

HIPAD - A Hybrid Interior-Point Alternating Direction algorithm for knowledge-based SVM and feature selection

Z. Qin¹, X. Tang², I. Akrotirianakis³, and A. Chakraborty³

¹ Columbia University, New York, NY,
zq2107@columbia.edu

² Lehigh University, Bethlehem, PA,
xit210@lehigh.edu

³ Siemens Corporation, Corporate Technology, Princeton, NJ,
(ioannis.akrotirianakis , amit.chakraborty)@siemens.com

Abstract. We consider classification tasks in the regime of scarce labeled training data in high dimensional feature space, where specific expert knowledge is also available. We propose a new hybrid optimization algorithm that solves the elastic-net support vector machine (SVM) through an alternating direction method of multipliers in the first phase, followed by an interior-point method for the classical SVM in the second phase. Both SVM formulations are adapted to knowledge incorporation. Our proposed algorithm addresses the challenges of automatic feature selection, high optimization accuracy, and algorithmic flexibility for taking advantage of prior knowledge. We demonstrate the effectiveness and efficiency of our algorithm and compare it with existing methods on a collection of synthetic and real-world data.

Keywords: Support Vector Machine, Alternating Direction Method of Multipliers, Interior Point Methods, Elastic Net, Domain Knowledge

1 Introduction

Classification tasks on data sets with large feature dimensions are very common in real-world machine learning applications. Typical examples include microarray data for gene selection and text documents for natural language processing. Despite the large number of features present in the data sets, usually only small subsets of the features are relevant to the particular learning tasks, and local correlation among the features is often observed. Hence, feature selection is required for good model interpretability. Popular classification techniques, such as support vector machine (SVM) and logistic regression, are formulated as convex optimization problems. An extensive literature has been devoted to optimization algorithms that solve variants of these classification models with sparsity regularization [12, 17]. Many of them are based on first-order (gradient-based) methods, mainly because the size of the optimization problem is very large. The advantage of first-order methods is that their computational and memory requirements at

each iteration are low and as a result they can handle the large optimization problems occurring in classification problems. Their major disadvantage is their slow convergence, especially when a good approximation of the feature support has been identified. Second-order methods exhibit fast local convergence, but their computational and memory requirements are much more demanding, since they need to store and invert the Newton matrix at every iteration. It is therefore very important to be able to intelligently combine the advantages of both the first and the second order optimization methods in such a way that the resulting algorithm can solve large classification problems efficiently and accurately. As we will demonstrate in this paper such combination is possible by taking advantage of the problem structure and the change in its size during the solution process. In addition, we will also show that our algorithmic framework is flexible enough to incorporate prior knowledge to improve classification performance.

1.1 Related Work

The above requirements demand three features from a learning algorithm: 1. it should be able to automatically select features which are possibly in groups and highly correlated; 2. it has to solve the optimization problem in the training phase efficiently and with high accuracy; and 3. the learning model needs to be flexible enough so that domain knowledge can be easily incorporated. Existing methods are available in the literature that meet some of the above requirements *individually*. For enforcing sparsity in the solution, efficient optimization algorithms such as that proposed in [9] can solve large-scale sparse logistic regression. On the other hand, the L_1 -regularization is unstable with the presence of highly correlated features - among a group of such features, essentially one of them is selected in a random manner. To handle local correlation among groups of features, the elastic-net regularization [22] has been successfully applied to SVM [20] and logistic regression [15]. However, incorporating domain knowledge into the logistic regression formulation is not straightforward. For SVM, including such knowledge in the optimization process has been demonstrated in [5]. Recently, an Alternating Direction Method of Multipliers (ADMM) has been proposed for the elastic-net SVM (ENSVM) [21]. ADMM is quick to find an approximate solution to the ENSVM problem, but it is known to converge very slowly to high accuracy optimal solutions [2]. The interior-point methods (IPM) for SVM are known to be able to achieve high accuracy in their solutions with a polynomial iteration complexity, and the dual SVM formulation is independent of the feature space dimensionality. However, the classical L_2 -norm SVM is not able to perform automatic feature selection. Although the elastic-net SVM can be formulated as a QP (in the primal form), its problem size grows substantially with the feature dimensionality. Due to the need to solve a Newton system in each iteration, the efficiency of IPM quickly deteriorates as the feature dimension becomes large.

1.2 Main Contributions

In this paper we propose a new hybrid algorithmic framework for SVM to address *all* of the above challenges and requirements *simultaneously*. Our framework combines the advantages of a first-order optimization algorithm (through the use of ADMM) and a second-order method (via IPM) to achieve both superior speed and accuracy. Through a novel algorithmic approach that is able to incorporate expert knowledge, our proposed framework is able to exploit domain knowledge to improve feature selection, and hence, prediction accuracy. Besides efficiency and generalization performance, we demonstrate through experiments on both synthetic and real data that our method is also more robust to inaccuracy in the supplied knowledge than existing approaches.

2 A Two-phase Hybrid Optimization Algorithm

As previously mentioned, for data sets with many features, the high dimensionality of the feature space still poses a computational challenge for IPM. Fortunately, many data sets of this kind are very sparse, and the resulting classifier w is also expected to be sparse, i.e. only a small subset of the features are expected to carry significant weights in classification. Naturally, it is ideal for IPM to train a classifier on the most important features only.

Inspired by the Hybrid Iterative Shrinkage (HIS) [16] algorithm for training large-scale sparse logistic regression classifiers, we propose a two-phase algorithm to shrink the feature space appropriately so as to leverage the high accuracy of IPM while maintaining efficiency. Specifically, we propose to solve an elastic-net SVM (ENSVM) or doubly-regularized SVM (DrSVM) [20] problem during the first phase of the algorithm. The elastic-net regularization performs feature selection with grouping effect and has been shown to be effective on data sets with many but sparse features and high local correlations [22]. This is the case for text classification, microarray gene expression, and fMRI data sets. The support of the weight vector w for ENSVM usually stabilizes well before the algorithm converges to the optimal solution. Taking advantage of that prospect, we can terminate the first phase of the hybrid algorithm early and proceed to solve a classical SVM problem with the reduced feature set in the second phase, using an IPM solver.

2.1 Solving the Elastic Net SVM using ADMM

SVM can be written in the regularized regression form as

$$\min_{\mathbf{w}, b} \frac{1}{N} \sum_{i=1}^N (1 - (y_i(\mathbf{x}_i^T \mathbf{w} + b)))_+ + \frac{\lambda}{2} \|\mathbf{w}\|_2^2, \quad (1)$$

where the first term is an averaged sum of the hinge losses and the second term is viewed as a ridge regularization on w . It is easy to see from this form that

the classical SVM does not enforce sparsity in the solution, and w is generally dense. The ENSVM adds an L_1 regularization on top of the ridge regularization term, giving

$$\min_{\mathbf{w}, b} \frac{1}{N} \sum_{i=1}^N (1 - y_i(\mathbf{x}_i^T \mathbf{w} + b))_+ + \lambda_1 \|\mathbf{w}\|_1 + \frac{\lambda_2}{2} \|\mathbf{w}\|_2^2. \quad (2)$$

Compared to the Lasso (L_1 -regularized regression) [18], the elastic-net has the advantage of selecting highly correlated features in groups (i.e. the grouping effect) while still enforcing sparsity in the solution. This is a particularly attractive feature for text document data, which is common in the hierarchical classification setting. Adopting the elastic-net regularization as in (2) brings the same benefit to SVM for training classifiers.

