
Feature Selection for Linear SVM with Provable Guarantees

Saurabh Paul

Computer Science Dept.
Rensselaer Polytechnic Institute
pauls2@rpi.edu

Malik Magdon-Ismail

Computer Science Dept.
Rensselaer Polytechnic Institute
magdon@cs.rpi.edu

Petros Drineas

Computer Science Dept.
Rensselaer Polytechnic Institute
drinep@cs.rpi.edu

Abstract

We give two *provably* accurate feature-selection techniques for the linear SVM. The algorithms run in deterministic and randomized time respectively. Our algorithms can be used in an unsupervised or supervised setting. The supervised approach is based on sampling features from support vectors. We prove that the margin in the feature space is preserved to within ϵ -relative error of the margin in the full feature space in the worst-case. In the unsupervised setting, we also provide worst-case guarantees of the radius of the minimum enclosing ball, thereby ensuring comparable generalization as in the full feature space and resolving an open problem posed in Dasgupta et al. (2007). We present extensive experiments on real-world datasets to support our theory and to demonstrate that our methods are competitive and often better than prior state-of-the-art, for which there are no known provable guarantees.

1 Introduction

The linear Support Vector Machine (SVM) is a popular classification method (Cristianini and Shawe-Taylor, 2000). Few theoretical results exist for feature selection with SVMs. Empirically, numerous feature selection techniques work well (e.g. Guyon et al. (2002); Fung and Mangasarian (2004)). We present a *deterministic* and a *randomized* feature selection technique for the linear SVM with a *provable worst-case performance guarantee* on the margin. The feature selection is unsupervised if features are selected obliviously to the data labels; otherwise, it is supervised.

Appearing in Proceedings of the 18th International Conference on Artificial Intelligence and Statistics (AISTATS) 2015, San Diego, CA, USA. JMLR: W&CP volume 38. Copyright 2015 by the authors.

Our algorithms can be used in an unsupervised or supervised setting. In the unsupervised setting, our algorithm selects a number of features proportional to the rank of the data and preserves both the margin and radius of minimum enclosing ball to within ϵ -relative error in the worst-case, thus resolving an open problem posed in Dasgupta et al. (2007). In the supervised setting, our algorithm selects $O(\#\text{support vectors})$ features using only the set of support vectors, and preserves the margin for the support vectors to within ϵ -relative error in the worst-case.

SVM basics. The training data has n points $\mathbf{x}_i \in \mathbb{R}^d$, with respective labels $y_i \in \{-1, +1\}$ for $i = 1 \dots n$. For linearly separable data, the primal SVM learning problem constructs a hyperplane \mathbf{w}^* which maximizes the geometric *margin* (the minimum distance of a data point to the hyperplane), while separating the data. For non-separable data the “soft” 1-norm margin is maximized. The dual lagrangian formulation of the soft 1-norm SVM reduces to the following quadratic program:

$$\begin{aligned} \max_{\alpha_i} : & \sum_{i=1}^n \alpha_i - \frac{1}{2} \sum_{i,j=1}^n \alpha_i \alpha_j y_i y_j \mathbf{x}_i^T \mathbf{x}_j \\ \text{subject to:} & \sum_{i=1}^n y_i \alpha_i = 0; \quad 0 \leq \alpha_i \leq C, \quad i = 1 \dots n. \end{aligned} \tag{1}$$

The regularizer C is part of the input and the hyperplane classifier can be constructed from the α_i . The out-of-sample performance is related to the VC -dimension of the resulting “*fat*”-separator. Assuming that the data lie in a ball of radius B , and that the hypothesis set consists of hyperplanes of width γ (the margin), then the VC -dimension is $O(B^2/\gamma^2)$ (Vapnik (1998)). Thus, by the VC -bound (Vapnik and Chervonenkis (1971)), the out-of-sample error is bounded by the in-sample error and a term monotonic in B^2/γ^2 .

Our Contributions. We give two provably accurate feature selection techniques for linear SVM in both unsupervised and supervised settings with worst-case performance guarantees on the margin. We use the single set spectral sparsification technique (BSS)

¹ from Batson et al. (2009) as our deterministic algorithm (the algorithm runs in deterministic time, hence the name ‘deterministic’) and leverage-score sampling (Dasgupta et al. (2007)) as the randomized algorithm. We give a new simple method of extending unsupervised feature selection to supervised in the context of SVMs by running the unsupervised technique on the support vectors. This allows us to select only $O(\#\text{support vectors})$ features for the deterministic algorithm ($\tilde{O}(\#\text{support vectors})$ features for the randomized algorithm, where \tilde{O} hides the logarithmic factors) while still preserving the margin on the support vectors. Since the support vectors are a sufficient statistic for learning a linear SVM, preserving the margin on the support vectors should be enough for learning on all the data with the sampled feature set.

More formally, let γ^* be the optimal margin for the support vector set (which is also the optimal margin for all the data). The optimal margin γ^* is obtained by solving the SVM optimization problem using all the features. For a suitably chosen number of features r , let $\tilde{\gamma}^*$ be the optimal margin obtained by solving the SVM optimization problem using the support vectors in the sampled feature space. We prove that the margin is preserved to within ϵ -relative error: $\tilde{\gamma}^{*2} \geq (1 - \epsilon)\gamma^{*2}$. For the deterministic algorithm, the number of features required is $r = O(\#\text{support vectors}/\epsilon^2)$, whereas the randomized algorithm requires $r = \tilde{O}(\#\text{support vectors}/\epsilon^2)$ features to be selected.

In the unsupervised setting, by running our algorithm on all the data, instead of only the support vectors, we get a stronger result statistically, but using more features. The deterministic algorithm requires $O(\rho/\epsilon^2)$ features to be selected, while the randomized algorithm requires $O(\rho/\epsilon^2 \log(\rho/\epsilon^2))$ features to be selected. Again, defining $\tilde{\gamma}^*$ as the optimal margin obtained by solving the SVM optimization problem using all the data in the sampled feature space, we prove that $\tilde{\gamma}^{*2} \geq (1 - \epsilon)\gamma^{*2}$. We can now also prove that the data radius is preserved, $\tilde{B}^2 \leq (1 + \epsilon)B^2$. This means that B^2/γ^{*2} is preserved to within ϵ -relative error, which means that the generalisation error is also preserved to within ϵ -relative error. The rank of the data is the effective dimension of the data, and one can construct this many combinations of pure features that preserve the geometry of the SVM exactly. What makes our result non-trivial is that we select this many pure features and preserve the geometry of the SVM to within ϵ -relative error.

On the practical side, we provide an efficient heuristic

¹We call single-set spectral sparsification BSS after the authors Batson, Spielman and Srivatsava

for our supervised feature selection using BSS which allows our algorithm to scale-up to large datasets. While the main focus of this paper is theoretical, we compare both supervised and unsupervised versions of feature selection using single-set spectral sparsification and leverage-score sampling with the corresponding supervised and unsupervised forms of Recursive Feature Elimination (RFE) (Guyon et al. (2002)), LPSVM (Fung and Mangasarian (2004)), uniform sampling and rank-revealing QR factorization (RRQR) based method of column selection. Feature selection based on the single-set spectral sparsification and leverage-score sampling technique is competitive and often better than RFE and LPSVM, and none of the prior art comes with provable performance guarantees in either the supervised or unsupervised setting.

