

Data-Dependent Complexity of First-Order Methods for Binary Classification*

Matthew Hough[†] Stephen A. Vavasis[‡]

December 3, 2025

Abstract

Large-scale problems in data science are often modeled with optimization, and the optimization model is usually solved with first-order methods that may converge at a sublinear rate. Therefore, it is of interest to terminate the optimization algorithm as soon as the underlying data science task is accomplished. We consider FISTA for solving two binary classification problems: the ellipsoid separation problem (ESP), and the soft-margin support-vector machine (SVM). For the ESP, we cast the dual second-order cone program into a form amenable to FISTA and show that the FISTA residual converges to the infimal displacement vector of the primal-dual hybrid gradient (PDHG) algorithm, that directly encodes a separating hyperplane. We further derive a data-dependent iteration upper bound scaling as $\mathcal{O}(1/\delta_{\mathcal{A}}^2)$, where $\delta_{\mathcal{A}}$ is the minimal perturbation that destroys separability. For the SVM, we propose a strongly-concave perturbed dual that admits efficient FISTA updates under a linear time projection scheme, and with our parameter choices, the objective has small condition number, enabling rapid convergence. We prove that, under a reasonable data model, early-stopped iterates identify well-classified points and yield a hyperplane that exactly separates them, where the accuracy required of the dual iterate is governed by geometric properties of the data. In particular, the proposed early-stopping criteria diminish the need for hard-to-select tolerance-based stopping conditions. Our numerical experiments on ESP instances derived from MNIST data and on soft-margin SVM benchmarks indicate competitive run-times and substantial speedups from stopping early.

1 Introduction

Binary classification is a central task in machine learning with numerous applications. Given data separated into two classes, the task is to construct a decision boundary called

*Supported in part by a Discovery Grant from the Natural Sciences and Engineering Research Council (NSERC) of Canada.

[†]Department of Combinatorics & Optimization, University of Waterloo, 200 University Ave. W., Waterloo, ON, N2L 3G1, Canada, mhough@uwaterloo.ca.

[‡]Department of Combinatorics & Optimization, University of Waterloo, 200 University Ave. W., Waterloo, ON, N2L 3G1, Canada, vavasis@uwaterloo.ca.

a *classifier* that separates the classes completely, or with minimal misclassifications. Since many binary classification models, including the support vector machine (SVM), can be modeled as convex optimization problems, first-order methods are a popular choice for training these models, especially in large-scale settings. Typically, algorithms for binary classification are guaranteed to yield a classifier under the assumption that numerical convergence of the underlying optimization problem is achieved. In practice, however, high-accuracy optimization is often unnecessary to correctly classify most points: a hyperplane that correctly classifies “well-separated” or “easy to classify” points may emerge after only a small number of iterations. One would expect that the ability to obtain such an approximate separator in few iterations depends on the *classifiability* of the data, i.e. the extent to which the data contain points that are robustly separable from the opposite class. In such cases, the computational effort required to find a useful classifier is not solely determined by worst-case complexity bounds of the optimization problem, but can instead be characterized in terms of *data-dependent complexity*, where the number of iterations required to obtain a classifier scales with geometric properties of the dataset rather than purely algorithmic constants.

The work herein investigates data-dependent complexity in two related binary classification tasks. The first is the ellipsoid separation problem (ESP), in which each class is represented by a finite family of ellipsoids, and the task is to determine a separating hyperplane that completely separates one class of ellipsoids from the other. This problem can be seen as an extension of classic binary classification, where the locations of the data are known only up to an ellipsoid, and a separating hyperplane must be found which separates all possible locations of the data. Modeling data uncertainty in this way is a form of robust classification, and was considered by Shivaswamy et al. in [26]. In the ellipsoid separation setting, we show that the FISTA algorithm [4] applied to a second-order cone program (SOCP), which models the dual of the ESP, yields iterates whose residual vectors encode a separating hyperplane. Furthermore, we prove that the ESP can be solved in $\mathcal{O}(1/\delta_A^2)$ iterations, where δ_A represents the smallest perturbation to the ellipsoid centers, such that the ellipsoids are no longer separable. The second problem we consider is the soft-margin support vector machine (SVM) for non-separable data introduced in [10]. Here, a perfect separator does not exist, but many points may still be classifiable. We propose a perturbed dual formulation of the soft-margin SVM that is strongly concave and converges quickly in FISTA due to the objective having a condition number $\kappa = L/\mu \leq 513$. We prove that, under a suitable data model, an approximate solution to the perturbed dual can be used to identify all well-classified points, and furthermore we show how one can construct a separating hyperplane from this approximate dual solution that can separate the well-classified points exactly. The degree of approximation required of the dual solution is also itself dependent on the geometry of the data.

For both the ESP and soft-margin SVM, we perform numerical experiments that show considerable speedup can be achieved using our methods. In particular, our method for approximately solving soft-margin support vector machines performs similarly to LIBSVM [9] and LIBLINEAR [13], and on average outperforms Pegasos [25] in terms of runtime in seconds and accuracy on select problems from the UCI Machine Learning Dataset [18].

Binary classification has been extensively studied in the contexts of convex optimiza-

tion and machine learning, most notably through the development of the soft-margin support vector machine by Cortes and Vapnik [10]. Various specialized algorithms to solve the SVM have been proposed over time. Osuna, Freund, and Girosi [22] developed a decomposition method for the dual SVM, which partitions the quadratic program (QP) into smaller subproblems and makes training feasible even when the number of support vectors is large. Platt then refined this work in his development of Sequential Minimal Optimization (SMO) [23], which instead decomposes the QP into the smallest possible subproblems, allowing them to be solved analytically. Around the same time, Joachims independently developed SVMlight [17], a more general decomposition framework than Platt's. Further refinements to the decomposition ideas of Osuna et al., Platt, and Joachims were proposed in [14], which forms the basis of the widely used LIBSVM library for solving SVMs [9]. In parallel, primal-based algorithms were also developed, most notably Pegasos [25], which employs stochastic subgradient descent to solve the primal SVM efficiently on very large-scale datasets. Shortly after, the LIBLINEAR solver was developed specifically for large-scale linear classification [13]. Beyond algorithmic advances, there is also a line of research analyzing the data-dependent complexity of algorithms. These ideas have been investigated recently in [12, 1, 20, 21, 27], but are related to earlier work done by Freund and Vera [15], and Cucker and Peña [11] on convergence rates for optimization algorithms that depend on the conditioning of the problem instance. However, all these works focus on the optimization problem itself, without an underlying data science problem such as binary classification.

The rest of this paper is separated into two main sections. The first, Section 2, covers the ESP, where in Section 2.1 we model the ESP, and in Section 2.2 we discuss using FISTA to solve the ESP. Section 2.3 builds on the theory of the Primal-Dual Hybrid Gradient algorithm to show how to obtain a separating hyperplane from the iterates of the FISTA formulation, and Section 2.4 introduces a data-dependent bound on the number of FISTA iterates required to find an approximate separator. The second main section is Section 3, which covers using FISTA to approximately solve the soft-margin SVM problem. Section 3.1 discusses how FISTA can be used to solve the standard SVM problem by considering its dual. In this section it is observed that the problem of projecting onto the constraint set of the dual SVM is a continuous quadratic knapsack problem, which can be solved in linear time. The section ends by noting that the dual SVM may have multiple minimizers, motivating Section 3.2, which introduces a perturbed dual SVM formulation and subsequently finds its dual to be a perturbed version of the standard primal SVM problem. Relevant theory is also introduced in this section that will be relied upon later. Sections 3.3 and 3.4 introduce the assumptions and algorithmic parameters we use to obtain the results of this main section. Section 3.5 details how well-classified points can be identified from an approximate dual solution, while Section 3.6 contains the main result of Section 3: how FISTA stopped early can obtain a hyperplane that separates the well-classified points. Implementation details and numerical experiments for both the ESP and SVM are provided in Section 4, before the conclusion of the paper. Details not essential to the readability of the paper are left to the appendix.

The contributions of this paper are both theoretical and practical. We show that two binary classification problems, namely the ellipsoid separation problem and the soft-

margin support vector machine, can be solved using FISTA. Moreover, we develop data-dependent complexity results for both the ESP and SVM that show that a good approximate separator can be obtained from the FISTA iterates when stopping early, giving rise to data-dependent stopping conditions that diminish the need for tolerance-based stopping conditions that may be difficult to select. In numerical experiments, our FISTA implementation for the SVM problem is competitive with widely used algorithms for solving SVMs, especially on large-scale problems, making it a candidate for routine use on large-scale instances of SVM.

2 Ellipsoid separation

2.1 Modeling the ellipsoid separation problem

In the ellipsoid separation problem, we are given $j + l$ ellipsoids $C_1, \dots, C_j, D_1, \dots, D_l$ in $\mathbb{R}^d$ and we want to know if there exists a hyperplane $\{\mathbf{z} \in \mathbb{R}^d : \mathbf{w}^T \mathbf{z} = s\}$ such that $\mathbf{w}^T \mathbf{z} < s$ for all $\mathbf{z} \in C_1 \cup \dots \cup C_j$ and $\mathbf{w}^T \mathbf{z} > s$ for all $\mathbf{z} \in D_1 \cup \dots \cup D_l$.

Denote the ellipsoids as

$$\begin{aligned} C_i &= \{\mathbf{z} : \|A_i^{-1}(\mathbf{z} - \mathbf{c}_i)\| \leq 1\}, i = 1, \dots, j, \\ D_i &= \{\mathbf{z} : \|B_i^{-1}(\mathbf{z} - \mathbf{d}_i)\| \leq 1\}, i = 1, \dots, l, \end{aligned}$$

where A_i and B_i are $d \times d$ nonsingular matrices, and $\mathbf{c}_i$ and $\mathbf{d}_i$ are d -vectors that represent the centers of the ellipsoids. We make the assumption throughout this section that

$$\max \left\{ \max_i \{\|\mathbf{c}_i\|, \|A_i\|\}, \max_j \{\|\mathbf{d}_j\|, \|B_j\|\} \right\} = 1. \quad (1)$$

Let us begin with a lemma that characterizes the containment of an ellipsoid in a halfspace.

Lemma 2.1. *Consider the ellipsoid $E := \{\mathbf{z} : \|A^{-1}(\mathbf{z} - \mathbf{c})\| \leq 1\}$, where $A \in \mathbb{R}^{d \times d}$, $\mathbf{c} \in \mathbb{R}^d$, and $s \in \mathbb{R}$. Let $\mathbf{w} \in \mathbb{R}^d \setminus \{\mathbf{0}\}$. Then E lies in a halfspace $H := \{\mathbf{z} : \mathbf{w}^T \mathbf{z} \leq -s\}$ if and only if $-s \geq \|A^T \mathbf{w}\| + \mathbf{c}^T \mathbf{w}$. Further, $E \subseteq \text{int}(H)$ if and only if this inequality is strict.*

Proof. See Section A of the appendix. □

For each $i = 1, \dots, j$, we want the ellipsoids C_i to be contained in the halfspace defined by $(\mathbf{w}, -s)$, i.e. $\{\mathbf{z} : \|A_i^{-1}(\mathbf{z} - \mathbf{c}_i)\| \leq 1\} \subseteq \{\mathbf{z} : \mathbf{w}^T \mathbf{z} \leq -s\}$. We have shown this is equivalent to requiring $\|A_i^T \mathbf{w}\| \leq -s - \mathbf{c}_i^T \mathbf{w}$ and the reader may observe that this is equivalent to writing

$$\begin{pmatrix} s + \mathbf{c}_i^T \mathbf{w} \\ A_i^T \mathbf{w} \end{pmatrix} \in -C_2^{d+1},$$

where C_2^{d+1} is the second-order cone. Similarly, we wish for all $i = 1, \dots, l$ that $\{\mathbf{z} : \|B_i^{-1}(\mathbf{z} - \mathbf{d}_i)\| \leq 1\} \subseteq \{\mathbf{z} : \mathbf{w}^T \mathbf{z} \geq t\}$, or equivalently

$$\begin{pmatrix} t - \mathbf{d}_i^T \mathbf{w} \\ B_i^T \mathbf{w} \end{pmatrix} \in -C_2^{d+1},$$

The problem of finding two disjoint halfspaces parameterized by $(\mathbf{w}, -s)$ and $(\mathbf{w}, t)$ such that the ellipsoids C_i for $i = 1, \dots, j$ are contained in the former halfspace and the ellipsoids D_i for $i = 1, \dots, l$ are contained in the latter, can be written as a homogeneous second-order cone program (SOCP):

$$\begin{aligned} \max_{s, t, \mathbf{w}} \quad & s + t \\ \text{s.t.} \quad & \mathcal{A}^T \begin{pmatrix} s \\ t \\ \mathbf{w} \end{pmatrix} \preceq_{\mathcal{C}} \mathbf{0}, \end{aligned} \tag{2}$$

where the matrix $\mathcal{A}^T \in \mathbb{R}^{(j+l)(d+1) \times (d+2)}$ is

$$\mathcal{A}^T := \begin{pmatrix} \bar{A}_1^T \\ \vdots \\ \bar{A}_j^T \\ \bar{B}_1^T \\ \vdots \\ \bar{B}_l^T \end{pmatrix},$$

where

$$\bar{A}_i^T := \begin{pmatrix} 1 & 0 & \mathbf{c}_i^T \\ \mathbf{0} & \mathbf{0} & A_i^T \end{pmatrix}, \quad \bar{B}_i^T := \begin{pmatrix} 0 & 1 & -\mathbf{d}_i^T \\ \mathbf{0} & \mathbf{0} & B_i^T \end{pmatrix},$$

and the notation $\mathbf{a} \preceq_{\mathcal{C}} \mathbf{b}$ means that $\mathbf{a} - \mathbf{b} \in -\mathcal{C}$, where

$$\mathcal{C} := \underbrace{C_2^{d+1} \times \dots \times C_2^{d+1}}_{(j+l) \text{ times}}.$$

Due to homogeneity, the objective does not yield a finite maximum; rather, feasibility with $s + t > 0$ certifies that the halfspaces are disjoint and a separator exists. Otherwise, the optimum is 0, certifying that no separating hyperplane exists.

2.2 FISTA for ellipsoid separation

It is not apparent how to apply FISTA [4] directly to (2), since there is no closed-form expression to project onto the constraint set $\{(s, t, \mathbf{w}) : \mathcal{A}^T(s, t, \mathbf{w}) \preceq_{\mathcal{C}} \mathbf{0}\}$. It turns out that the dual SOCP admits a constraint set that can be efficiently projected onto.

