
Fast Epigraphical Projection-based Incremental Algorithms for Wasserstein Distributionally Robust Support Vector Machine

Jiajin Li

Department of Systems Engineering and Engineering Management
The Chinese University of Hong Kong
jjli@se.cuhk.edu.hk

Caihua Chen*

School of Management and Engineering
Nanjing University
chchen@nju.edu.cn

Anthony Man-Cho So

Department of Systems Engineering and Engineering Management
The Chinese University of Hong Kong
manchoso@se.cuhk.edu.hk

Abstract

Wasserstein **D**istributionally **R**obust Optimization (DRO) is concerned with finding decisions that perform well on data that are drawn from the worst-case probability distribution within a Wasserstein ball centered at a certain nominal distribution. In recent years, it has been shown that various DRO formulations of learning models admit tractable convex reformulations. However, most existing works propose to solve these convex reformulations by general-purpose solvers, which are not well-suited for tackling large-scale problems. In this paper, we focus on a family of Wasserstein distributionally robust support vector machine (DRSVM) problems and propose two novel epigraphical projection-based incremental algorithms to solve them. The updates in each iteration of these algorithms can be computed in a highly efficient manner. Moreover, we show that the DRSVM problems considered in this paper satisfy a Hölderian growth condition with explicitly determined growth exponents. Consequently, we are able to establish the convergence rates of the proposed incremental algorithms. Our numerical results indicate that the proposed methods are orders of magnitude faster than the state-of-the-art, and the performance gap grows considerably as the problem size increases.

1 Introduction

Wasserstein distance-based distributionally robust optimization (DRO) has recently received significant attention in the machine learning community. This can be attributed to its ability to improve generalization performance by robustifying the learning model against unseen data [13, 22]. The DRO approach offers a principled way to regularize empirical risk minimization problems and provides a transparent probabilistic interpretation of a wide range of existing regularization techniques; see, e.g., [4, 10, 22] and the references therein. Moreover, many representative distributionally robust learning models admit equivalent reformulations as tractable convex programs via strong

*Corresponding author

duality [22, 12, 18, 27, 11]. Currently, a standard approach to solving these reformulations is to use off-the-shelf solvers such as YALMIP or CPLEX. However, these general-purpose solvers do not scale well with the problem size. Such a state of affairs greatly limits the use of the DRO methodology in machine learning applications and naturally motivates the study of the algorithmic aspects of DRO.

In this paper, we aim to design fast iterative methods for solving a family of Wasserstein distributionally robust support vector machine (DRSVM) problems. The SVM is one of the most frequently used classification methods and has enjoyed notable empirical successes in machine learning and data analysis [25, 23]. However, even for this seemingly simple learning model, there are very few works addressing the development of fast algorithms for its Wasserstein DRO formulation, which takes the form $\inf_w \{ \frac{c}{2} \|w\|_2^2 + \sup_{\mathbb{Q} \in B_\epsilon^p(\hat{\mathbb{P}}_n)} \mathbb{E}_{(x,y) \sim \mathbb{Q}} [\ell_w(x,y)] \}$ and can be reformulated as

$$\min_{w,\lambda} \lambda \epsilon + \frac{1}{n} \sum_{i=1}^n \max \{ 1 - w^T z_i, 1 + w^T z_i - \lambda \kappa, 0 \} + \frac{c}{2} \|w\|_2^2, \text{ s.t. } \|w\|_q \leq \lambda; \quad (1)$$

see [12, Theorem 2] and [22, Theorem 3.11]. Problem (1) arises from the vanilla soft-margin SVM model. Here, $\frac{c}{2} \|w\|_2^2$ is the regularization term with $c \geq 0$; $x \in \mathbb{R}^d$ denotes a feature vector and $y \in \{-1, +1\}$ is the associated binary label; $\ell_w(x,y) = \max\{1 - yw^T x, 0\}$ is the hinge loss w.r.t. the feature-label pair (x,y) and learning parameter $w \in \mathbb{R}^d$; $\{(\hat{x}_i, \hat{y}_i)\}_{i=1}^n$ are n training samples independently and identically drawn from an unknown distribution \mathbb{P}^* on the feature-label space $\mathcal{Z} = \mathbb{R}^d \times \{+1, -1\}$ and $z_i = \hat{x}_i \odot \hat{y}_i$; $\hat{\mathbb{P}}_n = \frac{1}{n} \sum_{i=1}^n \delta_{(\hat{x}_i, \hat{y}_i)}$ is the empirical distribution associated with the training samples; $B_\epsilon^p(\hat{\mathbb{P}}_n) = \{\mathbb{Q} \in \mathcal{P}(\mathcal{Z}) : W_p(\mathbb{Q}, \hat{\mathbb{P}}_n) \leq \epsilon\}$ is the ambiguity set defined on the space of probability distributions $\mathcal{P}(\mathcal{Z})$ centered at the empirical distribution $\hat{\mathbb{P}}_n$ and has radius $\epsilon \geq 0$ w.r.t. the ℓ_p norm-induced Wasserstein distance

$$W_p(\mathbb{Q}, \hat{\mathbb{P}}_n) = \inf_{\Pi \in \mathcal{P}(\mathcal{Z} \times \mathcal{Z})} \left\{ \int_{\mathcal{Z} \times \mathcal{Z}} d_p(\xi, \xi') \Pi(d\xi, d\xi') : \Pi(d\xi, \mathcal{Z}) = \mathbb{Q}(d\xi), \Pi(\mathcal{Z}, d\xi') = \hat{\mathbb{P}}_n(d\xi') \right\},$$

where $\xi = (x,y) \in \mathcal{Z}$, $\frac{1}{p} + \frac{1}{q} = 1$, and $d_p(\xi, \xi') = \|x - x'\|_p + \frac{\kappa}{2} |y - y'|$ is the transport cost between two data points $\xi, \xi' \in \mathcal{Z}$ with $\kappa \geq 0$ representing the relative emphasis between feature mismatch and label uncertainty. In particular, the larger the κ , the more reliable are the labels; see [22, 15] for further details. Intuitively, if the ambiguity set $B_\epsilon^p(\hat{\mathbb{P}}_n)$ contains the ground-truth distribution \mathbb{P}^* , then the estimator w^* obtained from an optimal solution to (1) will be less sensitive to unseen feature-label pairs.

In the works [12, 18], the authors proposed cutting surface-based methods to solve the ℓ_p -DRSVM problem (1). However, in their implementation, they still need to invoke off-the-shelf solvers for certain tasks. Recently, researchers have proposed to use stochastic (sub)gradient descent to tackle a class of Wasserstein DRO problems [5, 24]. Nevertheless, the results in [5, 24] do not apply to the ℓ_p -DRSVM problem (1), as they require $\kappa = \infty$; i.e., the labels are error-free. Moreover, the transport cost d_p does not satisfy the strong convexity-type condition in [5, Assumption 1] or [24, Assumption A]. On another front, the authors of [15] introduced an ADMM-based first-order algorithmic framework to deal with the Wasserstein distributionally robust logistic regression problem. Though the framework in [15] can be extended to handle the ℓ_p -DRSVM problem (1), it has two main drawbacks. First, under the framework, the optimal λ^* of problem (1) is found by an one-dimensional search, where each update involves fixing λ to a given value and solving for the corresponding optimal $w^*(\lambda)$ (which we refer to as the w -subproblem). Since the number of w -subproblems that arise during the search can be large, the framework is computationally rather demanding. Second, the w -subproblem is solved by an ADMM-type algorithm, which involves both primal and dual updates. In order to establish fast (e.g., linear) convergence rate guarantee for the algorithm, one typically requires a regularity condition on the set of primal-dual optimal pairs of the problem at hand. Unfortunately, it is not clear whether the ℓ_p -DRSVM problem (1) satisfies such a primal-dual regularity condition.

To overcome these drawbacks, we propose two new epigraphical projection-based incremental algorithms for solving the ℓ_p -DRSVM problem (1), which tackle the variables (w, λ) jointly. We focus on the commonly used ℓ_1 , ℓ_2 , and ℓ_∞ norm-induced transport costs, which correspond to $q \in \{1, 2, \infty\}$. Our first algorithm is the incremental projected subgradient descent (ISG) method, whose efficiency inherits from that of the projection onto the epigraph $\{(w, \lambda) : \|w\|_q \leq \lambda\}$ of the ℓ_q

Table 1: Convergence rates of incremental algorithms for ℓ_p -DRSVM

q	c	Hölderian growth	Step size scheme	Convergence rate
$q = 1, \infty$	$c = 0$	Sharp [8, Theorem 3.5]	$\alpha_{k+1} = \rho\alpha_k, \rho \in (0, 1)$	$\mathcal{O}(\rho^k)$
$q = 1, \infty$	$c > 0$	QG [28, Proposition 6]	$\alpha_k = \frac{\gamma}{nk}, \gamma > 0$	$\mathcal{O}(\frac{1}{k})$
$q = 2$	$c = 0$	Sharp (BLR) Not Known	$\alpha_{k+1} = \rho\alpha_k, \rho \in (0, 1)$ $\alpha_k = \frac{\gamma}{n\sqrt{k}}, \gamma > 0$	$\mathcal{O}(\rho^k)$ $\mathcal{O}(\frac{1}{\sqrt{k}})$
$q = 2$	$c > 0$	QG (BLR) Not Known	$\alpha_k = \frac{\gamma}{nk}, \gamma > 0$ $\alpha_k = \frac{\gamma}{n\sqrt{k}}, \gamma > 0$	$\mathcal{O}(\frac{1}{k})$ $\mathcal{O}(\frac{1}{\sqrt{k}})$

BLR: The result holds under the assumption of bounded linear regularity (BLR) (see Definition 2).
 Not Known: Without BLR, it is not known whether the Hölderian growth condition holds.

norm (with $q \in \{1, 2, \infty\}$). The second is the incremental proximal point algorithm (IPPA). Although in general IPPA is less sensitive to the choice of initial step size and can achieve better accuracy than ISG [16], in the context of the ℓ_p -DRSVM problem (1), each iteration of IPPA requires solving the following subproblem, which we refer to as the *single-sample proximal point update*:

$$\min_{w, \lambda} \max \{1 - w^T z_i, 1 + w^T z_i - \lambda\kappa, 0\} + \frac{1}{2\alpha} (\|w - \bar{w}\|_2^2 + (\lambda - \bar{\lambda})^2), \quad \text{s.t. } \|w\|_q \leq \lambda. \quad (2)$$

Here, $\alpha > 0$ is the step size, $q \in \{1, 2, \infty\}$, and $\bar{w}, \bar{\lambda}$ are given. By carefully exploiting the problem structure, we develop exceptionally efficient solutions to (2). Specifically, we show in Section 3 that the optimal solution to (2) admits an analytic form when $q = 2$ and can be computed by a fast algorithm based on a parametric approach and a modified secant method (cf. [9]) when $q = 1$ or ∞ .

Next, we investigate the convergence behavior of the proposed ISG and IPPA when applied to problem (1). Our main tool is the following regularity notion:

Definition 1 (Hölderian growth condition [6]) A function $f : \mathbb{R}^m \rightarrow \mathbb{R}$ is said to satisfy a Hölderian growth condition on the domain $\Omega \subseteq \mathbb{R}^m$ if there exist constants $\theta \in [0, 1]$ and $\sigma > 0$ such that

$$\text{dist}(x, \mathcal{X}) \leq \sigma^{-1} (f(x) - f^*)^\theta, \quad \forall x \in \Omega, \quad (3)$$

where \mathcal{X} denotes the optimal set of $\min_{x \in \Omega} f(x)$ and f^* is the optimal value. The condition (3) is known as sharpness when $\theta = 1$ and quadratic growth (QG) when $\theta = \frac{1}{2}$; see, e.g., [7].

We show that for different choices of $q \in \{1, 2, \infty\}$ and $c \geq 0$, the DRSVM problem (1) satisfies either the sharpness condition or QG condition; see Table 1. With the exception of the case $q \in \{1, \infty\}$, where the sharpness (resp. QG) of (1) when $c = 0$ (resp. $c > 0$) essentially follows from [8, Theorem 3.5] (resp. [28, Proposition 6]), the results on the Hölderian growth of problem (1) are new. Consequently, by choosing step sizes that decay at a suitable rate, we establish, for the first time, the fast sublinear (i.e., $\mathcal{O}(\frac{1}{k})$) or linear (i.e., $\mathcal{O}(\rho^k)$) convergence rate of the proposed incremental algorithms when applied to the DRSVM problem (1); see Table 1.

Lastly, we demonstrate the efficiency of our proposed methods through extensive numerical experiments on both synthetic and real data sets. It is worth mentioning that our proposed algorithms can be easily extended to an asynchronous decentralized parallel setting and thus can further meet the requirements of large-scale applications.