Related Work. All the prior art is heuristic in that there are no performance guarantees; nevertheless, they have been empirically tested against each other. Our algorithm comes with provable bounds, and performs comparably or better in empirical tests. We give a short survey of the prior art: Guyon et al. (2002) and Rakotomamonjy (2003) proposed SVM based criteria to rank features based on the weights. Weston et al. (2000) formulated a combinatorial optimization problem to select features by minimising B^2/γ^2 . Weston et al. (2003) used the zero norm to perform error minimization and feature selection in one step. A Newton based method of feature selection using linear programming was given in Fung and Mangasarian (2004). Tan et al. (2010) formulated the ℓ_0 -norm Sparse SVM using mixed integer programming. Do et al. (2009b) proposed R-SVM which performs feature selection and ranking by optimizing the radius-margin bound with a scaling factor, and extend this work in Kalousis and Do (2013) using a quadratic optimization problem with quadratic constraints. Another line of work includes the doubly regularised Support Vector Machine (DrSVM) (Wang et al., 2006) which uses a mixture of ℓ_2 -norm and ℓ_1 -norm penalties to solve the SVM optimization problem and perform variable selection. Subsequent works on DrSVM involve reducing the computational bottleneck (Wang et al. (2008), Ye et al. (2011)). Gilad-Bachrach et al. (2004) formulate the margin as a function of set of features and score to sets of features according to the margin induced. Park et al. (2012) studied the Fisher consistency and oracle property of penalized SVM where the dimension of inputs is fixed and showed that their method is able to identify the right model in most cases.

Paul et al. (2013, 2014) used random projections on linear SVM and showed that the margin and data-radius are preserved. However, this is different from our work, since they used linear combinations of features and we select pure features.

BSS and leverage-score sampling have been used to select features for k -means (Boutsidis and Magdon-Ismail (2013); Boutsidis et al. (2009)), regularized least-squares classifier (Dasgupta et al. (2007); Paul and Drineas (2014)). Our work further expands this research into sparsification algorithms for machine learning.

2 Background

Notation: $\mathbf{A}, \mathbf{B}, \dots$ denote matrices and $\boldsymbol{\alpha}, \mathbf{b}, \dots$ denote column vectors; \mathbf{e}_i (for all $i = 1 \dots n$) is the standard basis, whose dimensionality will be clear from context; and \mathbf{I}_n is the $n \times n$ identity matrix. The Singular Value Decomposition (SVD) of a matrix $\mathbf{A} \in \mathbb{R}^{n \times d}$ of rank $\rho \leq \min\{n, d\}$ is equal to $\mathbf{A} = \mathbf{U}\boldsymbol{\Sigma}\mathbf{V}^T$, where $\mathbf{U} \in \mathbb{R}^{n \times \rho}$ is an orthogonal matrix containing the left singular vectors, $\boldsymbol{\Sigma} \in \mathbb{R}^{\rho \times \rho}$ is a diagonal matrix containing the singular values $\sigma_1 \geq \sigma_2 \geq \dots \sigma_\rho > 0$, and $\mathbf{V} \in \mathbb{R}^{d \times \rho}$ is a matrix containing the right singular vectors. The spectral norm of \mathbf{A} is $\|\mathbf{A}\|_2 = \sigma_1$.

Matrix Sampling Formalism: Let \mathbf{A} be the data matrix consisting of n points and d dimensions, $\mathbf{S} \in \mathbb{R}^{d \times r}$ be a matrix such that $\mathbf{AS} \in \mathbb{R}^{n \times r}$ contains r columns of \mathbf{A} (\mathbf{S} is a sampling matrix as it samples r columns of \mathbf{A}). Let $\mathbf{D} \in \mathbb{R}^{r \times r}$ be the diagonal matrix such that $\mathbf{ASD} \in \mathbb{R}^{n \times r}$ rescales the columns of \mathbf{A} that are in \mathbf{AS} . We will replace the sampling and re-scaling matrices by a single matrix $\mathbf{R} \in \mathbb{R}^{d \times r}$, where $\mathbf{R} = \mathbf{SD}$ first samples and then rescales r columns of \mathbf{A} .

Let \mathbf{X} be a generic data matrix in d dimensions whose rows are data vectors \mathbf{x}_i^T , and let \mathbf{Y} be the diagonal label matrix whose diagonal entries are the labels, $\mathbf{Y}_{ii} = y_i$. Let $\boldsymbol{\alpha} = [\alpha_1, \alpha_2, \dots, \alpha_n] \in \mathbb{R}^n$ be the vector of lagrange multipliers to be determined by solving eqn. (2). In matrix form, the SVM dual optimization problem is

$$\begin{aligned} \max_{\boldsymbol{\alpha}} : \quad & \mathbf{1}^T \boldsymbol{\alpha} - \frac{1}{2} \boldsymbol{\alpha}^T \mathbf{Y} \mathbf{X} \mathbf{X}^T \mathbf{Y} \boldsymbol{\alpha} \\ \text{subject to:} \quad & \mathbf{1}^T \mathbf{Y} \boldsymbol{\alpha} = 0; \quad \mathbf{0} \leq \boldsymbol{\alpha} \leq \mathbf{C}. \end{aligned} \quad (2)$$

(In the above, $\mathbf{1}$, $\mathbf{0}$, \mathbf{C} are vectors with the implied constant entry.) When the data and label matrices contain all the data, we will emphasize this using the notation $\mathbf{X}^{\text{tr}} \in \mathbb{R}^{n \times d}$, $\mathbf{Y}^{\text{tr}} \in \mathbb{R}^{n \times n}$. Solving (2) with these full data matrices gives a solution $\hat{\boldsymbol{\alpha}}^*$. The data \mathbf{x}_i for which $\hat{\alpha}_i^* > 0$ are support vectors and we denote by $\mathbf{X}^{\text{sv}} \in \mathbb{R}^{p \times d}$, $\mathbf{Y}^{\text{sv}} \in \mathbb{R}^{p \times p}$ the data and label matrices containing only the p support vectors. Solving (2) with $(\mathbf{X}^{\text{tr}}, \mathbf{Y}^{\text{tr}})$ or $(\mathbf{X}^{\text{sv}}, \mathbf{Y}^{\text{sv}})$ result in the same classifier. Let $\boldsymbol{\alpha}^*$ be the solution to (2) for the support vector data. The optimal separating hyperplane is $\mathbf{w}^* = (\mathbf{X}^{\text{tr}})^T \mathbf{Y}^{\text{tr}} \hat{\boldsymbol{\alpha}}^* = (\mathbf{X}^{\text{sv}})^T \mathbf{Y}^{\text{sv}} \boldsymbol{\alpha}^*$, where \mathbf{X}^{sv} is the support vector matrix. The geometric margin of the optimal separating hyperplane is

$\gamma^* = 1 / \|\mathbf{w}^*\|_2$, where $\|\mathbf{w}^*\|_2^2 = \sum_{i=1}^n \alpha_i^*$. The data radius is $B = \min_{\mathbf{x}^*} \max_{\mathbf{x}_i} \|\mathbf{x}_i - \mathbf{x}^*\|_2$.