To approximately solve problem (2), we adopt the alternating direction method of multipliers (ADMM) for elastic-net SVM recently proposed in [21]. ADMM has a long history dating back to the 1970s [6]. Recently, it has been successfully applied to problems in machine learning [2]. ADMM is a special case of the inexact augmented Lagrangian (IAL) method [14] for the structured unconstrained problem

$$\min_x F(x) \equiv f(x) + g(Ax), \quad (3)$$

where both functions $f(\cdot)$ and $g(\cdot)$ are convex. We can decouple the two functions by introducing an auxiliary variable y and convert problem (3) into an equivalent constrained optimization problem

$$\min_{x, y} f(x) + g(y), \quad s.t. \ Ax = y. \quad (4)$$

This technique is often called variable-splitting [3]. The IAL method approximately minimizes in each iteration the augmented Lagrangian of (4) defined by $\mathcal{L}(x, y, \gamma) := f(x) + g(y) + \gamma^T(y - Ax) + \frac{\mu}{2} \|Ax - y\|_2^2$, followed by an update to the Lagrange multiplier $\gamma \leftarrow \gamma - \mu(Ax - y)$. The IAL method is guaranteed to converge to the optimal solution of (3), as long as the subproblem of approximately minimizing the augmented Lagrangian is solved with an increasing accuracy [14]. ADMM can be viewed as a practical implementation of IAL, where the subproblem is solved approximately by minimizing $\mathcal{L}(x, y; \gamma)$ with respect to x and y *alternatingly once*. Eckstein and Bertsekas [4] established the convergence of ADMM for the case of two-way splitting. Now applying variable-splitting and ADMM to problem (2), [21] introduced auxiliary variables (\mathbf{a}, \mathbf{c}) and linear constraints so that the non-smooth hinge loss and L_1 -norm in the objective function are decoupled, making it easy to optimize over each of the variables. Specifically, problem (2) is transformed into an equivalent constrained form

$$\begin{aligned} \min_{\mathbf{w}, b, \mathbf{a}, \mathbf{c}} \quad & \frac{1}{N} \sum_{i=1}^N (a_i)_+ + \lambda_1 \|\mathbf{c}\|_1 + \frac{\lambda_2}{2} \|\mathbf{w}\|_2^2 \\ s.t. \quad & \mathbf{a} = \mathbf{e} - Y(X\mathbf{w} + b\mathbf{e}) \quad \text{and} \quad \mathbf{c} = \mathbf{w} \end{aligned} \quad (5)$$

where \mathbf{x}_i^T is the i -th row of X , and $Y = \text{diag}(\mathbf{y})$. The augmented Lagrangian

$$\begin{aligned} \mathcal{L}(\mathbf{w}, b, \mathbf{a}, \mathbf{c}, \gamma_1, \gamma_2) := & \frac{1}{N} \sum_{i=1}^N a_i + \lambda_1 \|\mathbf{c}\|_1 + \frac{\lambda_2}{2} \|\mathbf{w}\|_2^2 + \gamma_1^T (\mathbf{e} - Y(X\mathbf{w} + b\mathbf{e}) - \mathbf{a}) \\ & + \gamma_2^T (\mathbf{w} - \mathbf{c}) + \frac{\mu_1}{2} \|\mathbf{e} - Y(X\mathbf{w} + b\mathbf{e}) - \mathbf{a}\|_2^2 + \frac{\mu_2}{2} \|\mathbf{w} - \mathbf{c}\|_2^2 \quad (6) \end{aligned}$$

is then minimized with respect to (\mathbf{w}, b) , \mathbf{a} , and \mathbf{c} sequentially in each iteration, followed by an update to the Lagrange multipliers γ_1 and γ_2 . The original problem is thus decomposed into three subproblems consisting of computing the proximal operator of the hinge loss function (with respect to \mathbf{a}), solving a special linear system (with respect to (\mathbf{w}, b)), and performing a soft-thresholding operation (with respect to \mathbf{c}), which can all be done in an efficient manner. Due to lack of space in the paper, we have included the detailed solution steps in the Appendix (see Algorithm A.1 ADMM-ENSVM), where we define by $\mathcal{S}_\lambda(\cdot)$ the proximal operator associated with the hinge loss

$$\mathcal{S}_\lambda(\omega) = \begin{cases} \omega - \lambda, & \omega > \lambda; \\ 0, & 0 \leq \omega \leq \lambda; \\ \omega, & \omega < 0. \end{cases}$$

and $\mathcal{T}_\lambda(\omega) = \text{sgn}(\omega) \max\{0, |\omega| - \lambda\}$ is the shrinkage operator.

2.2 SVM via Interior-Point Method

Interior Point Methods enjoy fast convergence rates for a wide class of QP problems. Their theoretical polynomial convergence ($O(n \log \frac{1}{\epsilon})$) was first established by Mizuno [13]. In addition, Andersen *et al* [1] showed that the number of iterations needed by IPMs to converge is $O(\log n)$, which demonstrates that their computational effort increases in a slower rate than the size of the problem.

Both the primal and the dual SVM are QP problems. The primal formulation of SVM [19] is defined as

$$\begin{aligned} \text{(SVM-P)} \quad & \min_{\mathbf{w}, b, \xi, \mathbf{s}} \frac{1}{2} \mathbf{w}^T \mathbf{w} + c \mathbf{e}^T \xi \\ & s.t. \quad y_i (\mathbf{w}^T \mathbf{x}_i - b) + \xi_i - s_i = 1, \quad i = 1, \dots, N, \\ & \quad \mathbf{s} \geq 0, \xi \geq 0. \end{aligned}$$

whereas the dual SVM has the form

$$\begin{aligned} \text{(SVM-D)} \quad & \min_{\alpha} \frac{1}{2} \alpha^T Q \alpha - \mathbf{e}^T \alpha \\ & s.t. \quad \mathbf{y}^T \alpha = 0, \quad \text{and} \quad 0 \leq \alpha_i \leq c, \quad i = 1, \dots, N, \end{aligned}$$

where $Q_{ij} = y_i y_j \mathbf{x}_i^T \mathbf{x}_j = \bar{X} \bar{X}^T$. By considering the KKT conditions of (SVM-P) and (SVM-D), the optimal solution is given by $\mathbf{w} = \bar{X}^T \alpha = \sum_{i \in SV} \alpha_i y_i \mathbf{x}_i$,

where SV is the set of sample indices corresponding to the support vectors. The optimal bias term b can be computed from the complementary slackness condition $\alpha_i(y_i(\mathbf{x}_i^T \mathbf{w} + b) - 1 + \xi_i) = 0$.

