
Compressed Support Vector Machines

Zhixiang (Eddie) Xu
xuzx@cse.wustl.edu

Jacob R. Gardner
gardner.jake@gmail.com

Stephen Tyree
swtyree@wustl.edu

Kilian Q. Weinberger
kilian@wustl.edu

Department of Computer Science & Engineering
Washington University in St. Louis
St. Louis, MO, USA

Abstract

Support vector machines (SVM) can classify data sets along highly non-linear decision boundaries because of the kernel-trick. This expressiveness comes at a price: During test-time, the SVM classifier needs to compute the kernel inner-product between a test sample and all support vectors. With large training data sets, the time required for this computation can be substantial. In this paper, we introduce a post-processing algorithm, which *compresses* the learned SVM model by reducing and optimizing support vectors. We evaluate our algorithm on several medium-scaled real-world data sets, demonstrating that it maintains high test accuracy while reducing the test-time evaluation cost by several orders of magnitude—in some cases from hours to seconds. It is fair to say that most of the work in this paper was previously been invented by Burges and Schölkopf almost 20 years ago. For most of the time during which we conducted this research, we were unaware of this prior work. However, in the past two decades, computing power has increased drastically, and we can therefore provide empirical insights that were not possible in their original paper.

1 Introductions

Support Vector Machines (SVM) are arguably one of the great success stories of machine learning and have been used in many real world applications, including email spam classification [7], face recognition [12] and gene selection [11]. In real world applications, the evaluation cost (in terms of memory and CPU) during test-time is of crucial importance. This is particularly prominent in settings with strong resource constraints (*e.g.* embedded devices, cell phones or tablets) or frequently repeated tasks (*e.g.* webmail spam classification, web-search ranking, face detection in uploaded images), which can be performed billions of times per day. Reducing the resource requirements to classify an input can reduce hardware costs, enable product improvements, and help curb power consumption.

Test-time cost is determined mainly by two components: classifier evaluation and feature extraction cost. Reducing feature extraction cost has recently obtained a significant amount of attention [3; 5; 15; 17; 19; 23; 25; 26]. These approaches reduce the test-time cost in scenarios where features are heterogeneous, extracted on-demand, and are significantly more expensive to compute than the classifier evaluation.

In this paper, we focus on the other common scenario where the classifier evaluation cost dominates the overall test-time cost. Specifically, we focus on kernel support vector machine (SVM) [20].

Kernel computation can be expensive because it is linear in the number of support vectors and, in addition, often requires expensive exponentiation (e.g. for the radial basis or χ^2 kernels). Previous work has reduced the classifier complexity by selecting few support vectors through budgeted training [6; 24] or with heuristic selection prior to learning [13].

We describe an approach that does not *select* support vectors from the training set, but instead *learns* them to match a pre-defined SVM decision boundary. Given an existing SVM model with m support vectors, it learns $r \ll m$ “artificial support vectors”, which are not originally part of the training set. The resulting model is a standard SVM classifier (thus can be saved, for example, in a LibSVM [4] compatible file). Relative to the original model, it has comparable accuracy, but it is up to several orders of magnitudes smaller and faster to evaluate. We refer to our algorithm as Compressed Vector Machine (CVM) and demonstrate on eight real-world data sets of various size and complexity that it achieves unmatched accuracy vs. test-time cost trade offs.

2 Related Work

Burges and Schölkopf [2] invented Compressed Vector Machines long before us. While we conducted our research, we were not aware of their work until very late during the final stages of paper writing. We still consider our perspective and additional experiments valuable and decided to post our results as a techreport. However we do want to emphasize that all academic credit should go to them who were clearly ahead of us.

Reducing test-time cost has recently attracted much attention. Much work [3; 5; 10; 15; 17; 19; 23; 25] focuses on scenarios where features are extracted on-demand and the extraction cost dominates the overall test-time cost. Their objective is to minimize the feature extraction cost.

Model compression was pioneered by [1]. Our work was inspired by their vision, however it differs substantially, as we do not focus on ensembles of classifiers and instead learn a model compressor explicitly for SVMs. More recently, [26] introduces an algorithm to reduce the test-time cost specifically for the SVM classifier. However, similar to the approaches mentioned above, they focus on learning a new representations consisting of cheap non-linear features for linear SVMs.

[6] propose an algorithm to limit the memory usage for kernel based online classification. Different from our approach, their algorithm is not a post-process procedure, and instead they modify the kernel function directly to limit the amount of memory the algorithm uses. Similar to [6], [24] also focusses on online kernel SVM, and attacks primarily the training time complexity.