Our goal is to study how the SVM performs when run in the sampled feature space. Let \mathbf{X} , \mathbf{Y} be data and label matrices (such as those above) and $\mathbf{R} \in \mathbb{R}^{d \times r}$ a sampling and rescaling matrix which selects r columns of \mathbf{X} . The transformed dataset into the r selected features is $\tilde{\mathbf{X}} = \mathbf{X}\mathbf{R}$, and the SVM optimization problem in this feature space becomes

$$\begin{aligned} \max_{\hat{\boldsymbol{\alpha}}} : \quad & \mathbf{1}^T \hat{\boldsymbol{\alpha}} - \frac{1}{2} \hat{\boldsymbol{\alpha}}^T \mathbf{Y} \mathbf{X} \mathbf{R} \mathbf{R}^T \mathbf{X}^T \mathbf{Y} \hat{\boldsymbol{\alpha}}, \\ \text{subject to:} \quad & \mathbf{1}^T \mathbf{Y} \hat{\boldsymbol{\alpha}} = 0; \quad \mathbf{0} \leq \hat{\boldsymbol{\alpha}} \leq \mathbf{C}. \end{aligned} \quad (3)$$

For the supervised setting, we select features from the support vector matrix and so we set $\mathbf{X} = \mathbf{X}^{\text{sv}}$ and $\mathbf{Y} = \mathbf{Y}^{\text{sv}}$ and we select $r_1 \ll d$ features using \mathbf{R} . For the unsupervised setting, we select features from the full data matrix and so we set $\mathbf{X} = \mathbf{X}^{\text{tr}}$ and $\mathbf{Y} = \mathbf{Y}^{\text{tr}}$ and we select $r_2 \ll d$ features using \mathbf{R} .

3 Our main tools

In this section, we describe our main tools of feature selection from the numerical linear algebra literature, namely single-set spectral sparsification and leverage-score sampling.

Single-set Spectral Sparsification. The BSS algorithm (Batson et al. (2009)) is a deterministic greedy technique that selects columns one at a time. The algorithm samples r columns in deterministic time, hence the name deterministic sampling. Consider the input matrix as a set of d column vectors $\mathbf{V}^T = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_d]$, with $\mathbf{v}_i \in \mathbb{R}^\ell$ ($i = 1, \dots, d$). Given ℓ and $r > \ell$, we iterate over $\tau = 0, 1, 2, \dots, r-1$. Define the parameters $L_\tau = \tau - \sqrt{r\ell}$, $\delta_L = 1$, $\delta_U = (1 + \sqrt{\ell/r}) / (1 - \sqrt{\ell/r})$ and $U_\tau = \delta_U (\tau + \sqrt{\ell r})$. For $U, L \in \mathbb{R}$ and $\mathbf{A} \in \mathbb{R}^{\ell \times \ell}$ a symmetric positive definite matrix with eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_\ell$, define $\Phi(L, \mathbf{A}) = \sum_{i=1}^\ell \frac{1}{\lambda_i - L}$ and $\hat{\Phi}(U, \mathbf{A}) = \sum_{i=1}^\ell \frac{1}{U - \lambda_i}$ as the lower and upper potentials respectively. These potential functions measure how far the eigenvalues of \mathbf{A} are from the upper and lower barriers U and L respectively. We define $\mathcal{L}(\mathbf{v}, \delta_L, \mathbf{A}, L)$ and $\mathcal{U}(\mathbf{v}, \delta_U, \mathbf{A}, U)$ as follows:

$$\begin{aligned} \mathcal{L}(\mathbf{v}, \delta_L, \mathbf{A}, L) &= \frac{\mathbf{v}^T (\mathbf{A} - (L + \delta_L) \mathbf{I}_\ell)^{-2} \mathbf{v}}{\Phi(L + \delta_L, \mathbf{A}) - \Phi(L, \mathbf{A})} - \mathbf{v}^T (\mathbf{A} - (L + \delta_L) \mathbf{I}_\ell)^{-1} \mathbf{v} \\ \mathcal{U}(\mathbf{v}, \delta_U, \mathbf{A}, U) &= \frac{\mathbf{v}^T ((U + \delta_U) \mathbf{I}_\ell - \mathbf{A})^{-2} \mathbf{v}}{\hat{\Phi}(U, \mathbf{A}) - \hat{\Phi}(U + \delta_U, \mathbf{A})} + \mathbf{v}^T ((U + \delta_U) \mathbf{I}_\ell - \mathbf{A})^{-1} \mathbf{v}. \end{aligned}$$

At every iteration, there exists an index i_τ and a weight $t_\tau > 0$ such that, $t_\tau^{-1} \leq \mathcal{L}(\mathbf{v}_i, \delta_L, \mathbf{A}, L)$

and $t_\tau^{-1} \geq \mathcal{U}(\mathbf{v}_i, \delta_U, \mathbf{A}, U)$. Thus, there will be at most r columns selected after τ iterations. The running time of the algorithm is dominated by the search for an index i_τ satisfying $\mathcal{U}(\mathbf{v}_i, \delta_U, \mathbf{A}_\tau, U_\tau) \leq \mathcal{L}(\mathbf{v}_i, \delta_L, \mathbf{A}_\tau, L_\tau)$ and computing the weight t_τ . One needs to compute the upper and lower potentials $\hat{\Phi}(U, \mathbf{A})$ and $\Phi(L, \mathbf{A})$ and hence the eigenvalues of \mathbf{A} . Cost per iteration is $O(\ell^3)$ and the total cost is $O(r\ell^3)$. For $i = 1, \dots, d$, we need to compute \mathcal{L} and \mathcal{U} for every \mathbf{v}_i which can be done in $O(d\ell^2)$ for every iteration, for a total of $O(rd\ell^2)$. Thus total running time of the algorithm is $O(rd\ell^2)$. We include the algorithm in the supplementary material. We present the following lemma for the single-set spectral sparsification algorithm. The proof is included in the supplementary material.

Lemma 1. *Given $\mathbf{V} \in \mathbb{R}^{d \times \ell}$ satisfying $\mathbf{V}^T \mathbf{V} = \mathbf{I}_\ell$ and $r > \ell$, we can deterministically construct sampling and rescaling matrices $\mathbf{S} \in \mathbb{R}^{d \times r}$ and $\mathbf{D} \in \mathbb{R}^{r \times r}$ such that for $\mathbf{R} = \mathbf{SD}$, $\left\| \mathbf{V}^T \mathbf{V} - \mathbf{V}^T \mathbf{R} \mathbf{R}^T \mathbf{V} \right\|_2 \leq 3\sqrt{\ell/r}$*

Leverage-Score Sampling. Our randomized feature selection method is based on importance sampling or the so-called leverage-score sampling of Dasgupta et al. (2007). Let \mathbf{V} be the top- ρ right singular vectors of the training set \mathbf{X} . A carefully chosen probability distribution of the form

$$p_i = \frac{\|\mathbf{V}_i\|_2^2}{n}, \text{ for } i = 1, 2, \dots, d, \quad (4)$$

i.e. proportional to the squared Euclidean norms of the rows of the right-singular vectors is constructed. Select r rows of \mathbf{V} in i.i.d trials and re-scale the rows with $1/\sqrt{p_i}$. The time complexity is dominated by the time to compute the SVD of \mathbf{X} .

Lemma 2. *Let $\epsilon \in (0, 1/2]$ be an accuracy parameter and $\delta \in (0, 1)$ be the failure probability. Given $\mathbf{V} \in \mathbb{R}^{d \times \ell}$ satisfying $\mathbf{V}^T \mathbf{V} = \mathbf{I}_\ell$. Let $\tilde{p} = \min\{1, rp_i\}$, let p_i be as Eqn. 4 and let $r = O\left(\frac{n}{\epsilon^2} \log\left(\frac{n}{\epsilon^2 \sqrt{\delta}}\right)\right)$. Construct the sampling and rescaling matrix \mathbf{R} . Then with probability at least 0.99, $\left\| \mathbf{V}^T \mathbf{V} - \mathbf{V}^T \mathbf{R} \mathbf{R}^T \mathbf{V} \right\|_2 \leq \epsilon$.*

4 Theoretical Analysis

Though our feature selection algorithms are relatively simple, we show that running the linear SVM in the selected feature space results in a classifier with provably comparable margin to the SVM classifier obtained from the full feature space. Our main results are in Theorems 1 & 2. We state the theorems for BSS, but similar theorems can be stated for leverage-score sampling. We include them in the Supplementary material.