Whether to solve (SVM-P) or (SVM-D) for a given data set depends on its dimensions as well as its sparsity. Even if X is a sparse matrix, Q in (SVM-D) is still likely to be dense, whereas the Hessian matrix in (SVM-P) is the identity. The primal problem (SVM-P), however, has a larger variable dimension and more constraints. It is often argued that one should solve (SVM-P) when the number of features is smaller than the number of samples, whereas (SVM-D) should be solved when the number of features is less than that of the samples. Since in the second phase of Algorithm 2.1 we expect to have identified a small number of promising features, we have decided to solve (SVM-D) by using IPM. Solving (SVM-D) is realized through the OOQP [7] software package that implements a primal-dual IPM for convex QP problems.

2.3 The Two-phase Algorithm

Let us keep in mind that the primary objective for the first phase is to appropriately reduce the feature space dimensionality without impacting the final prediction accuracy. As we mentioned above, the suitability of ADMM for the first phase depends on whether the support of the feature vector converges quickly or not. On an illustrative dataset from [21], which has 50 samples with 300 features each, ADMM converged in 558 iterations. The output classifier \mathbf{w} contained only 13 non-zero features, and the feature support converged in approximately 50 iteration (see Figure 1 in the Appendix for illustrative plots showing the early convergence of ADMM). Although the remaining more than 500 iterations are needed by ADMM in order to satisfy the optimality criteria, they do not offer any additional information regarding the feature selection process. Hence, it is important to monitor the change in the signs and indices of the support and terminate the first phase promptly. In our implementation, we adopt the criterion used in [16] and monitor the relative change in the iterates as a surrogate of the change in the support, i.e.

$$\frac{\|\mathbf{w}^{k+1} - \mathbf{w}^k\|}{\max(\|\mathbf{w}^k\|, 1)} < \epsilon_{tol}. \quad (7)$$

We have observed in our experiments that when the change over the iterates is small, the evolution of the support indices stabilizes too.

Upon starting the second phase, it is desirable for IPM to warm-start from the corresponding sub-vector of the solution returned by ADMM. It should also be noted that we apply IPM during the second phase to solve the classical L_2 -regularized SVM (1), instead of the ENSVM (2) in the first phase. There are two main reasons for this decision. First, although ENSVM can be reformulated as a QP, the size of the problem is larger than the classical SVM due to the additional linear constraints introduced by the L_1 -norm. Second, since we have already identified (approximately) the feature support in the first phase of the

algorithm, enforcing sparsity in the reduced feature space becomes less critical. The two-phase algorithm is summarized in Algorithm 2.1.

Algorithm 2.1 HIPAD (Hybrid Interior Point and Alternating Direction method)

- 1: Given $\mathbf{w}^0, b^0, \mathbf{a}^0, \mathbf{c}^0, \mathbf{u}^0$, and \mathbf{v}^0 .
 - 2: **PHASE 1: ADMM for ENSVM**
 - 3: $(\tilde{\mathbf{w}}^{\text{ADMM}}, b^{\text{ADMM}}) \leftarrow \text{ADMM-ENSVM}(\mathbf{w}^0, b^0, \mathbf{a}^0, \mathbf{c}^0, \mathbf{u}^0, \mathbf{v}^0)$
 - 4: **PHASE 2: IPM for SVM**
 - 5: $\tilde{\mathbf{w}} \leftarrow$ non-zero components of $\tilde{\mathbf{w}}^{\text{ADMM}}$
 - 6: $(\tilde{X}, \tilde{Y}) \leftarrow$ sub-matrices of (X, Y) corresponding to the support of $\tilde{\mathbf{w}}^{\text{ADMM}}$
 - 7: $(\mathbf{w}, b) \leftarrow \text{SVM-IPM}(\tilde{X}, \tilde{Y}, \tilde{\mathbf{w}})$, through (SVM-D).
 - 8: **return** (\mathbf{w}, b)
-

3 Domain Knowledge Incorporation

Very often, we have prior domain knowledge for specific classification tasks. Domain knowledge is most helpful when the training data does not form a comprehensive representation of the underlying unknown population, resulting in poor generalization performance of SVM on the unseen data from the same population. This often arises in situations where labeled training samples are scarce, while there is an abundance of unlabeled data.

For high dimensional data, ENSVM performs feature selection along with training to produce a simpler model and to achieve better prediction accuracy. However, the quality of the feature selection depends entirely on the training data. In pathological cases, it is very likely that the feature support identified by ENSVM does not form a good representation of the population. Hence, when domain knowledge about certain features is available, we should take it into consideration during the training phase and include the relevant features in the feature support should them be deemed important for classification.

In this section, we explore and propose a new approach to achieve this objective. We consider domain knowledge in the form of class-membership information associated with features. We can incorporate such information (or enforce such rules) in SVM by adding equivalent linear constraints to the SVM QP problem (KSVM) [5, 10]. To be specific, we can model the above information with the linear implication

$$B\mathbf{x} \leq \mathbf{d} \quad \Rightarrow \quad \mathbf{w}^T \mathbf{x} + b \geq 1, \quad (8)$$

where $B \in \mathbb{R}^{k_1 \times m}$ and $\mathbf{d} \in \mathbb{R}^{k_1}$. It is shown in [5] that by utilizing the non-homogeneous Farkas theorem of the alternative, (8) can be transformed into the following equivalent system of linear inequalities with a solution \mathbf{u}

$$B^T \mathbf{u} + \mathbf{w} = \mathbf{0}, \quad \mathbf{d}^T \mathbf{u} - b + 1 \leq 0, \quad \mathbf{u} \geq \mathbf{0}. \quad (9)$$

Similarly, for the linear implication for the negative class membership we have:

$$D\mathbf{x} \leq \mathbf{g} \Rightarrow \mathbf{w}^T \mathbf{x} + b \leq -1, \quad D \in \mathbb{R}^{k_2 \times m}, g \in \mathbb{R}^{k_2}, \quad (10)$$

which can be represented by the set of linear constraints in \mathbf{v}

$$D^T \mathbf{v} - \mathbf{w} = \mathbf{0}, \quad \mathbf{g}^T \mathbf{v} + b + 1 \leq 0, \quad \mathbf{v} \geq \mathbf{0}. \quad (11)$$

Hence, to incorporate the domain knowledge represented by (8) and (10) into SVM, Fung *et al* [5] simply add the linear constraints (9) and (11) to (SVM-P). Their formulation, however, increases both the variable dimension and the number of linear constraints by at least $2m$, where m is the number of features in the classification problem we want to solve. This is clearly undesirable when m is large, which is the scenario that we consider in this paper.

