

Scaling Up Sparse Support Vector Machines by Simultaneous Feature and Sample Reduction

Bin Hong*

HONGBINZJU@GMAIL.COM

*State Key Lab of CAD&CG, College of Computer Science, Zhejiang University
Hangzhou, 310058, China*

Weizhong Zhang*

ZHANGWEIZHONGZJU@GMAIL.COM

*Tencent AI Lab, Shenzhen, 518000, China
State Key Lab of CAD&CG, College of Computer Science, Zhejiang University
Hangzhou, 310058, China*

Wei Liu

WL2223@COLUMBIA.EDU

Tencent AI Lab, Shenzhen, 518000, China

Jieping Ye

JPYE@UMICH.EDU

*Department of Computational Medicine and Bioinformatics, University of Michigan
Ann Arbor, MI, 48104-2218, USA.*

Deng Cai†

DENGCAI@CAD.ZJU.EDU.CN

Xiaofei He

XIAOFEIHE@CAD.ZJU.EDU.CN

*State Key Lab of CAD&CG, College of Computer Science, Zhejiang University
Hangzhou, 310058, China*

Jie Wang

JIEWANGUSTC@GMAIL.COM

*Department of Electronic Engineering and Information Science
University of Science and Technology of China, Hefei, 230026, China*

Editor: Sanjiv Kumar

Abstract

Sparse support vector machine (SVM) is a popular classification technique that can simultaneously learn a small set of the most interpretable features and identify the support vectors. It has achieved great successes in many real-world applications. However, for large-scale problems involving a huge number of samples and ultra-high dimensional features, solving sparse SVMs remains challenging. By noting that sparse SVMs induce sparsities in both feature and sample spaces, we propose a novel approach, which is based on accurate estimations of the primal and dual optima of sparse SVMs, to simultaneously identify the inactive features and samples that are guaranteed to be irrelevant to the outputs. Thus, we can remove the identified inactive samples and features from the training phase, leading to substantial savings in the computational cost without sacrificing the accuracy. Moreover, we show that our method can be extended to multi-class sparse support vector machines. To the best of our knowledge, the proposed method is the *first static* feature and sample reduction method for sparse SVMs and multi-class sparse SVMs. Experiments on both synthetic and real data sets demonstrate that our approach significantly outperforms state-of-the-art methods and the speedup gained by our approach can be orders of magnitude.

Keywords: screening, SVM, dual problem, optimization, classification

*. The first two authors contribute equally. †. Corresponding author.

1. Introduction

Sparse support vector machine (SVM) (Bi et al., 2003; Wang et al., 2006) is a powerful technique that can simultaneously perform classification by margin maximization and variable selection via ℓ_1 -norm penalty. The last few years have witnessed many successful applications of sparse SVMs, such as text mining (Joachims, 1998; Yoshikawa et al., 2014), bioinformatics (Narasimhan and Agarwal, 2013), and image processing (Mohr and Obermayer, 2004; Kotsia and Pitas, 2007). Many algorithms (Hastie et al., 2004; Fan et al., 2008; Catanzaro et al., 2008; Hsieh et al., 2008; Shalev-Shwartz et al., 2011) have been proposed to efficiently solve sparse SVM problems. However, the applications of sparse SVMs to large-scale learning problems, which involve a huge number of samples and extremely high-dimensional features, still remain challenging.

An emerging technique, called *screening* (El Ghaoui et al., 2012), has been shown to be promising in accelerating the training processes of large-scale sparse learning models. The essential idea of screening is to quickly identify the zero coefficients in the sparse solutions without solving any optimization problems such that the corresponding features or samples—that are called *inactive* features or samples—can be removed from the training phase. Then, we only need to perform optimization on the reduced data sets instead of the full data sets, leading to a substantial saving in the computational cost. Here, we need to emphasize that screening differs greatly from feature selection, although they look similar at the first glance. To be precise, screening is devoted to accelerating the training processes of many sparse models including Lasso, sparse SVMs, etc., while feature selection is the goal of these models. In the past few years, many screening methods are proposed for a large set of sparse learning techniques, such as Lasso (Tibshirani et al., 2012; Xiang and Ramadge, 2012; Wang et al., 2015), group Lasso (Ndiaye et al., 2016), ℓ_1 -regularized logistic regression (Wang et al., 2014b), and SVMs (Ogawa et al., 2013). Most of the existing methods are in the same framework, i.e., estimating the optimum of the dual (resp. primal) problem and then developing the screening rules based on the estimations to infer which components of the primal (resp. dual) optimum are zero from the KKT conditions. The main differences among them are the techniques they use to develop their optima estimations and rules and the different sparse models they focus on. For example, SAFE (El Ghaoui et al., 2012) estimates the dual optimum by calculating the optimal value’s lower bound of the dual problem. The Lasso screening method (Wang et al., 2015) estimates the optimum based on the non-expensiveness (Bauschke et al., 2011) of the projection operator by noting that its dual problem boils down to finding the projection of a given point on a convex set. Moreover, a screening method is called *static* if it triggers its screening rules before training, and *dynamic* if during the training process. Empirical studies indicate that screening methods can lead to orders of magnitude of speedup in computation time.

However, most existing screening methods study either feature screening or sample screening individually and their applications have very different scenarios. Specifically, to achieve better performance (say, in terms of speedup), we favor feature screening methods when the number of features p is much larger than the sample size n , while sample screening methods are preferable when $n \gg p$. Note that there is another class of sparse learning techniques, like sparse SVMs, which induce sparsities in both feature and sample spaces. All these screening methods tend to be helpless in accelerating the training process of these

models with large n and p . We cannot address this problem by simply combining the existing feature and sample screening methods either. The reason is that they could mistakenly discard relevant data as they are specifically designed for different sparse models. Recently, Shibagaki et al. (2016) considered this issue and proposed a method to simultaneously identify the inactive features and samples in a dynamic manner, and they trigger their testing rule when there is a sufficient decrease in the duality gap during the training process. Thus, the method in Shibagaki et al. 2016 can discard more and more inactive features and samples as the training proceeds and one has small-scale problems to solve in the late stage of the training process. Nevertheless, the overall speedup is limited as the problems’ size can be large in the early stage of the training process. To be specific, the method in Shibagaki et al. 2016 depends heavily on the duality gap during the training process. The duality gap in the early stage can always be large, which makes the dual and primal estimations inaccurate and finally results in ineffective screening rules. Hence, it is essentially solving a large problem in the early stage. This will be verified in the experimental comparisons in Section 6 and similar results can also be found in the recent work (Massias et al., 2018), which merely focuses on Lasso.

In this paper, to address the limitations in the dynamic screening method, we propose a novel screening method that can **S**imultaneously identify **I**nactive **F**eatures and **S**amples (SIFS) for sparse SVMs in a static manner, and we only need to perform SIFS once *before* (instead of during) training. Thus, we only need to run the training algorithm on small-scale problems. The major technical challenge in developing SIFS is that we need to accurately estimate the primal and dual optima. The more accurate the estimations are, the more effective SIFS is in detecting inactive features and samples. Thus, our major technical contribution is a novel framework, which is based on the strong convexity of the primal and dual problems of sparse SVMs for deriving accurate estimations of the primal and dual optima (see Section 3). Another appealing feature of SIFS is the so-called *synergistic effect* (Shibagaki et al., 2016). Specifically, the proposed SIFS consists of two parts, i.e., **I**nactive **F**eature **S**creening (IFS) and **I**nactive **S**ample **S**creening (ISS). We show that discarding inactive features (resp. samples) identified by IFS (resp. ISS) leads to a more accurate estimation of the primal (resp. dual) optimum, which in turn dramatically enhances the capability of ISS (resp. IFS) in detecting inactive samples (resp. features). Thus, SIFS applies IFS and ISS in an alternating manner until no more inactive features and samples can be identified, leading to much better performance in scaling up large-scale problems than the application of ISS or IFS individually. Moreover, SIFS (see Section 4) is safe in the sense that the detected features and samples are guaranteed to be absent from the sparse representations. To the best of our knowledge, SIFS is the first static screening rule for sparse SVMs, which is able to simultaneously detect inactive features and samples. Experiments (see Section 6) on both synthetic and real data sets demonstrate that SIFS significantly outperforms the state-of-the-art (Shibagaki et al., 2016) in improving the training efficiency of sparse SVMs and the speedup can be orders of magnitude.

To demonstrate the flexibility of our proposed method SIFS, we extend its idea to multi-class sparse support vector machine (see Section 5), which also induces sparsities in both feature and sample spaces. Although multi-class sparse SVM has a very different structure and is more complex than sparse SVM, we will see that its dual problem has some similar properties (especially the strong convexity) with those of sparse SVM. Recall that

SIFS we developed for sparse SVMs is mainly based on the strong convexity of the primal and dual problems. Therefore, the idea of SIFS is also applicable for multi-class sparse SVMs. Experimental results show that the speedup gained by SIFS in multi-class sparse SVMs can also be orders of magnitude.

For the convenience of presentation, we postpone the detailed proofs of all the theorems in this paper to the appendix. At last, we should point out that this journal paper is an extension of our own previous work (Zhang et al., 2017) published at the International Conference on Machine Learning (ICML) 2017.

Notations: Let $\|\cdot\|_1$, $\|\cdot\|$, and $\|\cdot\|_\infty$ be ℓ_1 , ℓ_2 , and ℓ_∞ norms, respectively. We denote the inner product between vectors \mathbf{x} and \mathbf{y} by $\langle \mathbf{x}, \mathbf{y} \rangle$, and the i -th component of \mathbf{x} by $[\mathbf{x}]_i$. Let $[p] = \{1, 2, \dots, p\}$ for a positive integer p . Given a subset $\mathcal{J} := \{j_1, \dots, j_k\}$ of $[p]$, let $|\mathcal{J}|$ be the cardinality of \mathcal{J} . For a vector \mathbf{x} , let $[\mathbf{x}]_{\mathcal{J}} = ([\mathbf{x}]_{j_1}, \dots, [\mathbf{x}]_{j_k})^\top$. For a matrix \mathbf{X} , let $[\mathbf{X}]_{\mathcal{J}} = (\mathbf{x}_{j_1}, \dots, \mathbf{x}_{j_k})$ and $_{\mathcal{J}}[\mathbf{X}] = ((\mathbf{x}^{j_1})^\top, \dots, (\mathbf{x}^{j_k})^\top)^\top$, where \mathbf{x}^i and \mathbf{x}_j are the i^{th} row and j^{th} column of \mathbf{X} , respectively. For a scalar t , we denote $\max\{0, t\}$ by $[t]_+$. Let $\mathbf{e}_k \in \mathbb{R}^K$ be the index vector, that is, $[\mathbf{e}_k]_i = 1$ if $i = k$, otherwise $[\mathbf{e}_k]_i = 0$. At last, we denote the set of nonnegative real numbers as \mathbb{R}_+ .

2. Basics and Motivations

In this section, we briefly review some basics of sparse SVMs and then motivate SIFS via the KKT conditions. Specifically, we focus on an ℓ_1 -regularized SVM with a smoothed hinge loss, which takes the form of

$$\min_{\mathbf{w} \in \mathbb{R}^p} P(\mathbf{w}; \alpha, \beta) = \frac{1}{n} \sum_{i=1}^n \ell(1 - \langle \bar{\mathbf{x}}_i, \mathbf{w} \rangle) + \frac{\alpha}{2} \|\mathbf{w}\|^2 + \beta \|\mathbf{w}\|_1, \quad (\text{P}^*)$$

where $\mathbf{w} \in \mathbb{R}^p$ is the parameter vector to be estimated, $\{\mathbf{x}_i, y_i\}_{i=1}^n$ is the training set, $\mathbf{x}_i \in \mathbb{R}^p$, $y_i \in \{-1, +1\}$, $\bar{\mathbf{x}}_i = y_i \mathbf{x}_i$, α and β are two positive parameters, and $\ell(\cdot) : \mathbb{R} \rightarrow \mathbb{R}_+$ is the smoothed hinge loss, i.e.,

$$\ell(t) = \begin{cases} 0, & \text{if } t < 0, \\ \frac{t^2}{2\gamma}, & \text{if } 0 \leq t \leq \gamma, \\ t - \frac{\gamma}{2}, & \text{if } t > \gamma, \end{cases} \quad (1)$$

where $\gamma \in (0, 1)$.

Remark 1 *We use the smoothed hinge loss instead of the vanilla one in order to make the objective of the Lagrangian dual problem of (P*) strongly convex, which is needed in developing our accurate optima estimations. We should point out that the smoothed hinge loss is a pretty good approximation to the vanilla one with strong theoretical guarantees. See Section 5 in Shalev-Shwartz and Zhang 2016 for the details.*

We present the Lagrangian dual problem of problem (P*) and the KKT conditions in the following theorem, which play a fundamentally important role in developing our screening rules. We will not provide its proof in the appendix since it follows from Fenchel Duality (see Corollary 31.2.1 in Rockafellar, 1970) and a similar result can be found in Shibagaki et al. 2016.

Theorem 2 (Rockafellar, 1970) *Let $\bar{\mathbf{X}} = (\bar{\mathbf{x}}_1, \bar{\mathbf{x}}_2, \dots, \bar{\mathbf{x}}_n)$ and $\mathcal{S}_\beta(\cdot)$ be the soft-thresholding operator (Hastie et al., 2015), i.e., $[\mathcal{S}_\beta(\mathbf{u})]_i = \text{sign}([\mathbf{u}]_i)(|[\mathbf{u}]_i| - \beta)_+$. Then, for problem (\mathbf{P}^*) , the followings hold:*

(i) *the dual problem of (\mathbf{P}^*) is*

$$\min_{\theta \in [0,1]^n} D(\theta; \alpha, \beta) = \frac{1}{2\alpha} \left\| \mathcal{S}_\beta \left(\frac{1}{n} \bar{\mathbf{X}} \theta \right) \right\|^2 + \frac{\gamma}{2n} \|\theta\|^2 - \frac{1}{n} \langle \mathbf{1}, \theta \rangle, \quad (\mathbf{D}^*)$$

where $\mathbf{1} \in \mathbb{R}^n$ is a vector with all components equal to 1;

(ii) *denote the optima of (\mathbf{P}^*) and (\mathbf{D}^*) by $\mathbf{w}^*(\alpha, \beta)$ and $\theta^*(\alpha, \beta)$, respectively, then,*

$$\mathbf{w}^*(\alpha, \beta) = \frac{1}{\alpha} \mathcal{S}_\beta \left(\frac{1}{n} \bar{\mathbf{X}} \theta^*(\alpha, \beta) \right), \quad (\text{KKT-1})$$

$$[\theta^*(\alpha, \beta)]_i = \begin{cases} 0, & \text{if } 1 - \langle \bar{\mathbf{x}}_i, \mathbf{w}^*(\alpha, \beta) \rangle < 0; \\ 1, & \text{if } 1 - \langle \bar{\mathbf{x}}_i, \mathbf{w}^*(\alpha, \beta) \rangle > \gamma; \\ \frac{1}{\gamma}(1 - \langle \bar{\mathbf{x}}_i, \mathbf{w}^*(\alpha, \beta) \rangle), & \text{otherwise.} \end{cases} \quad (\text{KKT-2})$$

According to KKT-1 and KKT-2, we define 4 index sets:

$$\begin{aligned} \mathcal{F} &= \left\{ j \in [p] : \frac{1}{n} |[\bar{\mathbf{X}} \theta^*(\alpha, \beta)]_j| \leq \beta \right\}, \\ \mathcal{R} &= \{ i \in [n] : 1 - \langle \mathbf{w}^*(\alpha, \beta), \bar{\mathbf{x}}_i \rangle < 0 \}, \\ \mathcal{E} &= \{ i \in [n] : 1 - \langle \mathbf{w}^*(\alpha, \beta), \bar{\mathbf{x}}_i \rangle \in [0, \gamma] \}, \\ \mathcal{L} &= \{ i \in [n] : 1 - \langle \mathbf{w}^*(\alpha, \beta), \bar{\mathbf{x}}_i \rangle > \gamma \}, \end{aligned}$$

which imply that

$$(i) \ i \in \mathcal{F} \Rightarrow [\mathbf{w}^*(\alpha, \beta)]_i = 0, \quad (ii) \ \begin{cases} i \in \mathcal{R} \Rightarrow [\theta^*(\alpha, \beta)]_i = 0, \\ i \in \mathcal{L} \Rightarrow [\theta^*(\alpha, \beta)]_i = 1. \end{cases} \quad (\text{R})$$

Thus, we call the j -th feature *inactive* if $j \in \mathcal{F}$. The samples in \mathcal{E} are the so-called support vectors and we call the samples in \mathcal{R} and \mathcal{L} *inactive* samples.

Suppose that we are given subsets of \mathcal{F} , \mathcal{R} , and \mathcal{L} . Then by the rules in (R), we can see that many coefficients of $\mathbf{w}^*(\alpha, \beta)$ and $\theta^*(\alpha, \beta)$ are known. Thus, we may have much fewer unknowns to solve and the problem size can be dramatically reduced. We formalize this idea in Lemma 3.

Lemma 3 *Given index sets $\hat{\mathcal{F}} \subseteq \mathcal{F}$, $\hat{\mathcal{R}} \subseteq \mathcal{R}$, and $\hat{\mathcal{L}} \subseteq \mathcal{L}$, the followings hold:*

(i) $[\mathbf{w}^*(\alpha, \beta)]_{\hat{\mathcal{F}}} = 0$, $[\theta^*(\alpha, \beta)]_{\hat{\mathcal{R}}} = 0$, $[\theta^*(\alpha, \beta)]_{\hat{\mathcal{L}}} = 1$;

(ii) *let $\hat{\mathcal{D}} = \hat{\mathcal{R}} \cup \hat{\mathcal{L}}$, $\hat{\mathbf{G}}_1 = \frac{1}{|\hat{\mathcal{D}}^c|} [\bar{\mathbf{X}}]_{\hat{\mathcal{D}}^c}$, and $\hat{\mathbf{G}}_2 = \frac{1}{|\hat{\mathcal{D}}^c|} [\bar{\mathbf{X}}]_{\hat{\mathcal{L}}}$, where $\hat{\mathcal{F}}^c = [p] \setminus \hat{\mathcal{F}}$, $\hat{\mathcal{D}}^c = [n] \setminus \hat{\mathcal{D}}$, and $\hat{\mathcal{L}}^c = [n] \setminus \hat{\mathcal{L}}$. Then, $[\theta^*(\alpha, \beta)]_{\hat{\mathcal{D}}^c}$ solves the following scaled dual problem:*

$$\min_{\hat{\theta} \in [0,1]^{|\hat{\mathcal{D}}^c|}} \left\{ \frac{1}{2\alpha} \left\| \mathcal{S}_\beta \left(\frac{1}{n} \hat{\mathbf{G}}_1 \hat{\theta} + \frac{1}{n} \hat{\mathbf{G}}_2 \mathbf{1} \right) \right\|^2 + \frac{\gamma}{2n} \|\hat{\theta}\|^2 - \frac{1}{n} \langle \mathbf{1}, \hat{\theta} \rangle \right\}; \quad (\text{scaled-D}^*)$$

(iii) *suppose that $\theta^*(\alpha, \beta)$ is known, then,*

$$[\mathbf{w}^*(\alpha, \beta)]_{\hat{\mathcal{F}}^c} = \frac{1}{\alpha} \mathcal{S}_\beta \left(\frac{1}{n} [\bar{\mathbf{X}}]_{\hat{\mathcal{F}}^c} \theta^*(\alpha, \beta) \right).$$

Lemma 3 indicates that, if we can identify index sets $\hat{\mathcal{F}}$ and $\hat{\mathcal{D}}$ and the cardinalities of $\hat{\mathcal{F}}^c$ and $\hat{\mathcal{D}}^c$ are much smaller than the feature dimension p and the sample size n , we only need to solve problem (scaled-D*) that may be much *smaller* than problem (D*) to exactly recover the optima $\mathbf{w}^*(\alpha, \beta)$ and $\theta^*(\alpha, \beta)$ *without sacrificing any accuracy*.