Of particular relevance is [13], which, specifically reduces the SVM evaluation cost by reducing the number of support vectors. Heuristics are used to select a small subset of support vectors, up to a given budget, during training time, thus solving an approximate SVM optimization. In contrast, our method is a post-processing compression to the regular SVM. We begin from an exact SVM solution and compress the set of support vectors by choosing and optimizing over a small set of support vectors to approximate the optimal decision boundary. This post-processing optimization framework renders unmatched accuracy and cost performance. Similar approaches have successfully learned pseudo-inputs for compressed nearest neighbor classification sets [14] and sparse Gaussian process regression models [22].

3 Background

Let the data consist of input vectors $\{\mathbf{x}_1, \dots, \mathbf{x}_n\} \in \mathcal{R}^d$ and corresponding labels $\{y_1, \dots, y_n\} \in \{-1, +1\}$. For simplicity we assume binary classification in the following section, but our algorithm is easily extended to multi-class settings using one-vs-one [21], one-vs-all [18], or DAG [16] approaches, and results are included for several multi-class datasets.

Kernel support vector machines. SVMs are popular for their large margin enforcement, which leads to good generalization to unseen test data, and their formulation as a convex quadratic optimization problem, guaranteeing a globally optimal solution. Most importantly, the *kernel-trick* [20] may be employed to learn highly non-linear decision boundaries for data sets that are not linearly separable. Specifically, the kernel-trick maps the original feature space \mathbf{x}_i into a higher (possibly infinite) dimensional space $\phi(\mathbf{x}_i)$.

SVMs learn a hyperplane in this higher dimensional space by maximizing the margin $\frac{1}{\|\mathbf{w}\|}$ and penalizing training instances on the wrong side of the hyperplane,

$$\min_{\mathbf{w}, b} \|\mathbf{w}\| + C \sum_i^n \left(\max(1 - y_i \mathbf{w}^\top \phi(\mathbf{x}_i) + b, 0) \right)^2, \quad (1)$$

where b is the bias, and C trades-off regularization/margin and training accuracy. Note that we use the quadratic hinge loss penalty and thus (1) is differentiable. The power of the kernel trick is that the higher dimensional space $\phi(\mathbf{x}_i)$ never needs to be expressed explicitly, because (1) can be formulated in terms of inner products between input vectors. Let a matrix \mathbf{K} denote these inner products, where $\mathbf{K}_{ij} = \phi(\mathbf{x}_i)^\top \phi(\mathbf{x}_j)$, and \mathbf{K} is the training *kernel matrix*. The optimization in (1) can be then expressed in terms of kernel matrix \mathbf{K} in the dual form:

$$\max_{\alpha_1, \dots, \alpha_n} \sum_{i=1}^n \alpha_i - \frac{1}{2} \sum_{i,j=1}^n \alpha_i \alpha_j y_i y_j \mathbf{K}_{ij}, \quad \text{s.t.} \quad \sum_{i=1}^n \alpha_i y_i = 0 \text{ and } \alpha_i \geq 0, \quad (2)$$

where α_i are the Lagrange multipliers.

the classification rule $f(\cdot)$ for a test input \mathbf{x}_t can also be expressed by testing kernel $\tilde{\mathbf{K}}$ that consists of inner products between test inputs $\mathcal{E} = \{\mathbf{x}_t\}$ and support vectors $\mathcal{S} = \{\mathbf{x}_i | \alpha_i \neq 0\}$, $\tilde{\mathbf{K}}_{it} = \phi(\mathbf{x}_i)^\top \phi(\mathbf{x}_t)$, where

$$f(\phi(\mathbf{x}_t)) = \sum_{i=1}^n \alpha_i y_i \tilde{\mathbf{K}}_{it} + b. \quad (3)$$

Note that once testing kernel $\tilde{\mathbf{K}}$ is computed, generating the prediction is merely a linear combination, and thus the dominating cost is computing the testing kernel itself.

Least angle regression. LARS [8] is a widely used forward selection algorithm because of its simplicity and efficiency. Given input vectors \mathbf{x} , target labels \mathbf{y} , and the quadratic loss $\ell(\boldsymbol{\beta}) = (\mathbf{x}\boldsymbol{\beta} - \mathbf{y})^2$, LARS learns to approximate targets by building up the coefficient vector $\boldsymbol{\beta}$ in successive steps, starting from an all-zero vector. To minimize the loss function ℓ , LARS initially descends on a coordinate direction that has the largest gradient,

$$\beta_t = \operatorname{argmax}_{\beta_t} \frac{\partial \ell}{\partial \beta_t}. \quad (4)$$

The algorithm then incorporates this coordinate into its active set. After identifying the gradient direction, LARS selects the step size very carefully. Instead of too greedy or too tiny, LARS computes a step size that a new direction *outside* of the active set has the same maximum gradient as directions *in* the active set. LARS then include this new direction into the active set.