In order to avoid the above increase in the size of the optimization problem, we choose to penalize the quadratics $\|B^T \mathbf{u} + \mathbf{w}\|_2^2$ and $\|D^T \mathbf{v} - \mathbf{w}\|_2^2$ instead of their L_1 counterparts considered in [5]. By doing so the resulting problem is still a convex QP but with a much smaller size. Hence, we consider the following model for domain knowledge incorporation.

$$\begin{aligned} \text{(KSVM-P)} \quad & \min_{\mathbf{w}, b, \xi, \mathbf{u}, \mathbf{v}, \eta_u, \eta_v} \quad \frac{1}{2} \mathbf{w}^T \mathbf{w} + c \mathbf{e}^T \xi + \frac{\rho_1}{2} \|B^T \mathbf{u} + \mathbf{w}\|_2^2 \\ & + \rho_2 \eta_u + \frac{\rho_3}{2} \|D^T \mathbf{v} - \mathbf{w}\|_2^2 + \rho_4 \eta_v \\ \text{s.t.} \quad & y_i(\mathbf{w}^T \mathbf{x}_i + b) \geq 1 - \xi_i, \quad i = 1, \dots, N, \\ & \mathbf{d}^T \mathbf{u} - b + 1 \leq \eta_u, \\ & \mathbf{g}^T \mathbf{v} + b + 1 \leq \eta_v, \\ & (\xi, \mathbf{u}, \mathbf{v}, \eta_u, \eta_v) \geq \mathbf{0}. \end{aligned}$$

We are now ready to propose a novel combination of ENSVM and KSVM, and we will explain in the next section how the combined problem can be solved in our HIPAD framework. The main motivation behind this combination is to exploit domain knowledge to improve the feature selection, and hence, the generalization performance of HIPAD. To the best of our knowledge, this is the first method of this kind.

3.1 ADMM Phase

Our strategy for solving the elastic-net SVM with domain knowledge incorporation is still to apply the ADMM method. First, we combine problems (2) and (KSVM-P) and write the resulting optimization problem in an equivalent unconstrained form (by penalizing the violation of the inequality constraints through hinge losses in the objective function)

$$\text{(ENK-SVM)} \quad \min_{\mathbf{w}, b, \mathbf{u} \geq \mathbf{0}, \mathbf{v} \geq \mathbf{0}} F(\mathbf{w}, b, \mathbf{u}, \mathbf{v}),$$

where $F(\mathbf{w}, b, \mathbf{u}, \mathbf{v}) \equiv \frac{\lambda_2}{2} \|\mathbf{w}\|_2^2 + \lambda_1 \|\mathbf{w}\|_1 + \frac{1}{N} \sum_{i=1}^N (1 - y_i(x_i^T \mathbf{w} + b))_+ + \frac{\rho_1}{2} \|B^T \mathbf{u} + \mathbf{w}\|_2^2 + \rho_2 (\mathbf{d}^T \mathbf{u} - b + 1)_+ + \frac{\rho_3}{2} \|D^T \mathbf{v} - \mathbf{w}\|_2^2 + \rho_4 (\mathbf{g}^T \mathbf{v} + b + 1)_+$. We then apply variable-splitting to decouple the L_1 -norms and hinge losses and obtain the following equivalent constrained optimization problem:

$$\begin{aligned} \min_{\mathbf{w}, b, \mathbf{u}, \mathbf{v}, \mathbf{a}, \mathbf{c}, p, q} \quad & F(\mathbf{w}, b, \mathbf{u}, \mathbf{v}, \mathbf{a}, \mathbf{c}, p, q) \\ \text{s.t.} \quad & \mathbf{a} = \mathbf{e} - (\bar{X} \mathbf{w} + \mathbf{y}b), \quad \mathbf{c} = \mathbf{w}, \\ & q = \mathbf{d}^T \mathbf{u} - b + 1, \quad p = \mathbf{g}^T \mathbf{v} + b + 1, \\ & \mathbf{u} \geq \mathbf{0}, \quad \mathbf{v} \geq \mathbf{0}. \end{aligned} \quad (12)$$

with $F(\mathbf{w}, b, \mathbf{u}, \mathbf{v}, \mathbf{a}, \mathbf{c}, p, q) \equiv \frac{\lambda_2}{2} \|\mathbf{w}\|_2^2 + \lambda_1 \|\mathbf{c}\|_1 + \frac{1}{N} \mathbf{e}^T (\mathbf{a})_+ + \frac{\rho_1}{2} \|B^T \mathbf{u} + \mathbf{w}\|_2^2 + \rho_2 (q)_+ + \frac{\rho_3}{2} \|D^T \mathbf{v} - \mathbf{w}\|_2^2 + \rho_4 (p)_+$. As usual, we form the augmented Lagrangian \mathcal{L} of problem (12),

$$\begin{aligned} \mathcal{L} := & F(\mathbf{w}, b, \mathbf{u}, \mathbf{v}, \mathbf{a}, \mathbf{c}, p, q) + \gamma_1^T (\mathbf{e} - (\bar{X} \mathbf{w} + \mathbf{y}b) - \mathbf{a}) + \frac{\mu_1}{2} \|\mathbf{e} - (\bar{X} \mathbf{w} + \mathbf{y}b) - \mathbf{a}\|_2^2 \\ & + \gamma_2^T (\mathbf{w} - \mathbf{c}) + \frac{\mu_2}{2} \|\mathbf{w} - \mathbf{c}\|_2^2 + \gamma_3 (\mathbf{d}^T \mathbf{u} - b + 1 - q) + \frac{\mu_3}{2} \|\mathbf{d}^T \mathbf{u} - b + 1 - q\|_2^2 \\ & + \gamma_4 (\mathbf{g}^T \mathbf{v} + b + 1 - p) + \frac{\mu_4}{2} \|\mathbf{g}^T \mathbf{v} + b + 1 - p\|_2^2 \end{aligned}$$

and minimize \mathcal{L} with respect to $\mathbf{w}, b, \mathbf{c}, \mathbf{a}, p, q, \mathbf{u}, \mathbf{v}$ individually and in order. For the sake of readability, we do not penalize the non-negative constraints for \mathbf{u} and \mathbf{v} in the augmented Lagrangian.

Given $(\mathbf{a}^k, \mathbf{c}^k, p^k, q^k)$, solving for (\mathbf{w}, b) again involves solving a linear system

$$\begin{pmatrix} \kappa_1 I + \mu_1 X^T X & \mu_1 X^T \mathbf{e} \\ \mu_1 \mathbf{e}^T X & \mu_1 N + \kappa_2 \end{pmatrix} \begin{pmatrix} \mathbf{w}^{k+1} \\ b^{k+1} \end{pmatrix} = \begin{pmatrix} \mathbf{r}_w \\ \mathbf{r}_b \end{pmatrix}, \quad (13)$$

where $\kappa_1 = \lambda_2 + \mu_2 + \rho_1 + \rho_3, \kappa_2 = \mu_3 + \mu_4, \mathbf{r}_w = X^T Y \gamma_1^k + \mu_1 X^T Y (\mathbf{e} - \mathbf{a}^k) - \gamma_2^k + \mu_2 \mathbf{c}^k + \rho_3 D^T \mathbf{v}^k - \rho_1 B^T \mathbf{u}^k$ and $\mathbf{r}_b = \mathbf{e}^T Y \gamma_1^k + \mu_1 \mathbf{e}^T Y (\mathbf{e} - \mathbf{a}^k) + \gamma_3^k + \mu_3 (\mathbf{d}^T \mathbf{u}^k + 1 - q^k) - \gamma_4^k - \mu_4 (\mathbf{g}^T \mathbf{v}^k + 1 - p^k)$. Similar to solving the linear system in Algorithm A.1 ADMM-ENSVM, we can compute the solution to the above linear system through a few PCG iterations, taking advantage of the fact that the left-hand-side matrix is of low-rank.