However, we cannot directly apply Rules in (R) to identify subsets of \mathcal{F} , \mathcal{R} , and \mathcal{L} , as they require the knowledge of $\mathbf{w}^*(\alpha, \beta)$ and $\theta^*(\alpha, \beta)$ that are usually unavailable. Inspired by the idea in El Ghaoui et al. 2012, we can first estimate regions \mathcal{W} and Θ that contain $\mathbf{w}^*(\alpha, \beta)$ and $\theta^*(\alpha, \beta)$, respectively. Then, by denoting

$$\hat{\mathcal{F}} := \left\{ j \in [p] : \max_{\theta \in \Theta} \left\{ \left| \frac{1}{n} [\bar{\mathbf{X}}\theta]_j \right| \right\} \leq \beta \right\}, \quad (2)$$

$$\hat{\mathcal{R}} := \left\{ i \in [n] : \max_{\mathbf{w} \in \mathcal{W}} \{1 - \langle \mathbf{w}, \bar{\mathbf{x}}_i \rangle\} < 0 \right\}, \quad (3)$$

$$\hat{\mathcal{L}} := \left\{ i \in [n] : \min_{\mathbf{w} \in \mathcal{W}} \{1 - \langle \mathbf{w}, \bar{\mathbf{x}}_i \rangle\} > \gamma \right\}, \quad (4)$$

since it is easy to know that $\hat{\mathcal{F}} \subset \mathcal{F}$, $\hat{\mathcal{R}} \subset \mathcal{R}$, and $\hat{\mathcal{L}} \subset \mathcal{L}$, the rules in (R) can be relaxed as:

$$(i) \ j \in \hat{\mathcal{F}} \Rightarrow [\mathbf{w}^*(\alpha, \beta)]_j = 0, \quad (R1)$$

$$(ii) \ \begin{cases} i \in \hat{\mathcal{R}} \Rightarrow [\theta^*(\alpha, \beta)]_i = 0, \\ i \in \hat{\mathcal{L}} \Rightarrow [\theta^*(\alpha, \beta)]_i = 1. \end{cases} \quad (R2)$$

In view of Rules R1 and R2, we sketch the development of SIFS as follows.

Step 1: Derive estimations \mathcal{W} and Θ such that $\mathbf{w}^*(\alpha, \beta) \in \mathcal{W}$ and $\theta^*(\alpha, \beta) \in \Theta$, respectively.

Step 2: Develop SIFS by deriving the relaxed screening rules R1 and R2, i.e., by solving the optimization problems in Eqs. (2), (3) and (4).

3. Estimate the Primal and Dual Optima

In this section, we first show that the primal and dual optima admit closed-form solutions for specific values of α and β (Section 3.1). Then, in Sections 3.2 and 3.3, we present accurate estimations of the primal and dual optima, respectively. We would like to point out that in order to extend the optimum estimation results below to multi-class sparse SVM and avoid redundancy, we consider the optimum estimations for the primal and dual problems in more general forms.

3.1. Effective Intervals of Parameters α and β

Below we show two things. One is that if the value of β is sufficiently large, no matter what α is, the primal solution is 0. The other is for any β , if α is large enough, the primal and dual optima admit closed-form solutions.

Lemma 4 *Let $\beta_{\max} = \|\frac{1}{n} \bar{\mathbf{X}}\mathbf{1}\|_{\infty}$ and $\alpha_{\max}(\beta) = \frac{1}{1-\gamma} \max_{i \in [n]} \{\langle \bar{\mathbf{x}}_i, \mathcal{S}_{\beta}(\frac{1}{n} \bar{\mathbf{X}}\mathbf{1}) \rangle\}$. Then, (i) for $\alpha > 0$ and $\beta \geq \beta_{\max}$, we have*

$$\mathbf{w}^*(\alpha, \beta) = \mathbf{0}, \quad \theta^*(\alpha, \beta) = \mathbf{1};$$

(ii) for all $\alpha \in [\max\{\alpha_{\max}(\beta), 0\}, \infty) \cap (0, \infty)$, we have

$$\mathbf{w}^*(\alpha, \beta) = \frac{1}{\alpha} \mathcal{S}_\beta \left(\frac{1}{n} \bar{\mathbf{X}} \mathbf{1} \right), \quad \theta^*(\alpha, \beta) = \mathbf{1}. \quad (5)$$

By Lemma 4, we only need to consider the cases where $\beta \in (0, \beta_{\max}]$ and $\alpha \in (0, \alpha_{\max}(\beta)]$.

3.2. Primal Optimum Estimation

In Section 1, we mention that the proposed SIFS consists of IFS and ISS, which can identify the inactive features and samples, respectively. We also mention that an alternating application of IFS and ISS can improve the estimation accuracy of the primal and dual optimum estimations, which can in turn make ISS and IFS more effective in identifying inactive samples and features, respectively. We would like to show that discarding inactive features by IFS leads to a more accurate estimation of the primal optimum.

We consider the following general problem (g-P*):

$$\min_{\mathbf{w} \in \mathbb{R}^p} P(\mathbf{w}; \alpha, \beta) = L(\mathbf{w}) + \frac{\alpha}{2} \|\mathbf{w}\|^2 + \beta \|\mathbf{w}\|_1, \quad (\text{g-P}^*)$$

where $L(\mathbf{w}) : \mathbb{R}^p \rightarrow \mathbb{R}_+$ is smooth and convex. Exploiting the strong convexity of the objective function, we obtain the optimum estimation in the following lemma.

Lemma 5 *Suppose that the optimum $\mathbf{w}^*(\alpha_0, \beta_0)$ of problem (g-P*) at (α_0, β_0) with $\beta_0 \in (0, \beta_{\max}]$ and $\alpha_0 \in (0, \alpha_{\max}(\beta_0)]$ is known. Consider problem (g-P*) with parameters $\alpha > 0$ and β_0 . Let $\hat{\mathcal{F}}$ be the index set satisfying $[\mathbf{w}^*(\alpha, \beta_0)]_{\hat{\mathcal{F}}} = \mathbf{0}$ and define*

$$\mathbf{c} = \frac{\alpha_0 + \alpha}{2\alpha} [\mathbf{w}^*(\alpha_0, \beta_0)]_{\hat{\mathcal{F}}^c}, \quad (6)$$

$$r^2 = \frac{(\alpha_0 - \alpha)^2}{4\alpha^2} \|\mathbf{w}^*(\alpha_0, \beta_0)\|^2 - \frac{(\alpha_0 + \alpha)^2}{4\alpha^2} \|[\mathbf{w}^*(\alpha_0, \beta_0)]_{\hat{\mathcal{F}}}\|^2. \quad (7)$$

Then, the following holds:

$$[\mathbf{w}^*(\alpha, \beta_0)]_{\hat{\mathcal{F}}^c} \in \{\mathbf{w} : \|\mathbf{w} - \mathbf{c}\| \leq r\}.$$

For problem (P*), let $\hat{\mathcal{F}}$ be the index set of inactive features identified by previous IFS steps. We have $[\mathbf{w}^*(\alpha, \beta_0)]_{\hat{\mathcal{F}}} = \mathbf{0}$. Hence, we only need to find an estimation for $[\mathbf{w}^*(\alpha, \beta_0)]_{\hat{\mathcal{F}}^c}$. Since problem (P*) is a special case of problem (g-P*), given the reference solution $\mathbf{w}^*(\alpha_0, \beta_0)$ and the set $\hat{\mathcal{F}}$, from Lemma 5 we have:

$$[\mathbf{w}^*(\alpha, \beta_0)]_{\hat{\mathcal{F}}^c} \in \mathcal{W} := \{\mathbf{w} : \|\mathbf{w} - \mathbf{c}\| \leq r\}, \quad (8)$$

where \mathbf{c} and r are defined in Eqs. (6) and (7), respectively.

It shows that $[\mathbf{w}^*(\alpha, \beta_0)]_{\hat{\mathcal{F}}^c}$ lies in a ball of radius r centered at \mathbf{c} . Note that, before we perform IFS, the set $\hat{\mathcal{F}}$ is empty and the second term on the right hand side of Eq. (7) is thus 0. If we apply IFS multiple times alternatively with ISS, the set $\hat{\mathcal{F}}$ will be monotonically increasing. Thus, Eq. (7) implies that the radius will be monotonically decreasing, leading to a more accurate primal optimum estimation.

3.3. Dual Optimum Estimation

We consider a general dual problem (g-D*) below:

$$\min_{\theta \in [0,1]^n} D(\theta; \alpha, \beta) = \frac{1}{2\alpha} f_\beta(\theta) + \frac{\gamma}{2n} \|\theta\|^2 - \frac{1}{n} \langle \mathbf{v}, \theta \rangle, \quad (\text{g-D}^*)$$

where $\mathbf{v} \in \mathbb{R}^n$ and $f_\beta(\theta) : \mathbb{R}^n \rightarrow \mathbb{R}_+$ is smooth and convex. It is obvious that problem (g-D*) can be reduced to problem (D*) by letting $f_\beta(\theta) = \|\mathcal{S}_\beta(\frac{1}{n}\bar{\mathbf{X}}\theta)\|^2$ and $\mathbf{v} = \mathbf{1}$. The lemma below gives an optimum estimation of problem (g-D*) based on the strong convexity of its objective function.

Lemma 6 *Suppose that the optimum $\theta^*(\alpha_0, \beta_0)$ of problem (g-D*) with $\beta_0 \in (0, \beta_{\max}]$ and $\alpha_0 \in (0, \alpha_{\max}(\beta_0)]$ is known. Consider problem (g-D*) at (α, β_0) with $\alpha > 0$ and let $\hat{\mathcal{R}}$ and $\hat{\mathcal{L}}$ be two index sets satisfying $[\theta^*(\alpha, \beta_0)]_{\hat{\mathcal{R}}} = \mathbf{0}$ and $[\theta^*(\alpha, \beta_0)]_{\hat{\mathcal{L}}} = \mathbf{1}$, respectively. We denote $\hat{\mathcal{D}} = \hat{\mathcal{R}} \cup \hat{\mathcal{L}}$ and define*

$$\mathbf{c} = \frac{\alpha - \alpha_0}{2\gamma\alpha} [\mathbf{v}]_{\hat{\mathcal{D}}^c} + \frac{\alpha_0 + \alpha}{2\alpha} [\theta^*(\alpha_0, \beta_0)]_{\hat{\mathcal{D}}^c}, \quad (9)$$

$$\begin{aligned} r^2 = & \left(\frac{\alpha - \alpha_0}{2\alpha}\right)^2 \|\theta^*(\alpha_0, \beta_0) - \frac{1}{\gamma} \mathbf{v}\|^2 - \|\mathbf{1} - \frac{\alpha - \alpha_0}{2\gamma\alpha} [\mathbf{v}]_{\hat{\mathcal{L}}} - \frac{\alpha_0 + \alpha}{2\alpha} [\theta^*(\alpha_0, \beta_0)]_{\hat{\mathcal{L}}}\|^2 \\ & - \|\frac{\alpha - \alpha_0}{2\gamma\alpha} [\mathbf{v}]_{\hat{\mathcal{R}}} + \frac{\alpha_0 + \alpha}{2\alpha} [\theta^*(\alpha_0, \beta_0)]_{\hat{\mathcal{R}}}\|^2. \end{aligned} \quad (10)$$

Then, the following holds:

$$[\theta^*(\alpha, \beta_0)]_{\hat{\mathcal{D}}^c} \in \{\theta : \|\theta - \mathbf{c}\| \leq r\}.$$

We now turn back to problem (D*). As it is a special case of problem (g-D*), given the reference solution $\theta^*(\alpha_0, \beta_0)$ at (α_0, β_0) and the index sets of inactive samples identified by the previous ISS steps $\hat{\mathcal{R}}$ and $\hat{\mathcal{L}}$, using Lemma 6, we can obtain:

$$[\theta^*(\alpha, \beta_0)]_{\hat{\mathcal{D}}^c} \in \Theta := \{\theta : \|\theta - \mathbf{c}\| \leq r\}, \quad (11)$$

where \mathbf{c} and \mathbf{r} are defined by Eqs. (9) and (10) with $\mathbf{v} = \mathbf{1}$, respectively. Therefore, $[\theta^*(\alpha, \beta_0)]_{\hat{\mathcal{D}}^c}$ lies in the ball Θ . In view of Eq. (10), the index sets $\hat{\mathcal{L}}$ and $\hat{\mathcal{R}}$ monotonically increase and hence the last two terms on the right hand side of Eq. (10) monotonically increase when we perform ISS multiple times (alternating with IFS), which implies that the ISS steps can reduce the radius and thus improve the dual optimum estimation.

In addition, from both Lemmas 5 and 6, we can see that the radii of \mathcal{W} and Θ can be potentially large when α is very small, which may affect our estimation accuracy. We can sidestep this issue by letting the ratio α/α_0 be a constant. That is why we space the values of α equally at the logarithmic scale on the parameter value path in the experiments.

Remark 7 *To estimate the optima $\mathbf{w}^*(\alpha, \beta_0)$ and $\theta^*(\alpha, \beta_0)$ of problems (P*) and (D*) using Lemmas 5 and 6, we have a free reference solution pair $\mathbf{w}^*(\alpha_0, \beta_0)$ and $\theta^*(\alpha_0, \beta_0)$ with $\alpha_0 = \alpha_{\max}(\beta_0)$. The reason is that $\mathbf{w}^*(\alpha_0, \beta_0)$ and $\theta^*(\alpha_0, \beta_0)$ admit closed-form solutions in this setting (see Lemma 4).*

4. The Proposed SIFS Screening Rule

We first present the IFS and ISS rules in Sections 4.1 and 4.2, respectively. Then, in Section 4.3, we develop the SIFS screening rule by an alternating application of IFS and ISS.

4.1. Inactive Feature Screening (IFS)

Suppose that $\mathbf{w}^*(\alpha_0, \beta_0)$ and $\theta^*(\alpha_0, \beta_0)$ are known, we derive IFS to identify inactive features for problem (P*) at (α, β_0) by solving the optimization problem in Eq. (2) (see Section A.5 in the appendix):

$$s^i(\alpha, \beta_0) = \max_{\theta \in \Theta} \left\{ \frac{1}{n} |\langle [\bar{\mathbf{x}}^i]_{\hat{\mathcal{D}}^c}, \theta \rangle + \langle [\bar{\mathbf{x}}^i]_{\hat{\mathcal{L}}}, \mathbf{1} \rangle| \right\}, i \in \hat{\mathcal{F}}^c, \quad (12)$$

where Θ is given by Eq. (11) and $\hat{\mathcal{F}}$ and $\hat{\mathcal{D}} = \hat{\mathcal{R}} \cup \hat{\mathcal{L}}$ are the index sets of inactive features and samples that have been identified in previous screening processes, respectively. The next result shows the closed-form solution of problem (12).

Lemma 8 *Consider problem (12). Let \mathbf{c} and r be given by Eq. (9) and Eq. (10) with $\mathbf{v} = \mathbf{1}$. Then, for all $i \in \hat{\mathcal{F}}^c$, we have*

$$s^i(\alpha, \beta_0) = \frac{1}{n} (|\langle [\bar{\mathbf{x}}^i]_{\hat{\mathcal{D}}^c}, \mathbf{c} \rangle + \langle [\bar{\mathbf{x}}^i]_{\hat{\mathcal{L}}}, \mathbf{1} \rangle| + \|[\bar{\mathbf{x}}^i]_{\hat{\mathcal{D}}^c}\| r).$$

We are now ready to present the IFS rule.

Theorem 9 [Feature screening rule IFS] *Consider problem (P*). We suppose that $\mathbf{w}^*(\alpha_0, \beta_0)$ and $\theta^*(\alpha_0, \beta_0)$ are known. Then:*

(1) *the feature screening rule IFS takes the form of*

$$s^i(\alpha, \beta_0) \leq \beta_0 \Rightarrow [\mathbf{w}^*(\alpha, \beta_0)]_i = 0, \forall i \in \hat{\mathcal{F}}^c; \quad (\text{IFS})$$

(2) *we can update the index set $\hat{\mathcal{F}}$ by*

$$\hat{\mathcal{F}} \leftarrow \hat{\mathcal{F}} \cup \Delta \hat{\mathcal{F}} \text{ with } \Delta \hat{\mathcal{F}} = \{i : s^i \leq \beta_0, i \in \hat{\mathcal{F}}^c\}. \quad (13)$$

Recall that (see Lemma 6) previous sample screening results give us a tighter dual estimation, i.e., a smaller feasible region Θ for problem (12), which results in a smaller $s^i(\alpha, \beta_0)$. It finally brings about a more powerful feature screening rule IFS. This is the so called synergistic effect.

4.2. Inactive Sample Screening (ISS)

Similar to IFS, we derive ISS to identify inactive samples by solving the optimization problems in Eq. (3) and Eq. (4) (see Section A.7 in the appendix for details):

$$u_i(\alpha, \beta_0) = \max_{\mathbf{w} \in \mathcal{W}} \{1 - \langle [\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c}, \mathbf{w} \rangle\}, i \in \hat{\mathcal{D}}^c, \quad (14)$$

$$l_i(\alpha, \beta_0) = \min_{\mathbf{w} \in \mathcal{W}} \{1 - \langle [\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c}, \mathbf{w} \rangle\}, i \in \hat{\mathcal{D}}^c, \quad (15)$$

where \mathcal{W} is given by Eq. (8) and $\hat{\mathcal{F}}$ and $\hat{\mathcal{D}} = \hat{\mathcal{R}} \cup \hat{\mathcal{L}}$ are the index sets of inactive features and samples that have been identified in previous screening processes. We show that problems (14) and (15) admit closed-form solutions.

Lemma 10 *Consider problems (14) and (15). Let \mathbf{c} and r be given by Eq. (6) and Eq. (7). Then,*

$$\begin{aligned} u_i(\alpha, \beta_0) &= 1 - \langle [\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c}, \mathbf{c} \rangle + \|[\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c} \| r, i \in \hat{\mathcal{D}}^c, \\ l_i(\alpha, \beta_0) &= 1 - \langle [\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c}, \mathbf{c} \rangle - \|[\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c} \| r, i \in \hat{\mathcal{D}}^c. \end{aligned}$$

We are now ready to present the ISS rule.

Theorem 11 [Sample screening rule ISS] *Consider problem (D^*) and suppose that $\mathbf{w}^*(\alpha_0, \beta_0)$ and $\theta^*(\alpha_0, \beta_0)$ are known, then:*

(1) *the sample screening rule ISS takes the form of*

$$\begin{aligned} u_i(\alpha, \beta_0) < 0 &\Rightarrow [\theta^*(\alpha, \beta_0)]_i = 0, \\ l_i(\alpha, \beta_0) > \gamma &\Rightarrow [\theta^*(\alpha, \beta_0)]_i = 1, \end{aligned} \quad \forall i \in \hat{\mathcal{D}}^c; \quad (\text{ISS})$$

(2) *we can update the index sets $\hat{\mathcal{R}}$ and $\hat{\mathcal{L}}$ by*

$$\hat{\mathcal{R}} \leftarrow \hat{\mathcal{R}} \cup \Delta \hat{\mathcal{R}} \text{ with } \Delta \hat{\mathcal{R}} = \{i : u_i(\alpha, \beta_0) < 0, i \in \hat{\mathcal{D}}^c\}, \quad (16)$$

$$\hat{\mathcal{L}} \leftarrow \hat{\mathcal{L}} \cup \Delta \hat{\mathcal{L}} \text{ with } \Delta \hat{\mathcal{L}} = \{i : l_i(\alpha, \beta_0) > \gamma, i \in \hat{\mathcal{D}}^c\}. \quad (17)$$

The synergistic effect also exists here. Recall that (see Lemma 5), previous feature screening results lead to a smaller feasible region \mathcal{W} for problems (14) and (15), which results in smaller $u_i(\alpha, \beta_0)$ and larger $l_i(\alpha, \beta_0)$. It finally leads us to a more accurate sample screening rule ISS.

4.3. The Proposed SIFS Rule by An Alternating Application of IFS and ISS

In real applications, the optimal parameter values of α and β are usually unknown. To determine appropriate parameter values, common approaches, like cross validation and stability selection, need to solve the model over a grid of parameter values $\{(\alpha_{i,j}, \beta_j) : i \in [M], j \in [N]\}$ with $\beta_{\max} > \beta_1 > \dots > \beta_N > 0$ and $\alpha_{\max}(\beta_j) > \alpha_{1,j} > \dots > \alpha_{M,j} > 0$. This can be very time-consuming. Inspired by Strong Rule (Tibshirani et al., 2012) and SAFE (El Ghaoui et al., 2012), we develop a sequential version of SIFS in Algorithm 1. Specifically, given the primal and dual optima $\mathbf{w}^*(\alpha_{i-1,j}, \beta_j)$ and $\theta^*(\alpha_{i-1,j}, \beta_j)$ at $(\alpha_{i-1,j}, \beta_j)$, we apply SIFS to identify the inactive features and samples for problem (P^*) at $(\alpha_{i,j}, \beta_j)$. Then, we perform training on the reduced data set and solve the primal and dual optima at $(\alpha_{i,j}, \beta_j)$. We repeat this process until we solve problem (P^*) at all pairs of parameter values.