Introduce the dual variables $\lambda_1, \dots, \lambda_j, \mu_1, \dots, \mu_l \in \mathbb{R}$ and $\mathbf{p}_1, \dots, \mathbf{p}_j, \mathbf{q}_1, \dots, \mathbf{q}_l \in \mathbb{R}^d$, and define

$$\mathbf{x} := \begin{pmatrix} \lambda_1 \\ \mathbf{p}_1 \\ \vdots \\ \lambda_j \\ \mathbf{p}_j \\ \mu_1 \\ \mathbf{q}_1 \\ \vdots \\ \mu_l \\ \mathbf{q}_l \end{pmatrix}.$$

The dual SOCP is then

$$\begin{aligned} \min_{\mathbf{x}} \quad & 0 \\ \text{s.t.} \quad & \mathcal{A}\mathbf{x} = \mathbf{b} \\ & \mathbf{x} \succeq_{\mathcal{C}} \mathbf{0}, \end{aligned} \tag{3}$$

where $\mathbf{b} = (1, 1, \mathbf{0})$. Written explicitly, the constraint $\mathcal{A}\mathbf{x} = \mathbf{b}$ becomes

$$\begin{aligned} \lambda_1 + \dots + \lambda_j &= 1 \\ \mu_1 + \dots + \mu_l &= 1, \\ \sum_{i=1}^j (\lambda_i \mathbf{c}_i + A_i \mathbf{p}_i) - \sum_{i=1}^l (\mu_i \mathbf{d}_i - B_i \mathbf{q}_i) &= \mathbf{0}, \end{aligned} \tag{4}$$

while the constraint $\mathbf{x} \succeq_{\mathcal{C}} \mathbf{0}$ becomes

$$\begin{aligned} \|\mathbf{p}_i\| &\leq \lambda_i & \forall i = 1, \dots, j \\ \|\mathbf{q}_i\| &\leq \mu_i & \forall i = 1, \dots, l. \end{aligned}$$

The above constraints have a nice geometric interpretation: they characterize the points in $\text{conv}(C_1 \cup \dots \cup C_j) \cap \text{conv}(D_1 \cup \dots \cup D_l)$, so solving (3) is equivalent to finding a point in the above intersection of convex hulls. Strong duality holds for (3) due to the geometric fact that either a point exists in the above intersection or there exists a strict separating hyperplane for the two convex hulls. In this context, strong duality means that either the primal optimizer is 0 and the dual is feasible and attains the value 0 (there is no separating hyperplane), or the primal maximum is $+\infty$ and the dual is infeasible (there is a separating hyperplane).

In order to formulate (3) for FISTA, we use the fact that the objective is identically zero to move the constraint $\mathcal{A}\mathbf{x} = \mathbf{b}$ to the objective, obtaining the new problem:

$$\begin{aligned} \min_{\mathbf{x}} \quad & f(\mathbf{x}) \\ \text{s.t.} \quad & \mathbf{x} \succeq_{\mathcal{C}} \mathbf{0}, \end{aligned} \tag{5}$$

where $f(\mathbf{x}) := \frac{1}{2} \|\mathcal{A}\mathbf{x} - \mathbf{b}\|^2$. Our SOCP (5) is now in a form FISTA can solve: $\min f(\mathbf{x}) + i_{\mathcal{C}}(\mathbf{x})$, where we note that projection onto the second-order cone can be done in linear time (cf. [3, Example 6.37]).

2.3 From PDHG to FISTA: obtaining a separating hyperplane from the iterates

Primal-dual hybrid gradient (PDHG) [8] can also be applied to the primal-dual formulation of the ESP formed by (2) and (3), however its convergence rate in terms of function values is $\mathcal{O}(1/k)$, compared to FISTA which converges at a rate of $\mathcal{O}(1/k^2)$. Despite this, PDHG has rich theory that we make use of throughout Section 2. For the theoretical results from PDHG related to this section, we refer the reader to Section A.1 of the appendix. With the results of Section A.1, we may prove a theorem connecting the infimal displacement vector of PDHG to the residuals of FISTA when applied to (5):

Theorem 2.1. *Consider the FISTA formulation of the standard conic form problem where $\mathbf{c} = \mathbf{0}$, namely $\min\{f(\mathbf{x}) := \frac{1}{2}\|\mathcal{A}\mathbf{x} - \mathbf{b}\|^2 : \mathbf{x} \in \mathcal{C}\}$, and assume that $\text{Null}(\mathcal{A}) \cap \mathcal{C} = \{\mathbf{0}\}$ and $\mathbf{v} = (\mathbf{v}_R, \mathbf{v}_D)$ is the infimal displacement vector of PDHG applied to the ESP. Then the following hold:*

- (i) *There exists $\bar{\mathbf{x}} \in \mathcal{C}$ such that $\mathbf{b} - \mathcal{A}\bar{\mathbf{x}} = \mathbf{v}_D$.*
- (ii) *$\mathbf{v}_D$ is the unique minimal residual vector for FISTA, i.e. $\mathcal{S} := \{\mathbf{x} \in \mathcal{C} : \mathbf{b} - \mathcal{A}\mathbf{x} = \mathbf{v}_D\}$ is the set of optimizers for (5).*

Proof. Since $\text{Null}(\mathcal{A}) \cap \mathcal{C} = \{\mathbf{0}\}$, we have from Theorem A.1 that $\mathbf{v}_R = \mathbf{0}$ and $\mathbf{v} \in \text{Range}(\text{Id} - T)$, which means that there exists some $(\bar{\mathbf{x}}, \bar{\mathbf{y}})$ such that $\mathbf{v} = (\text{Id} - T)(\bar{\mathbf{x}}, \bar{\mathbf{y}})$. Referring to the PDHG update (29), where $\sigma > 0$ satisfies $\sigma < 1/\|\mathcal{A}\|^2$ and $\tau = 1$, we get

$$\bar{\mathbf{x}} = \text{proj}_{\mathcal{C}}(\bar{\mathbf{x}} - \sigma \mathcal{A}^T \bar{\mathbf{y}}) \quad (6)$$

$$\mathbf{v}_D = \bar{\mathbf{y}} - \bar{\mathbf{y}} - \mathcal{A}(2\text{proj}_{\mathcal{C}}(\bar{\mathbf{x}} - \sigma \mathcal{A}^T \bar{\mathbf{y}}) - \bar{\mathbf{x}}) + \mathbf{b}. \quad (7)$$

From (6), $\bar{\mathbf{x}} \in \mathcal{C}$, thus (7) reduces to $\mathbf{b} - \mathcal{A}\bar{\mathbf{x}} = \mathbf{v}_D$, which proves (i).

To prove (ii), use (i) to find some $\bar{\mathbf{x}}$ such that $\mathbf{b} - \mathcal{A}\bar{\mathbf{x}} = \mathbf{v}_D$. Then for any $\mathbf{u} \in \mathcal{C}$ we may write

$$\begin{aligned} (\mathbf{u} - \bar{\mathbf{x}})^T \mathcal{A}^T \mathbf{v}_D &= \mathbf{u}^T \mathcal{A}^T \mathbf{v}_D - \bar{\mathbf{x}}^T \mathcal{A}^T \mathbf{v}_D \\ &\leq -\mathbf{v}_D^T \mathcal{A} \bar{\mathbf{x}} \\ &= 0, \end{aligned}$$

where the inequality comes from Theorem A.1 (i) and the fact that $\mathbf{u} \in \mathcal{C}$, while the final equality comes from Theorem A.1 (iii). The inequality $(\mathbf{u} - \bar{\mathbf{x}})^T \mathcal{A}^T \mathbf{v}_D \leq 0$ is equivalent to the condition that $\mathcal{A}^T \mathbf{v}_D \in \mathcal{N}_{\mathcal{C}}(\bar{\mathbf{x}})$, which is itself equivalent to $-\nabla f(\bar{\mathbf{x}}) \in \mathcal{N}_{\mathcal{C}}(\bar{\mathbf{x}})$, where $\mathcal{N}_{\mathcal{C}}(\bar{\mathbf{x}})$ is the normal cone at the point $\bar{\mathbf{x}}$. It follows that $\bar{\mathbf{x}}$ is an optimal solution to $\min\{f(\mathbf{x}) : \mathbf{x} \in \mathcal{C}\}$, since f has full domain, is proper and convex, and $\mathcal{C}$ has nonempty interior (cf. [3, Theorem 3.67]).

To complete the proof of (ii), we note that $\frac{1}{2}\|\cdot\|^2$ is strongly convex, so any optimal solution must have the same value as $\mathbf{b} - \mathcal{A}\bar{\mathbf{x}}$, which we know is precisely $\mathbf{v}_D$. Hence, $\mathbf{v}_D$ is the unique minimal residual vector for the FISTA formulation. $\square$

Now, let

$$\mathbf{v}_D = \begin{pmatrix} s^* \\ t^* \\ \mathbf{w}^* \end{pmatrix}.$$

From Theorem A.1, we have $\mathcal{A}^T \mathbf{v}_D \in \mathcal{C}^\circ$, so $\mathcal{A}^T(s^*, t^*, \mathbf{w}^*) \preceq_{\mathcal{C}} \mathbf{0}$ and $\mathbf{v}_D$ is feasible for (2). Since $\mathbf{v}_D$ is the limiting residual vector of FISTA, the following lemma tells us that in the separable case of the ESP, $\mathbf{v}_D$ encodes a separating hyperplane.

Lemma 2.2 ([16, Theorem 6.19]). $\mathbf{v}_D \neq \mathbf{0} \iff s^* + t^* > 0$.

To obtain a separating hyperplane from the FISTA residual $\mathbf{v}_D = (s^*, t^*, \mathbf{w}^*)$, we simply choose any $s' \in (-s^*, t^*)$. The corresponding separating hyperplane is $\{\mathbf{z} : (\mathbf{w}^*)^T \mathbf{z} = s'\}$. It remains to deduce a separator from the residual of the FISTA iterates $\mathbf{v}_k := \mathbf{b} - \mathcal{A}\mathbf{x}_k$, and determine after how many iterations k is this separator valid.

Recall for $\mathbf{v}_D = (s^*, t^*, \mathbf{w}^*)$, we have

$$\begin{aligned} (\mathbf{w}^*)^T \mathbf{z} &\leq -s^*, \quad \forall \mathbf{z} \in C_1 \cup \dots \cup C_j, \\ (\mathbf{w}^*)^T \mathbf{z} &\geq t^*, \quad \forall \mathbf{z} \in D_1 \cup \dots \cup D_l. \end{aligned}$$

Define $\mathbf{v}_k := (s_k, t_k, \mathbf{w}_k)$ and let $m_k := (-s_k + t_k)/2$ so that we have a candidate separator $\{\mathbf{z} : \mathbf{w}_k^T \mathbf{z} = m_k\}$. We wish to determine when this candidate separates $C_1 \cup \dots \cup C_j$ from $D_1 \cup \dots \cup D_l$.

Theorem 2.2. *The residual $\mathbf{v}_k = (s_k, t_k, \mathbf{w}_k)$ defines an approximate separator $\{\mathbf{z} : \mathbf{w}_k^T \mathbf{z} = m_k\}$ provided*

$$\|\mathbf{v}_k - \mathbf{v}_D\| < \frac{\|\mathbf{v}_D\|^2}{\sqrt{18}}$$

Proof. For any $\mathbf{z} \in C_1 \cup \dots \cup C_j$, we have that $(\mathbf{w}^*)^T \mathbf{z} \leq -s^*$. This inequality holds true if and only if

$$\mathbf{w}_k^T \mathbf{z} \leq m_k + (\mathbf{w}_k - \mathbf{w}^*)^T \mathbf{z} - s^* - m_k.$$

Let

$$r := \|\mathbf{w}_k - \mathbf{w}^*\| \cdot \|\mathbf{z}\| - \frac{s^* - s_k}{2} + \frac{t^* - t_k}{2} - \frac{s^* + t^*}{2}.$$

By Cauchy-Schwarz, any $\mathbf{z} \in C_1 \cup \dots \cup C_j$ satisfies

$$\begin{aligned} \mathbf{w}_k^T \mathbf{z} &= m_k + (\mathbf{w}_k - \mathbf{w}^*)^T \mathbf{z} + \mathbf{w}^{*T} \mathbf{z} - m_k \\ &\leq m_k + \|\mathbf{w}_k - \mathbf{w}^*\| \cdot \|\mathbf{z}\| - s^* - m_k \\ &= m_k + r. \end{aligned}$$

It follows that if we have $r < 0$, then $\mathbf{w}_k^T \mathbf{z} \leq m_k$. Under the normalization assumption (1), we have that $\|\mathbf{z}\| \leq 2$, and it is easy to verify that $r < 0$ is implied by the condition

$$s^* + t^* > |s^* - s_k| + |t^* - t_k| + 4\|\mathbf{w}_k - \mathbf{w}^*\|. \quad (8)$$

The implication $\mathbf{z} \in D_1 \cup \dots \cup D_l \implies \mathbf{w}_k^T \mathbf{z} > m_k$ holds under the same assumption. Therefore, the satisfaction of the condition (8) implies that the residual $\mathbf{v}_k$ defines an approximate separator.

Now observe the bound

$$|s^* - s_k| + |t^* - t_k| + 4\|\mathbf{w}_k - \mathbf{w}^*\| \leq \sqrt{18}\|\mathbf{v}_k - \mathbf{v}_D\|$$

holds, so a new condition for $\mathbf{v}_k$ defining an approximate separator is

$$\|\mathbf{v}_k - \mathbf{v}_D\| \leq \frac{s^* + t^*}{\sqrt{18}}.$$

Applying Lemma A.1 completes the proof. $\square$

2.4 Obtaining a data-dependent bound on the number of FISTA iterations

In order to derive a data-dependent bound on the number of FISTA iterations needed to obtain a separating hyperplane, we must relate $\|\mathbf{v}_D\|$ to the geometry of the underlying data. This is done in the following theorem.