2 Epigraphical Projection-based Incremental Algorithms

In this section, we present our incremental algorithms for solving the ℓ_p -DRSVM problem. For simplicity, we focus on the case $c = 0$ in what follows. Our technical development can be extended to handle the general case $c \geq 0$ by noting that the subproblems corresponding to the cases $c = 0$ and $c > 0$ share the same structure.

To begin, observe that the ℓ_p -DRSVM problem (1) with $c = 0$ can be written compactly as

$$\min_{\|w\|_q \leq \lambda} \frac{1}{n} \sum_{i=1}^n f_i(w, \lambda), \quad (4)$$

where $f_i(w, \lambda) = \lambda\epsilon + \max\{1 - w^T z_i, 1 + w^T z_i - \lambda\kappa, 0\}$ is a piecewise affine function. Since problem (4) possesses the vanilla finite-sum structure with a single epigraphical projection constraint, a natural and widely adopted approach to tackling it is to use incremental algorithms. Roughly speaking, such algorithms select one mini-batch of component functions from the objective in (4) at a time based on a certain cyclic order and use the selected functions to update the current iterate. We shall focus on the following two incremental algorithms for solving the DRSVM problem (1). Here, k is the epoch index (i.e., the k -th time going through the cyclic order) and $\alpha_k > 0$ is the step size in the k -th epoch.

Incremental Mini-batch Projected Subgradient Algorithm (ISG)

$$(w_{i+1}^k, \lambda_{i+1}^k) = \text{proj}_{\{\|w\|_q \leq \lambda\}} [(w_i^k, \lambda_i^k) - \alpha_k g_i^k], \quad (5)$$

where g_i^k is a subgradient of $\frac{1}{|B_i|} \sum_{j \in B_i} f_j$ at (w_i^k, λ_i^k) and $B_i \subseteq \{1, \dots, n\}$ is the i -th mini-batch.

Incremental Proximal Point Algorithm (IPPA)

$$(w_{i+1}^k, \lambda_{i+1}^k) = \arg \min_{\|w\|_q \leq \lambda} \left\{ f_i(w, \lambda) + \frac{1}{2\alpha_k} (\|w - w_i^k\|_2^2 + (\lambda - \lambda_i^k)^2) \right\}, \quad (6)$$

where $(w_n^k, \lambda_n^k) = (w_0^{k+1}, \lambda_0^{k+1})$.

Now, a natural question is how to solve the subproblems (5) and (6) efficiently. As it turns out, the key lies in an efficient implementation of the ℓ_q norm epigraphical projection (with $q \in \{1, 2, \infty\}$). Indeed, such a projection appears explicitly in the ISG update (5) and, as we shall see later, plays a vital role in the design of fast iterative algorithms for the single-sample proximal point update (6). To begin, we note that the ℓ_2 norm epigraphical projection $\text{proj}_{\{\|w\|_2 \leq \lambda\}}$ has a well-known analytic solution; see [1, Theorem 3.3.6]. Next, the ℓ_1 norm epigraphical projection $\text{proj}_{\{\|w\|_1 \leq \lambda\}}$ can be found in linear time using the quick-select algorithm; see [26]. Lastly, the ℓ_∞ norm epigraphical projection $\text{proj}_{\{\|w\|_\infty \leq \lambda\}}$ can be computed in linear time via the Moreau decomposition

$$\text{proj}_{\{\|w\|_\infty \leq \lambda\}}(x, s) = (x, s) + \text{proj}_{\{\|w\|_1 \leq \lambda\}}(-x, -s).$$

From the above discussion, we see that the ISG update (5) can be computed efficiently. In the next section, we discuss how these epigraphical projections can be used to perform the single-sample proximal point update (6) in an efficient manner.

3 Fast Algorithms for Single-Sample Proximal Point Update (6)

Analytic solution for $q = 2$. We begin with the case $q = 2$. By combining the terms $\lambda\epsilon$ and $\frac{1}{2\alpha_k}(\lambda - \lambda_i^k)^2$ in (6), we see that the single-sample proximal point update takes the form (cf. (2))

$$\min_{w, \lambda} \underbrace{\max\{1 - w^T z_i, 1 + w^T z_i - \lambda\kappa, 0\}}_{h_i(w, \lambda)} + \frac{1}{2\alpha} (\|w - \bar{w}\|_2^2 + (\lambda - \bar{\lambda})^2), \quad \text{s.t. } \|w\|_2 \leq \lambda. \quad (7)$$

The main difficulty of (7) lies in the piecewise affine term h_i . To handle this term, let $h_{i,1}(w, \lambda) = 1 - w^T z_i$, $h_{i,2}(w, \lambda) = 1 + w^T z_i - \lambda\kappa$, and $h_{i,3}(w, \lambda) = 0$, so that $h_i = \max_{j \in \{1, 2, 3\}} h_{i,j}$. Observe that if (w^*, λ^*) is an optimal solution to (7), then there could only be one, two, or three affine pieces in h_i that are active at (w^*, λ^*) ; i.e., $\Gamma = |\{j : h_i(w^*, \lambda^*) = h_{i,j}(w^*, \lambda^*)\}| \in \{1, 2, 3\}$. This suggests that we can find (w^*, λ^*) by exhausting these possibilities. Due to space limitation, we only give an outline of our strategy here. The details can be found in the Appendix.

We start with the case $\Gamma = 1$. For $j = 1, 2, 3$, consider the following problem, which corresponds to the subcase where $h_{i,j}$ is the only active affine piece:

$$\min_{w, \lambda} h_{i,j}(w, \lambda) + \frac{1}{2\alpha} (\|w - \bar{w}\|_2^2 + (\lambda - \bar{\lambda})^2), \quad \text{s.t. } \|w\|_2 \leq \lambda. \quad (8)$$

Since $h_{i,j}$ is affine in (w, λ) , it is easy to verify that problem (8) reduces to an ℓ_2 norm epigraphical projection, which admits an analytic solution, say $(\hat{w}_j, \hat{\lambda}_j)$. If there exists a $j' \in \{1, 2, 3\}$ such

that $h_{i,j'}(\hat{w}_{j'}, \hat{\lambda}_{j'}) > h_{i,j}(\hat{w}_{j'}, \hat{\lambda}_{j'})$ for $j \neq j'$, then we know that $(\hat{w}_{j'}, \hat{\lambda}_{j'})$ is optimal for (7) and hence we can terminate the process. Otherwise, we proceed to the case $\Gamma = 2$ and consider, for $1 \leq j < j' \leq 3$, the following problem, which corresponds to the subcase where $h_{i,j}$ and $h_{i,j'}$ are the only two active affine pieces:

$$\min_{w,\lambda} h_{i,j}(w, \lambda) + \frac{1}{2\alpha} (\|w - \bar{w}\|_2^2 + (\lambda - \bar{\lambda})^2), \quad \text{s.t. } h_{i,j}(w, \lambda) = h_{i,j'}(w, \lambda), \quad \|w\|_2 \leq \lambda. \quad (9)$$

As shown in the Appendix (Proposition 6.2), the optimal solution to (9) can be found by solving a univariate quartic equation, which can be done efficiently. Now, let $(\hat{w}_{(j,j')}, \hat{\lambda}_{(j,j')})$ be the optimal solution to (9). If there exist j, j' with $1 \leq j < j' \leq 3$ such that $h_{i,j}(\hat{w}_{(j,j')}, \hat{\lambda}_{(j,j')}) = h_{i,j'}(\hat{w}_{(j,j')}, \hat{\lambda}_{(j,j')}) > h_{i,j''}(\hat{w}_{(j,j')}, \hat{\lambda}_{(j,j')})$ with $j'' \in \{1, 2, 3\} \setminus \{j, j'\}$, then $(\hat{w}_{(j,j')}, \hat{\lambda}_{(j,j')})$ is optimal for (7) and we can terminate the process. Otherwise, we proceed to the case $\Gamma = 3$. In this case, we consider the problem

$$\min_{w,\lambda} \frac{1}{2\alpha} (\|w - \bar{w}\|_2^2 + (\lambda - \bar{\lambda})^2), \quad \text{s.t. } h_{i,1}(w, \lambda) = h_{i,2}(w, \lambda) = h_{i,3}(w, \lambda), \quad \|w\|_2 \leq \lambda,$$

which reduces to

$$\min_w \frac{1}{2\alpha} \|w - \bar{w}\|_2^2, \quad \text{s.t. } w^T z_i = 1, \quad \|w\|_2 \leq \frac{2}{\kappa}. \quad (10)$$

It can be shown that problem (10) admits an analytic solution \hat{w} ; see the Appendix (Proposition 6.4). Then, the pair $(\hat{w}, \frac{2}{\kappa})$ is an optimal solution to (7).

Fast iterative algorithm for $q = 1$. The high-level idea is similar to that for the case $q = 2$; i.e., we systematically go through all valid subcollections of the affine pieces in h_i and test whether they can be active at the optimal solution to the single-sample proximal point update. The main difference here is that the subproblems arising from the subcollections do not necessarily admit analytic solutions. To overcome this difficulty, we propose a modified secant algorithm (cf. [9]) to search for the optimal dual multiplier of the subproblem and use it to recover the optimal solution to the original subproblem via ℓ_1 norm epigraphical projection. Again, we give an outline of our strategy here and relegate the details to the Appendix.

To begin, we rewrite the single-sample proximal point update (6) for the case $q = 1$ as

$$\begin{aligned} \min_{w,\lambda,\mu} \quad & \mu + \frac{1}{2\alpha} (\|w - \bar{w}\|_2^2 + (\lambda - \bar{\lambda})^2) \\ \text{s.t.} \quad & h_{i,j}(w, \lambda) \leq \mu, \quad j = 1, 2, 3; \quad \|w\|_1 \leq \lambda. \end{aligned} \quad (11)$$

For reason that would become clear in a moment, we shall not go through the cases $\Gamma = 1, 2, 3$ as before. Instead, consider first the case where $h_{i,3}$ is inactive. If $h_{i,1}$ is also inactive, then we consider the problem $\min_{w,\lambda} h_{i,2}(w, \lambda) + \frac{1}{2\alpha} (\|w - \bar{w}\|_2^2 + (\lambda - \bar{\lambda})^2)$, which, by the affine nature of $h_{i,2}$, is equivalent to an ℓ_1 norm epigraphical projection. If $h_{i,1}$ is active, then we consider the problem

$$\min_{w,\lambda} h_{i,1}(w, \lambda) + \frac{1}{2\alpha} (\|w - \bar{w}\|_2^2 + (\lambda - \bar{\lambda})^2), \quad \text{s.t. } h_{i,2}(w, \lambda) \leq h_{i,1}(w, \lambda), \quad \|w\|_1 \leq \lambda. \quad (12)$$

Note that $h_{i,2}$ can be active or inactive, and the constraint $h_{i,2}(w, \lambda) \leq h_{i,1}(w, \lambda)$ allows us to treat both possibilities simultaneously. Hence, we do not need to tackle them separately as we did in the case $q = 2$. Observe that problem (12) can be cast into the form

$$\min_{w,\lambda} \frac{1}{2\alpha} (\|w - \bar{w}\|_2^2 + (\lambda - \bar{\lambda})^2), \quad \text{s.t. } w^T z \leq a\lambda + b \quad (\leftarrow \sigma \geq 0), \quad \|w\|_1 \leq \lambda, \quad (13)$$

where, with an abuse of notation, we use $\bar{w} \in \mathbb{R}^d$, $\bar{\lambda} \in \mathbb{R}$ here again and caution the reader that they are different from those in (12), and $z = z_i$, $a = \frac{\kappa}{2}$, $b = 0$. Before we discuss how to solve the subproblem (13), let us note that it arises in the case where $h_{i,3}$ is active as well. Indeed, if $h_{i,3}$ is active and $h_{i,1}$ is inactive, then we have $z = z_i$, $a = \kappa$, $b = -1$, which corresponds to the constraint $h_{i,2}(w, \lambda) \leq h_{i,3}(w, \lambda)$ and covers the possibilities that $h_{i,2}$ is active and inactive. On the other hand, if $h_{i,3}$ is active and $h_{i,2}$ is inactive, then we have $z = -z_i$, $a = 0$, $b = -1$, which corresponds to the constraint $h_{i,1}(w, \lambda) \leq h_{i,3}(w, \lambda)$ and covers the possibilities that $h_{i,1}$ is active and inactive. The only remaining case is when $h_{i,1}, h_{i,2}, h_{i,3}$ are all active. In this case, we consider problem (10) with

$\|w\|_2 \leq \frac{2}{\kappa}$ replaced by $\|w\|_1 \leq \frac{2}{\kappa}$. As shown in the Appendix, such a problem can be tackled using the technique for solving (13). We go through the above cases sequentially and terminate the process if the solution to the subproblem in any one of the cases satisfies the optimality conditions of (11).