Note that we insert $\alpha_{0,j}$ into every sequence $\{\alpha_{i,j} : i \in [M]\}$ (see line 1 in Algorithm 1) to obtain a closed-form solution as the first reference solution. In this way, we can avoid solving problem at $(\alpha_{1,j}, \beta_j), j \in [N]$ directly (without screening), which is time consuming. At last, we would like to point out that the values $\{(\alpha_{i,j}, \beta_j) : i \in [M], j \in [N]\}$ in SIFS can be specified by users arbitrarily.

Algorithm 1 SIFS

```

1: Input:  $\beta_{\max} > \beta_1 > \dots > \beta_N > 0$  and  $\alpha_{\max}(\beta_j) = \alpha_{0,j} > \alpha_{1,j} > \dots > \alpha_{M,j} > 0$ .
2: for  $j = 1$  to  $N$  do
3:   Compute the first reference solution  $\mathbf{w}^*(\alpha_{0,j}, \beta_j)$  and  $\theta^*(\alpha_{0,j}, \beta_j)$  using the close-form
   formulas in Eq. (5).
4:   for  $i = 1$  to  $M$  do
5:     Initialization:  $\hat{\mathcal{F}} = \hat{\mathcal{R}} = \hat{\mathcal{L}} = \emptyset$ .
6:     repeat
7:       Run sample screening using rule ISS based on  $\mathbf{w}^*(\alpha_{i-1,j}, \beta_j)$ .
8:       Update the inactive sample sets  $\hat{\mathcal{R}}$  and  $\hat{\mathcal{L}}$ :
            $\hat{\mathcal{R}} \leftarrow \hat{\mathcal{R}} \cup \Delta\hat{\mathcal{R}}$  and  $\hat{\mathcal{L}} \leftarrow \hat{\mathcal{L}} \cup \Delta\hat{\mathcal{L}}$ .
9:       Run feature screening using rule IFS based on  $\theta^*(\alpha_{i-1,j}, \beta_j)$ .
10:      Update the inactive feature set  $\hat{\mathcal{F}}$ :
            $\hat{\mathcal{F}} \leftarrow \hat{\mathcal{F}} \cup \Delta\hat{\mathcal{F}}$ .
11:     until No new inactive features or samples are identified
12:     Compute  $\mathbf{w}^*(\alpha_{i,j}, \beta_j)$  and  $\theta^*(\alpha_{i,j}, \beta_j)$  by solving the scaled problem.
13:   end for
14: end for
15: Output:  $\mathbf{w}^*(\alpha_{i,j}, \beta_j)$  and  $\theta^*(\alpha_{i,j}, \beta_j)$ ,  $i \in [M]$ ,  $j \in [N]$ .
    
```

SIFS applies ISS and IFS in an alternating manner to reinforce their capability in identifying inactive samples and features. In Algorithm 1, we apply ISS first. Of course, we can also apply IFS first. The theorem below demonstrates that the orders have no impact on the performance of SIFS.

Theorem 12 *Given the optimal solutions $\mathbf{w}^*(\alpha_{i-1,j}, \beta_j)$ and $\theta^*(\alpha_{i-1,j}, \beta_j)$ at $(\alpha_{i-1,j}, \beta_j)$ as the reference solution pair at $(\alpha_{i,j}, \beta_j)$ for SIFS, we assume SIFS with ISS first stops after applying IFS and ISS for s times and denote the identified inactive features and samples as $\hat{\mathcal{F}}_s^A, \hat{\mathcal{R}}_s^A$, and $\hat{\mathcal{L}}_s^A$. Similarly, when we apply IFS first, the results are denoted as $\hat{\mathcal{F}}_t^B, \hat{\mathcal{R}}_t^B$, and $\hat{\mathcal{L}}_t^B$. Then, the followings hold:*

(1) $\hat{\mathcal{F}}_s^A = \hat{\mathcal{F}}_t^B, \hat{\mathcal{R}}_s^A = \hat{\mathcal{R}}_t^B$, and $\hat{\mathcal{L}}_s^A = \hat{\mathcal{L}}_t^B$.

(2) *with different orders of applying ISS and IFS, the difference between the times of ISS and IFS we need to apply in SIFS can never be larger than 1, that is, $|s - t| \leq 1$.*

4.4. Discussion

After developing the proposed method SIFS above, we now turn to the related work discussion in order to highlight the differences between SIFS and the existing methods and also the novelty of our method, although we have mentioned some of them in the introduction section. We divide the previous work into two parts: screening for sparse SVM and for other sparse learning models.

To the best of our knowledge, the screening method in Shibagaki et al. 2016 is the only method besides our SIFS, which can simultaneously identify the inactive features and samples for sparse SVM. There are mainly three big differences between SIFS and this method. First, the techniques used in the optima estimations of SIFS and Shibagaki

et al. 2016 are quite different. To be specific, the estimations in SIFS are developed by exploiting the reference solution pair and carefully studying the strong convexity of the objective functions and the optimum conditions of problems (P*) and (D*) at the current and reference parameter value pairs (see Lemmas 5 and 6 and also their proofs for the details). In contrast, Shibagaki et al. (2016) estimated the optima heavily based on the duality gap. As we mentioned in the introduction section, duality gap is usually large in the early stages and decreases gradually, which weakens its capability in the early stages and leads to a limited overall speedup. Second, algorithmically, Shibagaki et al. (2016) is dynamic while our method is static. In other words, Shibagaki et al. (2016) identifies the inactive features and samples during the training process, and in our method, we do this before the training process (see steps 6 to 11 in Algorithm 1), which means Shibagaki et al. (2016) needs to apply its screening rules for many times while we trigger SIFS only once. These two technical and algorithmic differences make our method outperform Shibagaki et al. (2016), which will be verified in the experimental section. At last, we theoretically prove that in the static scenario the orders of applying the feature and sample screening rules have no impact on the final performance, while Shibagaki et al. (2016) did not give the theoretical result accordingly in dynamic screening.

There are mainly three big differences between our SIFS and existing methods for other sparse learning models. First, these existing methods identify either features or samples individually and would be helpless in real applications involving a huge number of samples and extremely high-dimensional features, while SIFS identifies features and samples simultaneously. Second, technically, SIFS can incorporate the feature (resp. sample) screening results of the previous steps as the prior knowledge into the primal (reps. dual) optimum estimation to obtain a more accurate estimation. This is verified with a strong theoretical guarantee (see Lemmas 5 and 6). At last, we give a theoretical proof (see Sections 4.1 and 4.2) to show the existence of the synergistic effect between feature and sample screening for the first time in the static scenario.

Finally, we would like to point out that although the key ideas used in some of the screening methods including SIFS seem to be similar, they are developed specifically for the sparse models they focus on, which makes it nontrivial and even very difficult to reduce one method to another by simply resetting the loss and regularizer of the model. For example, SIFS cannot be reduced to Wang et al. (2014b) by setting $\alpha = 0$ and letting loss be the logistic loss, although both of their dual optimum estimations are based on the strong convexity of the objective. The reason is that they use the strong convexity in quite different ways due to their different dual problems including the feasible regions and the expressions of the dual objectives (see Theorem 2 in Wang et al., 2014b, Lemma 6 above, and their proofs for the details). Moreover, we cannot reduce SIFS to the methods (Ogawa et al., 2013; Wang et al., 2014a) for SVM by setting β to be 0, since they are based on the convexity of the objective while SIFS exploits the strong convexity.

5. SIFS for Multi-class Sparse SVMs

In this section, we consider extending SIFS to multi-class sparse SVMs. We will briefly review the basics of multi-class sparse SVMs and then derive a series of theorems as we did

for sparse SVMs above. Finally, we will present the detailed screening rule for multi-class sparse SVMs.

5.1. Basics of Multi-class Sparse SVMs

We focus on the ℓ_1 -regularized multi-class SVM with a smoothed hinge loss, which takes the form of

$$\min_{\mathbf{w} \in \mathbb{R}^{Kp}} P(\mathbf{w}; \alpha, \beta) = \frac{1}{n} \sum_{i=1}^n \ell_i(\mathbf{w}) + \frac{\alpha}{2} \|\mathbf{w}\|^2 + \beta \|\mathbf{w}\|_1, \quad (\text{m-P}^*)$$

where $\mathbf{w} = [\mathbf{w}_1; \mathbf{w}_2; \dots; \mathbf{w}_K] \in \mathbb{R}^{Kp}$ is the parameter vector to be estimated with $\mathbf{w}_k \in \mathbb{R}^p, k = 1, \dots, K$. The loss function $\ell_i(\mathbf{w}) = \sum_{k \neq y_i} \ell(\mathbf{w}_k^\top \mathbf{x}_i - \mathbf{w}_{y_i}^\top \mathbf{x}_i + 1)$, with $\{\mathbf{x}_i, y_i\}_{i=1}^n$ is the training data set of K classes, $\mathbf{x}_i \in \mathbb{R}^p, y_i \in \{1, \dots, K\}$, and $\ell(\cdot)$ is the smoothed hinge loss defined in Eq. (1).

The following theorem presents the Lagrangian dual problem of (m-P*) and the KKT conditions, which are essential for developing the screening rules.

Theorem 13 *For each sample (\mathbf{x}_i, y_i) , we define $\mathbf{u}_i = 1 - \mathbf{e}_{y_i} \in \mathbb{R}^K$ and $\mathbf{X}_i = [\mathbf{X}_i^1, \mathbf{X}_i^2, \dots, \mathbf{X}_i^K] \in \mathbb{R}^{Kp \times K}$, where $\mathbf{X}_i^k = \text{vec}(\mathbf{x}_i(\mathbf{e}_k - \mathbf{e}_{y_i})^\top) \in \mathbb{R}^{Kp}$. Let $\mathbf{u} = [\mathbf{u}_1; \dots; \mathbf{u}_n] \in \mathbb{R}^{Kn}$ and $\mathbf{X} = [\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_n] \in \mathbb{R}^{Kp \times Kn}$. Then, for the problem (m-P*), the followings hold:*

(i) *the dual problem of (m-P*) is*

$$\min_{\theta \in [0, 1]^{Kn}} D(\theta; \alpha, \beta) = \frac{1}{2\alpha} \left\| \mathcal{S}_\beta \left(\frac{1}{n} \mathbf{X} \theta \right) \right\|^2 + \frac{\gamma}{2n} \|\theta\|^2 - \frac{1}{n} \langle \mathbf{u}, \theta \rangle, \quad (\text{m-D}^*)$$

where $\theta = [\theta_1; \dots; \theta_n]$ with $\theta_i \in [0, 1]^K, i = 1, \dots, n$;

(ii) *denote the optima of problems (m-P*) and (m-D*) by $\mathbf{w}^*(\alpha, \beta)$ and $\theta^*(\alpha, \beta)$, respectively, then,*

$$\mathbf{w}^*(\alpha, \beta) = \frac{1}{\alpha} \mathcal{S}_\beta \left(-\frac{1}{n} \mathbf{X} \theta^*(\alpha, \beta) \right), \quad (\text{m-KKT-1})$$

$$[\theta_i^*(\alpha, \beta)]_k = \begin{cases} 0, & \text{if } \langle \mathbf{X}_i^k, \mathbf{w}^*(\alpha, \beta) \rangle + [\mathbf{u}_i]_k \leq 0; \\ 1, & \text{if } \langle \mathbf{X}_i^k, \mathbf{w}^*(\alpha, \beta) \rangle + [\mathbf{u}_i]_k \geq \gamma; \\ \frac{1}{\gamma} (\langle \mathbf{X}_i^k, \mathbf{w}^*(\alpha, \beta) \rangle + [\mathbf{u}_i]_k), & \text{otherwise;} \end{cases} \quad k = 1, 2, \dots, K. \quad (\text{m-KKT-2})$$

As we did in Section 2, we can also define 3 index sets here:

$$\begin{aligned} \mathcal{F} &= \left\{ (k, j) \in [K] \times [p] : \frac{1}{n} |[\mathbf{X} \theta^*(\alpha, \beta)]_{k,j}| \leq \beta \right\}, \\ \mathcal{R} &= \left\{ (i, k) \in [n] \times [K] : \langle \mathbf{X}_i^k, \mathbf{w}^*(\alpha, \beta) \rangle + [\mathbf{u}_i]_k \leq 0 \right\}, \\ \mathcal{L} &= \left\{ (i, k) \in [n] \times [K] : \langle \mathbf{X}_i^k, \mathbf{w}^*(\alpha, \beta) \rangle + [\mathbf{u}_i]_k \geq \gamma \right\}, \end{aligned}$$

which imply that

$$\begin{aligned} \text{(i)} \quad & (k, j) \in \mathcal{F} \Rightarrow [\mathbf{w}_k^*(\alpha, \beta)]_j = 0, \\ \text{(ii)} \quad & \begin{cases} (i, k) \in \mathcal{R} \Rightarrow [\theta_i^*(\alpha, \beta)]_k = 0, \\ (i, k) \in \mathcal{L} \Rightarrow [\theta_i^*(\alpha, \beta)]_k = 1. \end{cases} \end{aligned} \quad (\text{m-R})$$

Suppose that we are given three subsets of \mathcal{F} , \mathcal{R} , and \mathcal{L} , then we can infer the values of many coefficients of $\mathbf{w}^*(\alpha, \beta)$ and $\theta^*(\alpha, \beta)$ via Rule m-R. The lemma below shows that the rest coefficients of $\mathbf{w}^*(\alpha, \beta)$ and $\theta^*(\alpha, \beta)$ can be recovered by solving a small sized problem.

Lemma 14 *Given index sets $\hat{\mathcal{F}} \subseteq \mathcal{F}$, $\hat{\mathcal{R}} \subseteq \mathcal{R}$, and $\hat{\mathcal{L}} \subseteq \mathcal{L}$, the followings hold:*

- (i) $[\mathbf{w}^*(\alpha, \beta)]_{\hat{\mathcal{F}}} = 0$, $[\theta^*(\alpha, \beta)]_{\hat{\mathcal{R}}} = 0$, $[\theta^*(\alpha, \beta)]_{\hat{\mathcal{L}}} = 1$;
- (ii) let $\hat{\mathcal{D}} = \hat{\mathcal{R}} \cup \hat{\mathcal{L}}$, $\hat{\mathbf{G}}_1 = \hat{\mathcal{F}}^c[\mathbf{X}]_{\hat{\mathcal{D}}^c}$, and $\hat{\mathbf{G}}_2 = \hat{\mathcal{F}}^c[\mathbf{X}]_{\hat{\mathcal{L}}}$, where $\hat{\mathcal{F}}^c = [K] \times [p] \setminus \hat{\mathcal{F}}$, $\hat{\mathcal{D}}^c = [n] \times [K] \setminus \hat{\mathcal{D}}$, $\hat{\mathcal{L}}^c = [n] \times [K] \setminus \hat{\mathcal{L}}$ and $\hat{\mathbf{u}} = [\mathbf{u}]_{\hat{\mathcal{D}}^c}$, then, $[\theta^*(\alpha, \beta)]_{\hat{\mathcal{D}}^c}$ solves the following scaled dual problem:

$$\min_{\hat{\theta} \in [0, 1]^{|\hat{\mathcal{D}}^c|}} \left\{ \frac{1}{2\alpha} \left\| \mathcal{S}_\beta \left(\frac{1}{n} \hat{\mathbf{G}}_1 \hat{\theta} + \frac{1}{n} \hat{\mathbf{G}}_2 \mathbf{1} \right) \right\|^2 + \frac{\gamma}{2n} \|\hat{\theta}\|^2 - \frac{1}{n} \langle \hat{\mathbf{u}}, \hat{\theta} \rangle \right\}; \quad (\text{m-scaled-D}^*)$$

- (iii) suppose that $\theta^*(\alpha, \beta)$ is known, then,

$$[\mathbf{w}^*(\alpha, \beta)]_{\hat{\mathcal{F}}^c} = \frac{1}{\alpha} \mathcal{S}_\beta \left(-\frac{1}{n} \hat{\mathcal{F}}^c[\mathbf{X}] \theta^*(\alpha, \beta) \right).$$

Given two estimations \mathcal{W} and Θ for $\mathbf{w}^*(\alpha, \beta)$ and $\theta^*(\alpha, \beta)$, we can define three subsets of \mathcal{F} , \mathcal{R} , and \mathcal{L} below as we did in the binary case to relax Rule m-R into the applicable version:

$$\begin{aligned} \hat{\mathcal{F}} &= \left\{ (k, j) \in [K] \times [p] : \max_{\theta \in \Theta} \left\{ \frac{1}{n} |[\mathbf{X}\theta]_{k,j}| \right\} \leq \beta \right\}, \\ \hat{\mathcal{R}} &= \left\{ (i, k) \in [n] \times [K] : \max_{\mathbf{w} \in \mathcal{W}} \left\{ \langle \mathbf{X}_i^k, \mathbf{w} \rangle + [\mathbf{u}_i]_k \right\} \leq 0 \right\}, \\ \hat{\mathcal{L}} &= \left\{ (i, k) \in [n] \times [K] : \min_{\mathbf{w} \in \mathcal{W}} \left\{ \langle \mathbf{X}_i^k, \mathbf{w}^*(\alpha, \beta) \rangle + [\mathbf{u}_i]_k \right\} \geq \gamma \right\}. \end{aligned}$$

5.2. Estimate the Primal and Dual Optima

We first derive the effective intervals of the parameters α and β .

Lemma 15 *Let $\beta_{\max} = \|\frac{1}{n} \mathbf{X}\mathbf{u}\|_\infty$ and $\alpha_{\max}(\beta) = \frac{1}{1-\gamma} \max_{(i,k) \in [n] \times [K]} \left\{ \langle \mathbf{X}_i^k, \mathcal{S}_\beta(\frac{1}{n} \mathbf{X}\mathbf{u}) \rangle \right\}$. Then:*

- (i) for $\alpha > 0$ and $\beta \geq \beta_{\max}$, we have

$$\mathbf{w}^*(\alpha, \beta) = \mathbf{0}, \quad \theta^*(\alpha, \beta) = \mathbf{u};$$

- (ii) for all $\alpha \in [\max\{\alpha_{\max}(\beta), 0\}, \infty) \cap (0, \infty)$, we have

$$\mathbf{w}^*(\alpha, \beta) = \frac{1}{\alpha} \mathcal{S}_\beta \left(-\frac{1}{n} \mathbf{X}\mathbf{u} \right), \quad \theta^*(\alpha, \beta) = \mathbf{u}. \quad (18)$$

Hence, we only need to consider the cases with $\beta \in (0, \beta_{\max}]$ and $\alpha \in (0, \alpha_{\max}(\beta)]$.

Since problem (m-P*) is a special case of problem (g-P*) with $L(\mathbf{w}) = \frac{1}{m} \sum_{i=1}^n \ell_i(\mathbf{w})$, given the reference solution $\mathbf{w}^*(\alpha_0, \beta_0)$ and the index set $\hat{\mathcal{F}}$ of the inactive features identified by the previous IFS steps, we can apply Lemma 5 to obtain the estimation for $\mathbf{w}^*(\alpha, \beta_0)_{\hat{\mathcal{F}}^c}$:

$$[\mathbf{w}^*(\alpha, \beta_0)]_{\hat{\mathcal{F}}^c} \in \mathcal{W} := \{\mathbf{w} : \|\mathbf{w} - \mathbf{c}\| \leq r\}, \quad (19)$$

where \mathbf{c} and r are defined in Eqs. (6) and (7), respectively.

Moreover, noting that problem (m-D*) is also a special case of problem (g-D*), given the reference solution $\theta^*(\alpha_0, \beta_0)$ and the index sets of inactive samples identified by the previous ISS steps $\hat{\mathcal{R}}$ and $\hat{\mathcal{L}}$, we can obtain an estimation for $[\theta^*(\alpha, \beta_0)]_{\hat{\mathcal{D}}^c}$ from Lemma 6:

$$[\theta^*(\alpha, \beta_0)]_{\hat{\mathcal{D}}^c} \in \Theta := \{\theta : \|\theta - \mathbf{c}\| \leq r\}, \quad (20)$$

where \mathbf{c} and r are defined by Eqs. (9) and (10) with $\mathbf{v} = \mathbf{u}$, respectively.