In the following iterations, LARS gradient descends on a direction that maintains the same gradient for all directions in the active set. In other words, LARS descends following an equiangular direction of all directions in the active set. The algorithm then repeats computing step-size, including new directions into the active set, and descending on an equiangular directions. This process makes LARS very efficient, as after T iterations, LARS solution has exactly T directions in the active set, or equivalently, only T non-zero coefficients in $\boldsymbol{\beta}$.

4 Method

In this section, we detail the CVM approach to reduce the test-time SVM evaluation cost. We regard CVM as a post-processing compression to the original SVM solution. After solving an SVM, we obtain a set of support vectors $\mathcal{S} = \{\mathbf{x}_i | \alpha_i \neq 0\}$, and the corresponding Lagrange multipliers α_i . Given the original SVM solution, we can model the test-time evaluation cost explicitly.

Kernel SVM evaluation cost. Based on the prediction function (3) we can formulate the exact SVM classifier evaluation cost. Let e denote the cost of computing a test kernel entry $\tilde{\mathbf{K}}_{it}$ (*i.e.* kernel function of a test input \mathbf{x}_t and a support vector \mathbf{x}_i). We assume the computation cost is identical across all test inputs and all support vectors. As shown in (3), generating a prediction for a

testing input requires computing the kernel entry between the test input and all support vectors. The total evaluation cost is a function of the number of support vectors n_{sv} . After obtaining the kernel entries for a test point \mathbf{x}_t , prediction is simply linear combination of the kernel row $\tilde{\mathbf{K}}_t$ weighted by $\boldsymbol{\alpha}$. The cost of computing this linear combination is very low compared to the kernel computation, and therefore the total evaluation cost $c_e = n_{sv}e$. We aim to reduce the size of the support vector set n_{sv} without greatly affecting prediction accuracy.

Removing non-support vectors. Since the test-time evaluation cost is a function of the number of support vectors, the goal is to cherry-pick and optimize a subset of the optimal support vectors bounded in size by a user-specified compression ratio. We first note that all non-support vectors can be removed during this process without affecting the full SVM solution. If we define a design matrix $\hat{\mathbf{K}} \in \mathcal{R}^{n \times n}$, where $\hat{\mathbf{K}}_{ij} = y_i \mathbf{K}_{ij}$. The squared penalty SVM objective function in (1) can be expressed with Lagrange parameter $\boldsymbol{\alpha}$ and the kernel matrix \mathbf{K} :

$$\min_{\boldsymbol{\alpha}, b} \left(\max(\mathbf{1} - \hat{\mathbf{K}}\boldsymbol{\alpha} - yb, 0) \right)^2 + \boldsymbol{\alpha}^\top \mathbf{K}\boldsymbol{\alpha}. \quad (5)$$

Since (5) is a *strongly* convex function, and all non-support vectors have the corresponding Lagrange multiplier $\alpha_i = 0$, we can remove all non-support vectors from the optimization problem and the full SVM optimal solution stays the same.

To find an optimal subset of support vectors given the compression ratio, we re-train the SVM with only support vectors and a constraint on the number of support vectors. Note that $\boldsymbol{\alpha}$ are effectively the coefficients of support vectors, and we can efficiently control the number of support vectors by adding an l_0 norm on $\boldsymbol{\alpha}$. The optimization problem becomes

$$\begin{aligned} \min_{\boldsymbol{\alpha}, b} & \left(\mathbf{1} - \hat{\mathbf{K}}\boldsymbol{\alpha} - yb \right)^2 + \boldsymbol{\alpha}^\top \mathbf{K}\boldsymbol{\alpha} \\ \text{s.t.} & \|\boldsymbol{\alpha}\|_0 \leq \frac{1}{e} B_e, \end{aligned} \quad (6)$$

where B_e evaluation cost budget, and consequently, $\frac{1}{e} B_e$ is the desired number of support vectors based on the budget. Note that after removing non-support vectors, we obtain a condensed matrix $\hat{\mathbf{K}} \in \mathcal{R}^{n_{sv} \times n_{sv}}$.