Theorem 2.3. *Define $\delta_{\mathcal{A}}$ to be the smallest perturbation to the points $\mathbf{c}_i$ and $\mathbf{d}_i$ such that their ellipsoids are no longer separable. Then*

$$\delta_{\mathcal{A}} \leq 3\|\mathbf{v}_D\|.$$

Proof. See Section A.2 of the appendix. $\square$

Corollary 2.1. *The residual $\mathbf{v}_k = \mathbf{b} - \mathcal{A}\mathbf{x}_k = (s_k, t_k, \mathbf{w}_k)$ defines an approximate separator $\{\mathbf{z} : \mathbf{w}_k^T \mathbf{z} = m_k\}$, where $m_k = (-s_k + t_k)/2$, provided*

$$\|\mathbf{v}_k - \mathbf{v}_D\| < \frac{\delta_{\mathcal{A}}^2}{27\sqrt{2}}.$$

Proof. We have shown in Theorem 2.2 that $\mathbf{v}_k$ is an approximate separator provided $\|\mathbf{v}_k - \mathbf{v}_D\| < \|\mathbf{v}_D\|^2/\sqrt{18}$. We have also shown in Theorem 2.3 that we can lower bound $\|\mathbf{v}_D\|$ by the minimum perturbation $\delta_{\mathcal{A}}$ that makes the ESP no longer separable: $\|\mathbf{v}_D\| \geq \delta_{\mathcal{A}}/3$. Hence, to ensure $\|\mathbf{v}_k - \mathbf{v}_D\| < \|\mathbf{v}_D\|^2/\sqrt{18}$, we require that

$$\|\mathbf{v}_k - \mathbf{v}_D\| < \frac{\delta_{\mathcal{A}}^2}{27\sqrt{2}}.$$

$\square$

Theorem 2.4. *FISTA applied to the reformulation of the ESP (5) finds an approximate separating hyperplane in at most*

$$\frac{54\sqrt{2}\|\mathcal{A}\| \cdot \|\mathbf{x}_0 - \bar{\mathbf{x}}\|}{\delta_{\mathcal{A}}^2} - 1$$

iterations.

Proof. From Corollary 2.1, we know that $\mathbf{v}_k = (s_k, t_k, \mathbf{w}_k)$ encodes a separating hyperplane when $\|\mathbf{v}_k - \mathbf{v}_D\| < \delta_{\mathcal{A}}^2/27\sqrt{2}$. We want to know when $\mathbf{v}_k$ lies in a ball of radius $\delta_{\mathcal{A}}^2/27\sqrt{2}$ around $\mathbf{v}_D$. Recall (cf. [3, Theorem 10.34]) that for iterates $\mathbf{x}_k$ generated by FISTA and any $\bar{\mathbf{x}} \in \mathcal{S} = \{\mathbf{x} \in \mathcal{C} : \mathbf{b} - \mathcal{A}\mathbf{x} = \mathbf{v}_D\}$,

$$f(\mathbf{x}_k) - f(\bar{\mathbf{x}}) \leq \frac{2L\|\mathbf{x}_0 - \bar{\mathbf{x}}\|^2}{(k+1)^2},$$

where $\mathbf{x}_0$ is the starting point of the algorithm and L is the Lipschitz parameter of the gradient of f . One can verify that $L = \|\mathcal{A}\|^2$, so we obtain

$$f(\mathbf{x}_k) - f(\bar{\mathbf{x}}) \leq \frac{2\|\mathcal{A}\|^2\|\mathbf{x}_0 - \bar{\mathbf{x}}\|^2}{(k+1)^2}.$$

From Theorem A.2, we have $\|\mathbf{v}_k - \mathbf{v}_D\| \leq \sqrt{2}\sqrt{f(\mathbf{x}_k) - f(\bar{\mathbf{x}})}$, so applying this to the above we obtain the bound

$$\|\mathbf{v}_k - \mathbf{v}_D\| \leq \frac{2\|\mathcal{A}\|\|\mathbf{x}_0 - \bar{\mathbf{x}}\|}{k+1}.$$

It follows that we can guarantee that $\mathbf{v}_k$ lies in a ball of radius $\delta_{\mathcal{A}}^2/27\sqrt{2}$ around $\mathbf{v}_D$ when $2\|\mathcal{A}\|\|\mathbf{x}_0 - \bar{\mathbf{x}}\|/(k+1) \leq \delta_{\mathcal{A}}^2/27\sqrt{2}$. That is, when

$$k \geq \frac{54\sqrt{2}\|\mathcal{A}\|\|\mathbf{x}_0 - \bar{\mathbf{x}}\|}{\delta_{\mathcal{A}}^2} - 1. \quad (9)$$

□

The termination test (28) used in the numerical experiments for the ESP efficiently checks whether the algorithm has found a separator, ensuring termination when (9) is satisfied.

3 Support Vector Machines

The ellipsoid separation problem assumes the data are separable. We now suppose the data are not separable, but we still look for a hyperplane that separates the points well. This problem can be modeled as a soft-margin support vector machine (SVM).

In this section we assume the data are points rather than ellipsoids. Let the points be $\mathbf{c}_1, \dots, \mathbf{c}_j, \mathbf{d}_1, \dots, \mathbf{d}_l \in \mathbb{R}^d$. Let the (a priori unknown) separating hyperplane be $(\mathbf{w}, t)$, i.e. we want $\mathbf{w}^T \mathbf{c}_i < -t$ for $i = 1 \dots j$ and $\mathbf{w}^T \mathbf{d}_i > -t$ for $i = 1, \dots, l$. The soft-margin SVM is modeled as a convex quadratic programming problem:

$$\begin{aligned} \min_{\mathbf{s}, t, \mathbf{w}} \quad & \frac{1}{2}\|\mathbf{w}\|^2 + \gamma \sum_{i=1}^{j+l} s_i \\ \text{s.t.} \quad & \mathbf{c}_i^T \mathbf{w} + t \geq 1 - s_i, \quad i = 1, \dots, j, \\ & \mathbf{d}_i^T \mathbf{w} + t \leq -1 + s_{i+j}, \quad i = 1, \dots, l, \\ & \mathbf{s} \geq \mathbf{0}, \end{aligned} \quad (10)$$

where we say a point $\mathbf{c}_i$ or $\mathbf{d}_i$ is misclassified if its corresponding entry of $\mathbf{s}$ is positive, and $\gamma > 0$ is a hyperparameter that penalizes misclassification. Observe that (10) is bounded below by zero and always feasible, since one may always choose $\mathbf{s}$ to ensure the constraints are satisfied. In fact, we can rewrite (10) as unconstrained nonsmooth convex optimization by letting $\phi(\theta) := \gamma \max(1 - \theta, 0)$ and writing

$$\min_{\mathbf{w}, t} \frac{1}{2} \|\mathbf{w}\|^2 + \sum_{i=1}^j \phi(\mathbf{c}_i^T \mathbf{w} + t) + \sum_{i=1}^l \phi(-\mathbf{d}_i^T \mathbf{w} - t) =: f(\mathbf{w}, t). \quad (11)$$

Unlike in the ESP, there is no question of whether a separating hyperplane exists, but finding the optimal separating hyperplane requires convergence to $(\mathbf{w}^*, t^*)$. To develop the theory for only an approximate solution, we need to impose a data model. We return to this point in Section 3.3.

3.1 FISTA for support vector machines

The soft-margin SVM problem (10) is a sum of a convex smooth term and several convex nonsmooth terms. Unfortunately, there is no closed-form prox of the nonsmooth terms. We get around this problem by again passing to the dual, which for (10) is the following:

$$\begin{aligned} \max_{\mathbf{u}, \mathbf{v}} \quad & -\frac{1}{2} \left\| \sum_{i=1}^j u_i \mathbf{c}_i - \sum_{i=1}^l v_i \mathbf{d}_i \right\|^2 + \mathbf{e}^T \mathbf{u} + \mathbf{e}^T \mathbf{v} \\ \text{s.t.} \quad & \mathbf{e}^T \mathbf{u} = \mathbf{e}^T \mathbf{v}, \\ & \mathbf{0} \leq \mathbf{u} \leq \gamma \mathbf{e}, \\ & \mathbf{0} \leq \mathbf{v} \leq \gamma \mathbf{e}. \end{aligned} \quad (12)$$

Let $\Omega := \{(\mathbf{u}, \mathbf{v}) : \mathbf{e}^T \mathbf{u} = \mathbf{e}^T \mathbf{v}, \mathbf{0} \leq \mathbf{u} \leq \gamma \mathbf{e}, \mathbf{0} \leq \mathbf{v} \leq \gamma \mathbf{e}\}$. The projection $\text{proj}_\Omega(\mathbf{x}, \mathbf{y})$ is the solution to the following optimization problem:

$$\begin{aligned} \min_{\mathbf{u}, \mathbf{v}} \quad & \frac{1}{2} \|\mathbf{u} - \mathbf{x}\|^2 + \frac{1}{2} \|\mathbf{v} - \mathbf{y}\|^2 \\ \text{s.t.} \quad & \mathbf{e}^T \mathbf{u} = \mathbf{e}^T \mathbf{v}, \\ & \mathbf{0} \leq \mathbf{u} \leq \gamma \mathbf{e}, \\ & \mathbf{0} \leq \mathbf{v} \leq \gamma \mathbf{e}. \end{aligned} \quad (13)$$

After dropping constants, (13) can be rewritten in the form of a continuous quadratic knapsack problem:

$$\begin{aligned} \min_{\mathbf{r}} \quad & \frac{1}{2} \mathbf{r}^T \mathbf{I} \mathbf{r} - \mathbf{q}^T \mathbf{r} \\ \text{s.t.} \quad & \mathbf{e}^T \mathbf{r} = \mathbf{0}, \\ & \boldsymbol{\ell} \leq \mathbf{r} \leq \boldsymbol{\nu}, \end{aligned} \quad (14)$$

where $\mathbf{r} := (\mathbf{u}, -\mathbf{v})$, $\mathbf{q} := (\mathbf{x}, -\mathbf{y})$, $\boldsymbol{\ell} := (\mathbf{0}, -\gamma \mathbf{e})$, and $\boldsymbol{\nu} := (\gamma \mathbf{e}, \mathbf{0})$. An algorithm for exactly solving (14) in $\mathcal{O}(j + l)$ operations was introduced in [7], yielding a linear time method for projecting onto Ω .

3.1.1 Difficulties with identifying well-classified points

Since we have a linear time algorithm to compute $\text{proj}_\Omega(\mathbf{x}, \mathbf{y})$, FISTA can be applied efficiently to (12). In fact, the complementarity relation of the KKT (Karush-Kuhn-Tucker) conditions says that an optimal primal-dual pair for (10) and (12) must satisfy $s_i(\gamma - u_i) = 0$ for all $i \in [j]$ and $s_{j+i}(\gamma - v_i) = 0$ for all $i \in [l]$. Thus, if we have an optimal dual solution with $u_i < \gamma$, then $s_i = 0$ and $\mathbf{c}_i$ is correctly classified. The same applies for $\mathbf{v}$. A preliminary idea for identifying well-classified points is then to run FISTA on the dual of the soft-margin SVM problem (12) until the indices i such that $u_i < \gamma$ or $v_i < \gamma$ at the optimum are identified.

The difficulty with this approach is that the dual problem may have multiple minimizers. Furthermore, in some cases there exist minimizers where a well-classified point $\mathbf{c}_i$ or $\mathbf{d}_i$ satisfies $u_i = \gamma$ or $v_i = \gamma$. For an example of such a case, see Section B.1 in the Appendix. The convergence of the iterates of FISTA was recently proved in [5]; however, for convex problems with multiple minimizers, these results do not identify which minimizer FISTA converges to.

3.2 Solving the perturbed dual

To get around the problem of (12) having multiple minimizers, our approach is to make the dual (12) μ -strongly concave via a perturbation to the objective:

$$\begin{aligned} \max \quad & g(\mathbf{u}, \mathbf{v}) := -\frac{1}{2}\|C^T \mathbf{u} - D^T \mathbf{v}\|^2 + \mathbf{e}^T \mathbf{u} + \mathbf{e}^T \mathbf{v} - \frac{\mu}{2}\|\mathbf{u}\|^2 - \frac{\mu}{2}\|\mathbf{v}\|^2 \\ \text{s.t.} \quad & (\mathbf{u}, \mathbf{v}) \in \Omega, \end{aligned} \quad (15)$$

where $C = (\mathbf{c}_1^T, \dots, \mathbf{c}_j^T) \in \mathbb{R}^{j \times d}$ and $D = (\mathbf{d}_1^T, \dots, \mathbf{d}_l^T) \in \mathbb{R}^{l \times d}$. For the remainder of the section, we will focus on solving the formulation (15).

Theorem 3.1. *Suppose all points lie within a fixed distance R from the origin. The objective function $g(\mathbf{u}, \mathbf{v})$ defined in (15) is L -smooth, where $L = \mu + nR^2$.*

Proof. The gradient of g with respect to $\mathbf{u}$ is $-C(C^T \mathbf{u} - D^T \mathbf{v}) + \mathbf{e} - \mu \mathbf{u}$, while the gradient of g with respect to $\mathbf{v}$ is $D(C^T \mathbf{u} - D^T \mathbf{v}) + \mathbf{e} - \mu \mathbf{v}$. By letting

$$\mathbf{z} := \begin{pmatrix} \mathbf{u} & \mathbf{v} \end{pmatrix} \quad \text{and} \quad H^T := \begin{pmatrix} C^T & -D^T \end{pmatrix},$$

we may write

$$\nabla g(\mathbf{z}) = -HH^T \mathbf{z} + \mathbf{e} - \mu \mathbf{z},$$

and easily obtain the Hessian:

$$\nabla^2 g(\mathbf{z}) = -HH^T - \mu I.$$

It follows that

$$\begin{aligned} \|\nabla^2 g(\mathbf{z})\|_2 &= \|HH^T + \mu I\|_2 \\ &= \|H\|_2^2 + \mu \\ &\leq nR^2 + \mu, \end{aligned}$$

where in the final inequality we used Cauchy-Schwarz and that the norm of each row of H is bounded by R , under the assumption that all data lie within a fixed distance from the origin. $\square$

The next theorem gives a bound that allows us to monitor proximity to the optimizer $(\mathbf{u}^*, \mathbf{v}^*)$ of the perturbed dual (15) in terms of the projected gradient. Since the projected gradient is computed on each FISTA iteration, this theorem implies the projected gradient step is a low cost measure of convergence.