To minimize the augmented Lagrangian with respect to \mathbf{u} , we need to solve a convex quadratic problem with non-negative constraints

$$\min_{\mathbf{u} \geq \mathbf{0}} \quad \frac{\rho_1}{2} \|B^T \mathbf{u} + \mathbf{w}^{k+1}\|_2^2 + \gamma_3^k \mathbf{d}^T \mathbf{u} + \frac{\mu_3}{2} \|\mathbf{d}^T \mathbf{u} - b^{k+1} + 1 - q^k\|_2^2. \quad (14)$$

Solving problem (14) efficiently is crucial for the efficiency of the overall algorithm. We describe a novel way to do so. Introducing a slack variable \mathbf{s} and transferring the non-negative constraint on \mathbf{u} to \mathbf{s} , we decompose the problem into two parts which are easy to solve. Specifically, we reformulate (14) as

$$\begin{aligned} \min_{\mathbf{u}, \mathbf{s} \geq \mathbf{0}} \quad & \frac{\rho_1}{2} \|B^T \mathbf{u} + \mathbf{w}^{k+1}\|_2^2 + \gamma_3^k \mathbf{d}^T \mathbf{u} + \frac{\mu_3}{2} \|\mathbf{d}^T \mathbf{u} - b^{k+1} + 1 - q^k\|_2^2 \\ \text{s.t.} \quad & \mathbf{u} - \mathbf{s} = \mathbf{0}. \end{aligned}$$

Penalizing the linear constraint $\mathbf{u} - \mathbf{s} = \mathbf{0}$ in the new augmented Lagrangian, the new subproblem with respect to (\mathbf{u}, \mathbf{s}) is

$$\begin{aligned} \min_{\mathbf{u}, \mathbf{s} \geq \mathbf{0}} \quad & \frac{\rho_1}{2} \|B^T \mathbf{u} + \mathbf{w}^{k+1}\|_2^2 + \gamma_3^k \mathbf{d}^T \mathbf{u} \\ & + \frac{\mu_3}{2} \|\mathbf{d}^T \mathbf{u} - b^{k+1} + 1 - q^k\|_2^2 + \gamma_5^T (\mathbf{s} - \mathbf{u}) + \frac{\mu_5}{2} \|\mathbf{u} - \mathbf{s}\|_2^2. \end{aligned} \quad (15)$$

Given an $\mathbf{s}^k \geq \mathbf{0}$, we can compute \mathbf{u}^{k+1} by solving a $k_1 \times k_1$ linear system

$$(\rho_1 B B^T + \mu_3 \mathbf{d} \mathbf{d}^T + \mu_5 I) \mathbf{u}^{k+1} = \mathbf{r}_u, \quad (16)$$

where $\mathbf{r}_u = -\rho_1 B \mathbf{w}^{k+1} + \mu_3 \mathbf{d} b^{k+1} + \mu_3 \mathbf{d} (q^k - 1) - \mathbf{d} \gamma_3^k + \gamma_5 + \mu_5 \mathbf{s}^k$. We assume that B has full row-rank. This is a reasonable assumption since otherwise there is at least one redundant domain knowledge constraint and we can simply remove it. The number of domain knowledge constraints (k_1 and k_2) are usually small, so the system (16) can be solved exactly and efficiently by Cholesky factorization.

Solving for \mathbf{s}^{k+1} corresponding to \mathbf{u}^{k+1} is easy, observing that problem (15) is separable in the elements of \mathbf{s} . For each element s_i , the optimal solution to the one-dimensional quadratic problem with a non-negative constraint on s_i is given by $\max(0, u_i - \frac{(\gamma_5)_i}{\mu_5})$. Writing in the vector form, $\mathbf{s}^{k+1} = \max(\mathbf{0}, \mathbf{u}^{k+1} - \frac{\gamma_5^k}{\mu_5})$. Similarly, we solve for \mathbf{v}^{k+1} by introducing a non-negative slack variable \mathbf{t} and solve the linear system

$$(\rho_3 D D^T + \mu_4 \mathbf{g} \mathbf{g}^T + \mu_6 I) \mathbf{v}^{k+1} = \mathbf{r}_v, \quad (17)$$

where $\mathbf{r}_v = \rho_3 D \mathbf{w}^{k+1} - \mu_4 \mathbf{g} b^{k+1} - \mathbf{g} \gamma_4^k - \mu_4 \mathbf{g} (1 - p^k) + \gamma_6 + \mu_6 \mathbf{t}^k$, and $\mathbf{t}^{k+1} = \max(0, \mathbf{v}^{k+1} - \frac{\gamma_6^k}{\mu_6})$.

Now given $(\mathbf{w}^{k+1}, b^{k+1}, \mathbf{u}^{k+1}, \mathbf{v}^{k+1})$, the solutions for \mathbf{a} and \mathbf{c} are exactly the same as in Lines 4 and 5 of Algorithm 2.1, i.e.

$$\begin{aligned} \mathbf{a}^{k+1} &= \mathcal{S}_{\frac{1}{N\mu_1}} \left(\mathbf{e} + \frac{\gamma_1^k}{\mu_1} - Y(X \mathbf{w}^{k+1} + b^{k+1} \mathbf{e}) \right), \\ \mathbf{c}^{k+1} &= \mathcal{T}_{\frac{\lambda_1}{\mu_2}} \left(\frac{\gamma_2^k}{\mu_2} + \mathbf{w}^{k+1} \right). \end{aligned}$$

The subproblem with respect to q is

$$\begin{aligned} \min_q \quad & \rho_2(q)_+ - \gamma_3^k q + \frac{\mu_3}{2} \|\mathbf{d}^T \mathbf{u}^k - b^{k+1} + 1 - q\|_2^2 \equiv \\ & \rho_2(q)_+ + \frac{\mu_3}{2} \left\| q - \left(\mathbf{d}^T \mathbf{u}^k - b^{k+1} + 1 + \frac{\gamma_3^k}{\mu_3} \right) \right\|_2^2. \end{aligned} \quad (18)$$

The solution is given by a (one-dimensional) proximal operator associated with the hinge loss

$$q^{k+1} = \mathcal{S}_{\frac{\rho_2}{\mu_3}} \left(\mathbf{d}^T \mathbf{u}^k - b^{k+1} + 1 + \frac{\gamma_3^k}{\mu_3} \right). \quad (19)$$

Similarly, the subproblem with respect to p is

$$\min_p \quad \rho_4(p)_+ - \gamma_4^k p + \frac{\mu_4}{2} \|\mathbf{g}^T \mathbf{v}^k + b^{k+1} + 1 - p\|_2^2,$$

and the solution is given by

$$p^{k+1} = \mathcal{S}_{\frac{\rho_4}{\mu_4}} \left(\mathbf{g}^T \mathbf{v}^k + b^{k+1} + 1 + \frac{\gamma_4^k}{\mu_4} \right). \quad (20)$$

Due to lack of space in the paper, we summarize the detailed solution steps in the Appendix (see Algorithm A.2 ADMM-ENK)

Although there appears to be ten additional parameters (six ρ 's and four μ 's) in the ADMM method for ENK-SVM, we can usually set the ρ 's to the same value and do the same for the μ 's. Hence, in practice, there is only one additional parameter to tune, and our computational experience in Section 4.2 is that the algorithm is fairly insensitive to the μ 's and ρ 's.