Forming ordinary least squares problem. The current form of equation (6) can be made more amenable to optimization by rewriting the objective function as an ordinary least square problem. Expanding the squared term, simplifying, and fixing the bias term b (as it does not affect the solution dramatically), we re-format the objective function (6) into

$$\min_{\boldsymbol{\alpha}} (\mathbf{1} - yb)^\top (\mathbf{1} - yb) - 2\boldsymbol{\alpha}^\top \hat{\mathbf{K}}^\top (\mathbf{1} - yb) + \boldsymbol{\alpha}^\top (\hat{\mathbf{K}}^\top \hat{\mathbf{K}} + \mathbf{K})\boldsymbol{\alpha}. \quad (7)$$

We introduce two auxiliary variables Ω and β , where $\Omega^\top \Omega = \hat{\mathbf{K}}^\top \hat{\mathbf{K}} + \mathbf{K}$ and $\Omega^\top \beta = -\hat{\mathbf{K}}^\top (\mathbf{1} - yb)$. Because $\hat{\mathbf{K}}^\top \hat{\mathbf{K}} + \mathbf{K}$ is a symmetric matrix, we can compute its eigen-decomposition

$$\hat{\mathbf{K}}^\top \hat{\mathbf{K}} + \mathbf{K} = \mathbf{S}\mathbf{D}\mathbf{S}^\top, \quad (8)$$

where \mathbf{D} is the diagonal matrix of eigenvalues and \mathbf{S} is the orthonormal matrix of eigenvectors. Moreover, because the matrix $\hat{\mathbf{K}}^\top \hat{\mathbf{K}} + \mathbf{K}$ is positive semi-definite, we can further decompose $\mathbf{S}\mathbf{D}\mathbf{S}^\top$ into an inner product of two real matrices by taking the square root of \mathbf{D} . Let $\Omega = \sqrt{\mathbf{D}}\mathbf{S}^\top$, and we obtain a matrix Ω that satisfies $\Omega^\top \Omega = \hat{\mathbf{K}}^\top \hat{\mathbf{K}} + \mathbf{K}$. After computing Ω , we can readily compute $\beta = -(\Omega^\top)^{-1} \hat{\mathbf{K}}^\top (\mathbf{1} - yb)$, where $(\Omega^\top)^{-1} = \frac{1}{\sqrt{\mathbf{D}}} \mathbf{S}^\top$.

With the help of the two auxiliary variables, we convert (7), plus a constant term¹, into least squares format. Together with relaxation of the non-continuous l_0 - norm constraint to an l_1 -norm constraint, we obtain

$$\min_{\boldsymbol{\alpha}} (\Omega\boldsymbol{\alpha} + \beta)^2, \quad \text{s.t. } \|\boldsymbol{\alpha}\|_1 \leq \frac{1}{e} B_e. \quad (9)$$

¹ $(\mathbf{1} - yb)^\top (\hat{\mathbf{K}}(\hat{\mathbf{K}}^\top \hat{\mathbf{K}} + \mathbf{K})^{-1} \hat{\mathbf{K}}^\top - \mathbf{I}) (\mathbf{1} - yb)$

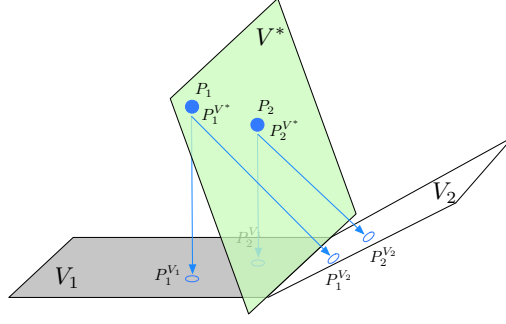


Figure 1: Illustration of searching for a subspace $V \in \mathcal{R}^2$ that best approximates predictions P_1 and P_2 of training instances in \mathcal{R}^3 space. Neither V_1 or V_2 , spanned by existing columns in the kernel matrix, is a good approximation. V^* spanned by kernel columns computed from two *artificial* support vectors is the optimal solution.

Compressing the support vector set. The squared loss and l_1 constraint in (9) lead naturally to the LARS algorithm. Given a budget B_e , we can determine the maximum size m of the compressed support vector set ($m = \frac{B_e}{\epsilon}$). Using LARS, we start from an empty support vector set and add m support vectors incrementally. Since adding a support vector is equivalent to activating a coefficient in α to a non-zero value, we can obtain m optimal support vectors by running LARS optimization in (9) exactly m steps, where each step activates one coefficient. The resulting solution gives the optimal set of m support vectors. We refer this intermediate step as *LARS-SVM*. Note that this step is crucial for the problem, as this LARS-SVM solution serves as a very good initialization for the next step, which is a non-convex optimization problem.