Now, let us return to the issue of solving (13). The main idea is to perform an one-dimensional search on the dual variable σ to find the optimal dual multiplier σ^* . Specifically, consider the problem

$$\min_{\|w\|_1 \leq \lambda} \frac{1}{2\alpha} (\|w - \bar{w}\|_2^2 + (\lambda - \bar{\lambda})^2) + \sigma(w^T z - a\kappa - b). \quad (14)$$


Let $(\hat{w}(\sigma), \hat{\lambda}(\sigma))$ be the optimal solution to (14) and define the function $p : \mathbb{R}_+ \rightarrow \mathbb{R}$ by $p(\sigma) = \hat{w}(\sigma)^T z - a\kappa - b$. Inspired by [17], we establish the following monotonicity property of p , which will be crucial to our development of an extremely efficient algorithm for solving (13) later.

Proposition 3.1 *If σ satisfies (i) $\sigma = 0$ and $p(\sigma) \leq 0$, or (ii) $p(\sigma) = 0$, then $(\hat{w}(\sigma), \hat{\lambda}(\sigma))$ is the optimal solution to (13). Moreover, p is continuous and monotonically non-increasing on \mathbb{R}_+ .*

In view of Proposition 3.1, we first check if $p(0) \leq 0$ via an ℓ_1 norm epigraphical projection. If not, then we search for the $\sigma^* \geq 0$ that satisfies $p(\sigma^*) = 0$ by the secant method, with some special modifications designed to speed up its convergence [9]. Let us now give a high-level description of our modified secant method. We refer the reader to the Appendix (Algorithm 1) for details.

At the beginning of a generic iteration of the method, we have an interval $[\sigma_l, \sigma_u]$ that contains σ^* , with $r_l = -p(\sigma_l) < 0$ and $r_u = -p(\sigma_u) > 0$. The initial interval can be found by considering the optimality conditions of (11) (i.e., $\sigma^* \in [0, 1]$). We then take a secant step to get a new point σ with $r = -p(\sigma)$ and perform the update on σ_l, σ_u as follows.

Suppose that $r > 0$. If σ lies in the left-half of the interval (i.e., $\sigma < \frac{\sigma_l + \sigma_u}{2}$), then we update σ_u to σ . Otherwise, we take an auxiliary secant step based on σ and σ_u to get a point σ' , and we update σ_u to $\max\{\sigma', 0.6\sigma_l + 0.4\sigma\}$. Such a choice ensures that the interval length is reduced by a factor of 0.6 or less. The case where $r < 0$ is similar, except that σ_l is updated. If $r = 0$, then by Proposition 3.1 we have found the optimal dual multiplier σ^* and hence can terminate.



Finally, for the case $q = \infty$, we can follow the same procedure as the case $q = 1$. The details can be found in the Appendix.

4 Convergence Rate Analysis of Incremental Algorithms

In this section, we study the convergence behavior of our proposed incremental methods ISG and IPPA. Our starting point is to understand the conditions under which the ℓ_p -DRSVM problem (1) possesses the Hölderian growth condition (3). Then, by determining the value of the growth exponent θ and using it to choose step sizes that decay at a suitable rate, we can establish the convergence rates of ISG and IPPA. To begin, let us consider problem (1) with $q \in \{1, \infty\}$. If $c = 0$, then problem (1) satisfies the sharpness (i.e., $\theta = 1$) condition. This follows essentially from [8, Theorem 3.5], as the objective of (1) has polyhedral epigraph and the constraint is polyhedral. On the other hand, if $c > 0$, then since the piecewise affine term in (1) has a polyhedral epigraph and the constraint is polyhedral, we can invoke [28, Proposition 6] and conclude that problem (1) satisfies the QG (i.e., $\theta = \frac{1}{2}$) condition.

Next, let us consider the case $q = 2$. From the above discussion, one may expect that similar conclusions hold for this case. However, as the following example shows, this case is more subtle and requires a more careful treatment.

Example 4.1 *Consider the problem $\min_{w, \lambda} 0.1\lambda + |1 - w_1|$, s.t. $\sqrt{w_1^2 + w_2^2} \leq \lambda$, which is an instance of (1) with $q = 2$, $c = 0$. It is easy to verify that the optimal solution is $w^* = (1, 0)$, $\lambda^* = 1$. Consider feasible points of the form $(w_1, w_2, \lambda) = (w_1, \sqrt{1 - w_1^2}, 1)$, which tend to (w^*, λ^*) as w_1 tends to 1. A simple calculation yields $\text{dist}((w_1, w_2, \lambda), (w^*, \lambda^*)) = \sqrt{2|1 - w_1|} = \omega(|1 - w_1|)$, which shows that the instance cannot satisfy the sharpness condition.*

As it turns out, it is still possible to establish the sharpness or QG condition for problem (1) with $q = 2$ under a well-known sufficient condition called *bounded linear regularity*. Let us begin with the definition.

Definition 2 (Bounded linear regularity [2, Definition 5.6]) *Let C_1, \dots, C_N be closed convex subsets of \mathbb{R}^d with a non-empty intersection C . We say that the collection $\{C_1, \dots, C_N\}$ is bounded linearly regular (BLR) if for every bounded subset \mathcal{B} of \mathbb{R}^d , there exists a constant $\kappa > 0$ such that*

$$\text{dist}(x, C) \leq \kappa \max_{i \in \{1, \dots, N\}} \text{dist}(x, C_i), \text{ for all } x \in \mathcal{B}.$$

Using the above definition, we can establish the following result; see the Appendix for the proof.

Proposition 4.2 *Consider problem (1) with $q = 2$. Let \mathcal{X} be the set of optimal solutions and $L_2^d = \{(w, \lambda) \in \mathbb{R}^d \times \mathbb{R} : \|w\|_2 \leq \lambda\}$ be the constraint set. Suppose that $\mathcal{X} \cap \text{ri}(L_2^d) \neq \emptyset$. Consequently, problem (1) satisfies the sharpness condition when $c = 0$ and the QG condition when $c > 0$.*

By combining Proposition 4.2 with an appropriate choice of step sizes, we obtain the following convergence rate guarantees for ISG and IPPA. The proof can be found in the Appendix.

Theorem 4.3 *Let $\{x^k = (w_0^k, \lambda_0^k)\}$ be the sequence of iterates generated by ISG or IPPA.*

- (1) *If problem (1) satisfies the sharpness condition, then by choosing the geometrically diminishing step sizes $\alpha_k = \alpha_0 \rho^k$ with $\alpha_0 \geq \frac{\sigma \text{dist}(x^0, \mathcal{X})}{2L^2 n}$ and $\sqrt{1 - \frac{\sigma^2}{2L^2}} \leq \rho < 1$, the sequence $\{x^k\}$ converges linearly to an optimal solution to (1); i.e., $\text{dist}(x^k, \mathcal{X}) \leq \mathcal{O}(\rho^k)$ for all $k \geq 0$.*
- (2) *If problem (1) satisfies the quadratic growth condition, then by choosing the polynomially decaying step sizes $\alpha_k = \frac{\gamma}{nk}$ with $\gamma > \frac{1}{2\sigma}$, the sequence $\{x^k\}$ converges to an optimal solution to (1) at the rate $\mathcal{O}(\frac{1}{\sqrt{k}})$ and $\{f(x^k) - f^*\}$ converges to zero at the rate $\mathcal{O}(\frac{1}{k})$.*
- (3) *(See [20, Proposition 2.10]) For the general convex problem (1), by choosing the step sizes $\alpha_k = \frac{\gamma}{n\sqrt{k}}$ with $\gamma > 0$, the sequence $\{\min_{0 \leq k \leq K} f(x^k) - f^*\}$ converges to zero at the rate $\mathcal{O}(\frac{1}{\sqrt{K}})$.*

5 Experiment Results

In this section, we present numerical results to demonstrate the efficiency of our proposed incremental methods. All simulations are implemented using MATLAB R2019b on a computer running Windows 10 with a 3.20 GHz, the Intel(R) Core(TM) i7-8700 processor, and 16 GB of RAM. To begin, we evaluate our two proposed incremental methods ISG and IPPA in different settings to corroborate our theoretical results in Section 4 and to better understand their empirical strengths and weaknesses. Based on this, we develop a hybrid algorithm that combines the advantages of both ISG and IPPA to further speed up the convergence in practice. Next, we compare the wall-clock time of our algorithms with GS-ADMM [15] and YALMIP (i.e., IPOPT) solver on real datasets. For sake of fairness, we only extend the first-order algorithmic framework (referred to as GS-ADMM) to tackle the ℓ_∞ -DRSVM problem. In fact, the faster inner solver conjugate gradient with an active set method can only tackle the ℓ_∞ case in [15]. The implementation details to reproduce all numerical results in this section are given in the Appendix. Our code is available at https://github.com/gerrili1996/Incremental_DRSVM.

5.1 Synthetic data: Different regularity conditions and their step size schemes

Our setup for the synthetic experiments is as follows. First, we generate the learning parameter w^* and feature vectors $\{x_i\}_{i=1}^n$ independently and identically (i.i.d) from the standard normal distribution $\mathcal{N}(0, I_d)$ and the noisy measurements $\{\xi_i\}_{i=1}^n$ i.i.d from $\mathcal{N}(0, \sigma^2 I_d)$ (e.g., $\sigma = 0.5$). Then, we compute the ground-truth labels $\{y_i\}_{i=1}^n$ by $y_i = \text{sign}(\langle w^*, x_i \rangle + \xi_i)$. Here, the model parameters are $n = 1000$, $d = 100$, $\kappa = 1$, $\epsilon = 0.1$. All the algorithmic parameters of ISG and IPPA have been fine-tuned via grid search for optimal performance. Recall from Theorem 4.3 that for instances satisfying the sharpness condition, the smaller shrinking rate ρ the algorithm can adopt, the faster its linear rate of convergence. The experiments results in Fig. 1(a–c) indicate that IPPA allows us

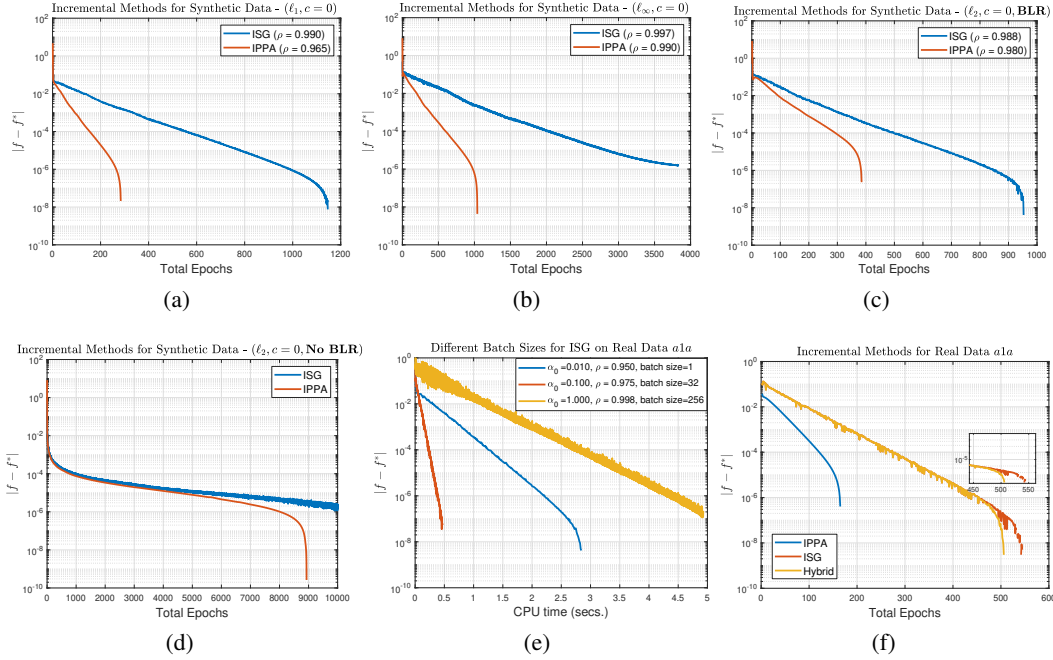


Figure 1: (a)–(d): Comparison between ISG and IPPA on both BLR and non-BLR instances generated from synthetic datasets. (e)–(f): Performance of ISG on different mini-batch sizes and performance of the hybrid algorithm on the a1a dataset.

to choose a more aggressive ρ when compared with ISG over all instances satisfying the sharpness condition. A similar phenomenon has also been observed in previous works; see, e.g., [16, Fig. 1]. Even for instances that do not satisfy the sharpness condition, IPPA performs better than ISG; see Fig. 1(d).