3.2 IPM Phase

The second phase for solving the knowledge-based SVM problem defined by (KSVM-P) follows the same steps as that described in section 2.2. Note that in the knowledge-based case we have decided to solve the primal problem. This decision was based on extensive numerical experiments with both the primal and dual formulation which revealed that the primal formulation is more efficient.

We found in our experiments that by introducing slack variables and transforming the above problem into a linearly equality-constrained QP, Phase 2 of HIPAD usually requires less time to solve.

3.3 HIPAD with domain knowledge incorporation

We formally state the new two-phase algorithm for the elastic-net KSVM in Algorithm 3.1.

Algorithm 3.1 HIPAD-ENK

- 1: Given $\mathbf{w}^0, b^0, \mathbf{a}^0, \mathbf{c}^0, \mathbf{u}^0, \mathbf{v}^0, p^0, q^0, \mathbf{s}^0 \geq \mathbf{0}, \mathbf{t}^0 \geq \mathbf{0}$.
 - 2: **PHASE 1: ADMM for ENK-SVM**
 - 3: $(\mathbf{w}, b, \mathbf{u}, \mathbf{v}) \leftarrow \text{ADMM-ENK}(\mathbf{w}^0, b^0, \mathbf{a}^0, \mathbf{c}^0, \mathbf{u}^0, \mathbf{v}^0, p^0, q^0, \mathbf{s}^0, \mathbf{t}^0)$
 - 4: **PHASE 2: IPM for KSVM**
 - 5: $\tilde{\mathbf{w}} \leftarrow$ non-zero components of \mathbf{w}
 - 6: $(\tilde{X}, \tilde{Y}) \leftarrow$ sub-matrices of (X, Y) corresponding to the support of \mathbf{w}
 - 7: $\eta_u^0 \leftarrow \mathbf{d}^T \mathbf{u} - b + 1$
 - 8: $\eta_v^0 \leftarrow \mathbf{g}^T \mathbf{v} + b + 1$
 - 9: $(\mathbf{w}, b) \leftarrow \text{SVM-IPM}(\tilde{X}, \tilde{Y}, \tilde{\mathbf{w}}, b, \mathbf{u}, \mathbf{v}, \eta_u^0, \eta_v^0)$
 - 10: **return** (\mathbf{w}, b)
-

4 Numerical results

We present our numerical experience with the two main algorithms proposed in this paper: HIPAD and its knowledge-based version HIPAD-ENK. We compare their performance with their non-hybrid counterparts, i.e., ADMM-ENSVM and ADMM-ENK, which use ADMM to solve the original SVM problem. The transition condition at the end of Phase 1 is specified in (7), with $\epsilon_{tol} = 10^{-2}$. The stopping criteria for ADMM are as follows: $\frac{|F^{k+1} - F^k|}{\max\{1, |F^k|\}} \leq \epsilon_1$, $\|\mathbf{a} - (\mathbf{e} - \bar{X}\mathbf{w} - \mathbf{y}b)\|_2 \leq \epsilon_1$, $\|\mathbf{c} - \mathbf{w}\|_2 \leq \epsilon_1$ and $\frac{\|\mathbf{w}^{k+1} - \mathbf{w}\|_2}{\|\mathbf{w}^k\|_2} \leq \epsilon_2$, with $\epsilon_1 = 10^{-5}$, and $\epsilon_2 = 10^{-3}$.

4.1 HIPAD vs ADMM

To demonstrate the practical effectiveness of HIPAD, we tested the algorithm on nine real data sets which are publicly available. **rcv1** [11] is a collection of manually categorized news wires from Reuters. The original multiple labels have been consolidated into two classes for binary classification. **real-sim** contains UseNet articles from four discussion groups, for simulated auto racing, simulated aviation, real autos, and real aviation. Both **rcv1** and **real-sim** have large feature dimensions but are highly sparse. The rest of the seven data sets are all dense data. **rcv1** and **real-sim** are subsets of the original data sets, where we randomly sampled 500 training instances and 1,000 test instances. **gisette** is a handwriting digit recognition problem from NIPS 2003 Feature Selection Challenge, and we also sub-sampled 500 instances for training. (For testing, we used the original test set of 1,000 instances.) **duke**, **leukemia**, and **colon-cancer** are data sets of gene expression profiles for breast cancer, leukemia, and colon cancer respectively. **fMRIa**, **fMRIb**, and **fMRIc** are functional MRI (fMRI) data of brain activities when the subjects are presented with pictures and text paragraphs. The data was compiled and made available by Tom Mitchell's neuroinformatics research group ⁴. Except the three fMRI data sets, all the other data sets and their references are available at the LIBSVM website ⁵.

The parameters of HIPAD, ADMM-ENSVM, and LIBSVM were selected through cross validation on the training data. We summarize the experiment results in Table 1. Clearly, HIPAD produced the best overall predication performance. In order to test the significance of the difference, we used the test statistic in [8] based on Friedman's χ_F^2 , and the results are significant at $\alpha = 0.1$. In terms of CPU-time, HIPAD consistently outperformed ADMM-ENSVM by several times on dense data. The feature support sizes selected by HIPAD were also very competitive or even better than the ones selected by ADMM-ENSVM. In most cases, HIPAD was able to shrink the feature space to below 10% of the original size.

⁴ <http://www.cs.cmu.edu/~tom/fmri.html>

⁵ <http://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets>

Data set	HIPAD			ADMM-ENSVM			LIBSVM
	Accuracy (%)	Support size	CPU (s)	Accuracy (%)	Support size	CPU (s)	accuracy (%)
rcv1	86.9	2,037	1.18	86.8	7,002	1.10	86.1
real-sim	94.0	2,334	0.79	93.9	2,307	0.31	93.4
gisetete	94.7	498	8.96	63.1	493	45.87	93.4
duke	90	168	1.56	90	150	5.52	80
leukemia	85.3	393	1.70	82.4	717	6.35	82.4
colon-cancer	84.4	195	0.45	84.4	195	1.34	84.4
fMRIa	90	157	0.25	90	137	2.17	60
fMRIb	90	45	0.23	90	680	0.75	90
fMRIc	90	321	0.14	90	659	1.58	90

Table 1. Experiment results of HIPAD and ADMM-ENSVM on real data. The best prediction accuracy for each data set is highlighted in bold.