Gradient support vectors. If we interpret α as coordinates and the corresponding columns in the kernel matrix \mathbf{K} as basis vectors, then these basis vectors span an $\mathcal{R}^{n_{sv}}$ space in which lie predictions of the original SVM model. In this compression algorithm, our goal is to find a lower dimensional subspace that supports good approximations of the original predictions. After running LARS for m iterations, we obtain m support vectors and their coefficients α , forming an \mathcal{R}^m subspace of the space spanned by the full kernel matrix.

We illustrate this lower dimensional approximation in Figure 1. Vectors P_1 and P_2 are predictions of two training points made in the full SVM solution space (\mathcal{R}^3 and spanned by three support vectors). We want to compress the model to two support vectors by looking for a subspace $V \in \mathcal{R}^2$ that supports the best approximations of these two predictions. Using existing support vectors as a basis, we can find subspaces V_1 and V_2 , each spanned by a pair of support vectors. The projections of P_1 and P_2 on plane V_1 ($P_1^{V_1}$ and $P_2^{V_1}$) are closest to the original predictions P_1 and P_2 , and thus V_1 is the better approximation. However, in this case, neither V_1 nor V_2 is a particularly good approximation. Suppose we remove the restriction of selecting a subspace spanned by *existing* basis vectors in the kernel matrix, instead optimizing the basis vectors to yield a more suitable subspace. In Figure 1, this is illustrated by the optimal subspace V^* which produces a better approximation to the target predictions.

Note that the basis vectors (columns of the kernel matrix) are parameterized by support vectors. By optimizing these underlying support vectors, we can search for a better low-dimensional subspace. If we denote \mathbf{K}_m as the training kernel matrix with only m columns corresponding to the support vectors chosen by LARS, and α_m as the coefficients of these support vectors, we can formulate the search for *artificial* support vector as an optimization problem. Specifically, we minimize a squared loss between approximate and full SVM predictions over all support vectors, and the parameters are support vectors.

$$\min_{(\mathbf{x}_1, \dots, \mathbf{x}_m)} \mathcal{L} = \left(\mathbf{K}_m \alpha_m - \mathbf{K} \alpha \right)^2, \quad (10)$$

where \mathbf{K}_{ij} is the kernel entry, and for simplicity, we use radial basis function (RBF) kernel function ($\mathbf{K}_{ij} = e^{-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|^2}{2\sigma^2}}$). However, other kernel functions are equally suitable. The unconstrained op-

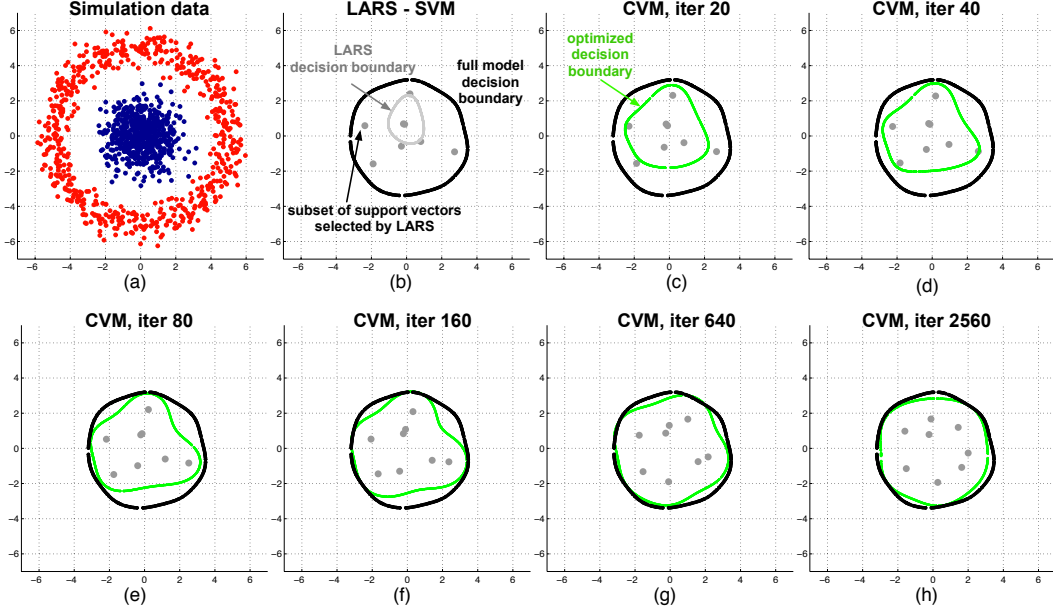


Figure 2: Illustration of each step of CVM on a synthetic data set. (a) Simulation inputs from two classes (red and blue). By design, the two classes are not linear separable. (b) Decision boundary formed by full SVM solution (black curve). A small subset of support vectors picked by LARS (gray points) and the compressed decision boundary formed by this subset of support vectors (gray curve). (c-h) Optimization iterations. The gradient support vectors are moved by the iterative optimization. The optimized decision boundary formed by gradient support vectors (green curve) gradually approaches the one formed by full SVM solution.