Theorem 3.2. *Let $\mathbf{r} := (\mathbf{u}, \mathbf{v}) \in \mathbb{R}^j \times \mathbb{R}^l$ and let*

$$(\mathbf{u}^+, \mathbf{v}^+) = \mathbf{r}^+ := \text{proj}_\Omega \left(\mathbf{r} + \frac{1}{L} \nabla g(\mathbf{r}) \right).$$

Then,

$$\|\mathbf{r}^+ - \mathbf{r}^*\| \leq \frac{2L}{\mu} \|\mathbf{r} - \mathbf{r}^+\|,$$

where $\mathbf{r}^ = (\mathbf{u}^*, \mathbf{v}^*)$ is the optimal solution of (15).*

Proof. Since $\mathbf{r}^*$ is optimal, there must exist some $\mathbf{z}^* \in \mathcal{N}_\Omega(\mathbf{r}^*)$ with $\mathbf{z}^* = \nabla g(\mathbf{r}^*)$. Also, from Lemma B.1, there exists some $\mathbf{z}^+ \in \mathcal{N}_\Omega(\mathbf{r}^+)$ such that $\|\nabla g(\mathbf{r}^+) - \mathbf{z}^+\| \leq 2L\|\mathbf{r} - \mathbf{r}^+\|$. By μ -strong convexity of $-g$,

$$-(\nabla g(\mathbf{r}^+) - \nabla g(\mathbf{r}^*))^T (\mathbf{r}^+ - \mathbf{r}^*) \geq \mu \|\mathbf{r}^+ - \mathbf{r}^*\|^2,$$

while by the definition of the normal cone, $(\mathbf{r}^+ - \mathbf{r}^*)^T \mathbf{z}^+ \geq 0$ and $(\mathbf{r}^+ - \mathbf{r}^*)^T \mathbf{z}^* \leq 0$. We may write $\nabla g(\mathbf{r}^*) = \mathbf{z}^+ + (\mathbf{z}^* - \mathbf{z}^+)$, which gives

$$\begin{aligned} -(\nabla g(\mathbf{r}^+) - \nabla g(\mathbf{r}^*))^T (\mathbf{r}^+ - \mathbf{r}^*) &= -(\nabla g(\mathbf{r}^+) - \mathbf{z}^+)^T (\mathbf{r}^+ - \mathbf{r}^*) + (\mathbf{z}^* - \mathbf{z}^+)^T (\mathbf{r}^+ - \mathbf{r}^*) \\ &\leq (\nabla g(\mathbf{r}^+) - \mathbf{z}^+)^T (\mathbf{r}^* - \mathbf{r}^+) \\ &\leq \|\nabla g(\mathbf{r}^+) - \mathbf{z}^+\| \cdot \|\mathbf{r}^+ - \mathbf{r}^*\| \\ &\leq 2L\|\mathbf{r} - \mathbf{r}^+\| \cdot \|\mathbf{r}^+ - \mathbf{r}^*\|. \end{aligned}$$

Combining the inequality obtained by strong concavity, and the above, we get

$$\mu \|\mathbf{r}^+ - \mathbf{r}^*\|^2 \leq 2L\|\mathbf{r} - \mathbf{r}^+\| \cdot \|\mathbf{r}^+ - \mathbf{r}^*\|,$$

which under the assumption that $\mathbf{r}^+ \neq \mathbf{r}^*$, gives

$$\|\mathbf{r}^+ - \mathbf{r}^*\| \leq \frac{2L}{\mu} \|\mathbf{r} - \mathbf{r}^+\|.$$

$\square$

In order to connect the perturbed dual back to the problem of finding a separating hyperplane, we take the dual of (15):

$$\begin{aligned} \min_{\mathbf{s}, \mathbf{t}, \mathbf{w}, \mathbf{z}} \quad & \frac{1}{2} \|\mathbf{w}\|^2 + \gamma \mathbf{e}^T \mathbf{s} + \frac{1}{2\mu} \|\mathbf{z}\|^2 \\ \text{s.t.} \quad & C\mathbf{w} + \mathbf{t}\mathbf{e} + \mathbf{z}_1 \geq \mathbf{e} - \mathbf{s}_1, \\ & D\mathbf{w} + \mathbf{t}\mathbf{e} - \mathbf{z}_2 \leq -\mathbf{e} + \mathbf{s}_2, \\ & \mathbf{s} \geq \mathbf{0}, \end{aligned} \tag{16}$$

where $\mathbf{s}_1 = \mathbf{s}(1 : j)$, $\mathbf{s}_2 = \mathbf{s}(j + 1 : j + l)$ and $\mathbf{z}_1 = \mathbf{z}(1 : j)$, $\mathbf{z}_2 = \mathbf{z}(j + 1 : j + l)$. The KKT conditions relating (15) and (16) are

$$\begin{aligned}
\mathbf{u} &= \mathbf{z}_1 / \mu \\
\mathbf{v} &= \mathbf{z}_2 / \mu \\
\mathbf{w} &= C^T \mathbf{u} - D^T \mathbf{v} \\
\mathbf{e}^T \mathbf{u} &= \mathbf{e}^T \mathbf{v} \\
\mathbf{z}_1 &= (1 - t) \mathbf{e} - \mathbf{s}_1 - C \mathbf{w} + \boldsymbol{\zeta}_1 \\
\mathbf{z}_2 &= (1 + t) \mathbf{e} - \mathbf{s}_2 + D \mathbf{w} + \boldsymbol{\zeta}_2 \\
\mathbf{s}_1^T (\mathbf{u} - \gamma \mathbf{e}) &= 0 \\
\mathbf{s}_2^T (\mathbf{v} - \gamma \mathbf{e}) &= 0 \\
\boldsymbol{\zeta}_1^T \mathbf{u} &= 0 \\
\boldsymbol{\zeta}_2^T \mathbf{v} &= 0 \\
\mathbf{0} &\leq \mathbf{u} \leq \gamma \mathbf{e} \\
\mathbf{0} &\leq \mathbf{v} \leq \gamma \mathbf{e} \\
\mathbf{s} &\geq \mathbf{0} \\
\boldsymbol{\zeta} &\geq \mathbf{0},
\end{aligned} \tag{17}$$

where $\boldsymbol{\zeta}_1 = \boldsymbol{\zeta}(1 : j)$ and $\boldsymbol{\zeta}_2 = \boldsymbol{\zeta}(j + 1 : j + l)$. If we suppose $\mathbf{w}, t$ are known in (16), then the entries of $\mathbf{s}$ and $\mathbf{z}$ separate. That is, if for some $i = 1, \dots, j$ we let $\theta = \mathbf{c}_i^T \mathbf{w} + t$ or for some $i = j + 1, \dots, l$, $\theta = -\mathbf{d}_i^T \mathbf{w} - t$, then the optimal s_i, z_i solve the problem

$$\psi(\theta) := \min \left\{ \gamma s + \frac{1}{2\mu} z^2 : \theta \geq 1 - s - z, s \geq 0 \right\}, \tag{18}$$

which can be written in closed form as the following differentiable approximation to hinge loss:

$$\psi(\theta) = \begin{cases} -\frac{1}{2}\mu\gamma^2 + \gamma(1 - \theta), & \theta \in (-\infty, 1 - \mu\gamma], \\ \frac{1}{2\mu}(1 - \theta)^2, & \theta \in [1 - \mu\gamma, 1], \\ 0, & \theta \in [1, +\infty). \end{cases} \tag{19}$$

In the context of (16), θ can be viewed as the signed margin of a point $\mathbf{p}_i$ with corresponding class label y_i , where $y_i = 1$ and $\mathbf{p}_i = \mathbf{c}_i$ for $i = 1, \dots, j$, and $y_i = -1$ and $\mathbf{p}_i = \mathbf{d}_i$ for $i = j + 1, \dots, l$:

$$\theta = y_i(\mathbf{w}^T \mathbf{p}_i + t).$$

With a closed-form solution to $\psi(\theta)$, we may now write (16) as an unconstrained convex optimization problem which will be helpful for our analysis.

$$\min_{\mathbf{w}, t} f(\mathbf{w}, t) := \frac{1}{2} \|\mathbf{w}\|^2 + \sum_{i=1}^j \psi(\mathbf{c}_i^T \mathbf{w} + t) + \sum_{i=1}^l \psi(-\mathbf{d}_i^T \mathbf{w} - t). \tag{20}$$

3.3 Assumptions on the data

We will assume throughout the rest of this section that the data form two clusters, each with a center, plus noisy points. Let $n := j + l$, and suppose $\sigma_1, \sigma_2 \in [1/2, 1]$ and $0 < \rho < \min(\sigma_1, \sigma_2)$. We assume that there are j_0 “well-classified” points in $B(\sigma_1 \mathbf{e}_1, \rho)$, there are l_0 “well-classified” points in $B(-\sigma_2 \mathbf{e}_1, \rho)$, and the n_{noise} noisy points are contained in the ball that encloses the two balls of *well-classified* points:

$$B\left(\frac{\sigma_1 - \sigma_2}{2} \mathbf{e}_1, \frac{\sigma_1 + \sigma_2}{2} + \rho\right).$$




Figure 1: An illustration of the assumption on the data, where we have two balls of well-classified points centered at -1 and $+1$, and a larger ball containing the two, which also includes the noisy points.

In our model we have $j_0 + l_0 + n_{\text{noise}} = n$, where $j_0 \geq n/3$, $l_0 \geq n/3$, and $n_{\text{noise}} \leq \nu n$ with $\nu > 0$. This assumption also implies that all points must lie within a fixed distance $R := \max(\sigma_1, \sigma_2) + \rho$ from the origin.

Now we have a data model, we can consider two useful questions:

1. Can well-classified points be identified with an approximate solution?
2. Can a (suboptimal) separator be identified from an approximate solution that separates the well-classified points?

3.4 Algorithmic parameters

In order to solve (15), we must select values for the misclassification penalty γ and the strong-concavity parameter μ . We fix $\gamma := 64/n$ and $\mu := n/128$, and observe that this makes $1 - \mu\gamma = 1/2$, ensuring that the breakpoints of the quadratic portion of $\psi(\theta)$ are $\theta \in [1/2, 1]$. Hence, if the signed margin $\theta \geq 1$, the corresponding point makes no

contribution to the objective, but as the margin decreases to $\theta \in [1/2, 1]$, the contribution begins to increase quadratically, before becoming linear when $\theta \leq 1/2$. This behavior is depicted in Figure 2. Our choice of γ follows the common convention of normalizing the hinge-loss term by the sample size, which ensures that the objective function remains well-scaled as n varies. Such normalization appears throughout the literature, for example in [24, Equation (15.4)] and [25, Equation (1)]. Another benefit of this choice of parameters is very fast convergence when using the strongly convex variant of FISTA (cf. [3, Chapter 7]). Indeed the strongly convex variant of FISTA converges at the rate $\mathcal{O}((1 - 1/\sqrt{\kappa})^k)$, where $\kappa = L/\mu$, and Theorem 3.1 gives $L = \mu + nR^2$, so

$$\kappa = 1 + 128R^2 \leq 513,$$

which is sufficiently small to ensure fast convergence of FISTA.




Figure 2: A plot of $\psi(\theta)$, where $n = 100$, $\gamma = 64/n$, $\mu = n/128$.

3.5 Identifying well-classified points

We say that a point is *properly classified* if its optimal value for (15) satisfies $u_i^* < \gamma$ or $v_i^* < \gamma$. The goal of this section is to prove that all well classified points are properly classified by FISTA provided that the distance $\|(\mathbf{u}_k, \mathbf{v}_k) - (\mathbf{u}^*, \mathbf{v}^*)\| \leq \Delta$, where Δ is a data-dependent quantity defined in (26). This result is stated precisely in Theorem 3.5. The termination test in this theorem depends on parameters that may be unknown a priori when the algorithm is run. In Section 4.2, we explain how to terminate FISTA in practice even when the parameters are unknown.

As in the unperturbed case, the complementarity relation of the KKT conditions (17) says that an optimal primal-dual pair for (15) and (16) must satisfy $s_i^*(\gamma - u_i^*) = 0$ for all $i \in [j]$ and $s_{j+i}^*(\gamma - v_i^*) = 0$ for all $i \in [l]$. This implies that for a properly classified point, $s_i^* = 0$.

The derivative of $\psi(\theta)$ for our choices of μ, γ is

$$\psi'(\theta) = \begin{cases} -\frac{64}{n}, & \theta \in (-\infty, 1/2), \\ \frac{128}{n}(\theta - 1), & \theta \in [1/2, 1], \\ 0, & \theta \in (1, +\infty), \end{cases} \quad (21)$$

which we use in the following theorem that connects the separator obtained from the primal optimizers $(\mathbf{w}^*, t^*)$ to the optimal dual variables. Here, we define $\boldsymbol{\psi}'_c(\mathbf{w}, t) := \sum_{i=1}^j \psi'(\mathbf{c}_i^T \mathbf{w} + t) \mathbf{e}_i$ and $\boldsymbol{\psi}'_d(\mathbf{w}, t) := \sum_{i=1}^l \psi'(-\mathbf{d}_i^T \mathbf{w} - t) \mathbf{e}_i$, i.e. $\boldsymbol{\psi}'_c(\mathbf{w}, t)$ is the vector with entries $\psi'(\mathbf{c}_i^T \mathbf{w} + t)$.

Theorem 3.3. *Suppose $(\mathbf{u}^*, \mathbf{v}^*)$ are optimal solutions to the perturbed dual SVM problem (15). Then*

$$\mathbf{u}^* = -\boldsymbol{\psi}'_c(\mathbf{w}^*, t^*), \text{ and } \mathbf{v}^* = -\boldsymbol{\psi}'_d(\mathbf{w}^*, t^*).$$

Proof. We will prove this result for $\mathbf{u}^*$ since the $\mathbf{v}^*$ case follows the same reasoning.

From the KKT conditions (17), we have $u_i^* = z_{i,1}^*/\mu = (1 - \theta_i^* - s_{i,1}^* + \zeta_{i,1}^*)/\mu$, where $\theta_i^* := \mathbf{c}_i^T \mathbf{w}^* + t^*$.

First, consider the case where $\theta_i^* \geq 1$. In this case, $\psi'(\theta_i^*) = 0$, and in the dual of the dual (16), $z_{i,1}^* = 0 = s_{i,1}^*$ must hold. It follows that $u_i^* = 0 = -\psi'(\theta_i^*)$.

