method for large-scale linear SVM,” in *Proceedings of the 25th International Conference on Machine Learning*. New York, NY: ACM, 2008, pp. 408–415.
- [7] Y. Grandvalet, “Least absolute shrinkage is equivalent to quadratic penalization,” in *Perspectives in Neural Computing*, vol. 1. Springer, 1998, pp. 201–206.
- [8] M. Schmidt, G. Fung, and R. Rosales, “Fast optimization methods for L1 regularization: A comparative study and two new approaches,” in *Proceedings of the 18th European Conference on Machine Learning*. Berlin, Heidelberg: Springer-Verlag, 2007, pp. 286–297.
- [9] J. H. Friedman and B. E. Popescu, “Gradient directed regularization for linear regression and classification,” Stanford University, Tech. Rep., 2004.
- [10] T. Joachims, “Making large-scale support vector machine learning practical,” in *Advances in Kernel Methods: Support Vector Learning*. Cambridge, MA, USA: MIT Press, 1999, pp. 169–184.
- [11] O. L. Mangasarian, “Exact 1-norm support vector machines via unconstrained convex differentiable minimization,” *Journal of Machine Learning Research*, vol. 7, pp. 1517–1530, 2006.
- [12] L. Breiman, “Better subset regression using the nonnegative garrote,” *Technometrics*, vol. 37, no. 4, pp. 373–384, 1995.
- [13] Y. Grandvalet and S. Canu, “Outcomes of the equivalence of adaptive ridge with least absolute shrinkage,” in *Advances in Neural Information Processing Systems*, D. A. C. Michael J. Kearns, Sara A. Solla, Ed. Cambridge, MA, USA: MIT Press, 1999, vol. 11, pp. 445–451.
- [14] A. Argyriou, T. Evgeniou, and M. Pontil, “Multi-task feature learning,” in *Advances in Neural Information Processing Systems*, B. Schölkopf, J. Platt, and T. Hoffman, Eds. Cambridge, MA: MIT Press, 2007, vol. 19, pp. 41–48.
- [15] T. Joachims, “Training linear SVMs in linear time,” in *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. New York, NY: ACM, 2006, pp. 217–226.
- [16] V. Franc and S. Sonnenburg, “Optimized cutting plane algorithm for support vector machines,” in *Proceedings of the 25th International Conference on Machine Learning*. New York, NY: ACM, 2008, pp. 320–327.
- [17] L. Bottou and O. Bousquet, “The tradeoffs of large scale learning,” in *Advances in Neural Information Processing Systems*, J. Platt, D. Koller, Y. Singer, and S. Roweis, Eds. NIPS Foundation, 2008, vol. 20, pp. 161–168.