timization problem (10) can be solved by conjugate gradient descent with respect to the chosen m support vectors. Since α 's are the coordinates with respect to the basis, we optimize α jointly with support vectors, which is faster than optimizing basis and solving coordinates alternatively. The gradients can be computed very efficiently using matrix operations. Since gradient descent on support vectors is equivalent to moving these support vectors in a continuous space, thereby generating m new support vectors, we refer to these newly generated support vectors as *gradient support vectors*. We denote this combined method of LARS-SVM and gradient support vectors as Compressed Vector Machine (CVM). Because the optimization problem in (10) is non-convex with respect to \mathbf{x}_i , we initialize our algorithm with the basis \mathbf{K}_m and coordinates α_m returned in the LARS-SVM solution.

In practice, it may be desirable to optimize both the SVM cost parameter C and any kernel parameters (e.g. σ^2 in the RBF kernel) for the final CVM model. Additionally, it may be preferable to optimize CVM constrained by the validation accuracy of the compressed model rather than the size of the support vector budget. Constrained Bayesian optimization [9] supports efficient constrained joint hyperparameter optimizations of this type. Additionally, the L1-penalized support vector selection in the LARS-SVM step may benefit from recent work on highly parallel Elastic Net solvers [27].

5 Results

In this section, we first demonstrate Compressed Vector Machine (CVM) on a synthetic data set to graphically illustrate each step in the algorithm. We then evaluate CVM on several medium-scale real-world data sets.

Synthetic data set. The data set contains 600 sample inputs from two classes (red and blue), where each input contains two features. The blue inputs are sampled from a Gaussian distribution with mean at the origin and variance 1, and red inputs are sampled from a noisy circle surrounding the blue inputs. As shown in Figure 2(a), by design the data set is not linearly separable. For simplicity,

we treat all inputs as training inputs. To evaluate CVM, we first learn an SVM with the RBF kernel from the full training set. We plot the resulting optimal decision boundary in Figure 2(b) with a black curve. In total, the full model has 80 support vectors.

To compress the model, we first select a subset of support vectors by solving LARS-SVM optimization (9). Specifically, we compress the model to 10% of its original size, 8 support vectors, by running LARS for 8 iterations. The 8 LARS-SVM support vectors are shown in Figure 2(b) as solid gray points, and the approximate LARS-SVM decision boundary is shown by the gray curve.

Since the subspace formed by 8 support vectors is heavily restricted by the discrete training input space, the approximation is poor. To overcome this problem, we search for a better subspace or basis in a continuous space, and perform gradient descent on support vectors by optimizing (10). In Figure 2(c-h), we illustrate the optimization with updated support vector locations and optimized decision boundaries as we gradually increase the number of iterations. The resulting *gradient support vectors* are shown as gray points and the new optimized decision boundaries formed from these new gradient support vectors are shown by green curves. After 2560 iterations, as shown in Figure 2(h), we can observe that the optimized decision boundary (green) is very close to the boundary captured in the full model (black). These optimized decision boundaries demonstrate that moving a small subset of support vectors in a continuous space can efficiently approximate the optimal decision boundary formed by full SVM solution, supporting effective SVM model compression.

Statistics	Pageblocks	Magic	Letters	20news	MNIST	DMOZ
#training exam.	4379	15216	16000	11269	60000	7184
#testing exam.	1094	3804	4000	7505	10000	1796
#features	10	10	16	200	784	16498
#classes	2	2	26	20	10	16

Table 1: Statistics of all six data sets.

Large real-world data sets. To evaluate the performance of CVM on real-world applications, we evaluate our algorithm on six data sets of varying size, dimensionality and complexity. Table 1 details the statistics of all six data sets. We use LibSVM [4] to train a regular RBF kernel SVM using regularization parameter C and RBF kernel width σ selected on a 20% validation split. For multi-class data sets, we use the one-vs-one multi-class scheme. The classification accuracy of test predictions from this SVM model serves as a baseline in Figure 3(full SVM).