4.2 Simulation for Knowledge Incorporation

We generated synthetic data to simulate the example presented at the beginning of Section 3 in the high dimensional feature space. Specifically, four groups of multi-variate Gaussians K_1, \dots, K_4 were sampled from $\mathcal{N}(\mu_1^+, \Sigma_1), \dots, \mathcal{N}(\mu_4^+, \Sigma_4)$ and $\mathcal{N}(\mu_1^-, \Sigma_1), \dots, \mathcal{N}(\mu_4^-, \Sigma_4)$ for four disjoint blocks of feature values $(\mathbf{x}_{K_1}, \dots, \mathbf{x}_{K_4})$. For positive class samples, $\mu_1^+ = 2, \mu_2^+ = 0.5, \mu_3^+ = -0.2, \mu_4^+ = -1$; for negative class samples, $\mu_1^- = -2, \mu_2^- = -0.5, \mu_3^- = 0.2, \mu_4^- = 1$. All the covariance matrices have 1 on the diagonal and 0.8 everywhere else. The training samples contain blocks K_2 and K_3 , while all four blocks are present in the test samples. A random fraction (5%-10%) of the remaining entries in all the samples are generated from the standard Gaussian distribution.

The training samples are apparently hard to separate because the values of blocks K_2 and K_3 for the two classes are close to each other. However, blocks K_1 and K_4 in the test samples are well-separated. Hence, if we are given information about these two blocks as general knowledge for the entire population, we could expect the resulting SVM classifier to perform much better on the test data. Since we know the mean values of the distributions from which the entries in K_1 and K_4 are generated, we can supply the following information about the relationship between the block sample means and class membership to the KSVM: $\frac{1}{L_1} \sum_{i \in K_1} x_i \geq 4 \Rightarrow \mathbf{x} \in A^+$, and $\frac{1}{L_4} \sum_{i \in K_4} x_i \geq 3 \Rightarrow \mathbf{x} \in A^-$ where L_j is the length of the K_j , $j = 1, \dots, 4$, A^+ and A^- represent the positive and negative classes, and the lowercase x_i denotes the i -th entry of the sample \mathbf{x} . Translating into the notation of (KSVM-P), we have

$$B = \left(\mathbf{0} \underbrace{-\frac{\mathbf{e}^T}{L_1}}_{K_1} \mathbf{0} \mathbf{0} \mathbf{0} \mathbf{0} \right), d = -4, \text{ and } D = \left(\mathbf{0} \mathbf{0} \mathbf{0} \mathbf{0} \underbrace{-\frac{\mathbf{e}^T}{L_4}}_{K_4} \mathbf{0} \right), g = -3.$$

The information given here is not precise, in that we are confident that a sample should belong to the positive (or negative) class only when the corresponding block sample mean well exceeds the distribution mean. This is consistent with

real-world situations, where the domain or expert knowledge tends to be conservative and often does not come in exact form.

We simulated two sets of synthetic data for ENK-SVM as described above, with $(N_{train} = 200, N_{test} = 400, m_{train} = 10,000)$ for **ksvm-s-10k** and $(N_{train} = 500, N_{test} = 1,000, m_{test} = 50,000)$ for **ksvm-s-50k**. The number of features in each of the four blocks (K_1, K_2, K_3, K_4) is 50 for both data sets. Clearly, HIPAD-ENK is very effective in terms of speed, feature selection, and prediction accuracy on these two data sets. Even though the features in blocks K_1 and K_4 are not discriminating in the training data, HIPAD-ENK was still able to identify all the 200 features in the four blocks correctly and exactly. This is precisely what we want to achieve as we explained at the beginning of Section 3. The expert knowledge not only helps rectify the separating hyperplane so that it generalizes better on the entire population, but also makes the training algorithm realize the significance of the features in blocks K_1 and K_4 .

Data set	HIPAD-ENK			ADMM-ENK			LIBSVM accuracy (%)
	Accuracy (%)	Support size	CPU (s)	Accuracy (%)	Support size	CPU (s)	
ksvm-s-10k	99	200	1.99	97.25	200	3.43	86.1
ksvm-s-50k	98.8	200	8.37	96.4	198	20.89	74.8

Table 2. Experiment results of HIPAD-ENK and ADMM-ENK on synthetic data. The best prediction accuracy for each data set is highlighted in bold.

5 Conclusion

We have proposed a two-phase hybrid optimization framework for solving the ENSVM, in which the first phase is solved by ADMM, followed by IPM in the second phase. In addition, we have proposed a knowledge-based extension of the ENSVM which can be solved by the same hybrid framework. Through a set of experiments, we demonstrated that our method has significant advantage over the existing method in terms of computation time and the resulting prediction accuracy. The algorithmic framework introduced in this paper is general enough and potentially applicable to other regularization-based classification or regression problems.

References

1. G. J. Andersen, A. C. Meszaros, and X. Xu. Implementation of interior point methods for large scale linear programming. In *In "Interior point methods in mathematical programming"*, T. Terlaky, editor, pages 189–252. Kluwer Academic Publishers, 1996.
2. S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.

3. P. Combettes and J. Pesquet. Proximal splitting methods in signal processing. *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*, pages 185–212, 2011.
4. J. Eckstein and D. Bertsekas. On the douglas-rachford splitting method and the proximal point algorithm for maximal monotone operators. *Mathematical Programming*, 55(1):293–318, 1992.
5. G. Fung, O. Mangasarian, and J. Shavlik. Knowledge-based support vector machine classifiers. *Advances in Neural Information Processing Systems*, pages 537–544, 2003.
6. D. Gabay and B. Mercier. A dual algorithm for the solution of nonlinear variational problems via finite element approximation. *Computers & Mathematics with Applications*, 2(1):17–40, 1976.
7. E. Gertz and S. Wright. Object-oriented software for quadratic programming. *ACM Transactions on Mathematical Software (TOMS)*, 29(1):58–81, 2003.
8. R. L. Iman and J. M. Davenport. Approximations of the critical region of the f-biostatistic. *Communications in Statistics-Theory and Methods*, 9(6):571–595, 1980.
9. K. Koh, S. Kim, and S. Boyd. An interior-point method for large-scale l_1 -regularized logistic regression. *Journal of Machine Learning Research*, 8(8):1519–1555, 2007.
10. F. Lauer and G. Bloch. Incorporating prior knowledge in support vector machines for classification: A review. *Neurocomputing*, 71(7-9):1578–1594, 2008.
11. D. Lewis, Y. Yang, T. Rose, and F. Li. Rcv1: A new benchmark collection for text categorization research. *The Journal of Machine Learning Research*, 5:361–397, 2004.
12. P. P. M. and P. E. Hansen. *Data Mining and Mathematical Programming*. American Mathematical Society, 2008.
13. S. Mizuno. Polynomiality of infeasible interior point algorithms for linear programming. *Mathematical Programming*, 67:109–119, 1994.
14. R. Rockafellar. The multiplier method of hestenes and powell applied to convex programming. *Journal of Optimization Theory and Applications*, 12(6):555–562, 1973.
15. S. Ryali, K. Supekar, D. Abrams, and V. Menon. Sparse logistic regression for whole-brain classification of fmri data. *NeuroImage*, 51(2):752–764, 2010.
16. J. Shi, W. Yin, S. Osher, and P. Sajda. A fast hybrid algorithm for large-scale l_1 -regularized logistic regression. *The Journal of Machine Learning Research*, 11:713–741, 2010.
17. N. S. Sra, S. and S. Wright. *Optimization for Machine Learning*. MIT Press, 2011.
18. R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1):267–288, 1996.
19. V. Vapnik. *The nature of statistical learning theory*. Springer Verlag, 2000.
20. L. Wang, J. Zhu, and H. Zou. The doubly regularized support vector machine. *Statistica Sinica*, 16(2):589–615, 2006.
21. G.-B. Ye, Y. Chen, and X. Xie. Efficient variable selection in support vector machines via the alternating direction method of multipliers. In *International Conference on Artificial Intelligence and Statistics*, pages 832–840, 2011.
22. H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005.