4.1 Margin is preserved by Supervised Feature Selection

Theorem 1 says that you get comparable margin from solving the SVM on the support vectors (equivalently all the data) and from solving the SVM on support vectors in a feature space with only $O(\#\text{support vectors})$ features.

Theorem 1. *Given $\epsilon \in (0, 1)$, run supervised BSS-feature selection on \mathbf{X}^{sv} with $r_1 = 36p/\epsilon^2$, to obtain the feature sampling and rescaling matrix \mathbf{R} . Let γ^* and $\tilde{\gamma}^*$ be the margins obtained by solving the SVM dual (2) with $(\mathbf{X}^{sv}, \mathbf{Y}^{sv})$ and $(\mathbf{X}^{sv} \mathbf{R}, \mathbf{Y}^{sv})$ respectively. Then, $\tilde{\gamma}^{*2} \geq (1 - \epsilon)\gamma^{*2}$.*

Proof. Let $\mathbf{X}^{\text{tr}} \in \mathbb{R}^{n \times d}$, $\mathbf{Y}^{\text{tr}} \in \mathbb{R}^{n \times n}$ be the feature matrix and class labels of the training set (as defined in Section 2) and let $\hat{\boldsymbol{\alpha}}^* = [\alpha_1^*, \alpha_2^*, \dots, \alpha_n^*]^T \in \mathbb{R}^n$ be the vector achieving the optimal solution for the problem of eqn. (2). Then,

$$Z_{opt} = \sum_{j=1}^n \hat{\alpha}_j^* - \frac{1}{2} \hat{\boldsymbol{\alpha}}^{*T} \mathbf{Y}^{\text{tr}} \mathbf{X}^{\text{tr}} (\mathbf{X}^{\text{tr}})^T \mathbf{Y}^{\text{tr}} \hat{\boldsymbol{\alpha}}^* \quad (5)$$

Let $p \leq n$ be the support vectors with $\hat{\alpha}_j > 0$. Let $\boldsymbol{\alpha}^* = [\alpha_1^*, \alpha_2^*, \dots, \alpha_p^*]^T \in \mathbb{R}^p$ be the vector achieving the optimal solution for the problem of eqn. (5). Let $\mathbf{X}^{\text{sv}} \in \mathbb{R}^{p \times d}$, $\mathbf{Y}^{\text{sv}} \in \mathbb{R}^{p \times p}$ be the support vector matrix and the corresponding labels respectively. Let $\mathbf{E} = \mathbf{V}^T \mathbf{V} - \mathbf{V}^T \mathbf{R} \mathbf{R}^T \mathbf{V}$. Then, we can write eqn (5) in terms of support vectors as,

$$\begin{aligned} Z_{opt} &= \sum_{i=1}^p \alpha_i^* - \frac{1}{2} \boldsymbol{\alpha}^{*T} \mathbf{Y}^{\text{sv}} \mathbf{X}^{\text{sv}} (\mathbf{X}^{\text{sv}})^T \mathbf{Y}^{\text{sv}} \boldsymbol{\alpha}^* \\ &= \sum_{i=1}^p \alpha_i^* - \frac{1}{2} \boldsymbol{\alpha}^{*T} \mathbf{Y}^{\text{sv}} \mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^T \mathbf{V} \boldsymbol{\Sigma} \mathbf{U}^T \mathbf{Y}^{\text{sv}} \boldsymbol{\alpha}^* \\ &= \sum_{i=1}^p \alpha_i^* - \frac{1}{2} \boldsymbol{\alpha}^{*T} \mathbf{Y}^{\text{sv}} \mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^T \mathbf{R} \mathbf{R}^T \mathbf{V} \boldsymbol{\Sigma} \mathbf{U}^T \mathbf{Y}^{\text{sv}} \boldsymbol{\alpha}^* \\ &\quad - \frac{1}{2} \boldsymbol{\alpha}^{*T} \mathbf{Y}^{\text{sv}} \mathbf{U} \boldsymbol{\Sigma} \mathbf{E} \boldsymbol{\Sigma} \mathbf{U}^T \mathbf{Y}^{\text{sv}} \boldsymbol{\alpha}^*. \end{aligned} \quad (6)$$

Let $\tilde{\boldsymbol{\alpha}}^* = [\tilde{\alpha}_1^*, \tilde{\alpha}_2^*, \dots, \tilde{\alpha}_p^*]^T \in \mathbb{R}^p$ be the vector achieving the optimal solution for the dimensionally-reduced SVM problem of eqn. (6) using $\tilde{\mathbf{X}}^{\text{sv}} = \mathbf{X}^{\text{sv}} \mathbf{R}$. Using the SVD of \mathbf{X}^{sv} ,

$$\tilde{Z}_{opt} = \sum_{i=1}^p \tilde{\alpha}_i^* - \frac{1}{2} \tilde{\boldsymbol{\alpha}}^{*T} \mathbf{Y}^{\text{sv}} \mathbf{U} \boldsymbol{\Sigma} \mathbf{V}^T \mathbf{R} \mathbf{R}^T \mathbf{V} \boldsymbol{\Sigma} \mathbf{U}^T \mathbf{Y}^{\text{sv}} \tilde{\boldsymbol{\alpha}}^*. \quad (7)$$

Since the constraints on $\boldsymbol{\alpha}^*$, $\tilde{\boldsymbol{\alpha}}^*$ do not depend on the data it is clear that $\tilde{\boldsymbol{\alpha}}^*$ is a feasible solution for the problem of eqn. (6). Thus, from the optimality of $\boldsymbol{\alpha}^*$,

and using eqn. (7), it follows that

$$\begin{aligned}
 Z_{opt} &= \sum_{i=1}^p \alpha_i^* - \frac{1}{2} \alpha^{*T} \mathbf{Y}^{\text{sv}} \mathbf{U} \Sigma \mathbf{V}^T \mathbf{R} \mathbf{R}^T \mathbf{V} \Sigma \mathbf{U}^T \mathbf{Y}^{\text{sv}} \alpha^* \\
 &\quad - \frac{1}{2} \alpha^{*T} \mathbf{Y}^{\text{sv}} \mathbf{U} \Sigma \mathbf{E} \mathbf{E} \Sigma \mathbf{U}^T \mathbf{Y}^{\text{sv}} \alpha^* \\
 &\geq \sum_{i=1}^p \tilde{\alpha}_i^* - \frac{1}{2} \tilde{\alpha}^{*T} \mathbf{Y}^{\text{sv}} \mathbf{U} \Sigma \mathbf{V}^T \mathbf{R} \mathbf{R}^T \mathbf{V} \Sigma \mathbf{U}^T \mathbf{Y}^{\text{sv}} \tilde{\alpha}^* \\
 &\quad - \frac{1}{2} \tilde{\alpha}^{*T} \mathbf{Y}^{\text{sv}} \mathbf{U} \Sigma \mathbf{E} \mathbf{E} \Sigma \mathbf{U}^T \mathbf{Y}^{\text{sv}} \tilde{\alpha}^* \\
 &= \tilde{Z}_{opt} - \frac{1}{2} \tilde{\alpha}^{*T} \mathbf{Y}^{\text{sv}} \mathbf{U} \Sigma \mathbf{E} \mathbf{E} \Sigma \mathbf{U}^T \mathbf{Y}^{\text{sv}} \tilde{\alpha}^*. \tag{8}
 \end{aligned}$$