Given the full SVM solution, we run CVM in two steps. First, we use LARS solve the optimization problem in (9) using all support vectors from the original SVM model. An initial compressed support vector set is selected with a target compressed size (*e.g.* 10 out of 500 support vectors). The selected support vectors serve as the second baseline in Figure 3(LARS-SVM). Second, we shift these support vectors in a continuous space by optimizing (10) w.r.t. the input support vectors and the corresponding Lagrange multipliers α , generating gradient support vectors. This final set of gradient support vectors constitutes the CVM model. To show the trend of accuracy/cost performance, we plot the classification accuracy for CVM after adding every 10 support vectors. Figure 3 shows the performance of CVM and the baselines on all six data sets.

Comparison with prior work. Figure 3 also shows a comparison of CVM with *Reduced-SVM* [13]. This algorithm takes an iterative two phase approach. First a set of support vectors is heuristically selected from random samples of the training set and added to the existing set of support vectors (initially empty). Then, the model weights are optimized by an SVM with the quadratic hinge loss. The algorithm alternates these two steps until the target number of support vectors is reached.

As shown in the Figure 3, CVM significantly improves over all baselines. Compared to the current state-of-the-art, Reduced-SVM, CVM has the capability of moving support vectors, generating a new basis, and learning a highly approximated basis to match the decision boundaries formed by the full SVM solution. It is this ability that distinguishes CVM from other algorithms when the evaluation budget is low. Across all data sets, CVM maintains close to the same accuracy as the full SVM with merely 10% of the support vectors.

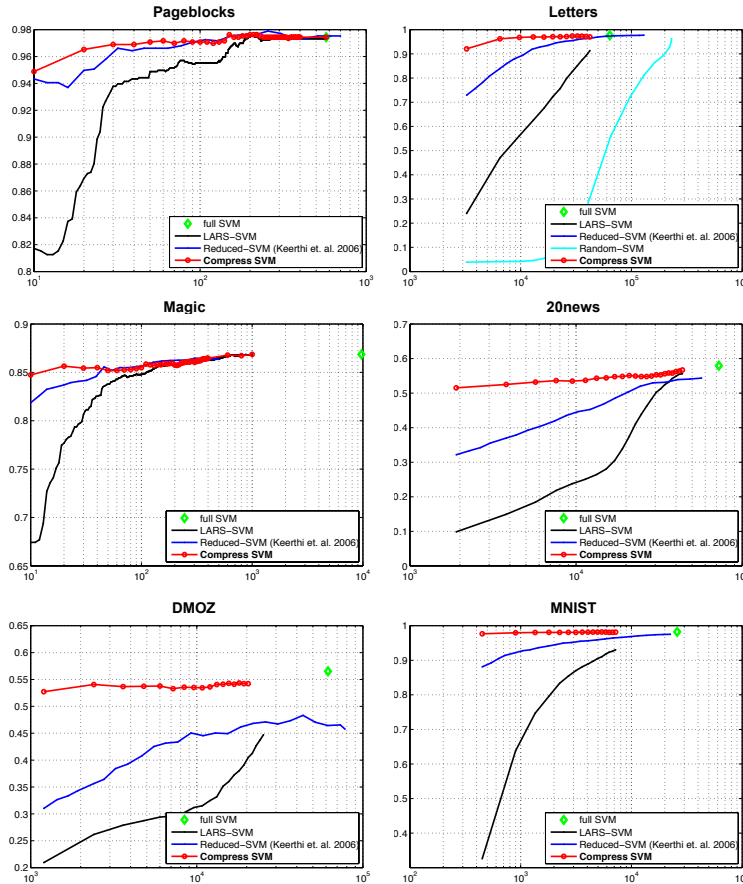


Figure 3: Accuracy versus number of support vectors (in log scale).

6 Acknowledgments

Most of this work was previously invented by Burges and Schölkopf [2] whose research was truly visionary at the time.

References

- [1] Cristian Bucilu, Rich Caruana, and Alexandru Niculescu-Mizil. Model compression. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 535–541. ACM, 2006.
- [2] Chris C. Burges and Bernhard Schölkopf. Improving the accuracy and speed of support vector machines. *Advances in neural information processing systems*, 9:375–381, 1997.
- [3] R. Busa-Fekete, D. Benbouzid, B. Kégl, et al. Fast classification using sparse decision dags. In *ICML*, 2012.
- [4] Chih-Chung Chang and Chih-Jen Lin. Libsvm: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(3):27, 2011.
- [5] M. Chen, Z. Xu, K. Q. Weinberger, and O. Chapelle. Classifier cascade for minimizing feature evaluation cost. In *AISTATS*, 2012.
- [6] Ofer Dekel, Shai Shalev-Shwartz, and Yoram Singer. The forgetron: A kernel-based perceptron on a budget. *SIAM Journal on Computing*, 37(5):1342–1372, 2008.
- [7] Harris Drucker, Donghui Wu, and Vladimir N Vapnik. Support vector machines for spam categorization. *Neural Networks, IEEE Transactions on*, 10(5):1048–1054, 1999.
- [8] Bradley Efron, Trevor Hastie, Iain Johnstone, and Robert Tibshirani. Least angle regression. *The Annals of statistics*, 32(2):407–499, 2004.