Nevertheless, IPPA can only handle one sample at a time. Thus, we are motivated to develop an approach that can combine the best features of both ISG and IPPA. Towards that end, observe from Fig. 1(e) that there is a tradeoff between the mini-batch size and the shrinking rate ρ , which means that there is an optimal mini-batch size for achieving the fastest convergence speed. Inspired by this, we propose to first apply the mini-batch ISG to obtain an initial point and then use IPPA in a local region around the optimal point to gain further speedup and get a more accurate solution. As shown in Fig. 1(f), such a hybrid algorithm is effective, thus confirming our intuition.

5.2 Efficiency of our incremental algorithms

Next, we demonstrate the efficiency of our proposed methods on the real datasets a1a–a9a, i j cnn1 downloaded from the LIBSVM². The results for ℓ_1 -DRSVM, which satisfies the sharpness condition, are shown in Table 2. Apparently, IPPA is slower than mini-batch ISG (i.e., M-ISG) in general but can obtain more accurate solutions. More importantly, the hybrid algorithm, which combines the advantages of both M-ISG and IPPA, has an excellent performance and achieves a well-balanced tradeoff between accuracy and efficiency. All of them are much faster than YALMIP. The results for ℓ_2 -DRSVM are reported in Table 3. As ISG is sensitive to hyper-parameters and has difficulty achieving the desired accuracy, we only present the results for IPPA. From the table, the superiority of IPPA over the solver is obvious.

To further demonstrate the efficiency of our proposed hybrid algorithm, we compare it with GS-ADMM [15] and YALMIP on ℓ_∞ -DRSVM, which again satisfies the sharpness condition. The results are shown in Table 4. The overall performance of our hybrid method dominates both GS-ADMM and YALMIP. Due to space limitation, we only present the results for the case $q \in \{1, 2, \infty\}$, $c = 0$. More numerical results can be found in the Appendix.

²<https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/binary.html>

Table 2: Wall-clock Time Comparison on UCI Real Dataset: ℓ_1 -DRSVM, $c = 0, \kappa = 1, \epsilon = 0.1$

	Objective Value				Wall-clock time (sec)			
	M-ISG	IPPA	Hybrid	YALMIP	M-ISG	IPPA	Hybrid	YALMIP
a1a	0.651090	0.651091	0.651090	0.651102	0.706	6.1242	1.560	12.221
a2a	0.670640	0.670640	0.670640	0.670652	0.717	7.040	1.720	9.695
a3a	0.662962	0.663093	0.662962	0.663060	1.800	21.242	3.740	11.854
a4a	0.674274	0.674274	0.674273	0.674274	3.764	25.980	4.664	16.638
a5a	0.660867	0.660867	0.660867	0.660869	2.026	24.752	24.752	24.207
a6a	0.654189	0.654189	0.654189	0.654194	2.277	26.127	2.509	39.311
a7a	0.656274	0.656274	0.656273	0.656411	2.528	33.094	2.799	60.046
a8a	0.650036	0.650036	0.650035	0.650081	3.004	41.249	3.729	94.377
a9a	0.642186	0.642186	0.642185	0.642596	2.285	35.554	3.063	155.980

Table 3: Wall-clock Time Comparison on UCI Real Dataset: ℓ_2 -DRSVM, $c = 0, \kappa = 1, \epsilon = 0.1$

	Objective Value		Wall-clock time (sec)		Regularity Condition
	IPPA	YALMIP	IPPA	YALMIP	
a1a	0.6339472	0.6338819	5.517	8.557	Not Known
a2a	0.6599856	0.6599099	9.355	11.989	Not Known
a3a	0.6443777	0.6442762	7.096	15.335	Not Known
a4a	0.6513987	0.6513899	14.162	23.122	Not Known
a5a	0.6484421	0.6484147	10.515	32.663	Not Known
a6a	0.6428831	0.6428806	15.195	67.695	Not Known
a7a	0.6459271	0.6462302	6.454	118.740	Not Known
a8a	0.6441057	0.6441057	27.242	161.000	Not Known
a9a	0.6389162	0.6437767	13.129	215.387	Sharpness
ijcnn	0.4781876	0.4781897	20.567	379.943	Sharpness

Table 4: Wall-clock Time Comparison on UCI Real Dataset: ℓ_∞ -DRSVM, $c = 0, \kappa = 1, \epsilon = 0.1$

	Hybrid	GS-ADMM	YALMIP		Hybrid	GS-ADMM	YALMIP
a1a	4.789	5.939	7.832	a6a	8.273	8.273	42.714
a2a	5.098	7.069	9.100	a7a	6.115	6.115	60.743
a3a	16.252	9.638	11.375	a8a	11.065	11.065	99.355
a4a	5.498	10.446	17.542	a9a	5.717	5.717	172.07
a5a	7.363	13.993	22.969	ijcnn	4.301	4.301	319.379

6 Conclusion and Future Work

In this paper, we developed two new and highly efficient epigraphical projection-based incremental algorithms to solve the Wasserstein DRSVM problem with ℓ_p norm-induced transport cost ($p \in \{1, 2, \infty\}$) and established their convergence rates. A natural future direction is to develop a mini-batch version of IPPA and extend our algorithms to the asynchronous decentralized parallel setting. Inspired by our paper, it would also be interesting to develop some new incremental/stochastic algorithms to tackle more general Wasserstein DRO problems; see, e.g., problem (11) in [19].

Acknowledgment Caihua Chen is supported in part by the National Natural Science Foundation of China (NSFC) projects 71732003, 11871269 and in part by the Natural Science Foundation of Jiangsu Province project BK20181259. Anthony Man-Cho So is supported in part by the CUHK Research Sustainability of Major RGC Funding Schemes project 3133236.

Broader Impact This work does not present any foreseeable societal consequence. A broader impact discussion is not applicable.

References

- [1] Heinz H. Bauschke. *Projection Algorithms and Monotone Operators*. PhD thesis, Simon Fraser University, 1996.
- [2] Heinz H Bauschke and Jonathan M Borwein. On projection algorithms for solving convex feasibility problems. *SIAM Review*, 38(3):367–426, 1996.
- [3] Dimitri P Bertsekas. Incremental proximal methods for large scale convex optimization. *Mathematical Programming*, 129(2):163–195, 2011.
- [4] Jose Blanchet, Yang Kang, and Karthyek Murthy. Robust Wasserstein profile inference and applications to machine learning. *Journal of Applied Probability*, 56(3):830–857, 2019.
- [5] Jose Blanchet, Karthyek Murthy, and Fan Zhang. Optimal transport based distributionally robust optimization: Structural properties and iterative schemes. *arXiv preprint arXiv:1810.02403*, 2018.
- [6] Jérôme Bolte, Trong Phong Nguyen, Juan Peypouquet, and Bruce W Suter. From error bounds to the complexity of first-order descent methods for convex functions. *Mathematical Programming*, 165(2):471–507, 2017.
- [7] J. Frédéric Bonnans and Alexander Shapiro. *Perturbation Analysis of Optimization Problems*. Springer Series in Operations Research. Springer–Verlag, New York, 2000.
- [8] James V Burke and Michael C Ferris. Weak sharp minima in mathematical programming. *SIAM Journal on Control and Optimization*, 31(5):1340–1359, 1993.
- [9] Yu-Hong Dai and Roger Fletcher. New algorithms for singly linearly constrained quadratic programs subject to lower and upper bounds. *Mathematical Programming*, 106(3):403–421, 2006.
- [10] Rui Gao, Xi Chen, and Anton J Kleywegt. Wasserstein distributional robustness and regularization in statistical learning. *arXiv preprint arXiv:1712.06050*, 2017.
- [11] Daniel Kuhn, Peyman Mohajerin Esfahani, Viet Anh Nguyen, and Soroosh Shafieezadeh-Abadeh. Wasserstein distributionally robust optimization: Theory and applications in machine learning. In *Operations Research & Management Science in the Age of Analytics*, pages 130–166. INFORMS, 2019.
- [12] Changhyeok Lee and Sanjay Mehrotra. A distributionally-robust approach for finding support vector machines. *Optimization Online*, 2015.
- [13] Jaeho Lee and Maxim Raginsky. Minimax statistical learning with Wasserstein distances. In *Advances in Neural Information Processing Systems*, pages 2687–2696, 2018.
- [14] Guoyin Li and Ting Kei Pong. Calculus of the exponent of Kurdyka–Łojasiewicz inequality and its applications to linear convergence of first-order methods. *Foundations of Computational Mathematics*, 18(5):1199–1232, 2018.
- [15] Jiajin Li, Sen Huang, and Anthony Man-Cho So. A first-order algorithmic framework for Wasserstein distributionally robust logistic regression. In *Advances in Neural Information Processing Systems*, pages 3939–3949, 2019.
- [16] Xiao Li, Zhihui Zhu, Anthony Man-Cho So, and Jason D Lee. Incremental methods for weakly convex optimization. *arXiv preprint arXiv:1907.11687*, 2019.
- [17] Meijiao Liu and Yong-Jin Liu. Fast algorithm for singly linearly constrained quadratic programs with box-like constraints. *Computational Optimization and Applications*, 66(2):309–326, 2017.
- [18] Fengqiao Luo and Sanjay Mehrotra. Decomposition algorithm for distributionally robust optimization using Wasserstein metric with an application to a class of regression models. *European Journal of Operational Research*, 278(1):20–35, 2019.

- [19] Peyman Mohajerin Esfahani and Daniel Kuhn. Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming*, 171(1-2):115–166, 2018.
- [20] Angelia Nedić and Dimitri Bertsekas. Convergence rate of incremental subgradient algorithms. In *Stochastic Optimization: Algorithms and Applications*, pages 223–264. Springer, 2001.
- [21] Angelia Nedić and Dimitri P Bertsekas. Incremental subgradient methods for nondifferentiable optimization. *SIAM Journal on Optimization*, 12(1):109–138, 2001.
- [22] Soroosh Shafieezadeh-Abadeh, Daniel Kuhn, and Peyman Mohajerin Esfahani. Regularization via mass transportation. *Journal of Machine Learning Research*, 20(103):1–68, 2019.
- [23] Manisha Singla, Debdas Ghosh, and KK Shukla. A survey of robust optimization based machine learning with special reference to support vector machines. *International Journal of Machine Learning and Cybernetics*, 11(7):1359–1385, 2020.
- [24] Aman Sinha, Hongseok Namkoong, and John Duchi. Certifying some distributional robustness with principled adversarial training. In *International Conference on Learning Representations*, 2018.
- [25] Johan AK Suykens and Joos Vandewalle. Least squares support vector machine classifiers. *Neural Processing Letters*, 9(3):293–300, 1999.
- [26] Po-Wei Wang, Matt Wytock, and Zico Kolter. Epigraph projections for fast general convex programming. In *International Conference on Machine Learning*, pages 2868–2877, 2016.
- [27] Wolfram Wiesemann, Daniel Kuhn, and Melvyn Sim. Distributionally robust convex optimization. *Operations Research*, 62(6):1358–1376, 2014.
- [28] Zirui Zhou and Anthony Man-Cho So. A unified approach to error bounds for structured convex optimization problems. *Mathematical Programming*, 165(2):689–728, 2017.

Appendix

This supplementary document is the appendix section of the paper titled “Fast Epigraphical Projection-based Incremental Algorithms for Wasserstein Distributionally Robust Support Vector Machine”. It is organized as follows. In Section A, we give the details of the algorithms for solving the subproblems (i.e., ℓ_q norm epigraphical projection and single-sample proximal point update (6)). In Section B, we prove the results in the section “Convergence Rate Analysis of Incremental Algorithms”. In Section C, we describe how to extend the algorithm in [15] (i.e., GS-ADMM) to tackle our ℓ_∞ -DRSVM problem. Subsequently, we provide additional experimental results to further demonstrate the effectiveness of our proposed method.

A: Algorithmic ingredients in ISG and IPPA

To begin, we provide a summary in Table 5, which aims to help the reader find the related algorithmic details as quickly as possible.