We now analyze the second term using standard sub-multiplicativity properties and $\mathbf{V}^T \mathbf{V} = \mathbf{I}$. Taking $\mathbf{Q} = \tilde{\alpha}^{*T} \mathbf{Y}^{\text{sv}} \mathbf{U} \Sigma$,

$$\begin{aligned}
 &\frac{1}{2} \tilde{\alpha}^{*T} \mathbf{Y}^{\text{sv}} \mathbf{U} \Sigma \mathbf{E} \mathbf{E} \Sigma \mathbf{U}^T \mathbf{Y}^{\text{sv}} \tilde{\alpha}^* \\
 &\leq \frac{1}{2} \|\mathbf{Q}\|_2 \|\mathbf{E}\|_2 \left\| \mathbf{Q}^T \right\|_2 \\
 &= \frac{1}{2} \|\mathbf{E}\|_2 \|\mathbf{Q}\|_2^2 \\
 &= \frac{1}{2} \|\mathbf{E}\|_2 \left\| \tilde{\alpha}^{*T} \mathbf{Y}^{\text{sv}} \mathbf{U} \Sigma \mathbf{V}^T \right\|_2^2 \\
 &= \frac{1}{2} \|\mathbf{E}\|_2 \left\| \tilde{\alpha}^{*T} \mathbf{Y}^{\text{sv}} \mathbf{X}^{\text{sv}} \right\|_2^2. \tag{9}
 \end{aligned}$$

Combining eqns. (8) and (9), we get

$$Z_{opt} \geq \tilde{Z}_{opt} - \frac{1}{2} \|\mathbf{E}\|_2 \left\| \tilde{\alpha}^{*T} \mathbf{Y}^{\text{sv}} \mathbf{X}^{\text{sv}} \right\|_2^2. \tag{10}$$

We now proceed to bound the second term in the right-hand side of the above equation. Towards that end, we bound the difference:

$$\begin{aligned}
 &\left| \tilde{\alpha}^{*T} \mathbf{Y}^{\text{sv}} \mathbf{U} \Sigma \left(\mathbf{V}^T \mathbf{R} \mathbf{R}^T \mathbf{V} - \mathbf{V}^T \mathbf{V} \right) \Sigma \mathbf{U}^T \mathbf{Y}^{\text{sv}} \tilde{\alpha}^* \right| \\
 &= \left| \tilde{\alpha}^{*T} \mathbf{Y}^{\text{sv}} \mathbf{U} \Sigma (-\mathbf{E}) \Sigma \mathbf{U}^T \mathbf{Y}^{\text{sv}} \tilde{\alpha}^* \right| \\
 &\leq \|\mathbf{E}\|_2 \left\| \tilde{\alpha}^{*T} \mathbf{Y}^{\text{sv}} \mathbf{U} \Sigma \right\|_2^2 \\
 &= \|\mathbf{E}\|_2 \left\| \tilde{\alpha}^{*T} \mathbf{Y}^{\text{sv}} \mathbf{U} \Sigma \mathbf{V}^T \right\|_2^2 \\
 &= \|\mathbf{E}\|_2 \left\| \tilde{\alpha}^{*T} \mathbf{Y}^{\text{sv}} \mathbf{X}^{\text{sv}} \right\|_2^2.
 \end{aligned}$$

We can rewrite the above inequality as

$$\begin{aligned}
 &\left| \left\| \tilde{\alpha}^{*T} \mathbf{Y}^{\text{sv}} \mathbf{X}^{\text{sv}} \mathbf{R} \right\|_2^2 - \left\| \tilde{\alpha}^{*T} \mathbf{Y}^{\text{sv}} \mathbf{X}^{\text{sv}} \right\|_2^2 \right| \\
 &\leq \|\mathbf{E}\|_2 \left\| \tilde{\alpha}^{*T} \mathbf{Y}^{\text{sv}} \mathbf{X}^{\text{sv}} \right\|_2^2 \\
 &\leq \frac{\|\mathbf{E}\|_2}{1 - \|\mathbf{E}\|_2} \left\| \tilde{\alpha}^{*T} \mathbf{Y}^{\text{sv}} \mathbf{X}^{\text{sv}} \mathbf{R} \right\|_2^2.
 \end{aligned}$$

Combining with eqn. (10), we get

$$Z_{opt} \geq \tilde{Z}_{opt} - \frac{1}{2} \left(\frac{\|\mathbf{E}\|_2}{1 - \|\mathbf{E}\|_2} \right) \left\| \tilde{\alpha}^{*T} \mathbf{Y}^{\text{sv}} \mathbf{X}^{\text{sv}} \mathbf{R} \right\|_2^2. \tag{11}$$

Now recall that $\mathbf{w}^{*T} = \alpha^{*T} \mathbf{Y}^{\text{sv}} \mathbf{X}^{\text{sv}}$, $\tilde{\mathbf{w}}^{*T} = \tilde{\alpha}^{*T} \mathbf{Y}^{\text{sv}} \mathbf{X}^{\text{sv}} \mathbf{R}$, $\|\mathbf{w}^*\|_2^2 = \sum_{i=1}^p \alpha_i^*$, and $\|\tilde{\mathbf{w}}^*\|_2^2 = \sum_{i=1}^p \tilde{\alpha}_i^*$. Then, the optimal solutions Z_{opt} and \tilde{Z}_{opt} can be expressed as follows:

$$Z_{opt} = \|\mathbf{w}^*\|_2^2 - \frac{1}{2} \|\mathbf{w}^*\|_2^2 = \frac{1}{2} \|\mathbf{w}^*\|_2^2. \tag{12}$$

$$\tilde{Z}_{opt} = \|\tilde{\mathbf{w}}^*\|_2^2 - \frac{1}{2} \|\tilde{\mathbf{w}}^*\|_2^2 = \frac{1}{2} \|\tilde{\mathbf{w}}^*\|_2^2. \tag{13}$$

Combining eqns. (11), (12) and (13), we get $\|\mathbf{w}^*\|_2^2 \geq \|\tilde{\mathbf{w}}^*\|_2^2 - \left(\frac{\|\mathbf{E}\|_2}{1 - \|\mathbf{E}\|_2} \right) \|\tilde{\mathbf{w}}^*\|_2^2 = \left(1 - \frac{\|\mathbf{E}\|_2}{1 - \|\mathbf{E}\|_2} \right) \|\tilde{\mathbf{w}}^*\|_2^2$.

Let $\gamma^* = \|\mathbf{w}^*\|_2^{-1}$ be the geometric margin of the problem of eqn. (6) and let $\tilde{\gamma}^* = \|\tilde{\mathbf{w}}^*\|_2^{-1}$ be the geometric margin of the problem of eqn. (7). Then, the above equation implies: $\gamma^{*2} \leq \left(1 - \frac{\|\mathbf{E}\|_2}{1 - \|\mathbf{E}\|_2} \right)^{-1} \tilde{\gamma}^{*2} \Rightarrow \tilde{\gamma}^{*2} \geq \left(1 - \frac{\|\mathbf{E}\|_2}{1 - \|\mathbf{E}\|_2} \right) \gamma^{*2}$. The result follows because $\|\mathbf{E}\|_2 \leq \epsilon/2$ by our choice of r , and so $\|\mathbf{E}\|_2 / (1 - \|\mathbf{E}\|_2) \leq \epsilon$. \square

4.2 Geometry is preserved by Unsupervised Feature Selection

In the unsupervised setting, the next theorem says that with a number of features proportional to the rank of the training data (which is at most the number of data points), we preserve B^2/γ^2 , thus ensuring comparable generalization error bounds (B is the radius of the minimum enclosing ball).