- [9] Jacob Gardner, Matt Kusner, Kilian Q. Weinberger, John Cunningham, and Zhixiang Xu. Bayesian optimization with inequality constraints. In Tony Jebara and Eric P. Xing, editors, *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 937–945. JMLR Workshop and Conference Proceedings, 2014.
- [10] A. Grubb and J. A. Bagnell. Speedboost: Anytime prediction with uniform near-optimality. In *AISTATS*, 2012.
- [11] Isabelle Guyon, Jason Weston, Stephen Barnhill, and Vladimir Vapnik. Gene selection for cancer classification using support vector machines. *Machine learning*, 46(1-3):389–422, 2002.
- [12] Bernd Heisele, Purdy Ho, and Tomaso Poggio. Face recognition with support vector machines: Global versus component-based approach. In *Computer Vision, 2001. ICCV 2001. Proceedings. Eighth IEEE International Conference on*, volume 2, pages 688–694. IEEE, 2001.
- [13] S Sathiya Keerthi, Olivier Chapelle, and Dennis DeCoste. Building support vector machines with reduced classifier complexity. *The Journal of Machine Learning Research*, 7:1493–1515, 2006.
- [14] Matt Kusner, Stephen Tyree, Kilian Q. Weinberger, and Kunal Agrawal. Stochastic neighbor compression. In Tony Jebara and Eric P. Xing, editors, *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 622–630. JMLR Workshop and Conference Proceedings, 2014.
- [15] L. Lefakis and F. Fleuret. Joint cascade optimization using a product of boosted classifiers. In *NIPS*, pages 1315–1323. 2010.
- [16] John C Platt, Nello Cristianini, and John Shawe-taylor. Large margin dags for multiclass classification. In *Advances in Neural Information Processing Systems 12*, 2000.
- [17] J. Pujara, H. Daumé III, and L. Getoor. Using classifier cascades for scalable e-mail classification. In *CEAS*, 2011.
- [18] Ryan Rifkin and Aldebaro Klautau. In defense of one-vs-all classification. *The Journal of Machine Learning Research*, 5:101–141, 2004.
- [19] M. Saberian and N. Vasconcelos. Boosting classifier cascades. In *NIPS*, pages 2047–2055. 2010.
- [20] B. Schölkopf and A.J. Smola. *Learning with kernels: Support vector machines, regularization, optimization, and beyond*. MIT press, 2001.
- [21] Cristianini Shawe-Taylor and Smola Schölkopf. The support vector machine. 2000.
- [22] Edward Snelson and Zoubin Ghahramani. Sparse gaussian processes using pseudo-inputs. In *Advances in Neural Information Processing Systems*, pages 1257–1264, 2005.
- [23] J. Wang and V. Saligrama. Local supervised learning through space partitioning. In *NIPS*, pages 91–99, 2012.
- [24] Zhuang Wang, Koby Crammer, and Slobodan Vucetic. Breaking the curse of kernelization: Budgeted stochastic gradient descent for large-scale svm training. *Journal of Machine Learning Research*, 13:3103–3131, 2012.
- [25] Zhixiang Xu, Matt Kusner, Minmin Chen, and Kilian Q. Weinberger. Cost-sensitive tree of classifiers. In Sanjoy Dasgupta and David McAllester, editors, *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, volume 28, pages 133–141. JMLR Workshop and Conference Proceedings, 2013.
- [26] Zhixiang (Eddie) Xu, Matt J. Kusner, Gao Huang, and Kilian Q. Weinberger. Anytime representation learning. In Sanjoy Dasgupta and David McAllester, editors, *ICML '13*, pages 1076–1084, 2013.
- [27] Quan Zhou, Wenlin Chen, Shiji Song, Jacob R. Gardner, Kilian Q. Weinberger, and Yixin Chen. A reduction of the elastic net to support vector machines with an application to gpu computing. In *Association for the Advancement of Artificial Intelligence (AAAI-15)*, 2015.