The next case is $\theta_i^* \in (1 - \mu\gamma, 1)$. Suppose for sake of contradiction that $u_i^* = 0$, then by the complementary slackness condition $s_{i,1}^*(u_i - \gamma) = 0$, $s_{i,1}^* = 0$, and $u_i^* = (1 - \theta_i^* + \zeta_{i,1}^*)/\mu$. But then

$$0 = u_i^* = (1 - \theta_i^* + \zeta_{i,1}^*)/\mu \iff \zeta_{i,1}^* = \theta_i^* - 1,$$

where $\theta_i^* - 1 < 0$. This is a contradiction because $\zeta_{i,1}^* \geq 0$. So $u_i^* > 0$, which from the other complementary slackness condition $u_i^* \zeta_{i,1}^* = 0$, means $\zeta_{i,1}^* = 0$. Now suppose that $u_i^* = \gamma$. This leads to a contradiction also, since

$$\gamma = u_i^* = (1 - \theta_i^* - s_{i,1}^*)/\mu \iff s_{i,1}^* = 1 - \mu\gamma - \theta_i^*,$$

but $1 - \mu\gamma - \theta_i^* < 0$ whereas $s_{i,1}^* \geq 0$. It follows that $u_i^* \in (0, \gamma)$, which implies by complementary slackness that $s_{i,1}^* = 0 = \zeta_{i,1}^*$, and

$$u_i^* = \frac{1}{\mu}(1 - \theta_i^*) = -\psi'(\theta_i^*).$$

Finally, consider the case where $\theta_i^* \leq 1 - \mu\gamma$. Suppose that $u_i^* < \gamma$. Then $s_{i,1}^* = 0$, and

$$\gamma > u_i^* = (1 - \theta_i^* + \zeta_{i,1}^*)/\mu \iff \mu\gamma > 1 - \theta_i^* + \zeta_{i,1}^*,$$

but $1 - \theta_i^* + \zeta_{i,1}^* \geq \mu\gamma$, which is a contradiction. It follows that

$$u_i^* = \gamma = -\psi'(\theta_i^*).$$

□

The following lemma says that, for iteration k , under the assumption that $\mathbf{c}_i$ is properly classified, once $u_{k,i}^+ := (\mathbf{u}_k^+)_i$ lies in an open ball of radius $\Delta/2$ around u_i^* , we can be sure that $u_{k,i} := (\mathbf{u}_k)_i$ is more than $\Delta/2$ away from its upper bound of γ . The exact same reasoning can be applied to $v_{k,i}^+ := (\mathbf{v}_k^+)_i$ for $i \in [l]$.

Lemma 3.1. *Suppose $\gamma - u_i^* = \Delta$ for some $\Delta > 0$ and $i \in [j]$. If $|u_{k,i}^+ - u_i^*| < \Delta/2$, then $\gamma - u_{k,i}^+ > \Delta/2$. Likewise, if $\gamma - v_i^* = \Delta$ for some $\Delta > 0$ and $i \in [l]$, then $\gamma - v_{k,i}^+ > \Delta/2$.*

Proof. Suppose $\gamma - u_i^* = \Delta$ for some $\Delta > 0$ and $i \in [j]$. Then

$$\begin{aligned} \frac{\Delta}{2} &> |u_{k,i}^+ - u_i^*| = |u_{k,i}^+ - \gamma + \Delta| \\ &\geq -|u_{k,i}^+ - \gamma| + |\Delta| \\ &= -(\gamma - u_{k,i}^+) + \Delta. \end{aligned}$$

The same argument applies for any v_i^* . $\square$

In fact, the two inequalities above in terms of $u_{k,i}^+$ can tell us that u_i^* is properly classified.

Lemma 3.2. *Let $\Delta > 0$ and $i \in [j]$. If $|u_{k,i}^+ - u_i^*| < \Delta/2$ and $\gamma - u_{k,i}^+ > \Delta/2$, then $u_i^* < \gamma$ and u_i^* is properly classified.*

Likewise, if $|v_{k,i}^+ - v_i^| < \Delta/2$ and $\gamma - v_{k,i}^+ > \Delta/2$, then v_i^* is properly classified for any $i \in [l]$.*

Proof. Suppose $|u_{k,i}^+ - u_i^*| < \Delta/2$ and $\gamma - u_{k,i}^+ > \Delta/2$. Then we may write

$$\begin{aligned} \frac{\Delta}{2} &> |u_{k,i}^+ - \gamma + \gamma - u_i^*| \\ &\geq |u_{k,i}^+ - \gamma| - |\gamma - u_i^*| \\ &> \frac{\Delta}{2} - |\gamma - u_i^*|, \end{aligned}$$

which implies $|\gamma - u_i^*| > 0$, so $u_i^* < \gamma$. The proof for the case of v_i is identical. $\square$

The conditions $|u_{k,i}^+ - u_i^*| < \Delta/2$ and $\gamma - u_{k,i}^+ > \Delta/2$ can be easily verified in FISTA. Recall Theorem 3.2 says that the size of the projected gradient step is an upper bound on $\|(\mathbf{u}_k^+, \mathbf{v}_k^+) - (\mathbf{u}^*, \mathbf{v}^*)\|$, so we may check upon each iteration if the size of the projected gradient step is less than $\Delta/2$ and then verify for which i we satisfy $\gamma - u_{k,i}^+ > \Delta/2$ and $\gamma - v_{k,i}^+ > \Delta/2$ to determine the properly classified points.

The auxiliary primal variables $(\mathbf{s}, \mathbf{z})$ are determined by $(\mathbf{w}, t)$ as the solutions to (18). In the remainder of this section we characterize $(\mathbf{w}^*, t^*)$ with the goal of finding bounds for $\mathbf{z}^*$ and subsequently expressing the optimal dual variables $\mathbf{u}^*$ and $\mathbf{v}^*$ via the identities $\mathbf{u}^* = \mathbf{z}_1^*/\mu$ and $\mathbf{v}^* = \mathbf{z}_2^*/\mu$.

Lemma 3.3. *If $\mathbf{c}_i^T \mathbf{w}^* + t^* \geq 1 - K$ for some $K \in [0, 1/2)$, then $u_i^* \leq 2K\gamma < \gamma$. Likewise, if $-\mathbf{d}_i^T \mathbf{w}^* - t^* \geq 1 - K$ for some $K \in [0, 1/2)$, then $v_i^* \leq 2K\gamma < \gamma$.*

Proof. Suppose $K \in [0, 1/2)$. From the KKT conditions of (18) it can be shown that $z_i^* = \max(1 - \theta, 0)$ for all $\theta \geq 1/2$, and in either case of $\theta = \mathbf{c}_i^T \mathbf{w}^* + t^* \geq 1 - K$ or $\theta = -\mathbf{d}_i^T \mathbf{w}^* - t^* \geq 1 - K$, $\theta \in (1/2, 1]$. We may now apply the KKT conditions (17) for the perturbed dual formulation (15), in particular the stationarity conditions $u_i^* = z_i^*/\mu = 128(1 - \theta)/n = 2K\gamma$ or $v_i^* = z_i^*/\mu = 128(1 - \theta)/n = 2K\gamma$. Since $K < 1/2$, $2K\gamma < \gamma$. $\square$

The above lemma is useful because it implies that if $\mathbf{c}_i^T \mathbf{w}^* + t^* \geq 1/2 + \Omega(1)$, then $u_i^* \leq \gamma - \Omega(1)$, certifying that $\mathbf{c}_i$ is properly classified. The case for $\mathbf{d}_i$ is symmetric. We now prove a bound on the norm of $\mathbf{w}^*$ in terms of the positive constant ξ . Notably, $\lim_{\rho, \nu \rightarrow 0} \xi = (1 - \sigma)/\sigma \in [0, 1]$, where $\sigma := \min(\sigma_1, \sigma_2)$.

Lemma 3.4. $\|\mathbf{w}^*\| \leq 1 + \xi$, where $\xi := \sqrt{(1 + \delta)^2 + 32\nu(11 + 8\delta)} - 1$ and $\delta := \frac{1}{\sigma - \rho} - 1$.

Proof. Recall the objective in (20):

$$f(\mathbf{w}, t) = \frac{1}{2}\|\mathbf{w}\|^2 + \sum_{i=1}^j \psi(\mathbf{c}_i^T \mathbf{w} + t) + \sum_{i=1}^l \psi(-\mathbf{d}_i^T \mathbf{w} - t)$$

and note that $\mathbf{w} = (1 + \delta)\mathbf{e}_1$, $t = 0$ is a valid solution, where the definitions $\sigma = \min(\sigma_1, \sigma_2)$ and $\rho < \sigma$ ensure $\delta > 0$. Under this choice of δ , we have for all $\mathbf{c}_i$ well-classified:

$$\begin{aligned} \mathbf{c}_i^T \mathbf{w} + t &= \mathbf{c}_i^T \mathbf{w} \\ &= (1 + \delta)\mathbf{c}_i^T \mathbf{e}_1 \\ &\geq (1 + \delta)(\sigma_1 - \rho)\mathbf{e}_1^T \mathbf{e}_1 \\ &= (1 + \delta)(\sigma_1 - \rho) \\ &= \frac{\sigma_1 - \rho}{\sigma - \rho} \\ &\geq 1, \end{aligned}$$

where we used that all well-classified points $\mathbf{c}_i$ lie in $B(\sigma_1 \mathbf{e}_1, \rho)$ and that $\rho < \min(\sigma_1, \sigma_2)$. The same argument can be made to see that $\mathbf{d}_i^T \mathbf{w} \leq -1$ for all $\mathbf{d}_i$ well-classified. Therefore, $f(\mathbf{w}, t)$ for this choice of $\mathbf{w}$ and t has contributions only from the term $\frac{1}{2}\|\mathbf{w}\|^2$ and the misclassified points, since $\psi(\theta) = 0$ for all well-classified points. For $\mathbf{c}_i$ not well-classified, we have that

$$\mathbf{c}_i^T \mathbf{w} + t = \mathbf{c}_i^T \mathbf{w} \geq -2(1 + \delta),$$

where $-2(1 + \delta) < -2$. The quadratic portion of ψ is bounded above by $64/n$, while the linear portion corresponding to the interval $(-\infty, 1/2)$ is always strictly larger than $64/n$ for $\theta < -2$. Hence, we may use the linear portion of ψ to obtain the upper bound

$$\psi(\mathbf{c}_i^T \mathbf{w} + t) \leq \psi(-2(1 + \delta)) \leq \frac{16}{n}(11 + 8\delta).$$

The argument for $\psi(-\mathbf{d}_i^T \mathbf{w} - t) \leq \frac{16}{n}(11+8\delta)$ for $\mathbf{d}_i$ not well-classified is similar. Therefore, the sum terms in (20) must be bounded as follows, since $n_{\text{noise}} \leq \nu n$.

$$\sum_{i=1}^j \psi(\mathbf{c}_i^T \mathbf{w} + t) + \sum_{i=1}^l \psi(-\mathbf{d}_i^T \mathbf{w} - t) \leq 16\nu(11+8\delta).$$

Moreover, $\frac{1}{2}\|\mathbf{w}\|^2 = \frac{1}{2}(1+\delta)^2$, so we have

$$f(\mathbf{w}, t) \leq \frac{1}{2}(1+\delta)^2 + 16\nu(11+8\delta),$$

and thus

$$f(\mathbf{w}^*, t^*) \leq f(\mathbf{w}, t) \leq \frac{1}{2}(1+\delta)^2 + 16\nu(11+8\delta).$$

We may now obtain a bound on $\|\mathbf{w}^*\|$ by noting that $\|\mathbf{w}^*\|^2 \leq 2f(\mathbf{w}^*, t^*)$:

$$\|\mathbf{w}^*\| \leq \sqrt{f(\mathbf{w}^*, t^*)} \leq \sqrt{(1+\delta)^2 + 32\nu(11+8\delta)}.$$

Letting $\xi := \sqrt{(1+\delta)^2 + 32\nu(11+8\delta)} - 1$, we may rewrite this bound as $\|\mathbf{w}^*\| \leq 1 + \xi$. $\square$

For the rest of this section we make the assumption that ρ, ν in our data model are chosen sufficiently small to ensure that

$$\frac{1 + \xi + 64\nu}{\sigma - \rho} \leq \frac{64}{3}. \quad (22)$$

Lemma 3.5. *Let $J \subseteq \{1, \dots, j\}$ be the well-classified points in the $\mathbf{c}$ -class, and $L \subseteq \{1, \dots, l\}$ be the well-classified points in the $\mathbf{d}$ -class. There exists an $i \in J$ such that*

$$\mathbf{c}_i^T \mathbf{w}^* + t^* \geq 1 - \frac{3(1 + \xi + 64\nu)}{128(\sigma_1 - \rho)}.$$

Likewise, there exists an $i \in L$ such that

$$-\mathbf{d}_i^T \mathbf{w}^* - t^* \geq 1 - \frac{3(1 + \xi + 64\nu)}{128(\sigma_2 - \rho)}.$$

Proof. We will prove the inequality for the $\mathbf{c}$ -points only, the proof for the $\mathbf{d}$ -points proceeds analogously.