Theorem 2. *Given $\epsilon \in (0, 1)$, run unsupervised BSS-feature selection on the full data \mathbf{X}^{tr} with $r_2 = O(\rho/\epsilon^2)$, where $\rho = \text{rank}(\mathbf{X}^{\text{tr}})$, to obtain the feature sampling and rescaling matrix \mathbf{R} . Let γ^* and $\tilde{\gamma}^*$ be the margins obtained by solving the SVM dual (2) with $(\mathbf{X}^{\text{tr}}, \mathbf{Y}^{\text{tr}})$ and $(\mathbf{X}^{\text{tr}} \mathbf{R}, \mathbf{Y}^{\text{tr}})$ respectively; and, let B and \tilde{B} be the radii for the data matrices \mathbf{X}^{tr} and $\mathbf{X}^{\text{tr}} \mathbf{R}$ respectively. Then,*

$$\frac{\tilde{B}^2}{\tilde{\gamma}^{*2}} \leq \frac{(1 + \epsilon) B^2}{(1 - \epsilon) \gamma^{*2}} = (1 + O(\epsilon)) \frac{B^2}{\gamma^{*2}}.$$

Proof. (sketch) The proof has two parts. First, as in Theorem 1 we prove that $\tilde{\gamma}^{*2} \geq (1 - \epsilon) \cdot \gamma^{*2}$. This proof is almost identical to the proof of Theorem 1 (replacing $(\mathbf{X}^{\text{sv}}, \mathbf{Y}^{\text{sv}})$ with $(\mathbf{X}^{\text{tr}}, \mathbf{Y}^{\text{tr}})$), and so we omit it. Second, we prove that $\tilde{B}^2 \leq (1 + \epsilon) B^2$. We give the result (with proof) in the supplementary material as Theorem 3. The theorem follows by combining these two results. \square

5 Experiments

We compared BSS and leverage-score sampling with RFE (Guyon et al. (2002)), LPSVM (Fung and Mangasarian (2004)), rank-revealing QR factorization (RRQR), random feature selection and full-data

Table 1: Most frequently selected features using the synthetic dataset.

	$r_1 = 30$		$r_1 = 40$	
	$k = 40$	$k = 50$	$k = 40$	$k = 50$
BSS	40, 39, 34, 36, 37	50, 49, 48, 47, 44	40, 39, 34, 37, 36	50, 49, 48, 47, 44
Lvg	40, 39, 37, 36, 34	50, 49, 48, 47, 46	40, 39, 37, 31, 32	50, 49, 48, 31, 47
RFE	40, 39, 38, 37, 36	50, 49, 48, 47, 46	40, 39, 38, 37, 36	50, 49, 48, 47, 46
LPSVM	40, 39, 38, 37, 34	50, 49, 48, 43, 40	40, 39, 38, 37, 34	50, 49, 48, 43, 40
RRQR	40, 30, 29, 28, 27	50, 30, 29, 28, 27	40, 39, 38, 37, 36	50, 40, 39, 38, 37

without feature selection on synthetic and real-world datasets. For the supervised case, we first run SVM on the training set, then run a feature selection method on the support-vector set and recalibrate the model using the support vector-set. For unsupervised feature selection, we perform feature selection on the training set. For LPSVM, we were not able to control the number of features and report the out-of-sample error using the features output by the algorithm. We **did not** extrapolate the values of out-of-sample error for LPSVM. We repeated random feature selection and leverage-score sampling five times. We performed ten-fold cross-validation and repeated it ten times. For medium-scale datasets like TechTC-300 we do not perform approximate BSS. For large-scale datasets like Reuters-CCAT (Lewis et al. (2004)) we use the approximate BSS method as described in Section 5.4. We used LIBSVM (Chang and Lin (2011)) as our SVM solver for medium-scale datasets and LIBLINEAR (Fan et al. (2008)) for large-scale datasets. We do not report running times in our experiments, since feature selection is an offline-task. We implemented all our algorithms in MATLAB R2013b on an Intel i-7 processor with 16GB RAM. BSS and leverage-score sampling are better than LPSVM and RRQR and comparable to RFE on 49 TechTC-300 datasets.

5.1 BSS Implementation Issues

At every iteration, there can be multiple columns which satisfy the condition $\mathcal{U}(\mathbf{v}_i, \delta_U, \mathbf{A}_\tau, U_\tau) \leq \mathcal{L}(\mathbf{v}_i, \delta_L, \mathbf{A}_\tau, L_\tau)$. Batson et al. (2009) suggest picking any column which satisfies this constraint. Selecting a column naively leaves out important features required for classification. Therefore, we choose the column \mathbf{v}_i which has not been selected in previous iterations and whose Euclidean-norm is highest among the candidate set. Columns with zero Euclidean norm never get selected by the algorithm. In our implementation, we do not use the data center as one of the inputs (since computing the center involves solving a quadratic program).

5.2 Experiments on Supervised Feature Selection

Synthetic Data: We generate synthetic data as described in Bhattacharyya (2004), where we control

the number of relevant features in the dataset. The dataset has n data-points and d features. The class label y_i of each data-point was randomly chosen to be 1 or -1 with equal probability. The first k features of each data-point \mathbf{x}_i are the relevant features and are drawn from $y_i \mathcal{N}(-j, 1)$ distribution, where $\mathcal{N}(\mu, \sigma^2)$ is a random normal distribution with mean μ and variance σ^2 and j varies from 1 to k . The remaining $(d - k)$ features are chosen from a $\mathcal{N}(0, 1)$ distribution and are noisy features. By construction, among the first k features, the k th feature has the most discriminatory power, followed by $(k - 1)$ th feature and so on. We set n to 200 and d to 1000. We set k to 40 and 50 and ran two sets of experiments. We set the value of r_1 , i.e. the number of features selected, to 30 and 40 for all experiments. We performed ten-fold cross-validation and repeated it ten times. We used LIBSVM with default settings and set $C = 1$. We compared with the other methods. The mean out-of-sample error was 0 for all methods for both $k = 40$ and $k = 50$. Table 1 shows the set of five most frequently selected features by the different methods for one such synthetic dataset. The top features picked up by the different methods are the relevant features by construction and also have good discriminatory power. This shows that supervised BSS and leverage-score sampling are as good as any other method in terms of feature selection. We repeated our experiments on ten different synthetic datasets and each time, the five most frequently selected features were from the set of relevant features. Thus, by selecting only 3% -4% of all features, we show that we are able to obtain the most discriminatory features along with good out-of-sample error using BSS and leverage-score sampling.