5.3. Screening Rule SIFS

Given the optima $\mathbf{w}^*(\alpha_0, \beta_0)$ and $\theta^*(\alpha_0, \beta_0)$, to derive the feature screening rule IFS, we need to solve the optimization problem below first:

$$s^{(k,j)}(\alpha, \beta_0) = \max_{\theta \in \Theta} \left\{ \frac{1}{n} |\langle [\mathbf{X}_{k,j}]_{\hat{\mathcal{D}}^c}, \theta \rangle + \langle [\mathbf{X}_{k,j}]_{\hat{\mathcal{L}}}, \mathbf{1} \rangle \right\}, (k, j) \in \hat{\mathcal{F}}^c, \quad (21)$$

where $\mathbf{X}_{k,j}$ is the $((k-1)p + j)$ -th row of \mathbf{X} , Θ is given by Eq. (20), and $\hat{\mathcal{F}}$ and $\hat{\mathcal{D}} = \hat{\mathcal{R}} \cup \hat{\mathcal{L}}$ are the index sets of the inactive features and samples identified in the previous screening process. We notice that this problem has exactly the same form as problem (12). Hence, from Lemma 8 we can obtain its closed-form solution directly. Now, from Rule m-R, we can obtain our IFS rule below.

- The feature screening rule IFS takes the form of

$$s^{(k,j)}(\alpha, \beta_0) \leq \beta_0 \Rightarrow [\mathbf{w}_k^*(\alpha, \beta_0)]_j = 0, \forall (k, j) \in \hat{\mathcal{F}}^c. \quad (\text{IFS})$$

- The index set $\hat{\mathcal{F}}$ can be updated by

$$\hat{\mathcal{F}} \leftarrow \hat{\mathcal{F}} \cup \Delta \hat{\mathcal{F}} \text{ with } \Delta \hat{\mathcal{F}} = \{(k, j) : s^{(k,j)} \leq \beta_0, (k, j) \in \hat{\mathcal{F}}^c\}. \quad (22)$$

Similarly, we need to solve the following problems to derive our sample screening rule ISS:

$$u_{(i,k)}(\alpha, \beta_0) = \max_{\mathbf{w} \in \mathcal{W}} \{ \langle [\mathbf{X}_i^k]_{\hat{\mathcal{F}}^c}, \mathbf{w} \rangle + [\mathbf{u}_i]_k \}, (i, k) \in \hat{\mathcal{D}}^c, \quad (23)$$

$$l_{(i,k)}(\alpha, \beta_0) = \min_{\mathbf{w} \in \mathcal{W}} \{ \langle [\mathbf{X}_i^k]_{\hat{\mathcal{F}}^c}, \mathbf{w} \rangle + [\mathbf{u}_i]_k \}, (i, k) \in \hat{\mathcal{D}}^c, \quad (24)$$

where \mathcal{W} is given by Eq. (19). We find that they can be solved by Lemma 10 directly. Therefore, we can obtain our sample screening rule ISS below.

- The sample screening rule ISS takes the form of

$$\begin{aligned} u_{(i,k)}(\alpha, \beta_0) \leq 0 &\Rightarrow [\theta_i^*(\alpha, \beta_0)]_k = 0, \\ l_{(i,k)}(\alpha, \beta_0) \geq \gamma &\Rightarrow [\theta_i^*(\alpha, \beta_0)]_k = 1, \end{aligned} \quad \forall (i, k) \in \hat{\mathcal{D}}^c. \quad (\text{ISS})$$

- We can update the index sets $\hat{\mathcal{R}}$ and $\hat{\mathcal{L}}$ by

$$\hat{\mathcal{R}} \leftarrow \hat{\mathcal{R}} \cup \Delta \hat{\mathcal{R}} \text{ with } \Delta \hat{\mathcal{R}} = \{(i, k) : u_{(i,k)}(\alpha, \beta_0) \leq 0, (i, k) \in \hat{\mathcal{D}}^c\}, \quad (25)$$

$$\hat{\mathcal{L}} \leftarrow \hat{\mathcal{L}} \cup \Delta \hat{\mathcal{L}} \text{ with } \Delta \hat{\mathcal{L}} = \{(i, k) : l_{(i,k)}(\alpha, \beta_0) \geq \gamma, (i, k) \in \hat{\mathcal{D}}^c\}. \quad (26)$$

The same as we did sparse SVM, we can also develop SIFS to reinforce the capabilities of ISS and IFS by applying them alternatively. The framework of SIFS for solving the model over a grid of parameter values here is the same as that in the case of sparse SVM, i.e. Algorithm 1, except for the different rules IFS and ISS, and the updates of $\hat{\mathcal{R}}$, $\hat{\mathcal{L}}$, and $\hat{\mathcal{F}}$.

In this version of SIFS, the order of applying IFS and ISS also has no impact on the final performance. Since the form of the theorem and its proof are nearly the same as that of Theorem 12, to avoid redundancy, we omit them in this section.

6. Experiments

We evaluate SIFS on both synthetic and real data sets in terms of three measurements: *speedup*, *scaling ratio*, and *rejection ratio*. Speedup is given by the ratio of the running time of the solver without screening to that with screening.

For sparse SVMs, scaling ratio is defined as $1 - \frac{(n-\tilde{n})(p-\tilde{p})}{np}$, where \tilde{n} , \tilde{p} , n , and p are the numbers of inactive samples and features identified by SIFS, sample size, and feature dimension of the data set. From steps 6 to 11 in Algorithm 1, we can see that we trigger the rules ISS and IFS repetitively until no new features and samples are identified. We adopt the rejection ratios of the i -th triggering of ISS and IFS defined as $\frac{\tilde{n}_i}{n_0}$ and $\frac{\tilde{p}_i}{p_0}$ to evaluate their performances in each triggering, where \tilde{n}_i and \tilde{p}_i are the numbers of inactive samples and features identified in the i -th triggering of ISS and IFS. n_0 and p_0 are the numbers of inactive samples and features.

For multi-class sparse SVMs, we let scaling ratio be $1 - \frac{(Kn-\tilde{n})(Kp-\tilde{p})}{K^2np}$, where $\tilde{n} = |\hat{\mathcal{R}} \cup \hat{\mathcal{L}}|$, $\tilde{p} = |\hat{\mathcal{F}}|$, and n and p are the sample size and the feature dimension of the data set. The rejection ratios of the i -th triggering of ISS and IFS are $\frac{\tilde{n}_i}{n_0}$ and $\frac{\tilde{p}_i}{p_0}$, respectively, where $p_0 = |\mathcal{F}|$, $n_0 = |\mathcal{R} \cup \mathcal{L}|$, $\tilde{p}_i = |\Delta\hat{\mathcal{F}}|$ and $\tilde{n}_i = |\Delta\hat{\mathcal{R}} \cup \Delta\hat{\mathcal{L}}|$ with $\Delta\hat{\mathcal{F}}$, $\Delta\hat{\mathcal{R}}$, and $\Delta\hat{\mathcal{L}}$ are the increments of $\hat{\mathcal{F}}$, $\hat{\mathcal{R}}$ and $\hat{\mathcal{L}}$ in the i -th triggering of the rules IFS and ISS. Please see Eqs. (22), (25), and (26) for the detailed definitions of $\Delta\hat{\mathcal{F}}$, $\Delta\hat{\mathcal{R}}$, and $\Delta\hat{\mathcal{L}}$.

Recall that, we can integrate SIFS with any solvers for problem (P*). In this experiment, we use Accelerated Proximal Stochastic Dual Coordinate Ascent (AProx-SDCA) (Shalev-Shwartz and Zhang, 2016) as a baseline, as it is one of the state-of-the-arts. We choose the state-of-art screening method for sparse SVMs in Shibagaki et al. 2016 as another baseline only in the experiments of sparse SVMs, since it cannot be applied in multi-class sparse SVMs. As we mentioned in the introduction section that screening differs greatly from feature selection methods, it is not appropriate to make comparisons with feature selection methods.

For each data set, we solve problem (P*) at a grid of turning parameter values. Specifically, we first compute β_{\max} by Lemma 4 and then select 10 values of β that are equally spaced at the logarithmic scale of β/β_{\max} from 1 to 0.05. Then, for each value of β , we first compute $\alpha_{\max}(\beta)$ by Lemma 4 and then select 100 values of α that are equally spaced at the logarithmic scale of $\alpha/\alpha_{\max}(\beta)$ from 1 to 0.01. Thus, for each data set, we solve problem (P*) at 1,000 pairs of parameter values in total. We write the code in C++ along with Eigen library for numerical computations. We perform all the computations on a single core of Intel(R) Core(TM) i7-5930K 3.50GHz, 128GB MEM.

6.1. Simulation Studies with Spare SVMs

We evaluate SIFS on 3 synthetic data sets named syn1, syn2, and syn3 with sample and feature size $(n, p) \in \{(10000, 1000), (10000, 10000), (1000, 10000)\}$. We present each data point as $\mathbf{x} = [\mathbf{x}_1; \mathbf{x}_2]$ with $\mathbf{x}_1 \in \mathbb{R}^{0.02p}$ and $\mathbf{x}_2 \in \mathbb{R}^{0.98p}$. We use Gaussian distributions $\mathcal{G}_1 = N(\mathbf{u}, 0.75\mathbf{I})$, $\mathcal{G}_2 = N(-\mathbf{u}, 0.75\mathbf{I})$, and $\mathcal{G}_3 = N(0, 1)$ to generate the data points, where $\mathbf{u} = 1.5\mathbf{1}$ and $\mathbf{I} \in \mathbb{R}^{0.02p \times 0.02p}$ is the identity matrix. To be precise, \mathbf{x}_1 for positive and negative points are sampled from \mathcal{G}_1 and \mathcal{G}_2 , respectively. For each entry in \mathbf{x}_2 , it has chance $\eta = 0.02$ to be sampled from \mathcal{G}_3 and chance $1 - \eta$ to be 0.

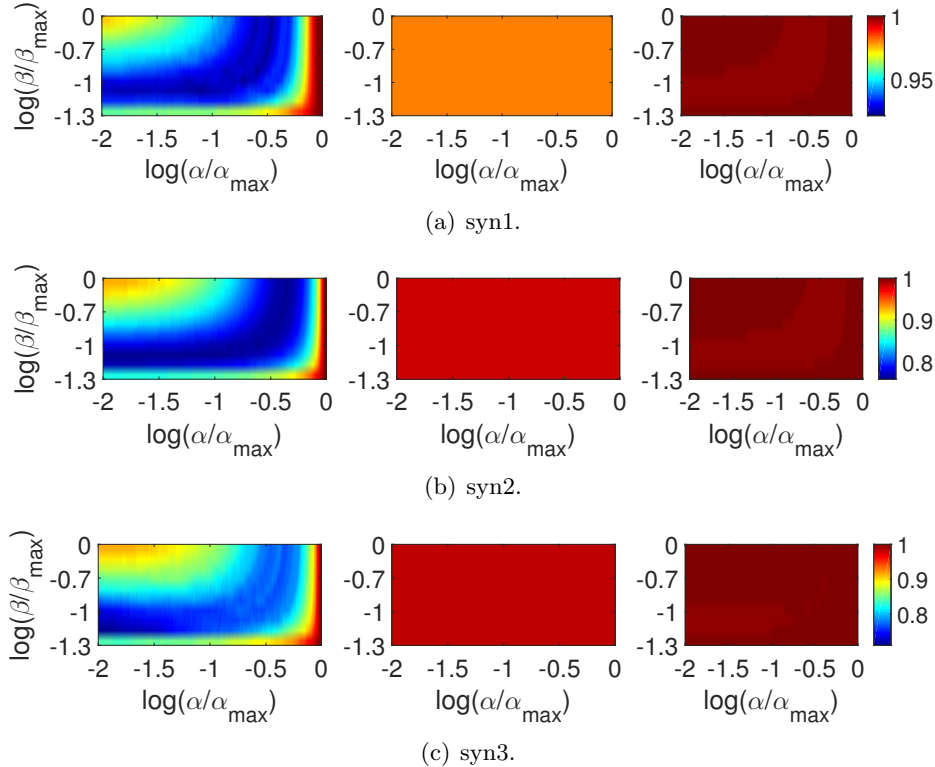


Figure 1: Scaling ratios of ISS, IFS, and SIFS (from left to right).

Fig. 1 shows the scaling ratios by ISS, IFS, and SIFS on the synthetic data sets at 1,000 parameter values. We can see that IFS is more effective in scaling problem size than ISS, with scaling ratios roughly 98% against 70 – 90%. Moreover, SIFS, which is an alternating application of IFS and ISS, significantly outperforms ISS and IFS, with scaling ratios roughly 99.9%. This high scaling ratios imply that SIFS can lead to a significant speedup.

Due to the space limitation, we only report the rejection ratios of SIFS on syn2. Other results can be found in the appendix. Fig. 2 shows that SIFS can identify most of the inactive features and samples. However, few features and samples are identified in the second and later triggerings of ISS and IFS. The reason may be that the task here is so simple that one triggering is enough.

Table 1 reports the running times of solver without and with IFS, ISS and SIFS for solving problem (P*) at 1,000 pairs of parameter values. We can see that SIFS leads to

significant speedups, that is, up to 76.8 times. Taking syn2 for example, without SIFS, the solver takes more than two hours to solve problem (P*) at 1,000 pairs of parameter values. However, combined with SIFS, the solver only needs less than three minutes for solving the same set of problems. From the theoretical analysis in Shalev-Shwartz and Zhang 2016 for AProx-SDCA, we can see that its computational complexity rises proportionately to the sample size n and the feature dimension p . From this theoretical result, we can see that the results in Fig. 1 are roughly consistent with the speedups we achieved shown in Table 1.

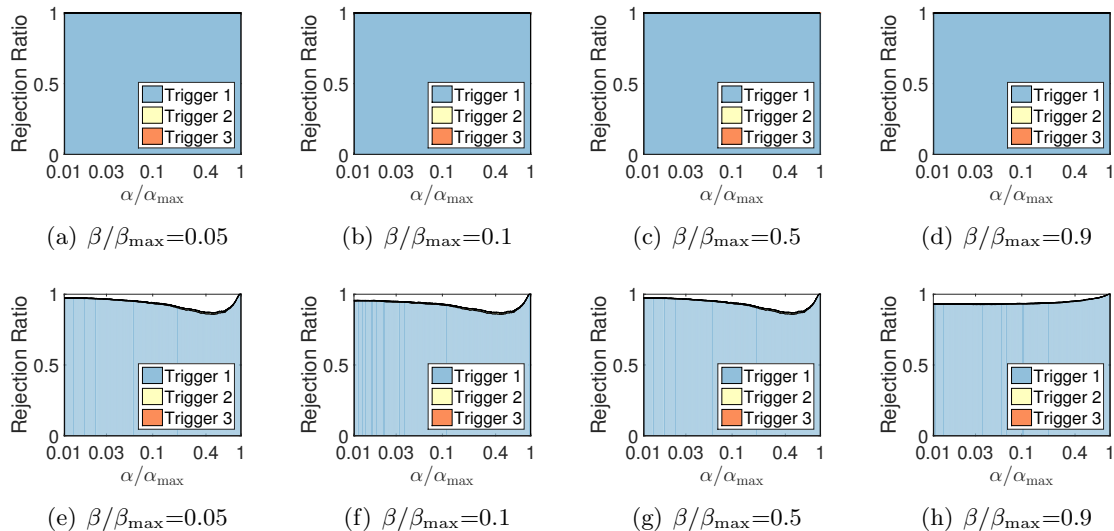


Figure 2: Rejection ratios of SIFS on syn 2 (first row: Feature Screening, second row: Sample Screening).

Data	Solver	ISS+Solver			IFS+Solver			SIFS+Solver		
		ISS	Solver	Speedup	IFS	Solver	Speedup	SIFS	Solver	Speedup
syn1	499.1	4.9	27.8	15.3	2.3	42.6	11.1	8.6	6.0	34.2
syn2	8749.9	24.9	1496.6	5.8	23.0	288.1	28.1	92.6	70.3	53.7
syn3	1279.7	2.0	257.1	4.9	2.2	33.4	36.0	7.2	9.5	76.8

Table 1: Running time (in seconds) for solving problem (P*) at 1,000 pairs of parameter values on three synthetic data sets.

6.2. Experiments with Sparse SVMs on Real Datasets

In this experiment, we evaluate the performance of SIFS on 5 large-scale real data sets: real-sim, rcv1-train, rcv1-test, url, and kddb, which are all collected from the project page of LibSVM (Chang and Lin, 2011). See Table 2 for a brief summary. We note that, the kddb data set has about 20 million samples with 30 million features.

Recall that, SIFS detects the inactive features and samples in a static manner, i.e., we perform SIFS only once before the training process and hence the size of the problem we need to perform optimization on is fixed. However, the method in Shibagaki et al. 2016 detects inactive features and samples in a dynamic manner (Bonney et al., 2014), i.e.,

they perform their method along with the training process and thus the size of the problem would keep decreasing during the iterations. Thus, comparing SIFS with this baseline in terms of the rejection ratios is inapplicable. We compare the performance of SIFS with the method in Shibagaki et al. 2016 in terms of speedup, i.e., the speedup gained by these two methods in solving problem (P*) at 1,000 pairs of parameter values. The code of the method in Shibagaki et al. 2016 is obtained from (<https://github.com/husk214/s3fs>).

Dataset	Feature size: p	Sample size: n	Classes
real-sim	20,958	72,309	2
rcv1-train	47,236	20,242	2
rcv1-test	47,236	677, 399	2
url	3,231,961	2,396,130	2
kddb	29,890,095	19,264,097	2

Table 2: Statistics of the binary real data sets.

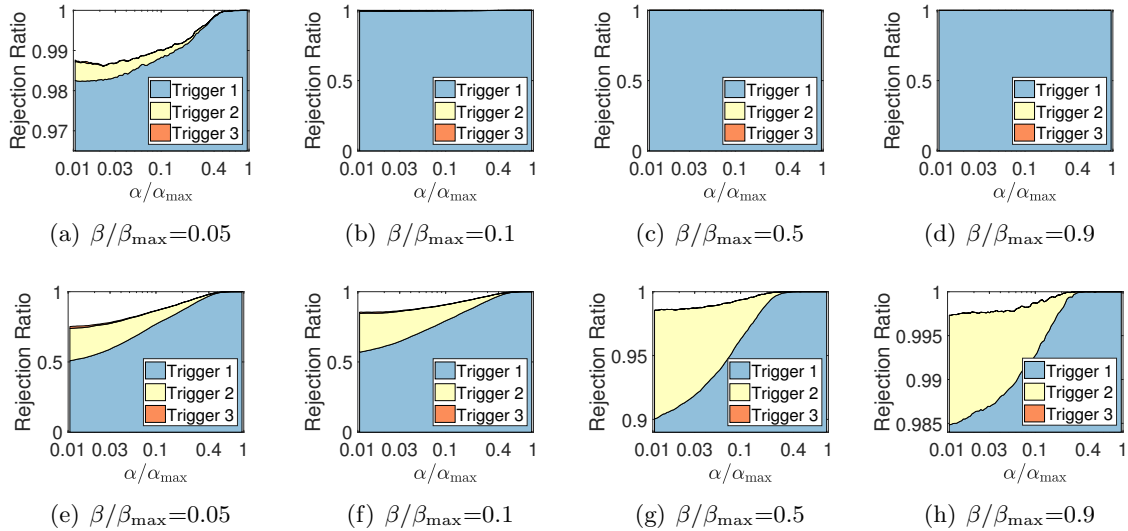


Figure 3: Rejection ratios of SIFS on real-sim when it performs **ISS first** (first row: Feature Screening, second row: Sample Screening).

Fig. 3 shows the rejection ratios of SIFS on the real-sim data set (other results are in the appendix). In Fig. 3, we can see that some inactive features and samples are identified in the 2nd and 3rd triggerings of ISS and IFS, which verifies the necessity of the alternating application of ISS and IFS. SIFS is efficient since it always stops in 3 times of triggering. In addition, most of (> 98%) the inactive features can be identified in the 1st triggering of IFS while identifying inactive samples needs to apply ISS two or more times. It may due to two reasons: 1) we run ISS first, which reinforces the capability of IFS due to the synergistic effect (see Sections 4.1 and 4.2), referring to the analysis below for further verification; 2) feature screening here may be easier than sample screening.