By optimality of $(\mathbf{w}^*, t^*)$, the following derivative must be zero:

$$\frac{d}{dw_1} f(\mathbf{w}^*, t^*) = w_1^* + \sum_{i=1}^j \psi'(\mathbf{c}_i^T \mathbf{w}^* + t^*) c_{1,i} - \sum_{i=1}^l \psi'(-\mathbf{d}_i^T \mathbf{w}^* - t^*) d_{i,1}.$$

Define

$$\begin{aligned} q &:= \max_{i \in J} \psi'(\mathbf{c}_i^T \mathbf{w}^* + t^*) c_{i,1}, \\ r &:= \max_{i \in L} (-\psi'(-\mathbf{d}_i^T \mathbf{w}^* - t^*) d_{i,1}). \end{aligned}$$

Note that in the definition of $\psi'(\theta)$, $-1/2 \leq \theta - 1 \leq 0$ on the interval $\theta \in [1/2, 1]$, so $|\psi'(\theta)| \leq 64/n$. Moreover, $|c_{i,1}| \leq 2$ and $|d_{i,1}| \leq 2$ for $\mathbf{c}_i$ and $\mathbf{d}_i$ not well-classified, allowing us to write the bound:

$$\begin{aligned}
0 &= \frac{d}{dw_1} f(\mathbf{w}^*, t^*) \\
&= w_1 + \sum_{i=1}^k \psi'(\mathbf{c}_i^T \mathbf{w}^* + t^*) c_{i,1} - \sum_{i=1}^l \psi'(-\mathbf{d}_i^T \mathbf{w}^* - t) d_{i,1} \\
&\leq (1 + \xi) + |J|q + |L|r + \frac{64n_{\text{noise}}}{n} \\
&\leq (1 + \xi) + |J|q + |L|r + 64\nu.
\end{aligned}$$

Here, the first term comes from Lemma 3.4, i.e. $w_1 \leq \|\mathbf{w}^*\| \leq 1 + \xi$, the second and third terms come from the well-classified points, while the last term comes from the ill-classified points. Rearranging the above inequality we obtain

$$|J|q + |L|r \geq -1 - \xi - 64\nu.$$

Since both terms on the left hand side are nonpositive (because $\psi'(\cdot) \leq 0$ and $\mathbf{c}_{i,1} \geq 0$ for well-classified $\mathbf{c}$ -points, and $d_{i,1} \leq 0$ for well-classified $\mathbf{d}$ -points), we can deduce two separate inequalities:

$$\begin{aligned}
|J|q &\geq -1 - \xi - 64\nu, \\
|L|r &\geq -1 - \xi - 64\nu.
\end{aligned}$$

The first inequality is rearranged to:

$$q \geq -\frac{1 + \xi + 64\nu}{|J|}.$$

Thus, for $i \in J$ giving the maximum in the definition of q ,

$$\psi'(\mathbf{c}_i^T \mathbf{w}^* + t^*) c_{i,1} \geq -\frac{1 + \xi + 64\nu}{|J|},$$

and since $c_{i,1} \geq \sigma_1 - \rho$, we must have that $\psi'(\mathbf{c}_i^T \mathbf{w}^* + t^*) c_{i,1} \leq \psi'(\mathbf{c}_i^T \mathbf{w}^* + t^*)(\sigma_1 - \rho)$, so

$$\psi'(\mathbf{c}_i^T \mathbf{w}^* + t^*) \geq -\frac{1 + \xi + 64\nu}{|J|(\sigma_1 - \rho)}. \quad (23)$$

By the assumption (22) on ρ, ν we can ensure that

$$\frac{1 + \xi + 64\nu}{\sigma_1 - \rho} \leq \frac{64}{3},$$

in which case the assumption that $|J| \geq n/3$ allows us to write

$$\begin{aligned}\psi'(\mathbf{c}_i^T \mathbf{w}^* + t^*) &\geq -\frac{3}{n} \left(\frac{1 + \xi + 64\nu}{\sigma_1 - \rho} \right) \\ &\geq -\frac{3}{n} \cdot \frac{64}{3} \\ &= -\frac{64}{n}.\end{aligned}$$

At the same time, we know that $\psi'(\mathbf{c}_i^T \mathbf{w}^* + t^*) \leq 0$, thus $\psi'(\mathbf{c}_i^T \mathbf{w}^* + t^*) \in (-64/n, 0]$ and we can be sure from the definition of ψ' in (21) that $\mathbf{c}_i^T \mathbf{w}^* + t^* \in (1/2, 1]$. This allows us to rewrite (23) as

$$\frac{128(\mathbf{c}_i^T \mathbf{w}^* + t^* - 1)}{n} \geq -\frac{1 + \xi + 64\nu}{|J|(\sigma_1 - \rho)}.$$

Multiply both sides by n , use the assumption on the data that $n/|J| \leq 3$, and divide by 128 to obtain

$$\mathbf{c}_i^T \mathbf{w}^* + t^* - 1 \geq -\frac{3(1 + \xi + 64\nu)}{128(\sigma_1 - \rho)}. \quad (24)$$

□

Corollary 3.1. *For any $i \in J$,*

$$\mathbf{c}_i^T \mathbf{w}^* + t^* \geq 1 - \frac{3(1 + \xi + 64\nu)}{128(\sigma - \rho)} - 3\rho(1 + \xi).$$

Likewise, for any $i \in L$

$$-\mathbf{d}_i^T \mathbf{w}^* - t^* \geq 1 - \frac{3(1 + \xi + 64\nu)}{128(\sigma - \rho)} - 3\rho(1 + \xi).$$

Proof. Again we prove this for $\mathbf{c}$ -points only because the proof for the $\mathbf{d}$ -points is analogous.

Let i be some other index in J . We may use Lemma 3.4 and (24) in Lemma 3.5 to write

$$\begin{aligned}\mathbf{c}_i^T \mathbf{w}^* + t^* - 1 &= c_{i',1} w_1^* + \mathbf{c}(2:d)^T \mathbf{w}^*(2:d) + t^* - 1 \\ &\geq (\sigma_1 - \rho) \mathbf{w}_1^* - \|\mathbf{c}(2:d)\| \cdot \|\mathbf{w}^*(2:d)\| + t^* - 1 \\ &\geq (\sigma_1 - \rho) \mathbf{w}_1^* - \rho(1 + \xi) + t^* - 1 \\ &= (\sigma_1 + \rho) w_1^* + t^* - 1 - \rho(1 + \xi) - 2\rho w_1^* \\ &\geq c_{i,1} w_1^* + t^* - 1 - \rho(1 + \xi) - 2\rho w_1^* \\ &\geq -\frac{3(1 + \xi + 64\nu)}{128(\sigma_1 - \rho)} - \rho(1 + \xi) - 2\rho w_1^* \\ &\geq -\frac{3(1 + \xi + 64\nu)}{128(\sigma - \rho)} - 3\rho(1 + \xi).\end{aligned}$$

□

Define $K(\rho, \nu) := 3(1 + \xi + 64\nu)/(128(\sigma - \rho)) + 3\rho(1 + \xi)$. Observe that as ρ, ν tend to zero, the right hand side in the inequality $K(\rho, \nu) \rightarrow 3/128\sigma^2 \leq 3/32 \leq 1/2 - \Omega(1)$. Hence, for each $i \in J$, a sufficiently small choice of ρ, ν will give $\mathbf{c}_i^T \mathbf{w}^* + t^* \geq 1 - K$ for $K \leq 1/2 - \Omega(1)$, and by Lemma 3.3, $u_i^* \leq \gamma - \Omega(1)$. Likewise for v_i^* where $i \in L$.

3.5.1 Identifying the well-classified points in FISTA

By Corollary 3.1 and Lemma 3.3, for ρ, ν sufficiently small per (22), the well-classified points at the optimum satisfy $u_i^*, v_i^* \leq 2K(\rho, \nu)\gamma$.

Theorem 3.4. *Suppose ρ, ν satisfy (22), and $K(\rho, \nu) < 1/2$. For all $i \in J$, $u_i^* \leq 2K(\rho, \nu)\gamma$ and for all $i \in L$, $v_i^* \leq 2K(\rho, \nu)\gamma$.*

Proof. This follows directly from Corollary 3.1 and Lemma 3.3. $\square$

If the data are such that $K(\rho, \nu) < 1/2$, then $\gamma - u_i^* \geq (1 - 2K(\rho, \nu))\gamma := \Delta_0$ for all $i \in J$ and $\gamma - v_i^* \geq \Delta_0$ for all $i \in L$. By Lemma 3.1, when $\|(\mathbf{u}_{k+1}, \mathbf{v}_{k+1}) - (\mathbf{u}^*, \mathbf{v}^*)\| < \Delta_0/2$, it must be the case that $\gamma - u_{k+1,i} > \Delta_0/2$ for all $i \in J$ and $\gamma - v_{k+1,i} > \Delta_0/2$ for all $i \in L$. But then we are in the case of Lemma 3.2, i.e. $\|(\mathbf{u}_{k+1}, \mathbf{v}_{k+1}) - (\mathbf{u}^*, \mathbf{v}^*)\| < \Delta_0/2$, and $\gamma - u_{k+1,i} > \Delta_0/2$ for all $i \in J$ and $\gamma - v_{k+1,i} > \Delta_0/2$ for all $i \in L$.

In the iterations of FISTA, we don't have access to $(\mathbf{u}^*, \mathbf{v}^*)$, however we do have the bound $\|(\mathbf{u}_{k+1}, \mathbf{v}_{k+1}) - (\mathbf{u}^*, \mathbf{v}^*)\| \leq \frac{2L}{\mu}\|(\mathbf{u}_k, \mathbf{v}_k) - (\mathbf{u}_{k+1}, \mathbf{v}_{k+1})\|$ from Theorem 3.2. If $\|(\mathbf{u}_k, \mathbf{v}_k) - (\mathbf{u}_{k+1}, \mathbf{v}_{k+1})\| < \Delta_0\mu/4L$, we can be sure that $\|(\mathbf{u}_{k+1}, \mathbf{v}_{k+1}) - (\mathbf{u}^*, \mathbf{v}^*)\| < \Delta_0/2$ as desired. It follows that we can find all the well-classified points in FISTA by checking the condition $\|(\mathbf{u}_k, \mathbf{v}_k) - (\mathbf{u}_{k+1}, \mathbf{v}_{k+1})\| < \Delta_0\mu/4L$, and once this condition is satisfied, finding all points such that $\gamma - u_{k+1,i} > \Delta_0/2$ and $\gamma - v_{k+1,i} > \Delta_0/2$. These points are properly classified and must include the well-classified points given by the sets J and L . Since the projected gradient is computed on each FISTA iteration, it is efficient to compute $\|(\mathbf{u}_k, \mathbf{v}_k) - (\mathbf{u}_{k+1}, \mathbf{v}_{k+1})\|$ on each iteration.

3.6 Obtaining a hyperplane that separates the well-classified points

So far we have shown how to determine which points are well-classified from the dual problem (15), but it remains to show how to obtain a hyperplane from an approximate dual solution that separates the well-classified $\mathbf{c}$ -points from the well-classified $\mathbf{d}$ -points. In this section we make the assumption that $\mathbf{c}_i^T \mathbf{w}^* + t^* \geq 1 - K$ for all $i \in J$ and $-\mathbf{d}_i^T \mathbf{w}^* - t^* \geq 1 - K$ for all $i \in L$, where $K := K(\rho, \nu) \in [0, 1/2)$ and $\bar{K} \in [1/2, 1)$ is chosen such that

$$\bar{K} - 3K^2 > 3\rho(1 + \xi) \left(\frac{4}{3} + 2\nu \right) + 3\nu \left(\frac{9}{4} + 2\xi \right) + \frac{1}{2}. \quad (25)$$

The above requires that ρ, ν are sufficiently small. Note that the right hand side above tends to $1/2$ as $\rho, \nu \rightarrow 0$.

Suppose we stop early with $\|(\mathbf{u}_k, \mathbf{v}_k) - (\mathbf{u}^*, \mathbf{v}^*)\| \leq \Delta$, where

$$\Delta := \frac{\frac{1}{3}(\bar{K} - 3K^2) - \rho(1 + \xi) \left(\frac{4}{3} + 2\nu\right) - \nu \left(\frac{9}{4} + 2\xi\right) - \frac{1}{6}}{4\sqrt{2n} \left(1 + \frac{2}{3}\rho\right)}, \quad (26)$$

and we compute $\mathbf{w}_k := C^T \mathbf{u}_k - D^T \mathbf{v}_k$. We take t_k to be the optimal solution to the unconstrained optimization problem

$$\min_t f(\mathbf{w}_k, t), \quad (27)$$

Since $f(\mathbf{w}_k, t)$ is the sum of convex functions, it is convex. Moreover, $\psi(\theta)$ is Lipschitz continuous. This allows us to show our main result, that upon stopping early, we obtain a hyperplane from $(\mathbf{w}_k, t_k)$ that separates the well-classified points with a margin of $1 - \bar{K}$.

Theorem 3.5. *Let $\bar{K}$ and K be given according to (25). Suppose $\|(\mathbf{u}_k, \mathbf{v}_k) - (\mathbf{u}^*, \mathbf{v}^*)\| \leq \Delta$, $\mathbf{w}_k := C^T \mathbf{u}_k - D^T \mathbf{v}_k$, and t_k is the solution to (27). Then $(\mathbf{w}_k, t_k)$ encode a separating hyperplane for the well-classified points, in particular*

$$\begin{aligned} \mathbf{c}_i^T \mathbf{w}_k + t_k &\geq 1 - \bar{K}, \quad \forall i \in J, \\ \mathbf{d}_i^T \mathbf{w}_k + t_k &\leq -1 + \bar{K}, \quad \forall i \in L. \end{aligned}$$

Proof. See Theorem B.1 in Section B.3 of the appendix. $\square$

4 Implementation and numerical experiments

4.1 Ellipsoid Separation Problem

The FISTA update for the ESP formulation (5) is

$$\begin{aligned} \mathbf{x}_{k+1} &= \text{proj}_{\mathcal{C}} \left(\mathbf{y}_k + \frac{1}{L} \mathcal{A}^T (\mathbf{b} - \mathcal{A} \mathbf{y}_k) \right) \\ t_{k+1} &= \frac{1 + \sqrt{1 + 4t_k^2}}{2} \\ \mathbf{y}_{k+1} &= \mathbf{x}_{k+1} + \left(\frac{t_k - 1}{t_{k+1}} \right) (\mathbf{x}_{k+1} - \mathbf{x}_k) \\ \mathbf{v}_{k+1} &= \mathbf{b} - \mathcal{A} \mathbf{x}_{k+1}, \end{aligned}$$

where $\mathbf{x}_0$ is the chosen starting point, $\mathbf{y}_0 = \mathbf{x}_0$, $t_0 = 1$, and $L = \|\mathcal{A}\|_2^2$. Since at every iteration we compute $\mathbf{v}_{k+1} = (s_{k+1}, t_{k+1}, \mathbf{w}_{k+1})$, it is only a small amount of extra work at the end of each iteration to check if the hyperplane $\{\mathbf{z} \in \mathbb{R}^d : \mathbf{w}_{k+1}^T \mathbf{z} = m_{k+1} = (-s_{k+1} + t_{k+1})/2\}$ separates $C_1 \cup \dots \cup C_j$ from $D_1 \cup \dots \cup D_l$ by verifying that

$$\begin{aligned} \mathbf{c}_i^T \mathbf{w}_{k+1} + \|\mathbf{A}_i^T \mathbf{w}_{k+1}\| &< m_{k+1}, \quad \forall i \in [j] \\ \mathbf{d}_i^T \mathbf{w}_{k+1} - \|\mathbf{B}_i^T \mathbf{w}_{k+1}\| &> m_{k+1}, \quad \forall i \in [l], \end{aligned}$$

or that (28)

$$\begin{aligned} \mathbf{c}_i^T \mathbf{w}_{k+1} - \|\mathbf{A}_i^T \mathbf{w}_{k+1}\| &> m_{k+1}, \quad \forall i \in [j] \\ \mathbf{d}_i^T \mathbf{w}_{k+1} + \|\mathbf{B}_i^T \mathbf{w}_{k+1}\| &< m_{k+1}, \quad \forall i \in [l]. \end{aligned}$$

By Lemma 2.1, if either of these conditions succeed, $\{\mathbf{z} \in \mathbb{R}^d : \mathbf{w}_{k+1}^T \mathbf{z} = m_{k+1}\}$ is a separating hyperplane. Moreover, one of these conditions is guaranteed to be satisfied once (9) is satisfied.