TechTC-300: For our first real dataset, we use 49 datasets of TechTC-300 (Davidov et al. (2004)) which contain binary classification tasks. Each data matrix consists of 150-280 documents (the rows of the data matrix), and each document is described with respect to 10,000-50,000 words (features are columns of the matrix). We removed all words with at most four letters from the datasets. We set the parameter $C = 1$ in LIBSVM and used default settings. We tried different values of C for the full-dataset and the out-of-sample error averaged over 49 TechTC-300 docu-

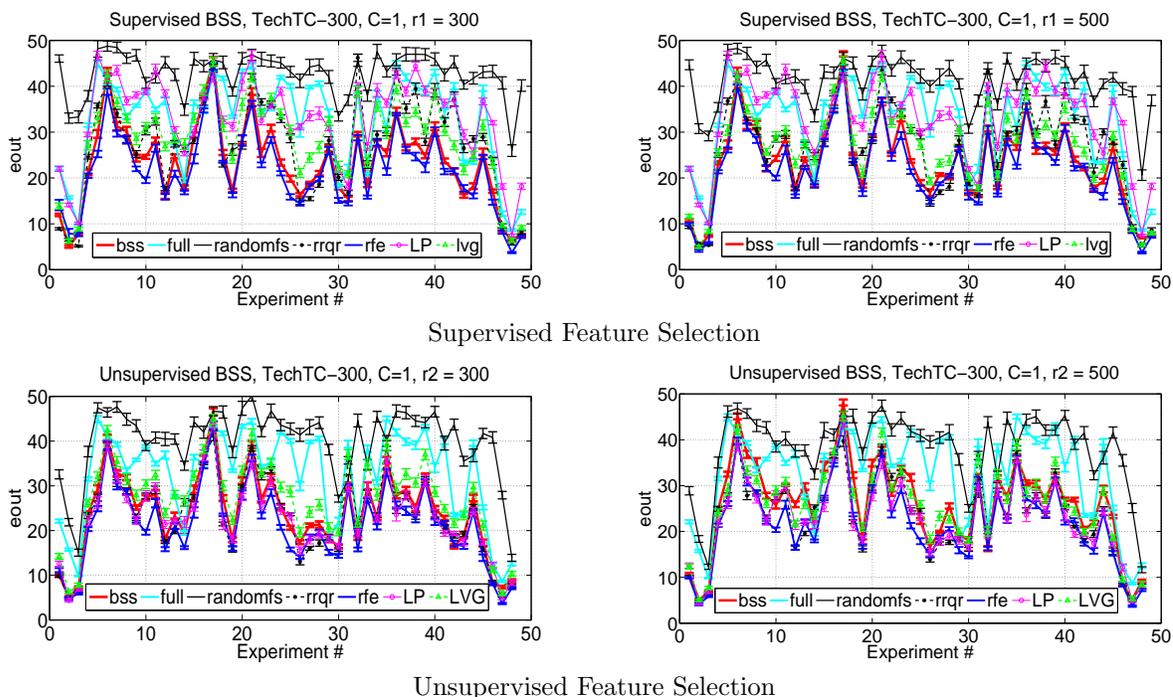


Figure 1: Plots of out-of-sample error (eout) of Supervised and Unsupervised BSS and leverage-score compared with other methods for 49 TechTC-300 documents averaged over ten ten-fold cross validation experiments. Vertical bars represent standard deviation.

Table 2: A subset of the TechTC matrices of our study

	id1	id2
(i)	Arts: Music: Styles: Opera	Arts: Education: Language: Reading Instructions
(ii)	Arts: Music: Styles: Opera	US Navy: Decommissioned Attack Submarines
(iii)	US: Michigan: Travel & Tourism	Recreation:Sailing Clubs: UK
(iv)	US: Michigan: Travel & Tourism	Science: Chemistry: Analytical: Products
(v)	US: Colorado: Localities: Boulder	Europe: Ireland: Dublin: Localities

Table 3: Frequently occurring terms of the five TechTC-300 datasets of Table 2 selected by supervised BSS and Leverage-score sampling.

	BSS	Leverage-score Sampling
(i)	reading, education, opera, frame	reading, opera, frame, spacer
(ii)	submarine, hullnumber, opera, tickets	hullnumber, opera, music, tickets
(iii)	michigan, vacation, yacht, sailing	sailing, yacht, michigan, vacation
(iv)	chemical, michigan, environmental, asbestos	travel, vacation, michigan, services, environmental
(v)	ireland, dublin, swords, boulder, colorado	ireland, boulder, swords, school, grade

Table 4: Results of Approximate BSS. CCAT (train / test): (23149 / 781265), $d=47236$. Mean and standard deviation (in parenthesis) of out-of-sample error (eout). Eout of full-data is 8.66 ± 0.54 .

Eout	r_1	BSS ($t = 128$)	BSS ($t = 256$)	RRQR	RFE	LPSVM
CCAT	1024	10.53 (0.59)	10.35 (0.64)	9.97 (0.62)	8.92 (0.57)	9.97 (0.55)
CCAT	2048	11.13 (0.66)	10.63 (0.62)	10.04 (0.66)	8.56 (0.54)	9.97 (0.55)

ments did not change much, so we report the results of $C = 1$. We set the number of features to 300, 400 and 500. Fig 1 shows the out-of-sample error for the 49 datasets for $r_1 = 300$ and $r_1 = 500$. We include the results of $r_1 = 400$ in the supplementary material. For the supervised feature selection, BSS is comparable to RFE and leverage-score sampling and better than RRQR, LPSVM, full-data and uniform sampling

in terms of out-of-sample error. For LPSVM, the number of selected features averaged over 49 datasets was greater than 500, but it performed worse than BSS and leverage-score sampling. Leverage-score sampling is comparable to BSS and better than RRQR, LPSVM, full-data and uniform sampling and slightly worse than RFE.

We list the most frequently occurring words selected

by supervised BSS and leverage-score for the $r_1 = 300$ case for five TechTC-300 datasets over 100 training sets. Table 2 shows the names of the five TechTC-300 document-term matrices. The words shown in Table 3 were selected in all cross-validation experiments for these five datasets. The words are closely related to the categories to which the documents belong, which shows that BSS and Leverage-score sampling select important features from the support-vector matrix. For example, for the document-pair (*ii*), where the documents belong to the category of “Arts:Music:Styles:Opera” and “US:Navy: Decommissioned Attack Submarines”, the BSS algorithm selects submarine, hullnumber, opera, tickets and Leverage-score sampling selects hullnumber, opera, music, tickets which are closely related to the two classes. Thus, we see that using only 2%-4% of all features we are able to obtain good out-of-sample error.

5.3 Experiments on Unsupervised Feature Selection

For the unsupervised feature selection case, we performed experiments on the same 49 TechTC-300 datasets and set r_2 to 300, 400 and 500. We include the results for $r_2 = 300$ and $r_2 = 500$ in Fig 1. For LPSVM, the number of selected features averaged over 49 datasets was close to 300. In the unsupervised case, BSS and leverage-score sampling are comparable to each other and also comparable to the other methods RRQR, LPSVM and RFE. These methods are better than random feature selection and full-data without feature selection. This shows that unsupervised BSS and leverage-score sampling are competitive feature selection algorithms.

Supervised feature selection is comparable to unsupervised feature selection for BSS, Leverage-score sampling and RFE, while unsupervised RRQR and LPSVM are better than their supervised versions. Running BSS (or leverage-score sampling) on the support-vector set is equivalent to running BSS (or leverage-score sampling) on the training data. However, RRQR and LPSVM are primarily used as unsupervised feature selection techniques and so they perform well in that setting. RFE is a heuristic based on SVM and running RFE on the support-vectors is equivalent to running RFE on the training data.