Table 5: Summary of all ingredients in ISG and IPPA

Cases	ISG epigraphical projection (5)	IPPA single sample update (6)
ℓ_2	closed-form; see Prop. 6.1	exhaust all seven cases; see Table 6 analytic form for subcases; see Prop. 6.2,6.4
ℓ_1	quick-select algorithm in linear time	exhaust all five cases; see Alg. 2
ℓ_∞	Moreau’s decomposition based on ℓ_1 case	modified secant alg. 1

Subproblems for $q = 2$. It is well known that the ℓ_2 norm epigraphical projection has a closed-form formula, which is given as follows:

Proposition 6.1 (Adopted from [1, Theorem 3.3.6]) Let $L_2^d = \{(x, s) \in \mathbb{R}^d \times \mathbb{R} : \|x\|_2 \leq s\}$. For any $(x, t) \in \mathbb{R}^d \times \mathbb{R}$, we have

$$\text{proj}_{L_2^d}(x, s) = \begin{cases} \left(\frac{\|x\|_2 + s}{2\|x\|_2}, \frac{\|x\|_2 + s}{2} \right) & \|x\|_2 \geq |s|, \\ (0, 0) & s < \|x\|_2 < -s, \\ (x, s) & \|x\|_2 \leq s. \end{cases} \quad (15)$$

Recall that the ℓ_2 single-sample proximal point subproblem (7) takes the form

$$\min_{w, \lambda} \max \underbrace{\{1 - w^T z_i, 1 + w^T z_i - \lambda\kappa, 0\}}_{h_i(w, \lambda)} + \frac{1}{2\alpha} (\|w - \bar{w}\|_2^2 + (\lambda - \bar{\lambda})^2), \quad \text{s.t. } \|w\|_2 \leq \lambda.$$

We start with $\Gamma = 1, j = 1$. Problem (8) can be written as

$$\min_{w, \lambda} \frac{1}{2\alpha} (\|w - \bar{w} - \alpha z_i\|_2^2 + (\lambda - \bar{\lambda})^2), \quad \text{s.t. } \|w\|_2 \leq \lambda,$$

whose optimal solution is given by $(\hat{w}, \hat{\lambda}) = \text{proj}_{L_2^d}(\bar{w} + \alpha z_i, \bar{\lambda})$. We further check whether $(\hat{w}, \hat{\lambda})$ satisfies the optimality condition of problem (7); i.e., $\{(w, \lambda) : h_{i,1}(w, \lambda) > h_{i,2}(w, \lambda), h_{i,1}(w, \lambda) > h_{i,3}(w, \lambda)\} = \{(w, \lambda) : w^T z < \min(\frac{\lambda\kappa}{2}, 1)\}$. The other two cases (i.e., $\Gamma = 1, j = 2, 3$) follow the same procedure. Then, we proceed to consider $\Gamma = 2$ (e.g., $(j, j') = (1, 2)$). Problem (9) is equivalent to

$$\min_{w, \lambda} \frac{1}{2\alpha} (\|w - \bar{w} - \alpha z_i\|_2^2 + (\lambda - \bar{\lambda})^2), \quad \text{s.t. } w^T z_i = \frac{\kappa}{2}\lambda, \|w\|_2 \leq \lambda.$$

We now derive an analytic solution for its prototypical form (16) in Proposition 6.2, which also covers the other two cases (i.e., $\Gamma = 2, (j, j') = (1, 3), (2, 3)$).

Proposition 6.2 Given $\bar{w} \in \mathbb{R}^d$ and $\bar{\lambda}, a, b \in \mathbb{R}$, consider the following optimization problem:

$$\begin{aligned} \min_{w, \lambda} \quad & \frac{1}{2} \|w - \bar{w}\|_2^2 + \frac{1}{2} (\lambda - \bar{\lambda})^2 \\ \text{s.t.} \quad & w^T z_i = a\lambda + b \leftarrow \mu_1, \\ & \|w\|_2 \leq \lambda \leftarrow \mu_2, \end{aligned} \tag{16}$$

where μ_1 and μ_2 are the associated dual multipliers. Then, the optimal solution (w^*, λ^*) to (16) is

$$\text{PPA_sub}(\bar{w}, \bar{\lambda}, a, b) \triangleq \left(\frac{\bar{w} - \mu_1^* z_i}{1 + 2\mu_2^*}, \frac{\bar{\lambda} + a\mu_1^*}{1 - 2\mu_2^*} \right),$$

where μ_1^* and μ_2^* are the optimal dual multipliers. In particular, we have

$$\begin{aligned} \mu_1^* &= \frac{(1 - 2\mu_2^*)\bar{w}^T z_i - a(1 + 2\mu_2^*)\bar{\lambda} - b(1 - 2\mu_2^*)(1 + 2\mu_2^*)}{(1 + 2\mu_2^*)a^2 + (1 - 2\mu_2^*)\|z_i\|_2^2}, \\ \mu_2^* &= \begin{cases} 0, & \text{if } \|\bar{w} - \mu_1^* z_i\|_2 \leq \bar{\lambda} + a\mu_1^*, \\ \hat{\mu}_2, & \text{otherwise} \end{cases} \end{aligned}$$

and $\hat{\mu}_2$ is the root of the following **quartic equation** satisfying $\hat{\mu}_2 > 0$ and $\lambda^* \geq 0$:

$$p_1\mu_2^4 + p_2\mu_2^3 + p_3\mu_2^2 + p_4\mu_2 + p_5 = 0.$$

Here, $A = \|\bar{w}\|_2^2$, $B = \|z_i\|_2^2$, $C = \bar{w}^T z_i$, and

$$\begin{aligned} p_1 &= a^2b^2 - Bb^2, \\ p_2 &= 4a^2b^2, \\ p_3 &= -4Bab\bar{\lambda} + 2BCa\bar{\lambda} - 2Ca^3\bar{\lambda} - 4Ca^2b + 2ABa^2 + 6a^2b^2 \\ &\quad + B^2\bar{\lambda}^2 + 2Bb^2 + BC^2 - Ba^2\bar{\lambda}^2 - C^2a^2 - Aa^4 - AB^2, \\ p_4 &= -8Bab\bar{\lambda} + 4BCa\bar{\lambda} - 2Ba^2\bar{\lambda}^2 - 4Ca^3\bar{\lambda} - 8Ca^2b - 2Aa^4 \\ &\quad - 2BC^2 + 2AB^2 + 4a^2b^2 + 2B^2\bar{\lambda}^2 + 2C^2a^2, \\ p_5 &= -4Bab\bar{\lambda} + 2BCa\bar{\lambda} - 2Ca^3\bar{\lambda} - 4Ca^2b - 2ABa^2 + a^2b^2 \\ &\quad + B^2\bar{\lambda}^2 + 3C^2a^2 + BC^2 - Ba^2\bar{\lambda}^2 - Bb^2 - Aa^4 - AB^2. \end{aligned} \tag{17}$$

Proof The Karush-Kuhn-Tucker (KKT) conditions of (16) are given by

$$\begin{cases} w^* - \bar{w} + \mu_1^* z_i + 2\mu_2^* w^* = 0, \\ \lambda^* - \bar{\lambda} - a\mu_1^* - 2\mu_2^* \lambda^* = 0, \\ w^{*T} z_i = a\lambda^* + b, \\ \|w^*\|_2 \leq \lambda^*, \\ \mu_2^* (\|w^*\|_2^2 - \lambda^{*2}) = 0, \\ \mu_2^* \geq 0. \end{cases} \tag{18}$$

Based on (18), we have

$$\begin{aligned} (1 + 2\mu_2^*)w^{*T} z_i - \bar{w}^T z_i + \mu_1^* \|z_i\|_2^2 = 0 &\Rightarrow w^{*T} z_i = \frac{\bar{w}^T z_i - \mu_1^* \|z_i\|_2^2}{1 + 2\mu_2^*}, \\ (1 - 2\mu_2^*)\lambda^* - \bar{\lambda} - a\mu_1^* = 0 &\Rightarrow (1 - 2\mu_2^*)(a\lambda^* + b) = a\bar{\lambda} + a^2\mu_1^* + b(1 - 2\mu_2^*). \end{aligned}$$

Plugging in $w^{*T} z_i = a\lambda + b$ yields

$$\begin{aligned} (1 - 2\mu_2^*) \frac{\bar{w}^T z_i - \mu_1^* \|z_i\|_2^2}{1 + 2\mu_2^*} &= a\bar{\lambda} + a^2\mu_1^* + b(1 - 2\mu_2^*) \\ \Rightarrow (1 - 2\mu_2^*)(\bar{w}^T z_i - \mu_1^* \|z_i\|_2^2) &= a\bar{\lambda}(1 + 2\mu_2^*) + a^2\mu_1^*(1 + 2\mu_2^*) + b(1 - 2\mu_2^*)(1 + 2\mu_2^*). \end{aligned}$$

Then,

$$\mu_1^* = \frac{(1 - 2\mu_2^*)\bar{w}^T z_i - a(1 + 2\mu_2^*)\bar{\lambda} - b(1 - 2\mu_2^*)(1 + 2\mu_2^*)}{(1 + 2\mu_2^*)a^2 + (1 - 2\mu_2^*)\|z_i\|_2^2}. \tag{19}$$

To handle the complementary slackness condition $\mu_2^* (\|w^*\|_2^2 - \lambda^{*2}) = 0$, we consider the following two cases:

- Case 1: If $\mu_2^* = 0$, then we have $\mu_1^* = \frac{\bar{w}^T z_i - a\bar{\lambda} - b}{a^2 + \|z_i\|_2^2}$ and hence

$$w^* = \bar{w} - \mu_1^* z_i, \quad \lambda^* = \bar{\lambda} + a\mu_1^*.$$

If $\|w^*\|_2 \leq \lambda^*$ does not hold, we go to Case 2.

- Case 2: If $\mu_2^* > 0$, then by incorporating $\|w^*\|_2 = \lambda^*$ into (19), we obtain the quartic equation

$$p_1\mu_2^4 + p_2\mu_2^3 + p_3\mu_2^2 + p_4\mu_2 + p_5 = 0,$$

whose coefficients p_1, \dots, p_5 are given in (17). Finally, μ_2^* is in effect the root of this quartic equation, which satisfies $\mu_2^* > 0$ and $\lambda^* \geq 0$. The optimal solution (w^*, λ^*) is then given by

$$w^* = \frac{\bar{w} - \mu_1^* z_i}{1 + 2\mu_2^*}, \quad \lambda^* = \frac{\bar{\lambda} + a\mu_1^*}{1 - 2\mu_2^*}.$$

□

Remark 6.3 *The KKT conditions are necessary and sufficient for optimality for problem (16). If there does not exist a KKT point $(w^*, \lambda^*, \mu_1^*, \mu_2^*)$ (i.e., no nonnegative roots for the quartic function), then the case $\Gamma = 2$ is not optimal for (7) and we proceed to other cases. For practical implementation, we apply the built-in function `roots` (`[p1, p2, p3, p4, p5]`) in MATLAB to get the roots of the quartic equation.*

Similarly, we check the corresponding optimality condition $\{(w, \lambda) : h_{i,1}(w, \lambda) = h_{i,2}(w, \lambda) > h_{i,3}(w, \lambda)\} = \{(w, \lambda) : \lambda\kappa - 1 < w^T z_i < 1, \lambda\kappa < 2\}$. The other two cases follow the same procedure. Lastly, we proceed to the case $\Gamma = 3$ and problem (9) in effect admits a closed-form update.