We tested our Julia implementation of FISTA for the ESP by separating ellipsoids derived from the digits 0 and 1 of the MNIST handwritten digits dataset [19]. A single digit from the MNIST dataset is a 28×28 matrix with entries in $[0, 255]$ representing pixel intensity, where 255 corresponds to a completely black pixel and 0 corresponds to a completely white pixel. In order to model an MNIST digit $\mathbf{x}$ as an ellipsoid, we first flatten $\mathbf{x}$ into a vector in $\mathbb{R}^{784}$ and normalize the entries to lie in $[0, 1]$. An ellipsoid is then constructed using $\mathbf{x}$ as the center:

$$E_{\mathbf{x}} = \{\mathbf{z} \in \mathbb{R}^{784} : \Sigma_{\mathbf{x}}^{-1}(\mathbf{z} - \mathbf{x})\},$$

where

$$\Sigma_{\mathbf{x}} = \text{diag}(\epsilon \mathbf{e} + \alpha(\mathbf{e} - \boldsymbol{\theta})),$$

and θ_i represents the “confidence” in a pixel being black or white:

$$\theta_i = |2x_i - 1|.$$

This model of an MNIST digit as an ellipsoid ensures that the ellipsoid stretches more in directions where the pixel color (black or white) becomes more uncertain, and contracts where confidence in the pixel color increases.

In our numerical experiments, we randomly select N digits from each class (0 or 1), varying N from 1 to 125. For each digit, we construct an ellipsoid using the process described above, with parameters $\epsilon = 10^{-10}$ and $\alpha = 1/2$. Before each experiment, the data are normalized according to (1). The results of these experiments can be seen in Table 1, which compares our method for solving the ellipsoid separation problem by early stopping to using FISTA with tolerance-based stopping: $\|\mathbf{y}_k - \text{proj}_{\mathcal{C}}(\mathbf{x} + \frac{1}{L}\mathcal{A}^T \mathbf{v}_k)\| \leq \text{tol}$, where $\text{tol} = 10^{-6}$. We also performed the same experiments on various unbalanced sets of ellipsoids, the results of which are listed in Appendix C, Table 3.

To study how the proposed stopping condition behaves as the sets of ellipsoids become better separated, we consider a simple two-dimensional ellipsoid separation problem. We first sample two random positive definite matrices $A_0, A_1 \in \mathbb{R}^{2 \times 2}$ of the form $A_j = \omega Q_j Q_j^T + \epsilon I$, where Q_j has i.i.d. standard normal entries, $\omega = 0.1$, and $\epsilon = 10^{-3}$. For a given distance d , we then define two ellipsoids $E_0(d)$ and $E_1(d)$ with centers $c_0 = (-d, 0)$ and $c_1 = (d, 0)$, and fixed shape matrices A_0 and A_1 . These ellipsoids are normalized by dividing both centers and shape matrices by a common factor so that the assumption (1) is satisfied. For each $d \in \{0.01, 0.02, \dots, 1.0\}$ we run FISTA on the corresponding ESP instance, starting from $\mathbf{x}_0 = \mathbf{0}$, using the proposed stopping condition with a maximum of 100 iterations. The results of this experiment are shown in Figure 3.

4.2 Support-Vector Machines

In order to compute an approximate separator for the SVM problem, we develop a stopping condition based on the results of Sections 3.2 and 3.5. In particular, from Theorem 3.2

Table 1: Comparing the number of iterations and time taken before finding a separating hyperplane in balanced data between early-stopping and tolerance-based stopping ($\text{tol} = 10^{-6}$).

Ellipsoids per class (N)	Columns in $\mathcal{A}$	FISTA (early-stopping)		FISTA (tolerance-based)	
		iterations	time(s)	iterations	time(s)
1	1570	1	0.172	531	0.372
5	7850	1	0.180	887	3.355
10	15700	2	0.193	988	7.522
25	39250	2	0.223	1374	26.662
50	78500	3	0.300	1627	66.916
75	117750	7	0.779	1686	118.945
100	157000	6	0.685	1818	166.617
125	196250	3	0.513	1856	256.162




Figure 3: Plot showing the number of iterations before obtaining a separating hyperplane vs the distance between the centers of two ellipsoids in two dimensions. A vertical red dashed line is used to mark the first distance at which the ESP was able to be solved.

$\|(\mathbf{u}_{k+1}, \mathbf{v}_{k+1}) - (\mathbf{u}^*, \mathbf{v}^*)\| \leq \frac{2L}{\mu} \|(\mathbf{u}_k, \mathbf{v}_k) - (\mathbf{u}_{k+1}, \mathbf{v}_{k+1})\|$, so set

$$\Delta := \frac{2L}{\mu} \|(\mathbf{u}_k, \mathbf{v}_k) - (\mathbf{u}_{k+1}, \mathbf{v}_{k+1})\|.$$

Clearly, $\|(\mathbf{u}_{k+1}, \mathbf{v}_{k+1}) - (\mathbf{u}^*, \mathbf{v}^*)\| < \Delta$ always holds, so by Lemma 3.2, if for any $i \in [j]$, $\gamma - u_{k+1,i} > \Delta$, then u_i^* is properly classified. Similarly, if for any $i \in [l]$, $\gamma - v_{k+1,i} > \Delta$, then v_i^* is properly classified. This means that throughout the FISTA iterations, we can keep a list T of the properly classified points $i \in n$, and on each iteration check if any more points should be added to T . To determine when to stop, we let p_{k+1} denote the percentage of points that have been properly classified at iteration $k+1$:

$$p_{k+1} = \frac{|T|}{n},$$

and let $\delta_{\min} > 0$ be some parameter that represents the minimum acceptable percentage change in properly classified points. Initially, we set two variables $p_{\text{best}} = 0$ and $\text{stale} = 0$, then at each iteration we verify $p_{k+1} > p_{\text{best}} + \delta_{\min}$. If the inequality is false, we increment stale by one, and if $\text{stale} \geq 2$, we stop the algorithm because it has stalled. If the inequality is true, we update p_{best} to p_{k+1} and reset stale to zero.

In all, we perform the following steps to solve the SVM:

1. Run FISTA on (15) with some choice of $\delta_{\min} > 0$ until it stalls as described above.
2. Compute

$$\mathbf{w}_{k+1} := C^T \mathbf{u}_{k+1} - D^T \mathbf{v}_{k+1}.$$

3. Solve the unconstrained optimization problem

$$t_{k+1} := \underset{t}{\operatorname{argmin}} f(\mathbf{w}_{k+1}, t)$$

by bisection.

By our choice of stopping condition, our theory guarantees that this method will produce a separator $(\mathbf{w}_{k+1}, t_{k+1})$ that separates all points in T .

We tested our implementation of this method on binary classification problems of various sizes chosen from the UCI Machine Learning repository [18]. Our numerical experiments used $\delta_{\min} = 10^{-4}$, and compared our FISTA implementation with three widely used SVM solvers: LIBSVM, LIBLINEAR, and Pegasos. All numerical experiments were written in Julia, so for LIBSVM and LIBLINEAR we used the LIBSVM.jl¹ and LIBLINEAR.jl² interfaces respectively, under the package defaults and with a linear kernel for LIBSVM. For Pegasos, we used the MLJ.jl library³ and set the maximum number of epochs to 500. All other parameters for Pegasos were left at their defaults.

¹<https://github.com/JuliaML/LIBSVM.jl>

²<https://github.com/JuliaML/LIBLINEAR.jl>

³<https://juliaai.github.io/MLJ.jl/stable/>




Figure 4: Plots depicting $\|\mathbf{y}_k - \text{proj}_\Omega(\mathbf{y}_k - \frac{1}{L}\nabla g(\mathbf{y}_k))\|$ for $k \in [0, 10000]$.

The results of the experiments are displayed in Table 2, which show that our method is competitive with LIBSVM and LIBLINEAR, and outperforms Pegasos. Notably, for the problem “a9a”, our method significantly outperformed the other solvers in time, while still achieving the highest accuracy. In this paper, the following definition of accuracy is used, where $\hat{\mathbf{y}} \in \mathbb{R}^n$ is the vector of predictions for each point and $\mathbf{y} \in \mathbb{R}^n$ is the vector of true labels:

$$\text{classification accuracy} = \frac{\text{correct classifications}}{\text{total classifications}} = \frac{|\{i : \hat{y}_i = y_i\}|}{n}.$$

Figure 4 plots the quantity $\|\mathbf{y}_k - \text{proj}_\Omega(\mathbf{y}_k - \frac{1}{L}\nabla g(\mathbf{y}_k))\|$ over the number of iterations for each problem, where $\mathbf{y}_k$ is the auxiliary sequence of FISTA. The plots include a horizontal dashed red line, which marks the value of $\|\mathbf{y}_k - \text{proj}_\Omega(\mathbf{y}_k - \frac{1}{L}\nabla g(\mathbf{y}_k))\|$ at which we stop early and obtain an approximate separator. These plots make it clear that a good separator can often be found without having to solve the underlying optimization problem to optimality.

Table 2: Classification accuracy and runtime for each algorithm and dataset.

Datasets	Samples	Features	FISTA		LIBSVM		LIBLINEAR		Pegasos	
			accu.	time(s)	accu.	time(s)	accu.	time(s)	accu.	time(s)
breast cancer	568	31	0.986	0.581	0.988	0.602	0.989	0.002	0.910	2.353
australian	589	14	0.858	0.577	0.856	0.029	0.872	0.005	0.779	0.350
a9a	22,696	123	0.847	1.919	0.847	146.434	0.741	2.695	0.786	30.992
flocking	24,014	2,400	0.999	46.005	0.999	38.174	1	7.867	0.795	677.247

5 Conclusion

We studied when FISTA applied to binary classification problems could be stopped early while still producing a useful classifier. For both the ellipsoid separation problem (ESP) and soft-margin support-vector machine problem (SVM), we showed how from the iterates of FISTA applied to the underlying optimization problem, we could extract a separator. Furthermore, we obtained data-dependent complexity results which showed that the effort to obtain a good separator scales with geometric properties of the data, rather than parameters of the underlying optimization problem.

For the ESP, we reformulated the dual second-order cone program so that we could apply FISTA, and showed that the FISTA residuals converge to the primal-dual hybrid gradient (PDHG) infimal displacement vector, which itself encodes a separating hyperplane. We showed that the number of iterations sufficient to produce a separating hyperplane scales as $\mathcal{O}(\delta_{\mathcal{A}}^{-2})$, where $\delta_{\mathcal{A}}$ is the smallest perturbation that destroys separability of the data. Our numerical experiments showed that as the classes of ellipsoids grew farther apart, the algorithm needed fewer iterations to find a separator, corroborating our data-dependent complexity theory.

For SVMs, we proposed a strongly concave perturbation of the standard soft-margin SVM dual, which ensures unique minimizers, and with our parameter choices, converges rapidly. Under a reasonable data model, we show that early-stopped iterates of FISTA applied to the perturbed dual identify well-classified points and can be used to efficiently compute a hyperplane. We obtain a data-dependent upper bound on the error $\|(\mathbf{u}_k, \mathbf{v}_k) - (\mathbf{u}^*, \mathbf{v}^*)\|$, after which the hyperplane computed from the FISTA iterates becomes a separating hyperplane for the well-classified points. In our experiments, our method for obtaining an approximate separating hyperplane performs well, and is competitive with LIBSVM and LIBLINEAR while outperforming Pegasos in runtime and accuracy.

Finally, we highlight two directions for future research. Specifically in the realm of binary classification, we selected parameter $\gamma = 64/n$ because this parameter worked well in both our theoretical developments and computational experiments. However, there could be data sets with either an exceptionally large or exceptionally small number of misclassified points in which another value for γ might be appropriate. Future research could extend our theory so that γ is selected in a data dependent manner. Looking beyond binary classification, there are countless problems in data science that are solved by optimization, both convex and nonconvex, and, to the best of our knowledge, few if

any have data dependent termination tests. Therefore, much work remains to be done in this arena.

References

- [1] David Applegate, Oliver Hinder, Haihao Lu, and Miles Lubin. Faster first-order primal-dual methods for linear programming using restarts and sharpness. *Mathematical Programming*, 201(1):133–184, Sep 2023.
- [2] J B Baillon, R E Bruck, and S Reich. On the Asymptotic Behavior of Nonexpansive Mappings and Semigroups in Banach Spaces. *Houston Journal of Mathematics*, 4(1):1–9, 1978.
- [3] Amir Beck. *First-Order Methods in Optimization*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2017.
- [4] Amir Beck and Marc Teboulle. A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009.
- [5] Radu Ioan Bot, Jalal Fadili, and Dang-Khoa Nguyen. The Iterates of Nesterov’s Accelerated Algorithm Converge in The Critical Regimes, 2025.
- [6] Ronald E Bruck and Simeon Reich. Nonexpansive Projections and Resolvents of Accretive Operators in Banach Spaces. *Houston Journal of Mathematics*, 3(4):459–470, 1977.
- [7] Peter Brucker. An $O(n)$ algorithm for quadratic knapsack problems. *Operations Research Letters*, 3(3):163–166, 1984.
- [8] Antonin Chambolle and Thomas Pock. A First-Order Primal-Dual Algorithm for Convex Problems with Applications to Imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, May 2011.
- [9] Chih-Chung Chang and Chih-Jen Lin. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27, 2011. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- [10] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine Learning*, 20(3):273–297, Sep 1995.
- [11] Felipe Cucker and Javier Peña. A primal-dual algorithm for solving polyhedral conic systems with a finite-precision machine. *SIAM Journal on Optimization*, 12(2):522–554, 2002.
- [12] Daniel Dadush, Sophie Huiberts, Bento Natura, and László A. Végh. A scaling-invariant algorithm for linear programming whose running time depends only on the constraint matrix, 2023.