Table 3 reports the running times of solver without and with the method in Shibagaki et al. 2016 and SIFS for solving problem (P*) at 1,000 pairs of parameter values on real data

sets. The speedup gained by SIFS is up to 300 times on real-sim, rcv1-train and rcv1-test. Moreover, SIFS significantly outperforms the method in Shibagaki et al. 2016 in terms of speedup—by about 30 to 40 times faster on the aforementioned three data sets. For data sets url and kddb, we do not report the results of the solver as the sizes of the data sets are huge and the computational cost is prohibitive. Instead, we can see that the solver with SIFS is about 25 times faster than the solver with the method in Shibagaki et al. 2016 on both data sets url and kddb. Let us take the data set kddb as an example. The solver with SIFS takes about 13 hours to solve problem (P*) for 1,000 pairs of parameter values, while the solver with the method in Shibagaki et al. 2016 needs 11 days to finish the same task.

Data Set	Solver	Method in Shibagaki et al. 2016+Solver			SIFS+Solver		
		Screen	Solver	Speedup	Screen	Solver	Speedup
real-sim	3.93E+04	24.10	4.94E+03	7.91	60.01	140.25	195.00
rcv1-train	2.98E+04	10.00	3.73E+03	7.90	27.11	80.11	277.10
rcv1-test	1.10E+06	398.00	1.35E+05	8.10	1.17E+03	2.55E+03	295.11
url	>3.50E+06	3.18E+04	8.60E+05	>4.00	7.66E+03	2.91E+04	> 100
kddb	>5.00E+06	4.31E+04	1.16E+06	>4.00	1.10E+04	3.6E+04	> 100

Table 3: Running time (in seconds) for solving problem (P*) at 1,000 pairs of parameter values on five real data sets.




Figure 4: Rejection ratios of SIFS on real-sim when it performs **IFS first** (first row: Feature Screening, second row: Sample Screening).

Verification of the Synergy Effect In Fig. 3, SIFS performs ISS (sample screening) first, while in Fig. 4, it performs IFS (feature screening) first. All the rejection ratios of the 1st triggering of IFS in Fig. 3(a)-3(d) where SIFS performs ISS first are much higher than (at least equal to) those in Fig. 4(a)-4(d) where SIFS performs IFS first. In turn, all the rejection ratios of the 1st triggering of ISS in Fig. 4(e)-4(h) where SIFS performs IFS

first are also much higher than those in Fig. 3(e)-3(h) where SIFS performs ISS first. This demonstrates that the screening result of ISS can reinforce the capability of IFS and vice versa, which is the so called synergistic effect. At last, in Fig. 4 and Fig. 3, we can see that the overall rejection ratios at the end of SIFS are exactly the same. Hence, no matter which rule (ISS or IFS) we perform first in SIFS, SIFS has the same screening performances in the end, which verifies the conclusion we present in Theorem 12.

Performance in Solving Single Problems In the experiments above, we solve problem (P*) at a grid of turning parameter values. This setting is meaningful and it arises naturally in various cases, such as cross validation (Kohavi, 1995) and feature selection (Meinshausen and Bühlmann, 2010; Yang et al., 2015). Therefore, the results above demonstrate that our method would be helpful in such real applications. We notice that sometimes one may be interested in the model at a specific parameter value pair. To this end, we now evaluate the performance of our method in solving a single problem. Due to the space limitation, we solve sparse SVM at 16 specific parameter values pairs on the real-sim data set for examples, To be precise, $\beta \in \{0.05, 0.10, 0.5, 0.9\}$ and each β has 4 values of α satisfying $\alpha/\alpha_{\max}(\beta) \in \{0.05, 0.1, 0.5, 0.9\}$. For each (α_i, β_i) , we construct a parameter value path $\{(\alpha_{j,i}, \beta_i), j = 0, \dots, M\}$ with $\alpha_{j,i}$ equally spaced at the logarithmic scale of $\alpha_{j,i}/\alpha_{\max}(\beta_i)$ between 1 and $\alpha_i/\alpha_{\max}(\beta_i)$, which implies that $\alpha_{\max}(\beta_i) = \alpha_{0,i} > \alpha_{1,i} > \dots > \alpha_{M,i} = \alpha_i$. Then we use AProx-SDCA integrated with SIFS to solve problem (P*) at the parameter value pairs on the path one by one and report the total time cost. It can be expected that for α_i far from $\alpha_{\max}(\beta_i)$, we need more points on the parameter value path in order to obtain higher rejection ratios and more significant speedups of SIFS and otherwise we need fewer ones. In this experiment, we fix $M = 50$ at each α_i for convenience. For comparison, we solve the problem at (α_i, β_i) using AProx-SDCA with random initializations, which would be faster than solving all the problems on the path. Table 4 reports the speedups achieved by SIFS at these parameter value pairs. It shows that SIFS can still obtain significant speedups. In addition, compared with the results in Tables 1 and 3, we can see that our method is much better at solving a problem sequence than solving a single problem.

$\alpha/\alpha_{\max}(\beta)$ \ β/β_{\max}	0.05	0.10	0.50	0.90
0.05	6.31	7.14	21.04	146.66
0.10	6.14	7.18	21.83	156.63
0.50	5.75	7.18	20.99	154.60
0.90	6.39	7.66	23.16	173.56

Table 4: The speedups achieved by SIFS in solving single problems

6.3. Experiments with Multi-class Sparse SVMs

We evaluate SIFS for multi-class Sparse SVMs on 3 synthetic data sets (syn-multi1, syn-multi2, and syn-multi3) and 2 real data sets (news20 and rcv1-multiclass). Their statistics are given in Table 5.

The synthetic data sets are generated in a similar way as we did in Section 6.1. Each of them has $K = 5$ classes and each class has n/K samples. The data points of them can be written as $\mathbf{x} = [\mathbf{x}_1; \mathbf{x}_2]$, where $\mathbf{x}_1 = [\mathbf{x}_1^1; \dots; \mathbf{x}_1^K] \in \mathbb{R}^{0.02p}$ with $\mathbf{x}_1^k \in \mathbb{R}^{0.02p/K}$ and $\mathbf{x}_2 \in \mathbb{R}^{0.98p}$. If \mathbf{x} belongs to the k -th class, then \mathbf{x}_1^k is sampled from a Gaussian distribution

$\mathcal{G} = N(\mathbf{u}, 0.75\mathbf{I})$ with $\mathbf{u} = 1.5\mathbf{1}$, $\mathbf{I} \in \mathbb{R}^{(0.02p/K) \times (0.02p/K)}$ and other components in \mathbf{x}_1 are from $N(0, 1)$. Each entry of \mathbf{x}_2 has chance $\eta = 0.2$ to be sampled from distribution $N(0, 1)$ and chance $1 - \eta$ to be 0.

Dataset	Feature size: p	Sample size: n	Classes
syn-multi1	1,000	10,000	5
syn-multi2	10,000	10,000	5
syn-multi3	10,000	1,000	5
news20	3,993	62,060	20
rcv1-multiclass	47,236	15,564	53

Table 5: Statistics of the multi-class data sets.

Fig. 5 shows the scaling ratios of ISS, IFS, and SIFS on the synthetic data sets at 1,000 pairs of parameter values. Similar to the results in sparse SVMs, SIFS totally surpasses IFS and ISS, with scaling ratios larger than 98% at any parameter value pair against 70 – 90%. Thus, we can expect significant speedups by integrating SIFS with the solver.

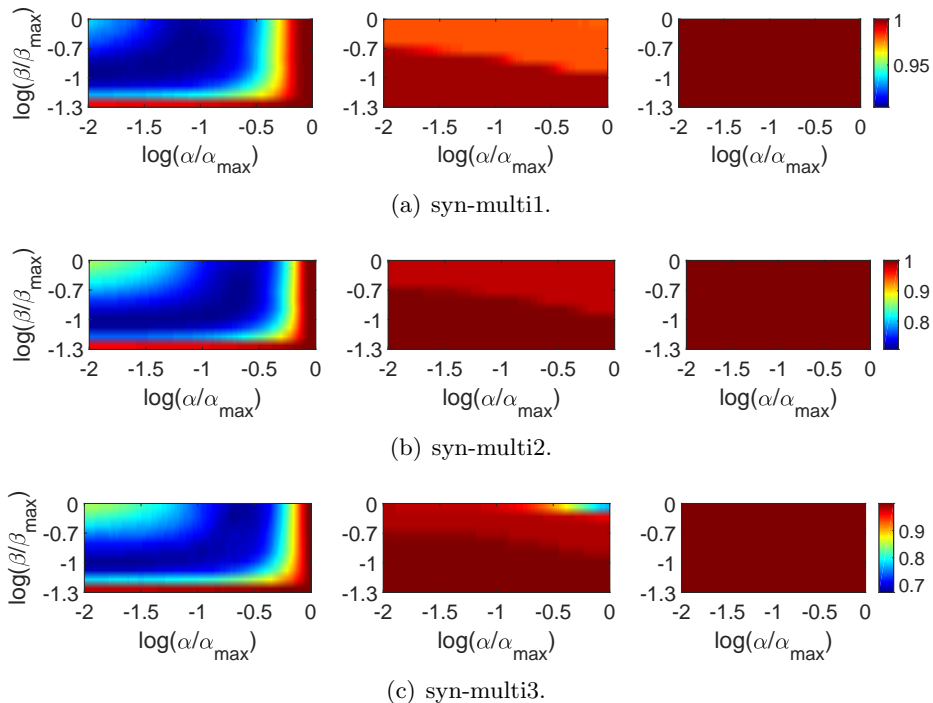


Figure 5: Scaling ratios of ISS, IFS, and SIFS (from left to right).

Fig. 6 and 7 present the rejection ratios of SIFS on the news20 data set (the results of other data sets can be found in the appendix). It indicates that, by alternatively applying IFS and ISS, we can finally identify most of the inactive samples ($> 95\%$) and features ($> 99\%$). The synergistic effect between IFS and ISS can also be found in these figures.

As a result of excellent performance of SIFS in identifying inactive features and samples, we observe significant speedups gained by SIFS in Table 6, which are up to 200 times on news20 and 300 times on syn-multi3. Specifically, on news20, the solver without any

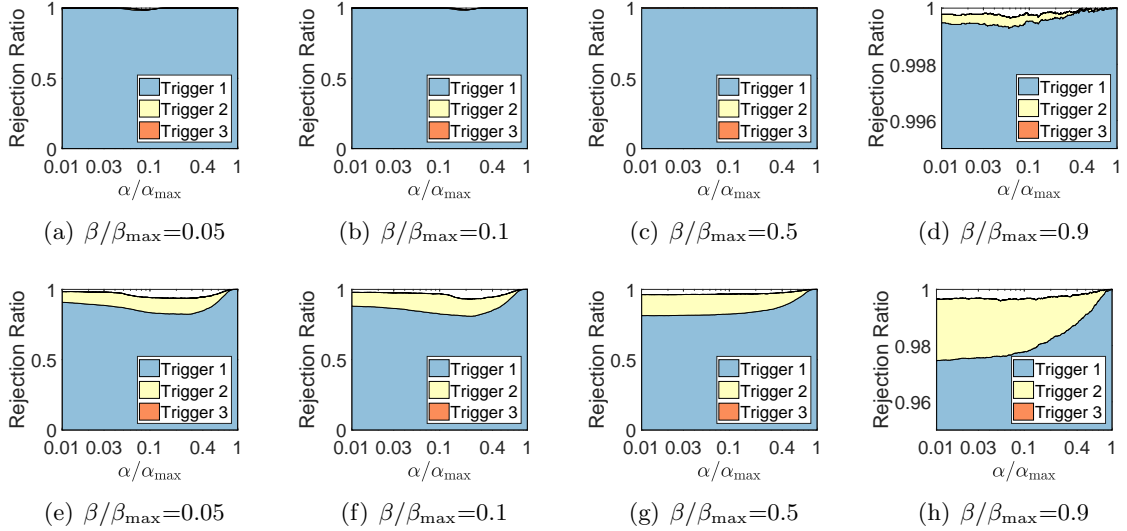


Figure 6: Rejection ratios of SIFS on syn-multi2 (first row: Feature Screening, second row: Sample Screening).

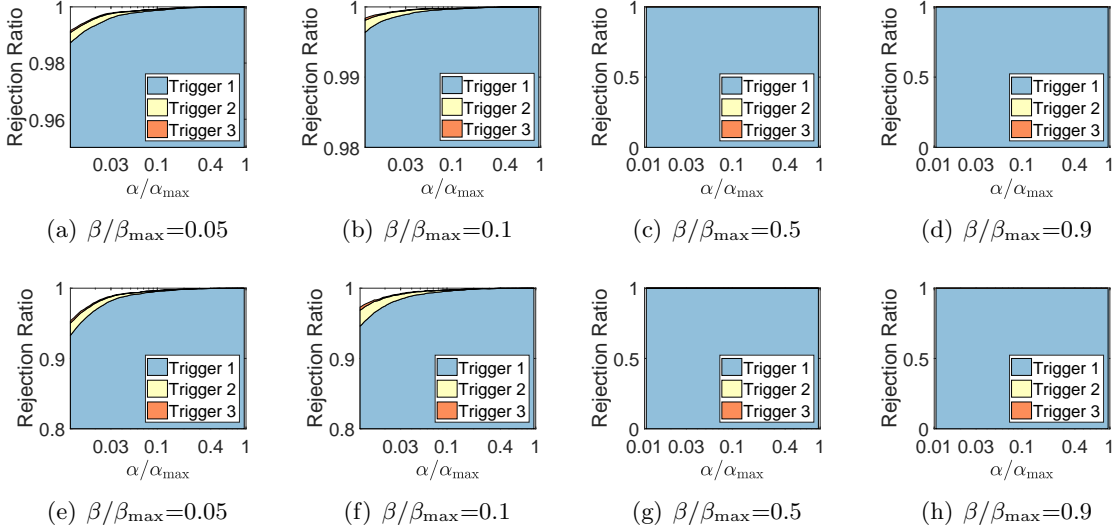


Figure 7: Rejection ratios of SIFS on news20 (first row: Feature Screening, second row: Sample Screening).

screening method takes about 97 hours to solve the problem at 1,000 pairs of parameter values; however, integrated with SIFS, the solver only needs less than 0.5 hour for the same task. Moreover, Table 6 also indicates that the speedup gained by SIFS is much more significant than that gained by IFS and ISS. This demonstrates again the great benefit of alternatively applying IFS and ISS in SIFS. At last, we notice that the speedup gained by

SIFS and IFS on syn-multi3 is far greater than those gained on syn-multi1 and syn-multi2. The reason is that the convergence rate of the solver in our problem may depend on the condition number of the data matrix. The dimension of \mathbf{w} in syn-multi3 is much larger than its sample size (50,000 over 1,000), which leads to a large condition number and a bad convergence rate. However, SIFS and IFS can remove the inactive features from the model, which can greatly improve the convergence rate of the solver by decreasing the condition number. This point can also be supported by the fact in the experiment on syn-multi3 that the solver with SIFS or IFS needs much fewer iterations to converge to the optima than that without SIFS or IFS.

Data	Solver	ISS+Solver			IFS+Solver			SIFS+Solver		
		ISS	Solver	Speedup	IFS	Solver	Speedup	SIFS	Solver	Speedup
syn-multi1	9009	72	1233	6.9	90	388	18.9	178	83	34.5
syn-multi2	105484	412	27661	3.8	920	2606	29.9	1936	893	37.3
syn-multi3	78373	44	18341	4.3	84	337	186.2	203	59	299.9
news20	350277	283	13950	24.6	1000	65175	5.3	1485	290	197.3
rcv1-multiclass	396205	793	44639	8.7	1620	33191	11.4	2850	1277	96.0

Table 6: Running time (in seconds) for solving problem (m-P*) at 1,000 pairs of parameter values on synthetic and real data sets.

7. Conclusions

In this paper, we proposed a novel data reduction method SIFS to simultaneously identify inactive features and samples for sparse SVMs. Our major contribution is a novel framework for an accurate estimation of the primal and dual optima based on strong convexity. To the best of our knowledge, the proposed SIFS is the first static screening method that is able to simultaneously identify inactive features and samples for sparse SVMs. An appealing feature of SIFS is that all detected features and samples are guaranteed to be irrelevant to the outputs. Thus, the model learned on the reduced data is exactly identical to that learned on the full data. To show the flexibility of the proposed SIFS, we extended it to multi-class sparse SVMs. The experimental results demonstrate that, for both sparse SVMs and multi-class sparse SVMs, SIFS can dramatically reduce the problem size and the resulting speedup can be orders of magnitude. We plan to generalize SIFS to more complicated models, e.g., SVMs with a structured sparsity-inducing penalty.

Acknowledgments

This work was supported by the National Basic Research Program of China (973 Program) under Grant 2013CB336500, National Natural Science Foundation of China under Grant 61233011, and National Youth Top-notch Talent Support Program.

Appendix A. Detailed Proofs and More Experimental Results

In this appendix, we first present the detailed proofs of all the theorems in the main paper and then report the rest experimental results which are omitted in the experimental section.

A.1. Proof for Lemma 3

Proof of Lemma 3:

- 1) It is the immediate conclusion of the analysis above.
- 2) After feature screening, the primal problem (P*) is scaled into:

$$\min_{\tilde{\mathbf{w}} \in \mathbb{R}^{|\hat{\mathcal{F}}^c|}} \frac{\alpha}{2} \|\tilde{\mathbf{w}}\|^2 + \beta \|\tilde{\mathbf{w}}\|_1 + \frac{1}{n} \sum_{i=1}^n \ell(1 - \langle [\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c}, \tilde{\mathbf{w}} \rangle). \quad (\text{scaled-}P^*-1)$$

Thus, we can easily derive out the dual problem of (scaled- P^* -1):

$$\min_{\tilde{\theta} \in [0, \alpha]^n} \tilde{D}(\tilde{\theta}; \alpha, \beta) = \frac{1}{2\alpha} \|\mathcal{S}_\beta(\frac{1}{n} \hat{\mathcal{F}}^c[\bar{\mathbf{X}}] \tilde{\theta})\|^2 + \frac{\gamma}{2n} \|\tilde{\theta}\|^2 - \frac{1}{n} \langle \mathbf{1}, \tilde{\theta} \rangle, \quad (\text{scaled-}D^*-1)$$

and the KKT conditions:

$$\begin{aligned} \tilde{\mathbf{w}}^*(\alpha, \beta) &= \frac{1}{\alpha} \mathcal{S}_\beta(\frac{1}{n} \hat{\mathcal{F}}^c[\bar{\mathbf{X}}] \tilde{\theta}^*(\alpha, \beta)), & (\text{scaled-KKT-1}) \\ [\tilde{\theta}^*(\alpha, \beta)]_i &= \begin{cases} 0, & \text{if } 1 - \langle [\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c}, \tilde{\mathbf{w}}^*(\alpha, \beta) \rangle < 0, \\ \frac{1}{\gamma} (1 - \langle [\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c}, \tilde{\mathbf{w}}^*(\alpha, \beta) \rangle), & \text{if } 0 \leq 1 - \langle [\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c}, \tilde{\mathbf{w}}^*(\alpha, \beta) \rangle \leq \gamma, \\ 1, & \text{if } 1 - \langle [\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c}, \tilde{\mathbf{w}}^*(\alpha, \beta) \rangle > \gamma, \end{cases} & (\text{scaled-KKT-2}) \end{aligned}$$

Then, it is obvious that $\tilde{\mathbf{w}}^*(\alpha, \beta) = [\mathbf{w}^*(\alpha, \beta)]_{\hat{\mathcal{F}}^c}$, since essentially, problem (scaled- P^* -1) can be derived by substituting 0 to the weights for the eliminated features in problem (P*) and optimizing over the rest weights.