Proposition 6.4 *Given $\bar{w} \in \mathbb{R}^d$ and $b, \lambda \in \mathbb{R}$, consider the following optimization problem:*

$$\begin{aligned} \min_w \quad & \frac{1}{2} \|w - \bar{w}\|_2^2 \\ \text{s.t.} \quad & w^T z_i = b \leftarrow \alpha, \\ & \|w\|_2 \leq \lambda \leftarrow \beta, \end{aligned} \tag{20}$$

where α and β are the associated dual multipliers. Then, the optimal solution w^* to (20) is

$$\mathcal{O}_{BH}(b, \lambda, \bar{w}) \triangleq \begin{cases} A, & \text{if } \|A\|_2 \leq \lambda, \\ \frac{1}{2\beta^* + 1} \left\{ A + \frac{2b\beta^*}{\|z_i\|_2^2} z_i \right\}, & \text{otherwise,} \end{cases}$$

where $A = \bar{w} - \frac{\bar{w}^T z_i - b}{\|z_i\|_2^2} z_i$, $B = \frac{2b}{\|z_i\|_2^2} z_i$, and β^* is the positive root of the following **quadratic** equation:

$$(4\lambda^2 - \|B\|_2^2)\beta^2 + (4\lambda^2 - 2A^T B)\beta + (\lambda^2 - \|A\|_2^2) = 0.$$

Proof The KKT conditions of (20) are given by

$$\begin{cases} w^* - \bar{w} + \alpha^* z_i + 2\beta^* w^* = 0, \\ w^{*T} z_i - b = 0, \\ \|w^*\|_2 \leq \lambda, \\ \beta^* (\|w^*\|_2^2 - \lambda^2) = 0, \\ \beta^* \geq 0. \end{cases} \tag{21}$$

Here, α^* and β^* are the optimal dual multipliers. On top of (21), we have

$$(1 + 2\beta^*)w^* - \bar{w} + \alpha^* z_i = 0 \quad \Rightarrow \quad (1 + 2\beta^*)w^{*T} z_i - \bar{w}^T z_i + \alpha^* \|z_i\|_2^2 = 0,$$

$$(1 + 2\beta^*)b - \bar{w}^T z_i + \alpha^* \|z_i\|_2^2 = 0 \quad \Rightarrow \quad \alpha^* = \frac{\bar{w}^T z_i - (1 + 2\beta^*)b}{\|z_i\|_2^2}.$$

Plugging in $w^* = \frac{\bar{w} - \alpha^* z_i}{1 + 2\beta^*}$ gives

$$w^* = \frac{1}{2\beta^* + 1} \left\{ \bar{w} - \frac{\bar{w}^T z_i - (1 + 2\beta^*)b}{\|z_i\|_2^2} z_i \right\}.$$

Similarly, to handle the complementary slackness condition, we consider the case $\beta^* = 0$. For this case, we check whether the condition $\|A\|_2 \leq \lambda$ holds. Otherwise, $\beta^* > 0$ and $\|w^*\|_2^2 = \lambda^2$. This is equivalent to finding the positive root of the quadratic function

$$(4\lambda^2 - \|B\|_2^2)\beta^{*2} + (4\lambda^2 - 2A^T B)\beta^* + (\lambda^2 - \|A\|_2^2) = 0.$$

The geometric interpretation of this case is illustrated in Fig. 2. Problem (20) seeks to find the projection onto the intersection of the Euclidean ball $\|w\|_2 \leq \lambda$ and the hyperplane $w^T z_i = b$. Observe that the projection of \bar{w} onto the hyperplane $w^T z_i = b$ is given by $A = \bar{w} - \frac{\bar{w}^T z_i - b}{\|z_i\|_2^2} z_i$. If $A \in \{w : \|w\|_2 \leq \lambda\}$ (i.e., red case), then A is the optimal solution to (20). Otherwise, we aim to find the point on the sphere (i.e., $\|w^*\|_2 = \lambda$) that is closer to A (i.e., blue case).




Figure 2: Projection onto the intersection of Euclidean ball and hyperplane.

□

Let us now summarize the optimal solutions of all sub-cases and the corresponding optimality conditions in Table 6.

Table 6: Summary of all sub-cases for ℓ_2 proximal point update (7)

Sub-cases	Optimal solution	Optimality condition
$\Gamma = 1, j = 1$	$\text{proj}_{L_2^d}(\bar{w} + \alpha z_i, \bar{\lambda})$	$w^{*T} z < \min(\frac{\lambda^* \kappa}{2}, 1)$
$\Gamma = 1, j = 2$	$\text{proj}_{L_2^d}(\bar{w} - \alpha z_i, \bar{\lambda} + \alpha \kappa)$	$w^{*T} z > \max(\frac{\lambda^* \kappa}{2}, \lambda^* \kappa - 1)$
$\Gamma = 1, j = 3$	$\text{proj}_{L_2^d}(\bar{w}, \bar{\lambda})$	$1 < w^{*T} z < \lambda^* \kappa - 1$ & $\lambda^* \kappa > 2$
$\Gamma = 2, (j, j') = (1, 2)$	$\text{PPA_sub}(\bar{w} + \alpha z_i, \bar{\lambda}, \frac{\kappa}{2}, 0)$	$\lambda^* \kappa - 1 < w^{*T} z < 1$ & $\lambda^* \kappa < 2$
$\Gamma = 2, (j, j') = (1, 3)$	$\text{PPA_sub}(\bar{w}, \bar{\lambda}, 0, 1)$	$w^{*T} z < \min(\frac{\lambda^* \kappa}{2}, \lambda^* \kappa - 1)$ & $\lambda^* \kappa > 2$
$\Gamma = 2, (j, j') = (2, 3)$	$\text{PPA_sub}(\bar{w}, \bar{\lambda}, \kappa, -1)$	$w^{*T} z > \max(\frac{\lambda^* \kappa}{2}, 1)$ & $\lambda^* \kappa > 2$
$\Gamma = 3$	$(\mathcal{O}_{BH}(1, \frac{2}{\kappa}, \bar{w}), \frac{2}{\kappa})$	$\lambda^* = \frac{2}{\kappa}$

Subproblems for $q = 1$. Consider the ℓ_1 norm epigraphical projection

$$\text{proj}_{L_1^d}(x, s) = \arg \min_{y, t} \left\{ \frac{1}{2} \|y - x\|_2^2 + \frac{1}{2} (t - s)^2, \text{ s.t. } \|y\|_1 \leq t \right\}, \quad (22)$$

where $L_1^d = \{(x, s) \in \mathbb{R}^d \times \mathbb{R} : \|x\|_1 \leq s\}$. Problem (22) is equivalent to finding the root of an one-dimensional piecewise linear equation. By inspecting the KKT conditions (with $\lambda \geq 0$

being the Lagrangian multiplier), we have $y = \text{sign}(x) \odot \max(|x| - \lambda, 0)$ (i.e., proximal operator for ℓ_1 norm) and $t = \lambda + s$. Combining this with the complementary slackness condition and $\lambda > 0$, the KKT conditions of (22) reduce to the piecewise linear root-finding problem $F(\lambda) = \sum_{i=1}^d \max(|x_i| - \lambda) - \lambda - s = 0$, which can be solved by the quick-select algorithm in linear time; see [26, Algorithm 2] for details. Otherwise, we have $\text{proj}_{L_1^d}(x, s) = (x, s)$.

Recall that the ℓ_1 single-sample proximal point subproblem (11) takes the form

$$\begin{aligned} \min_{w, \lambda, \mu} \quad & \mu + \frac{1}{2\alpha} (\|w - \bar{w}\|_2^2 + (\lambda - \bar{\lambda})^2) \\ \text{s.t.} \quad & h_{i,j}(w, \lambda) \leq \mu \quad (\leftarrow \sigma_j \geq 0), \quad j = 1, 2, 3; \quad \|w\|_1 \leq \lambda, \end{aligned}$$

where $\sigma_1, \sigma_2, \sigma_3 \geq 0$ are the corresponding dual multipliers.

- Case 1: $h_{i,1}, h_{i,3}$ are inactive. Then, problem (11) can be written as

$$\min_{w, \lambda} (1 + w^T z_i - \lambda \kappa) + \frac{1}{2\alpha} \|w - \bar{w}\|_2^2 + \frac{1}{2\alpha} (\lambda - \bar{\lambda})^2, \quad \text{s.t. } \|w\|_1 \leq \lambda.$$

Hence, we have $(w^*, \lambda^*) = \text{proj}_{L_1^d}(\bar{w} - \alpha z_i, \bar{\lambda} + \alpha \kappa)$.

- Case 2: $h_{i,1}$ is active and $h_{i,3}$ is inactive. Then, problem (12) can be reduced to

$$\min_{w, \lambda} \frac{1}{2\alpha} (\|w - \bar{w} - \alpha z_i\|_2^2 + (\lambda - \bar{\lambda})^2), \quad \text{s.t. } w^T z_i \leq \frac{\lambda \kappa}{2} \quad (\leftarrow \sigma_1 \geq 0), \quad \|w\|_1 \leq \lambda. \quad (23)$$

- Cases 3 and 4: ($h_{i,1}$ is inactive, $h_{i,3}$ is active) and ($h_{i,2}$ is inactive, $h_{i,3}$ is active). These two cases are similar to Case 2 and give rise to a problem of the form (13).

Now, let us demonstrate how to solve (13) efficiently. Recall that

$$\min_{w, \lambda} \frac{1}{2\alpha} (\|w - \bar{w}\|_2^2 + (\lambda - \bar{\lambda})^2), \quad \text{s.t. } w^T z \leq a\lambda + b \quad (\leftarrow \sigma \geq 0), \quad \|w\|_1 \leq \lambda.$$

Proposition 6.5 *Suppose that σ_1^* is the dual optimal solution to (23). Then, we have $\sigma_1^* \in [0, 1]$.*

Proof Based on the KKT conditions of (11), we have

$$1 - \sigma_1 - \sigma_2 - \sigma_3 = 0.$$

If the optimal solution to (23) is also optimal for (11), then we can match the two KKT systems. As $h_{i,3}$ is inactive for this case, we have $\sigma_3^* = 0$. This gives

$$\sigma_1^* + \sigma_2^* = 1, \sigma_1^*, \sigma_2^* \geq 0 \quad \Rightarrow \quad \sigma_1^* \in [0, 1].$$

□

Proposition 6.5 also holds for Cases 3 and 4. The analytic bound in Proposition 6.5 shows that σ^* can be efficiently found by an appropriate search strategy. Next, recall from (14) that

$$\begin{aligned} (\hat{w}(\sigma), \hat{\lambda}(\sigma)) &= \arg \min_{\|w\|_1 \leq \lambda} \frac{1}{2\alpha} (\|w - \bar{w}\|_2^2 + (\lambda - \bar{\lambda})^2) + \sigma(w^T z - a\lambda - b) \\ &= \text{proj}_{L_1^d}(\bar{w} - \sigma \alpha z, \bar{\lambda} + \sigma \alpha a). \end{aligned}$$

The following proposition establishes the monotonicity property of $\sigma \mapsto p(\sigma) = \hat{w}(\sigma)^T z - a\kappa - b$, which plays a vital role in our development of a fast algorithm for solving (13) later.

Proposition 6.6 *If σ satisfies (i) $\sigma = 0$ and $p(\sigma) \leq 0$, or (ii) $p(\sigma) = 0$, then $(\hat{w}(\sigma), \hat{\lambda}(\sigma))$ is the optimal solution to (13). Moreover, $p(\cdot)$ is continuous and monotonically non-increasing on \mathbb{R}_+ .*

Proof As $\text{proj}_{L_1^d}(\cdot, \cdot)$ is globally Lipschitz continuous, the function $(\hat{w}(\cdot), \hat{\lambda}(\cdot))$ is also globally Lipschitz continuous and further $p(\cdot)$ is continuous. Next, we prove the monotonicity property. Upon letting $h(\sigma) = \frac{1}{2\alpha} (\|\hat{w}(\sigma) - \bar{w}\|_2^2 + (\hat{\lambda}(\sigma) - \bar{\lambda})^2)$ and assuming that $0 \leq \sigma_1 < \sigma_2 \leq 1$, we have

$$\begin{aligned} h(\sigma_1) + \sigma_1 p(\sigma_1) &\leq h(\sigma_2) + \sigma_1 p(\sigma_2) \\ &= h(\sigma_2) + \sigma_2 p(\sigma_2) + (\sigma_1 - \sigma_2) p(\sigma_2) \\ &\leq h(\sigma_1) + \sigma_2 p(\sigma_1) + (\sigma_1 - \sigma_2) p(\sigma_2), \end{aligned}$$

which implies that $p(\sigma_1) \geq p(\sigma_2)$. □

Algorithm 1: A modified secant algorithm to solve the RHS of (13)—MSA($\bar{w}, \bar{\lambda}, z_i, a, b, \xi$)

Input: Tolerance Level ξ ;

if $p(0) \leq \xi$ **then** return $(\hat{w}(0), \hat{\lambda}(0), 0)$;

else

$\sigma_l = 0, r_l = -p(0)$; // Set the lower bound σ_l for σ^*

if $p(1) \geq 0$ **then** return -1 ;

 // $\sigma^* \in [0, 1]$; see Proposition 3.1

else

$\sigma_u = 1, r_u = -p(1)$; // Set the upper bound σ_u for σ^*

end

end

/ Secant Phase* **/*

$s = 1 - \frac{r_l}{r_u}, \sigma = \sigma_u - \frac{\sigma_u - \sigma_l}{s}$; calculate $r = -p(\sigma)$;

while $|r| > \xi$ **do**

 calculate $r = -p(\sigma)$; // ℓ_1 epigraph projection via the quick-select algorithm

if $r > 0$ **then**

if $s \leq 2$ **then** $\sigma_u = \sigma, r_u = r, s = 1 - \frac{r_l}{r_u}, \sigma = \sigma_u - \frac{\sigma_u - \sigma_l}{s}$;

else

$s = \max(\frac{r_u}{r} - 1, 0.1), \Delta\sigma = \frac{\sigma_u - \sigma}{s}, \sigma_u = \sigma, r_u = r$;