A Appendix

Algorithm A.1 ADMM-ENSVM

- 1: Given $\mathbf{w}^0, b^0, \mathbf{a}^0, \mathbf{c}^0, \gamma_1^0$, and γ_2^0 .
 - 2: **for** $k = 0, 1, \dots, K - 1$ **do**
 - 3: $(\mathbf{w}^{k+1}, b^{k+1}) \leftarrow$ PCG solution of:
$$\begin{pmatrix} (\lambda_2 + \mu_2)I + \mu_1 X^T X & \mu_1 X^T \mathbf{e} \\ \mu_1 \mathbf{e}^T X & \mu_1 N \end{pmatrix} \begin{pmatrix} \mathbf{w}^{k+1} \\ b^{k+1} \end{pmatrix} = \begin{pmatrix} X^T Y \gamma_1^k - \mu_1 X^T Y (\mathbf{a}^k - \mathbf{e}) - \gamma_2^k + \mu_2 \mathbf{c}^k \\ \mathbf{e}^T Y \gamma_1^k - \mu_1 \mathbf{e}^T Y (\mathbf{a}^k - \mathbf{e}) \end{pmatrix}$$
 - 4: $\mathbf{a}^{k+1} \leftarrow \mathcal{S}_{\frac{1}{N\mu_1}} \left(\mathbf{e} + \frac{\gamma_1^k}{\mu_1} - Y(X\mathbf{w}^{k+1} + b^{k+1}\mathbf{e}) \right)$
 - 5: $\mathbf{c}^{k+1} \leftarrow \mathcal{T}_{\frac{\lambda_1}{\mu_2}} \left(\frac{\gamma_2^k}{\mu_2} + \mathbf{w}^{k+1} \right)$
 - 6: $\gamma_1^{k+1} \leftarrow \gamma_1^k + \mu_1 (\mathbf{e} - Y(X\mathbf{w}^{k+1} + b^{k+1}\mathbf{e}) - \mathbf{a}^{k+1})$
 - 7: $\gamma_2^{k+1} \leftarrow \gamma_2^k + \mu_2 (\mathbf{w}^{k+1} - \mathbf{c}^{k+1})$
 - 8: **end for**
 - 9: **return** (\mathbf{w}^K, b^K)
-

Algorithm A.2 ADMM-ENK

- 1: Given $\mathbf{w}^0, b^0, \mathbf{a}^0, \mathbf{c}^0, \mathbf{u}^0, \mathbf{v}^0, p^0, q^0, \mathbf{s}^0 \geq \mathbf{0}, \mathbf{t}^0 \geq \mathbf{0}, \gamma_i^0$, and the parameters $\lambda_1, \lambda_2, \rho_i, i = 1, \dots, 6$.
 - 2: **for** $k = 0, 1, \dots, K - 1$ **do**
 - 3: $(\mathbf{w}^{k+1}, b^{k+1}) \leftarrow$ PCG solution of the structured linear system (13)
 - 4: $\mathbf{u}^{k+1} \leftarrow$ the solution of the linear system (16); $\mathbf{s}^{k+1} \leftarrow \max(\mathbf{0}, \mathbf{u}^{k+1} - \frac{\gamma_5^k}{\mu_5})$
 - 5: $\mathbf{v}^{k+1} \leftarrow$ the solution of the linear system (17); $\mathbf{t}^{k+1} \leftarrow \max(\mathbf{0}, \mathbf{v}^{k+1} - \frac{\gamma_6^k}{\mu_6})$
 - 6: $\mathbf{a}^{k+1} \leftarrow \mathcal{S}_{\frac{1}{N\mu_1}} \left(\mathbf{e} + \frac{\gamma_1^k}{\mu_1} - Y(X\mathbf{w}^{k+1} + b^{k+1}\mathbf{e}) \right)$; $\mathbf{c}^{k+1} \leftarrow \mathcal{T}_{\frac{\lambda_1}{\mu_2}} \left(\frac{\gamma_2^k}{\mu_2} + \mathbf{w}^{k+1} \right)$
 - 7: $q^{k+1} \leftarrow \mathcal{S}_{\frac{\rho_2}{\mu_3}} \left(\mathbf{d}^T \mathbf{u}^k - b^{k+1} + 1 + \frac{\gamma_3^k}{\mu_3} \right)$; $p^{k+1} \leftarrow \mathcal{S}_{\frac{\rho_4}{\mu_4}} \left(\mathbf{g}^T \mathbf{v}^k + b^{k+1} + 1 + \frac{\gamma_4^k}{\mu_4} \right)$
 - 8: $\gamma_1^{k+1} \leftarrow \gamma_1^k + \mu_1 (\mathbf{e} - Y(X\mathbf{w}^{k+1} + b^{k+1}\mathbf{e}) - \mathbf{a}^{k+1})$; $\gamma_2^{k+1} \leftarrow \gamma_2^k + \mu_2 (\mathbf{w}^{k+1} - \mathbf{c}^{k+1})$
 - 9: $\gamma_3^{k+1} \leftarrow \gamma_3^k + \mu_3 (\mathbf{d}^T \mathbf{u}^{k+1} - b^{k+1} + 1 - q^{k+1})$; $\gamma_4^{k+1} \leftarrow \gamma_4^k + \mu_4 (\mathbf{g}^T \mathbf{v}^{k+1} + b^{k+1} + 1 - p^{k+1})$
 - 10: $\gamma_5^{k+1} \leftarrow \gamma_5^k + \mu_5 (\mathbf{s}^{k+1} - \mathbf{u}^{k+1})$; $\gamma_6^{k+1} \leftarrow \gamma_6^k + \mu_6 (\mathbf{t}^{k+1} - \mathbf{v}^{k+1})$
 - 11: **end for**
 - 12: **return** (\mathbf{w}^K, b^K)
-




Fig. 1. Illustration of the early convergence (in approximately 50 iterations) of the feature support for ADMM.