Since the solutions $\mathbf{w}^*(\alpha, \beta)$ and $\theta^*(\alpha, \beta)$ satisfy the conditions (KKT-1) and (KKT-2) and $\langle [\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c}, \tilde{\mathbf{w}}^*(\alpha, \beta) \rangle = \langle \bar{\mathbf{x}}_i, \mathbf{w}^*(\alpha, \beta) \rangle$ for all i , we know that $\tilde{\mathbf{w}}^*(\alpha, \beta)$ and $\theta^*(\alpha, \beta)$ satisfy the conditions (scaled-KKT-1) and (scaled-KKT-2). Thus they are the solutions of problems (scaled- P^* -1) and (scaled- D^* -1). Then, due to the uniqueness of the solution of problem (scaled- D^* -1), we have

$$\theta^*(\alpha, \beta) = \tilde{\theta}^*(\alpha, \beta). \quad (27)$$

From 1) we have $[\tilde{\theta}^*(\alpha, \beta)]_{\hat{\mathcal{R}}^c} = 0$ and $[\tilde{\theta}^*(\alpha, \beta)]_{\hat{\mathcal{L}}^c} = 1$. Therefore, from the dual problem (scaled- D^*), we can see that $[\theta^*(C, \alpha)]_{\hat{\mathcal{D}}^c}$ can be recovered from the following problem:

$$\min_{\hat{\theta} \in [0, 1]^{|\hat{\mathcal{D}}^c|}} \frac{1}{2\alpha} \|\mathcal{S}_\beta(\frac{1}{n} \hat{\mathbf{G}}_1 \hat{\theta} + \frac{1}{n} \hat{\mathbf{G}}_2 \mathbf{1})\|^2 + \frac{\gamma}{2n} \|\hat{\theta}\|^2 - \frac{1}{n} \langle \mathbf{1}, \hat{\theta} \rangle.$$

Since $[\tilde{\theta}^*(\alpha, \beta)]_{\hat{\mathcal{D}}^c} = [\theta^*(\alpha, \beta)]_{\hat{\mathcal{D}}^c}$, the proof is therefore completed. \blacksquare

A.2. Proof for Lemma 4

Proof of Lemma 4:

(i) We prove this lemma by verifying that the solutions $\mathbf{w}^*(\alpha, \beta) = \mathbf{0}$ and $\theta^*(\alpha, \beta) = \mathbf{1}$ satisfy the conditions (KKT-1) and (KKT-2).

Firstly, since $\beta \geq \beta_{\max} = \|\frac{1}{n}\bar{\mathbf{X}}\mathbf{1}\|_{\infty}$, we have $\mathcal{S}_{\beta}(\frac{1}{n}\bar{\mathbf{X}}\mathbf{1}) = 0$. Thus $\mathbf{w}^*(\alpha, \beta) = \mathbf{0}$ and $\theta^*(\alpha, \beta) = \mathbf{1}$ satisfy the condition (KKT-1).

Then, for all $i \in [n]$, we have

$$1 - \langle \bar{\mathbf{x}}_i, \mathbf{w}^*(\alpha, \beta) \rangle = 1 - 0 > \gamma.$$

Thus, $\mathbf{w}^*(\alpha, \beta) = \mathbf{0}$ and $\theta^*(\alpha, \beta) = \mathbf{1}$ satisfy the condition (KKT-2). Hence, they are the solutions of the primal problem (P*) and the dual problem (D*), respectively.

(ii) Similar to the proof of (i), we prove this by verifying that the solutions $\mathbf{w}^*(\alpha, \beta) = \frac{1}{\alpha}\mathcal{S}_{\beta}(\frac{1}{n}\bar{\mathbf{X}}\theta^*(\alpha, \beta))$ and $\theta^*(\alpha, \beta) = \mathbf{1}$ satisfy the conditions (KKT-1) and (KKT-2).

1. **Case 1:** $\alpha_{\max}(\beta) \leq 0$. Then for all $\alpha > 0$, we have

$$\begin{aligned} & \min_{i \in [n]} \{1 - \langle \bar{\mathbf{x}}_i, \mathbf{w}^*(\alpha, \beta) \rangle\} \\ &= \min_{i \in [n]} \left\{1 - \frac{1}{\alpha} \langle \bar{\mathbf{x}}_i, \mathcal{S}_{\beta}(\frac{1}{n}\bar{\mathbf{X}}\theta^*(\alpha, \beta)) \rangle\right\} = \min_{i \in [n]} \left\{1 - \frac{1}{\alpha} \langle \bar{\mathbf{x}}_i, \mathcal{S}_{\beta}(\frac{1}{n}\bar{\mathbf{X}}\mathbf{1}) \rangle\right\} \\ &= 1 - \frac{1}{\alpha} \max_{i \in [n]} \langle \bar{\mathbf{x}}_i, \mathcal{S}_{\beta}(\frac{1}{n}\bar{\mathbf{X}}\mathbf{1}) \rangle = 1 - (1 - \gamma) \frac{1}{\alpha} \alpha_{\max}(\beta) \\ &\geq 1 > \gamma. \end{aligned}$$

Then, $\mathcal{L} = [n]$, and $\mathbf{w}^*(\alpha, \beta) = \frac{1}{\alpha}\mathcal{S}_{\beta}(\frac{1}{n}\bar{\mathbf{X}}\theta^*(\alpha, \beta))$ and $\theta^*(\alpha, \beta) = \mathbf{1}$ satisfy the conditions (KKT-1) and (KKT-2). Hence, they are the optimal solution of the primal and dual problems (P*) and (D*).

2. **Case 2:** $\alpha_{\max}(\beta) > 0$. Then for any $\alpha \geq \alpha_{\max}(\beta)$, we have

$$\begin{aligned} & \min_{i \in [n]} \{1 - \langle \bar{\mathbf{x}}_i, \mathbf{w}^*(\alpha, \beta) \rangle\} \\ &= \min_{i \in [n]} \left\{1 - \frac{1}{\alpha} \langle \bar{\mathbf{x}}_i, \mathcal{S}_{\beta}(\frac{1}{n}\bar{\mathbf{X}}\theta^*(\alpha, \beta)) \rangle\right\} = \min_{i \in [n]} \left\{1 - \frac{1}{\alpha} \langle \bar{\mathbf{x}}_i, \mathcal{S}_{\beta}(\frac{1}{n}\bar{\mathbf{X}}\mathbf{1}) \rangle\right\} \\ &= 1 - \frac{1}{\alpha} \max_{i \in [n]} \langle \bar{\mathbf{x}}_i, \mathcal{S}_{\beta}(\frac{1}{n}\bar{\mathbf{X}}\mathbf{1}) \rangle = 1 - (1 - \gamma) \frac{1}{\alpha} \alpha_{\max}(\beta) \geq 1 - (1 - \gamma) = \gamma. \end{aligned}$$

Thus, $\mathcal{E} \cup \mathcal{L} = [n]$, and $\mathbf{w}^*(\alpha, \beta) = \frac{1}{\alpha}\mathcal{S}_{\beta}(\frac{1}{n}\bar{\mathbf{X}}\theta^*(\alpha, \beta))$ and $\theta^*(\alpha, \beta) = \mathbf{1}$ satisfy the conditions (KKT-1) and (KKT-2). Hence, they are the optimal solution of the primal and dual problems (P*) and (D*).

The proof is complete. ■

A.3. Proof for Lemma 5

Proof of Lemma 5:

Due to the α -strong convexity of the objective $P(\mathbf{w}; \alpha, \beta)$, we have

$$\begin{aligned} P(\mathbf{w}^*(\alpha_0, \beta_0); \alpha, \beta_0) &\geq P(\mathbf{w}^*(\alpha, \beta_0); \alpha, \beta_0) + \frac{\alpha}{2} \|\mathbf{w}^*(\alpha_0, \beta_0) - \mathbf{w}^*(\alpha, \beta_0)\|^2, \\ P(\mathbf{w}^*(\alpha, \beta_0); \alpha_0, \beta_0) &\geq P(\mathbf{w}^*(\alpha_0, \beta_0); \alpha_0, \beta_0) + \frac{\alpha_0}{2} \|\mathbf{w}^*(\alpha_0, \beta_0) - \mathbf{w}^*(\alpha, \beta_0)\|^2, \end{aligned}$$

which are equivalent to

$$\begin{aligned} &\frac{\alpha}{2} \|\mathbf{w}^*(\alpha_0, \beta_0)\|^2 + \beta_0 \|\mathbf{w}^*(\alpha_0, \beta_0)\|_1 + L(\mathbf{w}^*(\alpha_0, \beta_0)) \\ &\geq \frac{\alpha}{2} \|\mathbf{w}^*(\alpha, \beta_0)\|^2 + \beta_0 \|\mathbf{w}^*(\alpha, \beta_0)\|_1 + L(\mathbf{w}^*(\alpha, \beta_0)) + \frac{\alpha}{2} \|\mathbf{w}^*(\alpha_0, \beta_0) - \mathbf{w}^*(\alpha, \beta_0)\|^2, \\ &\frac{\alpha_0}{2} \|\mathbf{w}^*(\alpha, \beta_0)\|^2 + \beta_0 \|\mathbf{w}^*(\alpha, \beta_0)\|_1 + L(\mathbf{w}^*(\alpha, \beta_0)) \\ &\geq \frac{\alpha_0}{2} \|\mathbf{w}^*(\alpha_0, \beta_0)\|^2 + \beta_0 \|\mathbf{w}^*(\alpha_0, \beta_0)\|_1 + L(\mathbf{w}^*(\alpha_0, \beta_0)) + \frac{\alpha_0}{2} \|\mathbf{w}^*(\alpha_0, \beta_0) - \mathbf{w}^*(\alpha, \beta_0)\|^2. \end{aligned}$$

Adding the above two inequalities together, we obtain

$$\begin{aligned} \frac{\alpha - \alpha_0}{2} \|\mathbf{w}^*(\alpha_0, \beta_0)\|^2 &\geq \frac{\alpha - \alpha_0}{2} \|\mathbf{w}^*(\alpha, \beta_0)\|^2 + \frac{\alpha_0 + \alpha}{2} \|\mathbf{w}^*(\alpha_0, \beta_0) - \mathbf{w}^*(\alpha, \beta_0)\|^2 \\ \Rightarrow \|\mathbf{w}^*(\alpha, \beta_0) - \frac{\alpha_0 + \alpha}{2\alpha} \mathbf{w}^*(\alpha_0, \beta_0)\|^2 &\leq \frac{(\alpha - \alpha_0)^2}{4\alpha^2} \|\mathbf{w}^*(\alpha_0, \beta_0)\|^2. \end{aligned} \quad (28)$$

Substituting the prior that $[\mathbf{w}^*(\alpha, \beta_0)]_{\hat{\mathcal{F}}} = 0$ into Eq. (28), we obtain

$$\begin{aligned} &\|[\mathbf{w}^*(\alpha, \beta_0)]_{\hat{\mathcal{F}}^c} - \frac{\alpha_0 + \alpha}{2\alpha} [\mathbf{w}^*(\alpha_0, \beta_0)]_{\hat{\mathcal{F}}^c}\|^2 \\ &\leq \frac{(\alpha - \alpha_0)^2}{4\alpha^2} \|\mathbf{w}^*(\alpha_0, \beta_0)\|^2 - \frac{(\alpha_0 + \alpha)^2}{4\alpha^2} \|[\mathbf{w}^*(\alpha_0, \beta_0)]_{\hat{\mathcal{F}}}\|^2. \end{aligned}$$

The proof is complete. ■

A.4. Proof for Lemma 6

Proof of Lemma 6: Firstly, we need to extend the definition of $D(\theta; \alpha, \beta)$ to \mathbb{R}^n :

$$\tilde{D}(\theta; \alpha, \beta) = \begin{cases} D(\theta; \alpha, \beta), & \text{if } \theta \in [0, 1]^n, \\ +\infty, & \text{otherwise.} \end{cases} \quad (29)$$

Due to the strong convexity of objective $\tilde{D}(\theta; \alpha, \beta)$, we have

$$\begin{aligned} \tilde{D}(\theta^*(\alpha_0, \beta_0), \alpha, \beta_0) &\geq \tilde{D}(\theta^*(\alpha, \beta_0), \alpha, \beta_0) + \frac{\gamma}{2n} \|\theta^*(\alpha_0, \beta_0) - \theta^*(\alpha, \beta_0)\|^2, \\ \tilde{D}(\theta^*(\alpha, \beta_0), \alpha_0, \beta_0) &\geq \tilde{D}(\theta^*(\alpha_0, \beta_0), \alpha_0, \beta_0) + \frac{\gamma}{2n} \|\theta^*(\alpha_0, \beta_0) - \theta^*(\alpha, \beta_0)\|^2. \end{aligned}$$

Since $\theta^*(\alpha_0, \beta_0), \theta^*(\alpha, \beta_0) \in [0, \mathbf{1}]^n$, the above inequalities are equivalent to

$$\begin{aligned}
 & \frac{1}{2\alpha} f_{\beta_0}(\theta^*(\alpha_0, \beta_0)) + \frac{\gamma}{2n} \|\theta^*(\alpha_0, \beta_0)\|^2 - \frac{1}{n} \langle \mathbf{1}, \theta^*(\alpha_0, \beta_0) \rangle \\
 & \geq \frac{1}{2\alpha} f_{\beta_0}(\theta^*(\alpha, \beta_0)) + \frac{\gamma}{2n} \|\theta^*(\alpha, \beta_0)\|^2 - \frac{1}{n} \langle \mathbf{1}, \theta^*(\alpha, \beta_0) \rangle + \frac{\gamma}{2n} \|\theta^*(\alpha_0, \beta_0) - \theta^*(\alpha, \beta_0)\|^2, \\
 & \frac{1}{2\alpha_0} f_{\beta_0}(\theta^*(\alpha, \beta_0)) + \frac{\gamma}{2n} \|\theta^*(\alpha, \beta_0)\|^2 - \frac{1}{n} \langle \mathbf{1}, \theta^*(\alpha, \beta_0) \rangle \\
 & \geq \frac{1}{2\alpha_0} f_{\beta_0}(\theta^*(\alpha_0, \beta_0)) + \frac{\gamma}{2n} \|\theta^*(\alpha_0, \beta_0)\|^2 - \frac{1}{n} \langle \mathbf{1}, \theta^*(\alpha_0, \beta_0) \rangle + \frac{\gamma}{2n} \|\theta^*(\alpha_0, \beta_0) - \theta^*(\alpha, \beta_0)\|^2.
 \end{aligned}$$

Adding the above two inequalities, we obtain

$$\begin{aligned}
 & \frac{\gamma(\alpha - \alpha_0)}{2n} \|\theta^*(\alpha_0, \beta_0)\|^2 - \frac{\alpha - \alpha_0}{n} \langle \mathbf{1}, \theta^*(\alpha_0, \beta_0) \rangle \\
 & \geq \frac{\gamma(\alpha - \alpha_0)}{2n} \|\theta^*(\alpha, \beta_0)\|^2 - \frac{\alpha - \alpha_0}{n} \langle \mathbf{1}, \theta^*(\alpha, \beta_0) \rangle + \frac{\gamma(\alpha_0 + \alpha)}{2n} \|\theta^*(\alpha_0, \beta_0) - \theta^*(\alpha, \beta_0)\|^2.
 \end{aligned}$$

That is equivalent to

$$\begin{aligned}
 & \|\theta^*(\alpha, \beta_0)\|^2 - \left\langle \frac{\alpha - \alpha_0}{\gamma\alpha} \mathbf{1} + \frac{\alpha_0 + \alpha}{\alpha} \theta^*(\alpha_0, \beta_0), \theta^*(\alpha, \beta_0) \right\rangle \\
 & \leq -\frac{\alpha_0}{\alpha} \|\theta^*(\alpha_0, \beta_0)\|^2 - \frac{\alpha - \alpha_0}{\gamma\alpha} \langle \mathbf{1}, \theta^*(\alpha_0, \beta_0) \rangle.
 \end{aligned} \tag{30}$$

That is

$$\|\theta^*(\alpha, \beta_0) - \left(\frac{\alpha - \alpha_0}{2\gamma\alpha} \mathbf{1} + \frac{\alpha_0 + \alpha}{2\alpha} \theta^*(\alpha_0, \beta_0) \right)\|^2 \leq \left(\frac{\alpha - \alpha_0}{2\alpha} \right)^2 \|\theta^*(\alpha_0, \beta_0) - \frac{1}{\gamma} \mathbf{1}\|^2. \tag{31}$$

Substituting the priors that $[\theta^*(\alpha, \beta_0)]_{\hat{\mathcal{R}}} = 0$ and $[\theta^*(\alpha, \beta_0)]_{\hat{\mathcal{L}}} = 1$ into Eq. (31), we have

$$\begin{aligned}
 & \|[\theta^*(\alpha, \beta_0)]_{\hat{\mathcal{D}}^c} - \left(\frac{\alpha - \alpha_0}{2\gamma\alpha} \mathbf{1} + \frac{\alpha_0 + \alpha}{2\alpha} [\theta^*(\alpha_0, \beta_0)]_{\hat{\mathcal{D}}^c} \right)\|^2 \\
 & \leq \left(\frac{\alpha - \alpha_0}{2\alpha} \right)^2 \|\theta^*(\alpha_0, \beta_0) - \frac{1}{\gamma} \mathbf{1}\|^2 - \left\| \frac{(2\gamma - 1)\alpha + \alpha_0}{2\gamma\alpha} \mathbf{1} - \frac{\alpha_0 + \alpha}{2\alpha} [\theta^*(\alpha_0, \beta_0)]_{\hat{\mathcal{L}}} \right\|^2 \\
 & \quad - \left\| \frac{\alpha - \alpha_0}{2\gamma\alpha} \mathbf{1} + \frac{\alpha_0 + \alpha}{2\alpha} [\theta^*(\alpha_0, \beta_0)]_{\hat{\mathcal{R}}} \right\|^2.
 \end{aligned}$$

The proof is complete. ■

A.5. Proof for Lemma 8

Before the proof of Lemma 8, we should prove that the optimization problem in (2) is equivalent to

$$s^i(\alpha, \beta_0) = \max_{\theta \in \Theta} \left\{ \frac{1}{n} |\langle [\bar{\mathbf{x}}^i]_{\hat{\mathcal{D}}^c}, \theta \rangle + \langle [\bar{\mathbf{x}}^i]_{\hat{\mathcal{L}}}, \mathbf{1} \rangle| \right\}, i \in \hat{\mathcal{F}}^c. \tag{32}$$

To avoid notational confusion, we denote the feasible region Θ and θ in Eq. (2) as $\tilde{\Theta}$ and $\tilde{\theta}$, respectively. Then,

$$\begin{aligned}
 & \max_{\tilde{\theta} \in \tilde{\Theta}} \left\{ \left| \frac{1}{n} [\bar{\mathbf{X}} \tilde{\theta}]_i \right| \right\} = \max_{\tilde{\theta} \in \tilde{\Theta}} \left\{ \frac{1}{n} |\bar{\mathbf{x}}^i \tilde{\theta}| \right\} \\
 &= \max_{\tilde{\theta} \in \tilde{\Theta}} \left\{ \frac{1}{n} \left| [\bar{\mathbf{x}}^i]_{\hat{\mathcal{D}}^c} [\tilde{\theta}]_{\hat{\mathcal{D}}^c} + [\bar{\mathbf{x}}^i]_{\hat{\mathcal{L}}} [\tilde{\theta}]_{\hat{\mathcal{L}}} + [\bar{\mathbf{x}}^i]_{\hat{\mathcal{R}}} [\tilde{\theta}]_{\hat{\mathcal{R}}} \right| \right\} \\
 &= \max_{\tilde{\theta} \in \tilde{\Theta}} \left\{ \frac{1}{n} \left| \langle [\bar{\mathbf{x}}^i]_{\hat{\mathcal{D}}^c}, [\tilde{\theta}]_{\hat{\mathcal{D}}^c} \rangle + \langle [\bar{\mathbf{x}}^i]_{\hat{\mathcal{L}}}, \mathbf{1} \rangle \right| \right\} \\
 &= \max_{\theta \in \Theta} \left\{ \frac{1}{n} \left| \langle [\bar{\mathbf{x}}^i]_{\hat{\mathcal{D}}^c}, \theta \rangle + \langle [\bar{\mathbf{x}}^i]_{\hat{\mathcal{L}}}, \mathbf{1} \rangle \right| \right\} = s^i(\alpha, \beta_0).
 \end{aligned}$$

The last two equations hold since $[\theta]_{\hat{\mathcal{L}}} = \mathbf{1}$, $[\theta]_{\hat{\mathcal{R}}} = 0$, and $[\theta]_{\hat{\mathcal{D}}^c} \in \Theta$.