$\sigma = \max(\sigma_u - \Delta\sigma, 0.6\sigma_l + 0.4\sigma_u), s = \frac{\sigma_u - \sigma_l}{\sigma_u - \sigma}$;

end

else

if $s \geq 2$ **then** $\sigma_l = \sigma, r_l = r, s = 1 - \frac{r_l}{r_u}, \sigma = \sigma_u - \frac{\sigma_u - \sigma_l}{s}$;

else

$s = \max(\frac{r_l}{r} - 1, 0.1), \Delta\sigma = \frac{\sigma - \sigma_l}{s}, \sigma_l = \sigma, r_l = r$;

$\sigma = \max(\sigma_l + \Delta\sigma, 0.6\sigma_u + 0.4\sigma_l), s = \frac{\sigma_u - \sigma_l}{\sigma_u - \sigma}$;

end

end

end

Algorithm 2: A fast algorithm based on parametric approach to solve (11)

Input: Tolerance Level ξ ; parameters $\bar{w}, \bar{\lambda}, z_i, \epsilon, \kappa$;

/ Case 1: $h_{i,1}, h_{i,3}$ are inactive* **/*

$(w^*, \lambda^*) = \text{proj}_{L_1^q}(\bar{w} - \alpha z_i, \bar{\lambda} + \alpha \kappa)$;

if $\langle w^*, z \rangle > \max(\lambda^* \kappa - 1, \frac{\lambda^* \kappa}{2})$ **then** return (w^*, λ^*) ;

// Check optimality

/ Case 2: $h_{i,1}$ is active; $h_{i,3}$ is inactive* **/*

$(w^*, \lambda^*, \sigma^*) = \text{MSA}(\bar{w} + \alpha z_i, \bar{\lambda}, z_i, \kappa/2, 0)$; // Apply the modified secant algorithm 1

if $\langle w^*, z \rangle < 1$ & $\sigma^* \in [0, 1]$ **then** return (w^*, λ^*) ;

/ Case 3: $h_{i,1}$ is inactive; $h_{i,3}$ is active* **/*

$(w^*, \lambda^*, \sigma^*) = \text{MSA}(\bar{w}, \bar{\lambda}, z_i, \kappa, -1)$;

if $\langle w^*, z \rangle > 1$ & $\sigma^* \in [0, 1]$ **then** return (w^*, λ^*) ;

/ Case 4: $h_{i,2}$ is inactive; $h_{i,3}$ is active* **/*

$(w^*, \lambda^*, \sigma^*) = \text{MSA}(\bar{w}, \bar{\lambda}, -z_i, 0, -1)$;

if $\lambda^* \kappa > 2$ & $\sigma^* \in [0, 1]$ **then** return (w^*, λ^*) ;

/ Case 5: $h_{i,1}, h_{i,2}, h_{i,3}$ are active* **/*

else

$w^* = \arg \min_w \{\|w - \bar{w}\|_2^2, \text{ s.t. } w^T z_i = 1, \|w\|_1 \leq \frac{2}{\kappa}\}, \lambda^* = \frac{\kappa}{2}$;

 // Apply the modified secant algorithm in [9] **/*

end

B: Convergence Rate Analysis of Incremental Algorithms

We now give a condition under which problem (1) with $q = 2$ satisfies the sharpness or quadratic growth (QG) property. Consider the following more general formulation of the ℓ_2 -DRSVM problem:

$$\min_{w, \lambda} \frac{c}{2} \|w\|_2^2 + \frac{1}{n} \sum_{i=1}^n f_i(w, \lambda) + \mathbb{I}_{\{(w, \lambda) \in L_2^d\}}, \quad (24)$$

where f_1, \dots, f_n are non-smooth convex functions with polyhedral epigraphs. Our condition is based on the following lemma:

Lemma 1 *Let C_1, \dots, C_N be closed convex subsets of \mathbb{R}^n , where C_{r+1}, \dots, C_N are polyhedral for some $r \in \{0, 1, \dots, N\}$. Suppose that*

$$\bigcap_{i=1}^r \text{ri}(C_i) \cap \bigcap_{i=r+1}^N C_i \neq \emptyset.$$

Then, the collection $\{C_1, \dots, C_N\}$ is boundedly linearly regular (BLR).

Proposition 6.7 *Consider problem (24). Let \mathcal{X} be the set of optimal solutions and $L_2^d = \{(w, \lambda) \in \mathbb{R}^d \times \mathbb{R} : \|w\|_2 \leq \lambda\}$ be the constraint set. Suppose that $\mathcal{X} \cap \text{ri}(L_2^d) \neq \emptyset$. Then, problem (1) satisfies the sharpness condition when $c = 0$ and the QG condition when $c > 0$.*

Proof Let $x = (w, \lambda)$, $h(x) = \frac{c}{2} \|w\|_2^2 + \frac{1}{n} \sum_{i=1}^n f_i(w, \lambda)$, and $g(x) = h(x) + \mathbb{I}_{\{x \in L_2^d\}}$. Consider the case where $c = 0$. The set \mathcal{X} can then be written as

$$\mathcal{X} = \{x : 0 \in \partial h(x) + \mathcal{N}_{L_2^d}(x)\},$$

where $\mathcal{N}_{L_2^d}(x)$ is the normal cone of L_2^d at x . As $\mathcal{X} \cap \text{ri}(L_2^d) \neq \emptyset$, we can find an $x^* \in \mathcal{X} \cap \text{ri}(L_2^d)$ that satisfies $0 \in \partial h(x^*)$. Thus, x^* is also an optimal solution to the unconstrained problem $\min_x h(x)$ and $h(x^*) = g(x^*)$.

Let \mathcal{X}_U denote the set of optimal solutions to the problem $\min_x h(x)$. It is not difficult to check that

$$\mathcal{X} = \mathcal{X}_U \cap L_2^d.$$

Since f_1, \dots, f_n have polyhedral epigraphs, by Lemma 1, the collection $\{\mathcal{X}_U, L_2^d\}$ is BLR. This implies that there exists a constant $\kappa > 0$ satisfying

$$\text{dist}(x, \mathcal{X}) = \text{dist}(x, \mathcal{X}_U \cap L_2^d) \leq \kappa \text{dist}(x, \mathcal{X}_U), \quad \forall x \in L_2^d.$$

Furthermore, the problem $\min_x h(x)$ enjoys the sharpness property; see [8, Corollary 3.6]. This gives

$$g(x) - g^* = h(x) - h^* \geq \sigma \text{dist}(x, \mathcal{X}_U) \geq \frac{\sigma}{\kappa} \text{dist}(x, \mathcal{X}), \quad \forall x \in L_2^d.$$

For $c > 0$, we note that the problem $\min_x h(x)$ can be regarded as one with a polyhedral convex regularizer; see [28, Section 4.2]. As such, it satisfies a proximal error bound (see [28, Proposition 6]) and hence the QG condition (see [14, Theorem 4.1]). It follows that

$$g(x) - g^* = h(x) - h^* \geq \sigma \text{dist}^2(x, \mathcal{X}_U) \geq \frac{\sigma}{\kappa^2} \text{dist}^2(x, \mathcal{X}), \quad \forall x \in L_2^d.$$

□

To derive the convergence rates of the incremental algorithms, we also need the following assumption.

Assumption 6.8 (Subgradient boundedness) *There exists a scalar $L > 0$ such that*

$$\|\nabla f_i(x)\| \leq L, \quad \forall \nabla f_i(x) \in \partial f_i(x), \quad i \in [n].$$

Lemma 2 (ISG; see [21, Lemma 2.1]) Suppose that Assumption 6.8 holds and $\{x^k = (w_0^k, \lambda_0^k)\}$ is a sequence generated by ISG. Then, for all y and $k \geq 0$, we have

$$\|x^{k+1} - y\|_2^2 \leq \|x^k - y\|_2^2 - 2\alpha_k n(f(x^k) - f(y)) + a_k^2 n^2 L^2.$$

Lemma 3 (IPPA) Suppose that Assumption 6.8 holds and $\{x^k = (w_0^k, \lambda_0^k)\}$ is a sequence generated by IPPA. Then, for all y and $k \geq 0$, we have

$$\|x^{k+1} - y\|_2^2 \leq \|x^k - y\|_2^2 - 2\alpha_k n(f(x^k) - f(y)) + a_k^2 n(n+1)L^2.$$

Proof Based on Proposition 1 in [3] with $x_i^k = (w_i^k, \lambda_i^k)$, we have

$$\|x_{i+1}^k - y\|_2^2 \leq \|x_i^k - y\|_2^2 - 2\alpha_k (f_{i+1}(x_{i+1}^k) - f_{i+1}(y)), \quad \forall i = 0, \dots, n-1.$$

Summing up,

$$\begin{aligned} \|x_n^k - y\|_2^2 &\leq \|x_0^k - y\|_2^2 - 2\alpha_k \sum_{i=0}^{n-1} (f_{i+1}(x_{i+1}^k) - f_{i+1}(y)) \\ &= \|x_0^k - y\|_2^2 - 2\alpha_k \sum_{i=0}^{n-1} (f_{i+1}(x_{i+1}^k) - f_{i+1}(x_0^k) + f_{i+1}(x_0^k) - f_{i+1}(y)) \\ &= \|x_0^k - y\|_2^2 - 2\alpha_k n(f(x_0^k) - f(y)) - 2\alpha_k \sum_{i=0}^{n-1} (f_{i+1}(x_{i+1}^k) - f_{i+1}(x_0^k)) \\ &\leq \|x_0^k - y\|_2^2 - 2\alpha_k n(f(x_0^k) - f(y)) + 2\alpha_k L \sum_{i=0}^{n-1} \|x_{i+1}^k - x_0^k\|_2 \\ &\leq \|x_0^k - y\|_2^2 - 2\alpha_k n(f(x_0^k) - f(y)) + 2\alpha_k^2 L^2 \sum_{i=0}^{n-1} (i+1) \\ &\leq \|x_0^k - y\|_2^2 - 2\alpha_k n(f(x_0^k) - f(y)) + a_k^2 n(n+1)L^2. \end{aligned}$$

□

Combining Lemmas 2 and 3, we have

$$\|x^{k+1} - y\|_2^2 \leq \|x^k - y\|_2^2 - 2\alpha_k n(f(x^k) - f(y)) + 2a_k^2 n^2 L^2.$$

Theorem 6.9 Let $\{x^k = (w_0^k, \lambda_0^k)\}$ be the sequence of iterates generated by ISG or IPPA.

- (1) If problem (1) satisfies the sharpness condition, then by choosing the geometrically diminishing step sizes $\alpha_k = \alpha_0 \rho^k$ with $\alpha_0 \geq \frac{\sigma \text{dist}(x^0, \mathcal{X})}{2L^2 n}$ and $\sqrt{1 - \frac{\sigma^2}{2L^2}} \leq \rho < 1$, the sequence $\{x^k\}$ converges linearly to an optimal solution to (1); i.e., $\text{dist}(x^k, \mathcal{X}) \leq \mathcal{O}(\rho^k)$ for all $k \geq 0$.
- (2) If problem (1) satisfies the quadratic growth condition, then by choosing the polynomially decaying step sizes $\alpha_k = \frac{\gamma}{nk}$ with $\gamma > \frac{1}{2\sigma}$, the sequence $\{x^k\}$ converges to an optimal solution to (1) at the rate $\mathcal{O}(\frac{1}{\sqrt{k}})$ and $\{f(x^k) - f^*\}$ converges to zero at the rate $\mathcal{O}(\frac{1}{k})$.
- (3) (See [20, Proposition 2.10]) For the general convex problem (1), by choosing the step sizes $\alpha_k = \frac{\gamma}{n\sqrt{k}}$ with $\gamma > 0$, the sequence $\{\min_{0 \leq k \leq K} f(x^k) - f^*\}$ converges to zero at the rate $\mathcal{O}(\frac{1}{\sqrt{K}})$.

Proof (1):

By the sharpness condition $f(x_k) - f^* \geq \sigma \text{dist}(x_k, \mathcal{X})$, we have

$$\text{dist}^2(x^{k+1}, \mathcal{X}) \leq \text{dist}^2(x^k, \mathcal{X}) - 2\alpha_k \sigma n \text{dist}(x^k, \mathcal{X}) + 2\alpha_k^2 L^2 n^2.$$

We now prove by induction that

$$\text{dist}(x^k, \mathcal{X}) \leq \frac{2\alpha_0 L^2 n}{\sigma} \rho^k.$$

The base case trivially holds, as $\text{dist}(x^0, \mathcal{X}) \leq \frac{2\alpha_0 L^2 n}{\sigma}$. For the inductive step, we compute

$$\begin{aligned}
\text{dist}^2(x^{k+1}, \mathcal{X}) &\leq \left(\frac{2\alpha_0 L^2 n}{\sigma} \rho^k \right)^2 - 2\alpha_k \sigma n \frac{2\alpha_0 L^2 n}{\sigma} \rho^k + 2\alpha_k^2 L^2 n^2 \\
&= \frac{4\alpha_0^2 L^4 n^2}{\sigma^2} \rho^{2k} - 2\alpha_0^2 L^2 n^2 \rho^{2k} \\
&= \frac{4\alpha_0^2 L^4 n^2}{\sigma^2} \rho^{2k} \left(1 - \frac{\sigma^2}{2L^2} \right) \\
&\leq \left(\frac{2\alpha_0 L^2 n}{\sigma} \right)^2 \rho^{2(k+1)}.
\end{aligned} \tag{25}$$

This completes the proof.