5.4 Approximate BSS

We describe a heuristic to make supervised BSS scalable to large-scale datasets. For datasets with large number of support vectors, we premultiply the support vector matrix \mathbf{X} with a random gaussian matrix $\mathbf{G} \in \mathbb{R}^{t \times p}$ to obtain $\hat{\mathbf{X}} = \mathbf{GX}$ and then use BSS to select features from the right singular vectors of $\hat{\mathbf{X}}$. The right singular vectors of $\hat{\mathbf{X}}$ closely approximates

the right singular vectors of \mathbf{X} . Hence the columns selected from $\hat{\mathbf{X}}$ will be approximately same as the columns selected from \mathbf{X} . We performed experiments on a subset of Reuters Corpus dataset, namely reuters-CCAT, which contains binary classification task. We used the L2-regularized L2-loss SVM formulation in the dual form in LIBLINEAR and set the value of C to 10. We experimented with different values of C on the full-dataset, and since there was small change in classification accuracy among the different values of C , we chose $C = 10$ for our experiments. We premultiplied the support vector matrix with a random gaussian matrix of size $t \times p$, where p is the number of support vectors and t was set to 128 and 256. We repeated our experiments five times using five different random gaussian matrices to get around the randomness. We set the value of r_1 in BSS to 1024 and 2048. LPSVM selects 1898 features for CCAT. Table 4 shows the results of our experiments. We observe that the out-of-sample error using approx-BSS is close to that of RRQR and comparable to RFE, LPSVM and full-data. The out-of-sample error of approx-BSS decreases with an increase in the value of t . This shows that we get a good approximation of the right singular vectors of the support vector matrix with an increase in number of projections.

6 Conclusions

Our simple method of extending an unsupervised feature selection method into a supervised one for SVM not only has a provable guarantee, but also works well empirically: BSS and leverage-score sampling are comparable and often better than prior state-of-the-art feature selection methods for SVM, and those methods don’t come with guarantees.

Our supervised sparsification algorithms only preserve the margin for the support vectors in the feature space. We do not make any claims about the margin of the full data in the feature space constructed from the support vectors. This appears challenging and it would be interesting to see progress made in this direction: can one choose $O(\#\text{support vectors})$ features for the full data set and obtain provable guarantees on the margin and data radius? There have been recent advances in approximate leverage-scores for large-scale datasets. A possible future work in this direction would be to see if those algorithms indeed work well with SVMs.

Acknowledgements

PD and SP are supported by NSF IIS-1447283 and IIS-1319280 respectively.

References

- J. Batson, D. Spielman, and N. Srivastava. Twice-ramanujan sparsifiers. In *Proceedings of the 41st annual ACM STOC*, pages 255–262, 2009.
- C. Bhattacharyya. Second order cone programming formulations for feature selection. *JMLR*, 5:1417–1433, 2004.
- C. Boutsidis and M. Magdon-Ismail. Deterministic feature selection for k -means clustering. *IEEE Transactions on Information Theory*, 59(9):6099–6110, 2013.
- C. Boutsidis, M. W. Mahoney, and P. Drineas. Unsupervised feature selection for the k -means clustering problem. *Advances in Neural Information Processing Systems (NIPS)*, 2009.
- C.-C. Chang and C.-J. Lin. Libsvm: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27, 2011. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- N. Cristianini and J. Shawe-Taylor. *Support Vector Machines and other kernel-based learning methods*. Cambridge University Press, 2000.
- A. Dasgupta, P. Drineas, B. Harb, V. Josifovski, and M. Mahoney. Feature selection methods for text classification. In *Proceedings of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 230–239, 2007.
- D. Davidov, E. Gabrilovich, and S. Markovitch. Parameterized generation of labeled datasets for text categorization based on a hierarchical directory. In *Proceedings of the 27th Annual International ACM SIGIR Conference*, pages 250–257, 2004. <http://techtc.cs.technion.ac.il/techtc300/techtc300.html>.
- H. Do, A. Kalousis, and M. Hilario. Feature weighting using margin and radius based error bound optimization in svms. In *European Conference on Machine Learning (ECML)*, pages 315–329, 2009b.
- R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. Liblinear: A library for large linear classification. *JMLR*, pages 1871–1874, 2008.
- M. G. Fung and O. Mangasarian. A feature selection newton method for support vector machine classification. *Comput. Optim. Appl.*, 28(2):185–202, 2004.
- R. Gilad-Bachrach, A. Navot, and N. Tishby. Margin based feature selection-theory and algorithms. In *Proceedings of the twenty-first international conference on Machine learning (ICML)*, page 43, 2004.
- I. Guyon, J. Weston, S. Barnhill, and V. Vapnik. Gene selection for cancer classification using support vector machines. *Machine Learning*, 46(1-3):389–422, 2002.
- A. Kalousis and H. Do. Convex formulations of radius-margin based support vector machines. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*, pages 169–177, 2013.
- D. D. Lewis, Y. Yang, T. G. Rose, and F. Li. Rcv1: A new benchmark collection for text categorization research. *JMLR*, pages 361–397, 2004.
- C. Park, K.-R. Kim, R. Myung, and J.-Y. Koo. Oracle properties of scad-penalized support vector machine. *Journal of Statistical Planning and Inference*, 142(8):2257–2270, 2012.
- S. Paul and P. Drineas. Deterministic feature selection for regularized least squares classification. In *Machine Learning and Knowledge Discovery in Databases (ECML-PKDD)*, volume 8725 of *LNCS*, pages 533–548, 2014.
- S. Paul, C. Boutsidis, M. Magdon-Ismail, and P. Drineas. Random projections for support vector machines. In *Sixteenth International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 498–506. JMLR W&CP 31, 2013.
- S. Paul, C. Boutsidis, M. Magdon-Ismail, and P. Drineas. Random projections for linear support vector machines. *ACM Transactions on Knowledge Discovery from Data*, 8(4):22:1–22:25, 2014.
- A. Rakotomamonjy. Variable selection using svm based criteria. *JMLR*, 3:1357–1370, 2003.
- M. Rudelson and R. Vershynin. Sampling from large matrices: An approach through geometric functional analysis. *J. ACM*, 54(4), 2007.
- M. Tan, L. Wang, and I. Tsang. Learning sparse svm for feature selection on very high dimensional datasets. In *Proceedings of the 27th International Conference on Machine Learning (ICML)*, pages 1047–1054, 2010.
- V. Vapnik. *Statistical Learning Theory. Theory of Probability and its Applications*, 16:264–280, 1998.
- V. Vapnik and A. Chervonenkis. On the Uniform Convergence of Relative Frequencies of Events to their Probabilities. *Theory of Probability and its Applications*, 16:264–280, 1971.
- L. Wang, J. Zhu, and H. Zou. The doubly regularized support vector machine. *Statistica Sinica*, 16(2):589, 2006.
- L. Wang, J. Zhu, and H. Zou. Hybrid huberized support vector machines for microarray classification and gene selection. *Bioinformatics*, 24(3):412–419, 2008.
- J. Weston, S. Mukherjee, O. Chapelle, M. Pontil, T. Poggio, and V. Vapnik. Feature selection for svms. In *NIPS*, volume 12, pages 668–674, 2000.
- J. Weston, A. Elisseeff, B. Schölkopf, and M. Tipping. Use of the zero norm with linear models and kernel methods. *JMLR*, 3:1439–1461, 2003.
- G. Ye, Y. Chen, and X. Xie. Efficient variable selection in support vector machines via the alternating direction method of multipliers. In *AISTATS*, pages 832–840, 2011.