Proof of Lemma 8:

$$\begin{aligned}
 s^i(\alpha, \beta_0) &= \max_{\theta \in B(\mathbf{c}, r)} \left\{ \frac{1}{n} \left| \langle [\bar{\mathbf{x}}^i]_{\hat{\mathcal{D}}^c}, \theta \rangle + \langle [\bar{\mathbf{x}}^i]_{\hat{\mathcal{L}}}, \mathbf{1} \rangle \right| \right\} \\
 &= \max_{\eta \in B(\mathbf{0}, r)} \left\{ \frac{1}{n} \left| \langle [\bar{\mathbf{x}}^i]_{\hat{\mathcal{D}}^c}, \mathbf{c} \rangle + \langle [\bar{\mathbf{x}}^i]_{\hat{\mathcal{L}}}, \mathbf{1} \rangle + \langle [\bar{\mathbf{x}}^i]_{\hat{\mathcal{D}}^c}, \eta \rangle \right| \right\} \\
 &= \frac{1}{n} \left(\left| \langle [\bar{\mathbf{x}}^i]_{\hat{\mathcal{D}}^c}, \mathbf{c} \rangle + \langle [\bar{\mathbf{x}}^i]_{\hat{\mathcal{L}}}, \mathbf{1} \rangle \right| + \| [\bar{\mathbf{x}}^i]_{\hat{\mathcal{D}}^c} \| r \right).
 \end{aligned}$$

The last equality holds since $-\| [\bar{\mathbf{x}}^i]_{\hat{\mathcal{D}}^c} \| r \leq \langle [\bar{\mathbf{x}}^i]_{\hat{\mathcal{D}}^c}, \eta \rangle \leq \| [\bar{\mathbf{x}}^i]_{\hat{\mathcal{D}}^c} \| r$.

The proof is complete. ■

A.6. Proof for Theorem 9

Proof of Theorem 9:

- (1) It can be obtained from the rule R1.
- (2) It is from the definition of $\hat{\mathcal{F}}$. ■

A.7. Proof for Lemma 10

Firstly, we need to point out that the optimization problems in Eqs. (3) and (4) are equivalent to the problems:

$$\begin{aligned}
 u_i(\alpha, \beta_0) &= \max_{\mathbf{w} \in \mathcal{W}} \{1 - \langle [\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c}, \mathbf{w} \rangle\}, i \in \hat{\mathcal{D}}^c, \\
 l_i(\alpha, \beta_0) &= \min_{\mathbf{w} \in \mathcal{W}} \{1 - \langle [\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c}, \mathbf{w} \rangle\}, i \in \hat{\mathcal{D}}^c.
 \end{aligned}$$

They follow from the fact that $[\mathbf{w}]_{\hat{\mathcal{F}}^c} \in \mathcal{W}$ and

$$\begin{aligned}
 & \{1 - \langle \mathbf{w}, \bar{\mathbf{x}}_i \rangle\} \\
 &= \{1 - \langle [\mathbf{w}]_{\hat{\mathcal{F}}^c}, [\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c} \rangle - \langle [\mathbf{w}]_{\hat{\mathcal{F}}}, [\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}} \rangle\} \\
 &= \{1 - \langle [\mathbf{w}]_{\hat{\mathcal{F}}^c}, [\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c} \rangle\} \text{ (since } [\mathbf{w}]_{\hat{\mathcal{F}}} = 0\text{)}.
 \end{aligned}$$

Proof of Lemma 10:

$$\begin{aligned}
 u_i(\alpha, \beta_0) &= \max_{\mathbf{w} \in B(\mathbf{c}, r)} \{1 - \langle [\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c}, \mathbf{w} \rangle\} \\
 &= \max_{\eta \in B(\mathbf{0}, r)} \{1 - \langle [\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c}, \mathbf{c} \rangle - \langle [\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c}, \eta \rangle\} \\
 &= 1 - \langle [\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c}, \mathbf{c} \rangle + \max_{\eta \in B(\mathbf{0}, r)} \{-\langle [\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c}, \eta \rangle\} \\
 &= 1 - \langle [\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c}, \mathbf{c} \rangle + \|[\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c}\| r.
 \end{aligned}$$

$$\begin{aligned}
 l_i(\alpha, \beta_0) &= \min_{\mathbf{w} \in B(\mathbf{c}, r)} \{1 - \langle [\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c}, \mathbf{w} \rangle\} \\
 &= \min_{\eta \in B(\mathbf{0}, r)} \{1 - \langle [\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c}, \mathbf{c} \rangle - \langle [\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c}, \eta \rangle\} \\
 &= 1 - \langle [\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c}, \mathbf{c} \rangle + \min_{\eta \in B(\mathbf{0}, r)} \{-\langle [\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c}, \eta \rangle\} \\
 &= 1 - \langle [\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c}, \mathbf{c} \rangle - \|[\bar{\mathbf{x}}_i]_{\hat{\mathcal{F}}^c}\| r.
 \end{aligned}$$

The proof is complete. ■

A.8. Proof for Theorem 11

Proof of Theorem 11:

- (1) It can be obtained from the rule R2.
- (2) It is from the definitions of $\hat{\mathcal{R}}$ and $\hat{\mathcal{L}}$. ■

A.9. Proof for Theorem 12

Proof of Theorem 12:

- (1) Given the reference solutions pair $\mathbf{w}^*(\alpha_{i-1, j}, \beta_j)$ and $\theta^*(\alpha_{i-1, j}, \beta_j)$, if we do ISS first in SIFS and apply ISS and IFS for infinite times. If after s times of triggering, no new inactive features or samples are identified, then we can denote the sequence of $\hat{\mathcal{F}}$, $\hat{\mathcal{R}}$, and $\hat{\mathcal{L}}$ as:

$$\hat{\mathcal{F}}_0^A = \hat{\mathcal{R}}_0^A = \hat{\mathcal{L}}_0^A = \emptyset \xrightarrow{ISS} \hat{\mathcal{F}}_1^A, \hat{\mathcal{R}}_1^A, \hat{\mathcal{L}}_1^A \xrightarrow{IFS} \hat{\mathcal{F}}_2^A, \hat{\mathcal{R}}_2^A, \hat{\mathcal{L}}_2^A \xrightarrow{ISS} \dots \hat{\mathcal{F}}_s^A, \hat{\mathcal{R}}_s^A, \hat{\mathcal{L}}_s^A \xrightarrow{IFS/ISS} \dots \quad (33)$$

$$\text{with } \hat{\mathcal{F}}_s^A = \hat{\mathcal{F}}_{s+1}^A = \hat{\mathcal{F}}_{s+2}^A = \dots, \hat{\mathcal{R}}_s^A = \hat{\mathcal{R}}_{s+1}^A = \hat{\mathcal{R}}_{s+2}^A = \dots \text{ and } \hat{\mathcal{L}}_s^A = \hat{\mathcal{L}}_{s+1}^A = \hat{\mathcal{L}}_{s+2}^A = \dots \quad (34)$$

In the same way, if we do IFS first in SIFS and no new inactive features or samples are identified after t times of triggering of ISS and IFS, then the sequence can be denoted as:

$$\hat{\mathcal{F}}_0^B = \hat{\mathcal{R}}_0^B = \hat{\mathcal{L}}_0^B = \emptyset \xrightarrow{IFS} \hat{\mathcal{F}}_1^B, \hat{\mathcal{R}}_1^B, \hat{\mathcal{L}}_1^B \xrightarrow{ISS} \hat{\mathcal{F}}_2^B, \hat{\mathcal{R}}_2^B, \hat{\mathcal{L}}_2^B \xrightarrow{IFS} \dots \hat{\mathcal{F}}_t^B, \hat{\mathcal{R}}_t^B, \hat{\mathcal{L}}_t^B \xrightarrow{IFS/ISS} \dots \quad (35)$$

$$\text{with } \hat{\mathcal{F}}_t^B = \hat{\mathcal{F}}_{t+1}^B = \hat{\mathcal{F}}_{t+2}^B = \dots, \hat{\mathcal{R}}_t^B = \hat{\mathcal{R}}_{t+1}^B = \hat{\mathcal{R}}_{t+2}^B = \dots \text{ and } \hat{\mathcal{L}}_t^B = \hat{\mathcal{L}}_{t+1}^B = \hat{\mathcal{L}}_{t+2}^B = \dots \quad (36)$$

We first prove that $\hat{\mathcal{F}}_k^B \subseteq \hat{\mathcal{F}}_{k+1}^A$, $\hat{\mathcal{R}}_k^B \subseteq \hat{\mathcal{R}}_{k+1}^A$ and $\hat{\mathcal{L}}_k^B \subseteq \hat{\mathcal{L}}_{k+1}^A$ hold for all $k \geq 0$ by induction.

1) When $k = 0$, the equalities $\hat{\mathcal{F}}_0^B \subseteq \hat{\mathcal{F}}_1^A, \hat{\mathcal{R}}_0^B \subseteq \hat{\mathcal{R}}_1^A$ and $\hat{\mathcal{L}}_0^B \subseteq \hat{\mathcal{L}}_1^A$ hold since $\hat{\mathcal{F}}_0^B = \hat{\mathcal{R}}_0^B = \hat{\mathcal{L}}_0^B = \emptyset$.

2) If $\hat{\mathcal{F}}_k^B \subseteq \hat{\mathcal{F}}_{k+1}^A, \hat{\mathcal{R}}_k^B \subseteq \hat{\mathcal{R}}_{k+1}^A$ and $\hat{\mathcal{L}}_k^B \subseteq \hat{\mathcal{L}}_{k+1}^A$ hold, by the synergistic effect of ISS and IFS, we have that $\hat{\mathcal{F}}_{k+1}^B \subseteq \hat{\mathcal{F}}_{k+2}^A, \hat{\mathcal{R}}_{k+1}^B \subseteq \hat{\mathcal{R}}_{k+2}^A$ and $\hat{\mathcal{L}}_{k+1}^B \subseteq \hat{\mathcal{L}}_{k+2}^A$ hold.

Thus, $\hat{\mathcal{F}}_k^B \subseteq \hat{\mathcal{F}}_{k+1}^A, \hat{\mathcal{R}}_k^B \subseteq \hat{\mathcal{R}}_{k+1}^A$ and $\hat{\mathcal{L}}_k^B \subseteq \hat{\mathcal{L}}_{k+1}^A$ hold for all $k \geq 0$.

Similar to the analysis in (1), we can also prove that $\hat{\mathcal{F}}_k^A \subseteq \hat{\mathcal{F}}_{k+1}^B, \hat{\mathcal{R}}_k^A \subseteq \hat{\mathcal{R}}_{k+1}^B$ and $\hat{\mathcal{L}}_k^A \subseteq \hat{\mathcal{L}}_{k+1}^B$ hold for all $k \geq 0$.

Combining (1) and (2), we can acquire

$$\hat{\mathcal{F}}_0^B \subseteq \hat{\mathcal{F}}_1^A \subseteq \hat{\mathcal{F}}_2^B \subseteq \hat{\mathcal{F}}_3^A \dots \quad (37)$$

$$\hat{\mathcal{F}}_0^A \subseteq \hat{\mathcal{F}}_1^B \subseteq \hat{\mathcal{F}}_2^A \subseteq \hat{\mathcal{F}}_3^B \dots \quad (38)$$

$$\hat{\mathcal{R}}_0^B \subseteq \hat{\mathcal{R}}_1^A \subseteq \hat{\mathcal{R}}_2^B \subseteq \hat{\mathcal{R}}_3^A \dots \quad (39)$$

$$\hat{\mathcal{R}}_0^A \subseteq \hat{\mathcal{R}}_1^B \subseteq \hat{\mathcal{R}}_2^A \subseteq \hat{\mathcal{R}}_3^B \dots \quad (40)$$

$$\hat{\mathcal{L}}_0^B \subseteq \hat{\mathcal{L}}_1^A \subseteq \hat{\mathcal{L}}_2^B \subseteq \hat{\mathcal{L}}_3^A \dots \quad (41)$$

$$\hat{\mathcal{L}}_0^A \subseteq \hat{\mathcal{L}}_1^B \subseteq \hat{\mathcal{L}}_2^A \subseteq \hat{\mathcal{L}}_3^B \dots \quad (42)$$

By the first equality of Eqs. (34), (37), and (38), we can obtain $\hat{\mathcal{F}}_p^A = \hat{\mathcal{F}}_t^B$. Similarly, we can obtain $\hat{\mathcal{R}}_p^A = \hat{\mathcal{R}}_t^B$ and $\hat{\mathcal{L}}_p^A = \hat{\mathcal{L}}_t^B$.

(2) If p is odd, then by Eqs. (37), (39), and (41), we have $\hat{\mathcal{F}}_s^A \subseteq \hat{\mathcal{F}}_{p+1}^B, \hat{\mathcal{R}}_s^A \subseteq \hat{\mathcal{R}}_{p+1}^B$, and $\hat{\mathcal{L}}_s^A \subseteq \hat{\mathcal{L}}_{p+1}^B$. Thus $q \leq p + 1$.

Else if p is even, then by Eqs. (38), (40), and (42), we have $\hat{\mathcal{F}}_s^A \subseteq \hat{\mathcal{F}}_{p+1}^B, \hat{\mathcal{R}}_s^A \subseteq \hat{\mathcal{R}}_{p+1}^B$, and $\hat{\mathcal{L}}_s^A \subseteq \hat{\mathcal{L}}_{p+1}^B$. Thus $q \leq p + 1$.

Doing the same analysis for q , we can obtain $p \leq q + 1$.

Hence, $|p - q| \leq 1$.

The proof is complete. ■

A.10. Proof for Theorem 13

Proof of Theorem 13:

We notice $(\mathbf{X}_i^k)^\top \mathbf{w} = \mathbf{w}_k^\top \mathbf{x}_i - \mathbf{w}_{y_i}^\top \mathbf{x}_i$. By denoting $\phi_i(\mathbf{a}) = \sum_{k=1}^K \ell([\mathbf{u}_i]_k + [\mathbf{a}]_k)$, it is easy to verify that $\ell_i(\mathbf{w}) = \phi_i(\mathbf{X}_i^\top \mathbf{w})$. Thus, the problem (m-P*) can be rewritten as:

$$\min_{\mathbf{w} \in \mathbb{R}^{Kp}} P(\mathbf{w}; \alpha, \beta) = \frac{1}{n} \sum_{i=1}^n \phi_i(\mathbf{X}_i^\top \mathbf{w}) + \frac{\alpha}{2} \|\mathbf{w}\|^2 + \beta \|\mathbf{w}\|_1.$$

Let $\mathbf{z}_i = \mathbf{X}_i^\top \mathbf{w}$. The primal problem (P*) is then equivalent to

$$\begin{aligned} \min_{\mathbf{w} \in \mathbb{R}^p, \mathbf{z} \in \mathbb{R}^n} & \frac{\alpha}{2} \|\mathbf{w}\|^2 + \beta \|\mathbf{w}\|_1 + \frac{1}{n} \sum_{i=1}^n \phi_i(\mathbf{z}_i), \\ \text{s.t. } & \mathbf{z}_i = \mathbf{X}_i^\top \mathbf{w}, i = 1, 2, \dots, n. \end{aligned}$$

The Lagrangian then becomes

$$\begin{aligned}
 L(\mathbf{w}, \mathbf{z}, \theta) &= \frac{\alpha}{2} \|\mathbf{w}\|^2 + \beta \|\mathbf{w}\|_1 + \frac{1}{n} \sum_{i=1}^n \phi_i(\mathbf{z}_i) + \frac{1}{n} \sum_{i=1}^n \langle \mathbf{X}_i^\top \mathbf{w} - \mathbf{z}_i, \theta_i \rangle \\
 &= \underbrace{\frac{\alpha}{2} \|\mathbf{w}\|^2 + \beta \|\mathbf{w}\|_1 + \frac{1}{n} \langle \mathbf{X}\theta, \mathbf{w} \rangle}_{:=f_1(\mathbf{w})} + \underbrace{\frac{1}{n} \sum_{i=1}^n (\phi_i(\mathbf{z}_i) - \langle \mathbf{z}_i, \theta_i \rangle)}_{:=f_2(\mathbf{z})}. \tag{43}
 \end{aligned}$$

We first consider the subproblem $\min_{\mathbf{w}} L(\mathbf{w}, \mathbf{z}, \theta)$:

$$\begin{aligned}
 0 \in \partial_{\mathbf{w}} L(\mathbf{w}, \mathbf{z}, \theta) &= \partial_{\mathbf{w}} f_1(\mathbf{w}) = \alpha \mathbf{w} + \frac{1}{n} \mathbf{X}\theta + \beta \partial \|\mathbf{w}\|_1 \Leftrightarrow \\
 \frac{1}{n} \mathbf{X}\theta &\in -\alpha \mathbf{w} - \beta \partial \|\mathbf{w}\|_1 \Rightarrow \mathbf{w} = -\frac{1}{\alpha} \mathcal{S}_\beta \left(\frac{1}{n} \mathbf{X}\theta \right). \tag{44}
 \end{aligned}$$

By substituting Eq. (44) into $f_1(\mathbf{w})$, we obtain

$$f_1(\mathbf{w}) = \frac{\alpha}{2} \|\mathbf{w}\|^2 + \beta \|\mathbf{w}\|_1 - \langle \alpha \mathbf{w} + \beta \partial \|\mathbf{w}\|_1, \mathbf{w} \rangle = -\frac{\alpha}{2} \|\mathbf{w}\|^2 = -\frac{1}{2\alpha} \|\mathcal{S}_\beta \left(\frac{1}{n} \mathbf{X}\theta \right)\|^2. \tag{45}$$

Then, we consider the problem $\min_{\mathbf{z}} L(\mathbf{w}, \mathbf{z}, \theta)$:

$$\begin{aligned}
 0 = \nabla_{[\mathbf{z}_i]_k} L(\mathbf{w}, \mathbf{z}, \theta) &= \nabla_{[\mathbf{z}_i]_k} f_2(\mathbf{z}) = \begin{cases} -\frac{1}{n} [\theta_i]_k, & \text{if } [\mathbf{z}_i]_k + [\mathbf{u}_i]_k < 0, \\ \frac{1}{\gamma n} ([\mathbf{z}_i]_k + [\mathbf{u}_i]_k) - \frac{1}{n} [\theta_i]_k, & \text{if } 0 \leq [\mathbf{z}_i]_k + [\mathbf{u}_i]_k \leq \gamma, \\ \frac{1}{n} - \frac{1}{n} [\theta_i]_k, & \text{if } [\mathbf{z}_i]_k + [\mathbf{u}_i]_k > \gamma. \end{cases} \\
 \Rightarrow [\theta_i]_k &= \begin{cases} 0, & \text{if } [\mathbf{z}_i]_k + [\mathbf{u}_i]_k < 0, \\ \frac{1}{\gamma} ([\mathbf{z}_i]_k + [\mathbf{u}_i]_k), & \text{if } 0 \leq [\mathbf{z}_i]_k + [\mathbf{u}_i]_k \leq \gamma, \\ 1, & \text{if } [\mathbf{z}_i]_k + [\mathbf{u}_i]_k > \gamma. \end{cases} \tag{46}
 \end{aligned}$$

Thus, we have

$$f_2(\mathbf{z}) = \begin{cases} -\frac{\gamma}{2n} \|\theta\|^2 + \frac{1}{n} \langle \mathbf{u}, \theta \rangle, & \text{if } [\theta_i]_k \in [0, 1], \forall i \in [n], k \in [K], \\ -\infty, & \text{otherwise.} \end{cases} \tag{47}$$

Combining Eqs.(43), (45), and (47), we obtain the dual problem:

$$\min_{\theta \in [0,1]^{Kn}} \frac{1}{2\alpha} \|\mathcal{S}_\beta \left(\frac{1}{n} \mathbf{X}\theta \right)\|^2 + \frac{\gamma}{2n} \|\theta\|^2 - \frac{1}{n} \langle \mathbf{u}, \theta \rangle.$$

■

A.11. Proof for Lemma 14

Lemma 14 can be proved quite similarly with Lemma 3. Therefore, we omit this proof here.

A.12. Proof for Lemma 15

Proof of Lemma 15:

(i) We prove this lemma by verifying that the solutions $\mathbf{w}^*(\alpha, \beta) = \mathbf{0}$ and $\theta^*(\alpha, \beta) = \mathbf{u}$ satisfy the conditions (m-KKT-1) and (m-KKT-2).