(2): By the quadratic growth condition $f(x_k) - f^* \geq \sigma \text{dist}^2(x_k, \mathcal{X})$, we have

$$\text{dist}^2(x^{k+1}, \mathcal{X}) \leq (1 - 2\alpha_k \sigma n) \text{dist}^2(x^k, \mathcal{X}) + 2\alpha_k^2 L^2 n^2.$$

Plugging in the corresponding step size scheme $\alpha_k = \frac{\gamma}{nk}$, we obtain

$$\text{dist}^2(x^{k+1}, \mathcal{X}) \leq \left(1 - \frac{2\gamma\sigma}{k} \right) \text{dist}^2(x^k, \mathcal{X}) + \frac{2\gamma^2 L^2}{k^2}.$$

We now prove by induction that

$$\text{dist}^2(x^k, \mathcal{X}) \leq \frac{B}{k},$$

where $B > 0$ is a given number. Indeed, we have

$$\begin{aligned}
\text{dist}^2(x^{k+1}, \mathcal{X}) &\leq \left(1 - \frac{2\gamma\sigma}{k} \right) \frac{B}{k} + \frac{2\gamma^2 L^2}{k^2} \\
&= \frac{B}{k+1} + \frac{B}{k(k+1)} - \frac{2\gamma\sigma B}{k^2} + \frac{2\gamma^2 L^2}{k^2} \\
&\leq \frac{B}{k+1} + \frac{B}{k^2} - \frac{2\gamma\sigma B}{k^2} + \frac{2\gamma^2 L^2}{k^2} \\
&= \frac{B}{k+1} + \frac{(1-2\gamma\sigma)B}{k^2} + \frac{2\gamma^2 L^2}{k^2} \\
&\leq \frac{B}{k+1},
\end{aligned}$$

where the last inequality holds if $\frac{(1-2\gamma\sigma)B}{k^2} + \frac{2\gamma^2 L^2}{k^2} < 0$. Hence, we have $B > \frac{2\gamma^2 L^2}{2\gamma\sigma-1}$ due to $\gamma > \frac{1}{2\sigma}$.

Combining this with the base case, we have $B > \max\left\{ \frac{2\gamma^2 L^2}{2\gamma\sigma-1}, \text{dist}^2(x^0, \mathcal{X}) \right\}$. \square

C: Additional Experimental Results

To begin with, we show how to extend the GS-ADMM framework in [15] to tackle our DRSVM problems, which serves as a baseline in Section 5. We reformulate problem (1) as

$$\begin{aligned}
\min_{w, s, \lambda} \quad & \lambda\epsilon + \frac{1}{n} \sum_{i=1}^n s_i \\
\text{s.t.} \quad & 1 - w^T z_i \leq s_i, i \in [n], \\
& 1 + w^T z_i + \lambda\kappa \leq s_i, i \in [n], \\
& s_i \geq 0, i \in [n], \\
& \|w\|_q \leq \lambda.
\end{aligned} \tag{26}$$

We follow the technique used in [15, Proposition 3.1].

Proposition 6.10 *Suppose that (w^*, λ^*, s^*) is a global minimizer of (26). Then, we have $\lambda^* \leq \lambda^U = \frac{1}{\epsilon}$.*

Proof For simplicity, we consider the case where $q = 2$. Since problem (26) satisfies the Managarian-Fromovitz Constraint Qualification (MFCQ), the KKT conditions are necessary and sufficient. Let $a_{ij} \geq 0, \forall j \in [3], i \in [N]$ and $\beta \geq 0$ be the dual variables. Then, we can write down the KKT conditions as follows:

$$\begin{cases} \sum_{i=1}^N (a_{i2} - a_{i1})z_i + 2\beta w = 0, \\ a_{i1} + a_{i2} + a_{i3} = \frac{1}{N}, \forall i \leq N, \\ \kappa \sum_{i=1}^N a_{i2} + 2\beta\lambda = \epsilon, \\ a_{i1}(1 - w^T z_i - s_i) = 0, \\ a_{i2}(1 + w^T z_i - \lambda\kappa - s_i) = 0, \\ a_{i3}s_i = 0, \\ \beta(\|w\|_2^2 - \lambda^2) = 0. \end{cases} \quad (27)$$

Based on (27), we have

$$\begin{aligned} \mathbf{0} &= \sum_{i=1}^N (a_{i2} - a_{i1})w^T z_i + 2\beta\|w\|_2^2 \\ &= \sum_{i=1}^N (a_{i2} - a_{i1})w^T z_i + 2\beta\lambda^2 = \sum_{i=1}^N (a_{i2} - a_{i1})w^T z_i + \lambda(\epsilon - \kappa \sum_{i=1}^N a_{i2}) \\ &= \lambda\epsilon + \sum_{i=1}^N a_{i2}(w^T z_i - \lambda\kappa) - \sum_{i=1}^N a_{i1}w^T z_i = \lambda\epsilon + \sum_{i=1}^N (a_{i2} + a_{i1})(s_i - 1). \end{aligned}$$

Thus, we have

$$\lambda = \frac{1}{\epsilon} \sum_{i=1}^N (a_{i2} + a_{i1})(1 - s_i) = \frac{1}{\epsilon} \sum_{i=1}^N \left(\frac{1}{N} - a_{i3}\right)(1 - s_i) \leq \frac{1}{\epsilon N} \sum_{i=1}^N (1 - s_i) \leq \frac{1}{\epsilon}.$$

□

Remark 6.11 Although we focus on the case where $q = 2$ in this proof, we can easily extend the techniques to study the case where $q \in \{1, \infty\}$. We just need to modify the blue part in (27). Specifically, observe that $\|w\|_1 \leq \lambda$ is equivalent to $Bw \leq \lambda e_{2^d}$, where B is the $2^d \times d$ matrix whose rows are all the possible arrangements of $+1$'s and -1 's. On the other hand, $\|w\|_\infty \leq \lambda$ is equivalent to $e_i^T w \leq \lambda, -e_i^T w \leq \lambda, \forall i \in [n]$.

Subsequently, we develop a standard ADMM algorithm to address the w -subproblem

$$\min_w \frac{1}{n} \sum_{i=1}^n \max \{1 - w^T z_i, 1 + w^T z_i - \lambda\kappa, 0\}, \text{ s.t. } \|w\|_q \leq \lambda.$$

We apply the operator splitting technique to reformulate it as

$$\begin{aligned} \min_{w,y} \quad & \frac{1}{n} \sum_{i=1}^n \max \{1 - y_i, 1 + y_i - \lambda\kappa, 0\} \\ \text{s.t.} \quad & Zw - y = 0, \\ & \|w\|_q \leq \lambda. \end{aligned}$$

Algorithm 3: ADMM for solving w -subproblem

Input: Choose value $(w^0, y^0, g^0) \in \mathbb{R}^d \times \mathbb{R}^n \times \mathbb{R}^n$;Initialized the penalty parameter ρ_0 and shrinking parameter $\gamma \geq 1$;**Output:** $\{(w^k, y^k, g^k)\}_{k=1}^K$ and function value sequences;**for each iteration do**

/* Accelerated projected gradient algorithm, see [15, Algorithm 5] */

$$w^{k+1} = \arg \min_{\|w\|_q \leq \lambda} \left\{ \frac{\rho_k}{2} \|Zw - y^k + \frac{g^k}{\rho_k}\|_2^2 \right\};$$

/* Closed-form update */

$$y^{k+1} = \arg \min_{y \in \mathbb{R}^n} \left\{ \frac{1}{n} \sum_{i=1}^n \max \{1 - y_i, 1 + y_i - \lambda \kappa, 0\} + \frac{\rho_k}{2} \|y - Zw^{k+1} - \frac{g^k}{\rho_k}\|_2^2 \right\};$$

/* Dual update */

$$g^{k+1} = g^k + \rho_k (Zw^{k+1} - y^{k+1});$$

$$\rho_{k+1} = \gamma \rho_k;$$

end

Single-sample proximal point update for $c > 0$ Recall that

$$\begin{aligned} \min_{w, \lambda} \quad & \frac{c}{2} \|w\|_2^2 + \max \{1 - w^T z_i, 1 + w^T z_i - \lambda \kappa, 0\} + \frac{1}{2\alpha} (\|w - \bar{w}\|_2^2 + (\lambda - \bar{\lambda})^2) \\ \text{s.t.} \quad & \|w\|_q \leq \lambda. \end{aligned}$$

Note that $\mu = \frac{\lambda}{\sqrt{1+\alpha c}}$ and $\kappa' = \kappa \sqrt{1+\alpha c}$. The above problem can be written as

$$\begin{aligned} \min_{w, \mu} \quad & \max \{1 - w^T z_i, 1 + w^T z_i - \mu \kappa', 0\} + \frac{1 + \alpha c}{2\alpha} \left(\left\| w - \frac{\bar{w}}{1 + \alpha c} \right\|_2^2 + (\mu - \bar{\mu})^2 \right) \\ \text{s.t.} \quad & \|w\|_q \leq \sqrt{1 + \alpha c} \mu. \end{aligned} \quad (28)$$

Indeed, problem (28) shares the same structure as problem (2).

Table 7: Wall-clock Time Comparison on UCI Real Dataset: ℓ_1 -DRSVM, $c = 1, \kappa = 1, \epsilon = 0.1$

	Objective Value				Wall-clock time (sec)			
	M-ISG	IPPA	Hybrid	YALMIP	M-ISG	IPPA	Hybrid	YALMIP
a1a	0.7871456	0.7871450	0.7871445	0.7871462	1.233	2.169	1.545	16.495
a2a	0.8002882	0.8002879	0.8002879	0.8003101	0.602	3.546	0.720	23.557
a3a	0.7826654	0.7826653	0.7826653	0.7826653	3.687	4.626	3.696	32.575
a4a	0.7929724	0.7929724	0.7929724	0.7929959	0.579	4.109	0.713	62.456
a5a	0.7852845	0.7852848	0.7852844	0.7853440	0.639	2.929	1.156	109.620
a6a	0.7764425	0.7764425	0.7764425	0.7767701	1.272	5.756	1.576	185.080
a7a	0.7822116	0.7822114	0.7822115	0.7827171	2.121	4.336	2.976	270.480
a8a	0.7805498	0.7805498	0.7805498	0.7836023	2.502	10.184	2.798	372.050
a9a	0.7767114	0.7767099	0.7767113	0.7791881	3.018	9.295	6.203	642.160

Table 8: Wall-clock Time Comparison on UCI Real Dataset: ℓ_∞ -DRSVM, $c = 1, \kappa = 1, \epsilon = 0.1$

	Objective Value				Wall-clock time (sec)			
	M-ISG	IPPA	Hybrid	YALMIP	M-ISG	IPPA	Hybrid	YALMIP
a1a	0.7853266	0.7853265	0.7853265	0.7853269	0.510	0.707	0.687	12.928
a2a	0.7987669	0.7987666	0.7987667	0.7987663	0.861	1.233	1.182	20.850
a3a	0.7810149	0.7810140	0.7810145	0.7810568	0.350	0.666	0.563	28.494
a4a	0.7913534	0.7913534	0.7913534	0.7913963	0.540	1.100	0.679	63.037
a5a	0.7836246	0.7836189	0.7836246	0.7836478	0.737	1.421	0.897	96.314
a6a	0.7748542	0.7748533	0.7748537	0.7759606	1.203	2.250	1.763	201.510
a7a	0.7806894	0.7806894	0.7806894	0.7811091	1.753	3.244	2.124	370.510
a8a	0.7789801	0.7789800	0.7789800	0.7850996	2.390	4.568	2.981	365.140
a9a	0.7750633	0.7750633	0.7750633	0.7776798	3.368	6.706	4.371	753.330

Remark 6.12 For ℓ_2 -DRSVM problems, it is worth mentioning that the $c > 0$ case is identical to the $c = 0$ case from a modeling perspective, which depends on the different robustness levels ϵ and κ .