Firstly, since $\beta \geq \beta_{\max} = \|\frac{1}{n}\mathbf{X}\mathbf{u}\|_{\infty}$, we have $\mathcal{S}_{\beta}(\frac{1}{n}\mathbf{X}\mathbf{u}) = \mathbf{0}$. Thus, $\mathbf{w}^*(\alpha, \beta) = \mathbf{0}$ and $\theta^*(\alpha, \beta) = \mathbf{u}$ satisfy the condition (m-KKT-1). Then, for all $i \in [n]$, we have

$$\langle \mathbf{X}_i^k, \mathbf{w}^*(\alpha, \beta) \rangle + [\mathbf{u}_i]_k = [\mathbf{u}_i]_k = \begin{cases} 1, & \text{if } k \neq y_i; \\ 0, & \text{if } k = y_i; \end{cases} \quad \forall i = 1, \dots, n, k = 1, \dots, K.$$

Thus, $\mathbf{w}^*(\alpha, \beta) = \mathbf{0}$ and $\theta^*(\alpha, \beta) = \mathbf{u}$ satisfy the condition (m-KKT-2). Hence, they are the solutions of the primal problem (m-P*) and the dual problem (m-D*), respectively.

(ii) Similar to the proof of (i), we prove this by verifying that the solutions $\mathbf{w}^*(\alpha, \beta) = \frac{1}{\alpha}\mathcal{S}_{\beta}(-\frac{1}{n}\mathbf{X}\mathbf{u})$ and $\theta^*(\alpha, \beta) = \mathbf{u}$ satisfy the conditions (m-KKT-1) and (m-KKT-2).

1. **Case 1:** $\alpha_{\max}(\beta) \leq 0$. Then for all $\alpha > 0$, we have

$$\begin{aligned} \text{if } k = y_i, \quad & \langle \mathbf{X}_i^k, \mathbf{w}^*(\alpha, \beta) \rangle + \mathbf{u}_i^k = \mathbf{u}_i^k = 0, \\ \text{if } k \neq y_i, \quad & \langle \mathbf{X}_i^k, \mathbf{w}^*(\alpha, \beta) \rangle + \mathbf{u}_i^k = 1 - \langle \mathbf{X}_i^k, \frac{1}{\alpha}\mathcal{S}_{\beta}(\frac{1}{n}\mathbf{X}\mathbf{u}) \rangle \\ & = 1 - \frac{1}{\alpha}\langle \mathbf{X}_i^k, \mathcal{S}_{\beta}(\frac{1}{n}\mathbf{X}\mathbf{u}) \rangle \geq 1 > \gamma. \end{aligned}$$

Thus, $\mathbf{w}^*(\alpha, \beta) = \frac{1}{\alpha}\mathcal{S}_{\beta}(-\frac{1}{n}\mathbf{X}\mathbf{u})$ and $\theta^*(\alpha, \beta) = \mathbf{u}$ satisfy the conditions (m-KKT-1) and (m-KKT-2). Hence, they are the optimal solution of problems (m-P*) and (m-D*).

2. **Case 2:** $\alpha_{\max}(\beta) > 0$. Then for any $\alpha \geq \alpha_{\max}(\beta)$, we have

$$\begin{aligned} \text{if } k = y_i, \quad & \langle \mathbf{X}_i^k, \mathbf{w}^*(\alpha, \beta) \rangle + \mathbf{u}_i^k = \mathbf{u}_i^k = 0, \\ \text{if } k \neq y_i, \quad & \langle \mathbf{X}_i^k, \mathbf{w}^*(\alpha, \beta) \rangle + \mathbf{u}_i^k = 1 - \langle \mathbf{X}_i^k, \frac{1}{\alpha}\mathcal{S}_{\beta}(\frac{1}{n}\mathbf{X}\mathbf{u}) \rangle \\ & = 1 - \frac{1}{\alpha}\langle \mathbf{X}_i^k, \mathcal{S}_{\beta}(\frac{1}{n}\mathbf{X}\mathbf{u}) \rangle \\ & \geq 1 - \frac{\alpha_{\max}(\beta)}{\alpha}(1 - \gamma) \geq 1 - (1 - \gamma) = \gamma. \end{aligned}$$

Thus, $\mathbf{w}^*(\alpha, \beta) = \frac{1}{\alpha}\mathcal{S}_{\beta}(-\frac{1}{n}\mathbf{X}\mathbf{u})$ and $\theta^*(\alpha, \beta) = \mathbf{u}$ satisfy the conditions (m-KKT-1) and (m-KKT-2). Hence, they are the optimal solutions of the primal and dual problems (m-P*) and (m-D*).

The proof is complete. ■

A.13. More Experimental Results

Below, we report the rejection ratios of SIFS on syn1 (Fig. 8), syn3 (Fig. 9), rcv1-train (Fig. 10), rcv1-test (Fig. 11), url (Fig. 12), kddb (Fig. 13), syn-multi1 (Fig. 14), syn-multi3 (Fig. 15), and rcv1-multiclass (Fig. 16), which are omitted in the main paper. In all the figures, the first row presents the rejection ratios of feature screening and the second row presents those of sample screening.

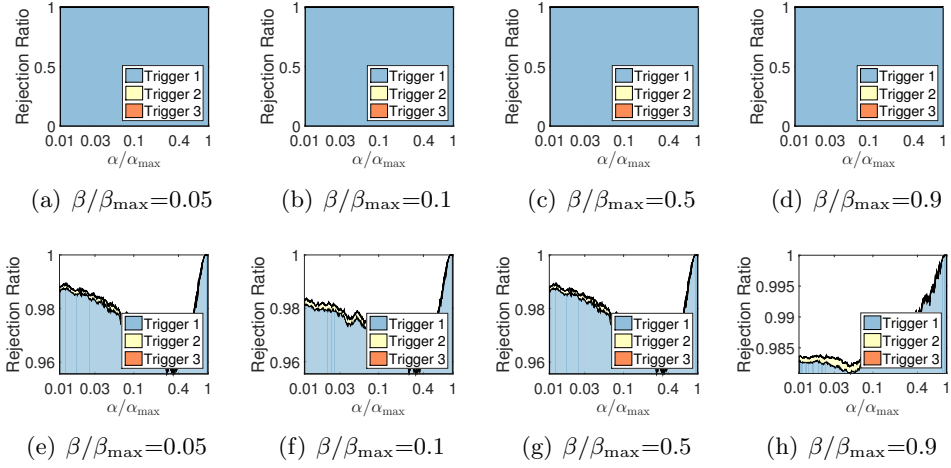


Figure 8: Rejection ratios of SIFS on syn1.

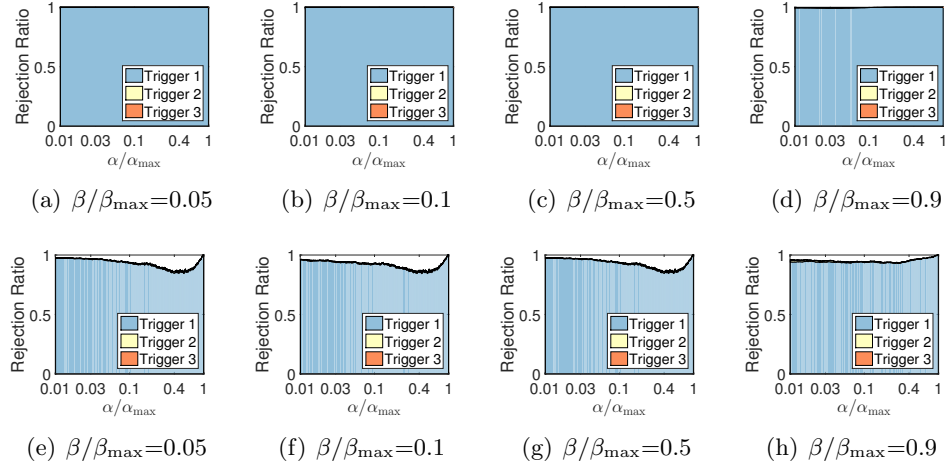


Figure 9: Rejection ratios of SIFS on syn3.




Figure 10: Rejection ratios of SIFS on rcv1-train.




Figure 11: Rejection ratios of SIFS on rcv1-test.

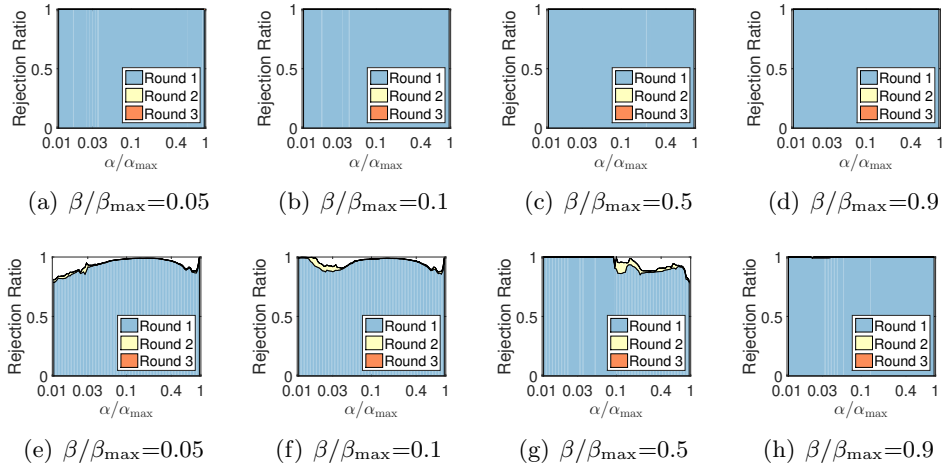


Figure 12: Rejection ratios of SIFS on url.

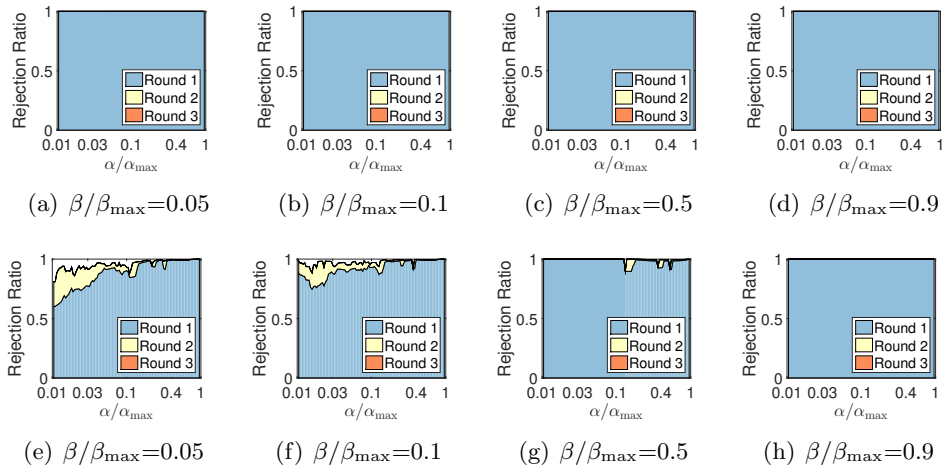


Figure 13: Rejection ratios of SIFS on kddb.




Figure 14: Rejection ratios of SIFS on syn-multi1.

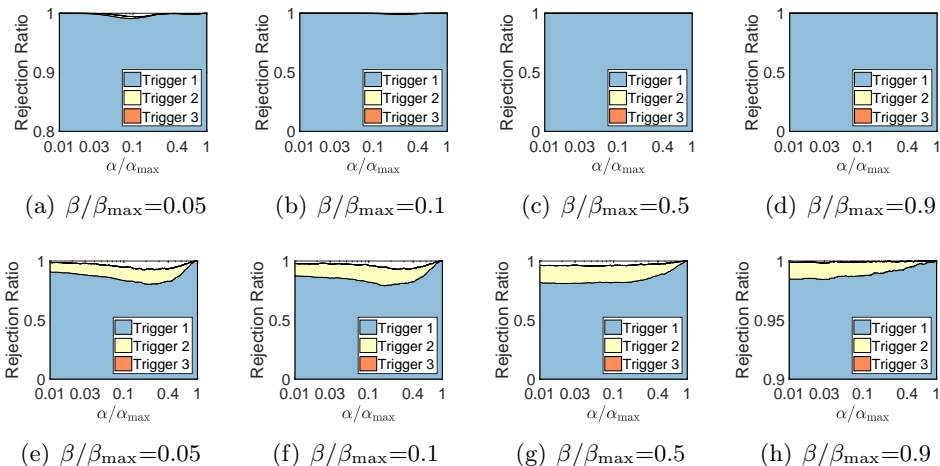


Figure 15: Rejection ratios of SIFS on syn-multi3.

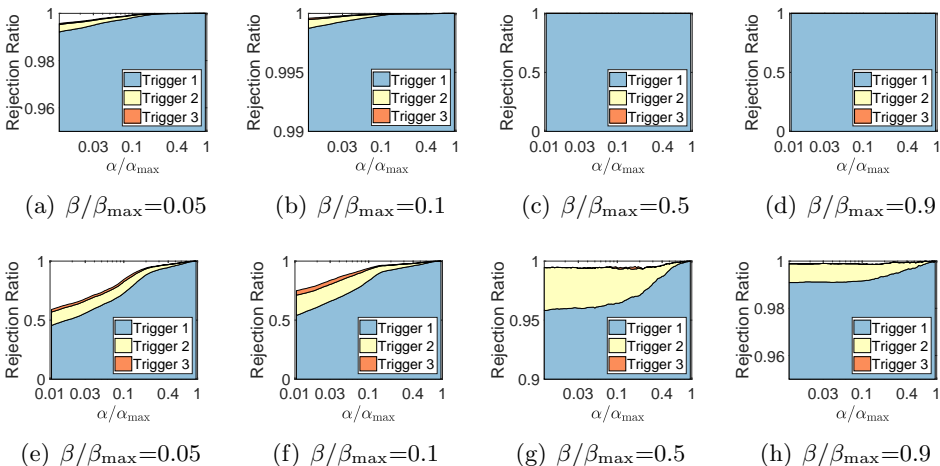


Figure 16: Rejection ratios of SIFS on rcv1-multiclass.

References

- Heinz H Bauschke, Patrick L Combettes, et al. *Convex analysis and monotone operator theory in Hilbert spaces*, volume 408. Springer, 2011.
- Jinbo Bi, Kristin Bennett, Mark Embrechts, Curt Breneman, and Minghu Song. Dimensionality reduction via sparse support vector machines. *The Journal of Machine Learning Research*, 3:1229–1243, 2003.
- Antoine Bonnefoy, Valentin Emiya, Liva Ralaivola, and Rémi Gribonval. A dynamic screening principle for the lasso. In *Signal Processing Conference (EUSIPCO), 2014 Proceedings of the 22nd European*, pages 6–10. IEEE, 2014.
- Bryan Catanzaro, Narayanan Sundaram, and Kurt Keutzer. Fast support vector machine training and classification on graphics processors. In *Proceedings of the 25th international conference on Machine learning*, pages 104–111. ACM, 2008.
- Chih-Chung Chang and Chih-Jen Lin. Libsvm: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(3):27, 2011.
- Laurent El Ghaoui, Vivian Viallon, and Tarek Rabbani. Safe feature elimination in sparse supervised learning. *Pacific Journal of Optimization*, 8:667–698, 2012.
- Rong-En Fan, Kai-Wei Chang, Cho-Jui Hsieh, Xiang-Rui Wang, and Chih-Jen Lin. Liblinear: A library for large linear classification. *The Journal of Machine Learning Research*, 9: 1871–1874, 2008.
- Trevor Hastie, Saharon Rosset, Robert Tibshirani, and Ji Zhu. The entire regularization path for the support vector machine. *The Journal of Machine Learning Research*, 5: 1391–1415, 2004.
- Trevor Hastie, Robert Tibshirani, and Martin Wainwright. *Statistical learning with sparsity: the lasso and generalizations*. CRC Press, 2015.
- Cho-Jui Hsieh, Kai-Wei Chang, Chih-Jen Lin, S Sathiya Keerthi, and Sellamanickam Sundararajan. A dual coordinate descent method for large-scale linear svm. In *Proceedings of the 25th international conference on Machine learning*, pages 408–415. ACM, 2008.
- Thorsten Joachims. *Text categorization with support vector machines: Learning with many relevant features*. Springer, 1998.
- Ron Kohavi. A study of cross-validation and bootstrap for accuracy estimation and model selection. In *International Joint Conference on Artificial Intelligence*, pages 1137–1145, 1995.
- Irene Kotsia and Ioannis Pitas. Facial expression recognition in image sequences using geometric deformation features and support vector machines. *Image Processing, IEEE Transactions on*, 16(1):172–187, 2007.

- Mathurin Massias, Joseph Salmon, and Alexandre Gramfort. Celer: a fast solver for the lasso with dual extrapolation. In *International Conference on Machine Learning*, pages 3321–3330, 2018.
- Nicolai Meinshausen and Peter Bühlmann. Stability selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(4):417–473, 2010.
- Johannes Mohr and Klaus Obermayer. A topographic support vector machine: Classification using local label configurations. In *Advances in Neural Information Processing Systems*, pages 929–936, 2004.
- Harikrishna Narasimhan and Shivani Agarwal. Svm pauc tight: a new support vector method for optimizing partial auc based on a tight convex upper bound. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 167–175. ACM, 2013.
- Eugene Ndiaye, Olivier Fercoq, Alexandre Gramfort, and Joseph Salmon. Gap safe screening rules for sparse-group lasso. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 388–396. Curran Associates, Inc., 2016.
- Kohei Ogawa, Yoshiki Suzuki, and Ichiro Takeuchi. Safe screening of non-support vectors in pathwise svm computation. In *Proceedings of the 30th International Conference on Machine Learning*, pages 1382–1390, 2013.
- Ralph Tyrell Rockafellar. *Convex analysis*. Princeton university press, 1970.
- Shai Shalev-Shwartz and Tong Zhang. Accelerated proximal stochastic dual coordinate ascent for regularized loss minimization. *Mathematical Programming*, 155(1-2):105–145, 2016.
- Shai Shalev-Shwartz, Yoram Singer, Nathan Srebro, and Andrew Cotter. Pegasos: Primal estimated sub-gradient solver for svm. *Mathematical programming*, 127(1):3–30, 2011.
- Atsushi Shibagaki, Masayuki Karasuyama, Kohei Hatano, and Ichiro Takeuchi. Simultaneous safe screening of features and samples in doubly sparse modeling. In *Proceedings of The 33rd International Conference on Machine Learning*, 2016.
- Robert Tibshirani, Jacob Bien, Jerome Friedman, Trevor Hastie, Noah Simon, Jonathan Taylor, and Ryan J Tibshirani. Strong rules for discarding predictors in lasso-type problems. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 74(2):245–266, 2012.
- Jie Wang, Peter Wonka, and Jieping Ye. Scaling svm and least absolute deviations via exact data reduction. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 523–531, 2014a.
- Jie Wang, Jiayu Zhou, Jun Liu, Peter Wonka, and Jieping Ye. A safe screening rule for sparse logistic regression. In *Advances in Neural Information Processing Systems*, pages 1053–1061, 2014b.

- Jie Wang, Peter Wonka, and Jieping Ye. Lasso screening rules via dual polytope projection. *Journal of Machine Learning Research*, 16:1063–1101, 2015.
- Li Wang, Ji Zhu, and Hui Zou. The doubly regularized support vector machine. *Statistica Sinica*, pages 589–615, 2006.
- Zhen James Xiang and Peter J Ramadge. Fast lasso screening tests based on correlations. In *Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on*, pages 2137–2140. IEEE, 2012.
- Tao Yang, Jie Wang, Qian Sun, Derrek P Hibar, Neda Jahanshad, Li Liu, Yalin Wang, Liang Zhan, Paul M Thompson, and Jieping Ye. Detecting genetic risk factors for alzheimer’s disease in whole genome sequence data via lasso screening. In *Biomedical Imaging (ISBI), 2015 IEEE 12th International Symposium on*, pages 985–989. IEEE, 2015.
- Yuya Yoshikawa, Tomoharu Iwata, and Hiroshi Sawada. Latent support measure machines for bag-of-words data classification. In *Advances in Neural Information Processing Systems*, pages 1961–1969, 2014.
- Weizhong Zhang, Bin Hong, Wei Liu, Jieping Ye, Deng Cai, Xiaofei He, and Jie Wang. Scaling up sparse support vector machine by simultaneous feature and sample reduction. In *Proceedings of The 34th International Conference on Machine Learning*, pages 4016–4025, 2017.