- [13] Rong-En Fan, Kai-Wei Chang, Cho-Jui Hsieh, Xiang-Rui Wang, and Chih-Jen Lin. Liblinear: A library for large linear classification. *J. Mach. Learn. Res.*, 9:1871–1874, June 2008.
- [14] Rong-En Fan, Pai-Hsuen Chen, and Chih-Jen Lin. Working set selection using second order information for training support vector machines. *Journal of Machine Learning Research*, 6(63):1889–1918, 2005.
- [15] Robert M. Freund and Jorge R. Vera. Condition-based complexity of convex optimization in conic linear form via the ellipsoid algorithm. *SIAM Journal on Optimization*, 10(1):155–176, 1999.
- [16] Tao Jiang, Walaa M. Moursi, and Stephen A. Vavasis. Range of the displacement operator of PDHG with applications to quadratic and conic programming, 2023.
- [17] T. Joachims. Making large-Scale SVM Learning Practical. LS8-Report 24, Universität Dortmund, LS VIII-Report, 1998.
- [18] Markelle Kelly, Rachel Longjohn, and Kolby Nottingham. The UCI Machine Learning Repository. <https://archive.ics.uci.edu>, 2025.
- [19] Yann LeCun, Corinna Cortes, and Christopher J. C. Burges. MNIST handwritten digit database. <https://web.archive.org/web/20220121031128/https://yann.lecun.org/exdb/mnist/>, 1998. Original: <https://yann.lecun.org/exdb/mnist/> (archived snapshot). Accessed: 2025-07-08.
- [20] Haihao Lu and Jinwen Yang. On the geometry and refined rate of primal-dual hybrid gradient for linear programming, 2023. arXiv:2307.03664 [math.OC].
- [21] Haihao Lu and Jinwen Yang. On the infimal sub-differential size of primal-dual hybrid gradient method and beyond, 2023. arXiv:2206.12061 [math.OC].
- [22] E. Osuna, R. Freund, and F. Girosi. An improved training algorithm for support vector machines. In *Neural Networks for Signal Processing VII. Proceedings of the 1997 IEEE Signal Processing Society Workshop*, pages 276–285, 1997.
- [23] John C. Platt. A Fast Algorithm for Training Support Vector Machines. Technical report, Microsoft Research, 1998.
- [24] Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, Cambridge, UK, 2014.
- [25] Shai Shalev-Shwartz, Yoram Singer, and Nathan Srebro. Pegasos: Primal Estimated sub-GrAdient Solver for SVM. In *Proceedings of the 24th International Conference on Machine Learning, ICML '07*, page 807–814, New York, NY, USA, 2007. Association for Computing Machinery.

- [26] Pannagadatta K. Shivaswamy, Chiranjib Bhattacharyya, and Alexander J. Smola. Second order cone programming approaches for handling missing and uncertain data. *Journal of Machine Learning Research*, 7(47):1283–1314, 2006.
- [27] Zikai Xiong and Robert M. Freund. The Role of Level-Set Geometry on the Performance of PDHG for Conic Linear Optimization, 2024.

A Detailed results and proofs of Section 2

Lemma 2.1. *Consider the ellipsoid $E := \{z : \|A^{-1}(z - c)\| \leq 1\}$, where $A \in \mathbb{R}^{d \times d}$, $c \in \mathbb{R}^d$, and $s \in \mathbb{R}$. Let $w \in \mathbb{R}^d \setminus \{0\}$. Then E lies in a halfspace $H := \{z : w^T z \leq -s\}$ if and only if $-s \geq \|A^T w\| + c^T w$. Further, $E \subseteq \text{int}(H)$ if and only if this inequality is strict.*

Proof. Suppose $\|A^T w\| + w^T c \leq -s$. For any $z \in E$, we get

$$\begin{aligned} w^T z &= w^T A A^{-1}(z - c) + w^T c \\ &= \|A^T w\| \|A^{-1}(z - c)\| + w^T c \\ &\leq \|A^T w\| + w^T c \\ &\leq -s, \end{aligned}$$

so $z \in H$. Now suppose $E \subseteq H$ and let $z := A A^T w / \|A^T w\| + c$. Then $A^{-1}(z - c) = A^T w / \|A^T w\|$ and clearly $z \in E \subseteq H$. It follows that

$$\begin{aligned} -s &\geq w^T z \\ &= w^T A A^{-1}(z - c) + w^T c \\ &= w^T A A^{-1}(A A^T w / \|A^T w\|) + w^T c \\ &= \|A^T w\| + w^T c. \end{aligned}$$

The proof that $E \subseteq \text{int}(H)$ if and only if the inequality is strict proceeds in the same way. $\square$

A.1 Related results from primal-dual hybrid gradient

PDHG applies to the saddle-point problem

$$\inf_{x \in \mathbb{R}^n} \sup_{y \in \mathbb{R}^m} \phi(x) - \psi^*(y) + y^T \mathcal{A}x,$$

where ψ^* denotes the Fenchel conjugate of ψ . Specializing to the standard conic form problem $\min\{c^T x : \mathcal{A}x = b, x \in \mathcal{C}\}$, we take $\phi(x) := c^T x + i_{\mathcal{C}}(x)$ and $\psi^*(y) := b^T y$, then the PDHG iteration for standard conic form is

$$\begin{pmatrix} x^+ \\ y^+ \end{pmatrix} = T \begin{pmatrix} x \\ y \end{pmatrix} := \begin{pmatrix} \text{proj}_{\mathcal{C}}(x - \sigma \mathcal{A}^T y - c) \\ y + \tau \mathcal{A}(2x^+ - x) - \tau b \end{pmatrix}, \quad (29)$$

where $\sigma, \tau > 0$ are parameters of PDHG. We make the assumption that $\sigma < 1/\|\mathcal{A}\|^2$, $\tau = 1$. If PDHG converges to (x, y) , then $(x, y) = T(x, y)$, i.e. $(x, y) \in \text{zer}(\text{Id} - T)$. If the problem is inconsistent, then $0 \notin \text{Range}(\text{Id} - T)$ and we define the *infimal displacement vector*

$$\begin{pmatrix} v_R \\ v_D \end{pmatrix} = \text{argmin} \{\|v\|_M : v \in \text{cl}(\text{Range}(\text{Id} - T))\},$$

where

$$M = \begin{pmatrix} \frac{1}{\sigma}I & -\mathcal{A}^T \\ -\mathcal{A} & \frac{1}{\tau}I \end{pmatrix} = \begin{pmatrix} \frac{1}{\sigma}I & -\mathcal{A}^T \\ -\mathcal{A} & I \end{pmatrix}.$$

Note that $\text{cl}(\text{Range}(\text{Id} - T))$ is a closed nonempty convex set, so the argmin above has a unique solution. It is a well-known fact (see [6] or [2]) that if the dual is infeasible, then the iterates $(\mathbf{x}_k, \mathbf{y}_k)$ do not converge. Instead, $(\mathbf{x}_{k+1}, \mathbf{y}_{k+1}) - (\mathbf{x}_k, \mathbf{y}_k) \rightarrow (\mathbf{v}_R, \mathbf{v}_D)$. We now recount some useful results from [16].

Lemma A.1. $s^* + t^* = \|\mathbf{v}_D\|^2$.

Proof. From [16, Lemma 6.8 (viii)], $\|\mathbf{v}_D\|^2 = \mathbf{b}^T \mathbf{v}_D$, and by definition $\mathbf{b} = (1, 1, \mathbf{0})$, so

$$s^* + t^* = \mathbf{b}^T \mathbf{v}_D = \|\mathbf{v}_D\|^2.$$

□

Theorem A.1 ([16]). *For the standard conic form problem $\min\{\mathbf{c}^T \mathbf{x} : \mathcal{A}\mathbf{x} = \mathbf{b}, \mathbf{x} \in \mathcal{C}\}$, we have the following:*

- (i) $\mathcal{A}^T \mathbf{v}_D \in \mathcal{C}^\circ$.
- (ii) If $\text{Null}(\mathcal{A}) \cap \mathcal{C} = \{\mathbf{0}\}$, then $\mathbf{v}_R = \mathbf{0}$.
- (iii) If $\text{Null}(\mathcal{A}) \cap \mathcal{C} = \{\mathbf{0}\}$, then $\mathbf{v}_D^T \mathcal{A}\mathbf{x} = \mathbf{0}$ for all $\mathbf{x}$ such that $\mathbf{b} - \mathcal{A}\mathbf{x} = \mathbf{v}_D$.
- (iv) If $\mathbf{c} = \mathbf{0}$ and $\text{Null}(\mathcal{A}) \cap \mathcal{C} = \{\mathbf{0}\}$, then $\mathbf{v} \in \text{Range}(\text{Id} - T)$, where $\mathbf{v}$ is the infimal displacement vector.

In the case of the ESP dual problem (3), both $\mathbf{c} = \mathbf{0}$ and $\text{Null}(\mathcal{A}) \cap \mathcal{C} = \{\mathbf{0}\}$, so all of the items in Theorem A.1 apply. We connected the results from PDHG to the FISTA formulation in Theorem 2.1, so we may now prove a theorem that bounds the distance between a residual vector $\mathbf{v}$ and $\mathbf{v}_D$ by the distance to the optimal function value. This theorem implies that an approximation to the infimal displacement vector converges at the rate $\mathcal{O}(1/k)$ under FISTA, since it is known that the function values of FISTA converge at the rate $\mathcal{O}(1/k^2)$ [3, Theorem 10.34].

Lemma A.2 ([16, Lemma 6.8 (iv)]). *Let $\mathbf{z} \in \mathcal{A}(\mathcal{C})$ and suppose $\text{Null}(\mathcal{A}) \cap \mathcal{C} = \{\mathbf{0}\}$. Then*

$$\mathbf{v}_D^T (\mathbf{v}_D - (\mathbf{b} - \mathbf{z})) \leq 0.$$

Lemma A.3. *For arbitrary vectors $\mathbf{a}, \mathbf{b}, \mathbf{c}$ of the same length, we have the following identity*

$$\|\mathbf{b} - \mathbf{a}\|^2 + 2(\mathbf{a} - \mathbf{c})^T (\mathbf{b} - \mathbf{a}) = \|\mathbf{b} - \mathbf{c}\|^2 - \|\mathbf{a} - \mathbf{c}\|^2$$

Theorem A.2. *Fix $\bar{\mathbf{x}}$ such that $\mathbf{b} - \mathcal{A}\bar{\mathbf{x}} = \mathbf{v}_D$. If $\mathbf{x} \in \mathcal{C}$ and $\mathbf{v} := \mathbf{b} - \mathcal{A}\mathbf{x}$, then*

$$\|\mathbf{v} - \mathbf{v}_D\|^2 \leq 2(f(\mathbf{x}) - f(\bar{\mathbf{x}})),$$

where $f(\mathbf{x}) := \frac{1}{2}\|\mathcal{A}\mathbf{x} - \mathbf{b}\|^2$.

Proof. First, $\bar{\mathbf{x}}$ exists by Theorem 2.1, so the statement is well-defined. By Lemma A.3, we may write

$$\|\mathcal{A}\mathbf{x} - \mathbf{b}\|^2 - \|\mathcal{A}\bar{\mathbf{x}} - \mathbf{b}\|^2 = \|\mathcal{A}\mathbf{x} - \mathcal{A}\bar{\mathbf{x}}\|^2 + 2(\mathcal{A}\bar{\mathbf{x}} - \mathbf{b})^T(\mathcal{A}\mathbf{x} - \mathcal{A}\bar{\mathbf{x}}).$$

Using that $\mathbf{b} - \mathcal{A}\bar{\mathbf{x}} = \mathbf{v}_D$, we may rewrite the above as

$$\|\mathcal{A}\mathbf{x} - \mathbf{b}\|^2 - \|\mathcal{A}\bar{\mathbf{x}} - \mathbf{b}\|^2 = \|\mathcal{A}\mathbf{x} - \mathcal{A}\bar{\mathbf{x}}\|^2 - 2\mathbf{v}_D^T(\mathbf{v}_D - (\mathbf{b} - \mathcal{A}\mathbf{x})).$$

We may now use Lemma A.2 to obtain the desired inequality:

$$\begin{aligned} \|\mathcal{A}\mathbf{x} - \mathbf{b}\|^2 - \|\mathcal{A}\bar{\mathbf{x}} - \mathbf{b}\|^2 &= \|\mathcal{A}\mathbf{x} - \mathcal{A}\bar{\mathbf{x}}\|^2 - 2\mathbf{v}_D^T(\mathbf{v}_D - (\mathbf{b} - \mathcal{A}\mathbf{x})) \\ &\geq \|\mathcal{A}\mathbf{x} - \mathcal{A}\bar{\mathbf{x}}\|^2 \\ &= \|\mathbf{v} - \mathbf{v}_D\|^2. \end{aligned}$$

□

A.2 Obtaining a data-dependent bound on the number of FISTA iterations

Our strategy is to perturb the data points $\mathbf{c}_1, \dots, \mathbf{c}_j$ to $\hat{\mathbf{c}}_1, \dots, \hat{\mathbf{c}}_j$, yielding a perturbation of $\mathcal{A}$ to $\hat{\mathcal{A}}$ such that there exists $\hat{\mathbf{x}} \in \mathcal{C}$ satisfying $\hat{\mathcal{A}}\hat{\mathbf{x}} = \mathbf{b}$.

Let $\mathbf{x}^* \in \mathcal{C}$ such that $\mathbf{b} - \mathcal{A}\mathbf{x}^* = \mathbf{v}_D$, which exists by Theorem 2.1 (i). We will say that

$$\mathbf{x}^* := \begin{pmatrix} \lambda_1^* \\ \mathbf{p}_1^* \\ \vdots \\ \lambda_j^* \\ \mathbf{p}_j^* \\ \mu_1^* \\ \mathbf{q}_1^* \\ \vdots \\ \mu_l^* \\ \mathbf{q}_l^* \end{pmatrix} \quad \text{and} \quad \hat{\mathbf{x}} := \begin{pmatrix} \hat{\lambda}_1 \\ \hat{\mathbf{p}}_1 \\ \vdots \\ \hat{\lambda}_j \\ \hat{\mathbf{p}}_j \\ \hat{\mu}_1 \\ \hat{\mathbf{q}}_1 \\ \vdots \\ \hat{\mu}_l \\ \hat{\mathbf{q}}_l \end{pmatrix},$$

where $\boldsymbol{\lambda}^* = (\lambda_1^*, \dots, \lambda_j^*)$, $\hat{\boldsymbol{\lambda}} = (\hat{\lambda}_1, \dots, \hat{\lambda}_j)$, $\boldsymbol{\mu}^* = (\mu_1^*, \dots, \mu_j^*)$, $\hat{\boldsymbol{\mu}} = (\hat{\mu}_1, \dots, \hat{\mu}_j)$, and define

$$\begin{aligned} \hat{\lambda} &:= \text{proj}_{\Delta^j}(\boldsymbol{\lambda}^*) \\ \hat{\mu} &:= \text{proj}_{\Delta^l}(\boldsymbol{\mu}^*) \\ \hat{\mathbf{p}}_i &:= \begin{cases} \frac{\hat{\lambda}_i}{\lambda_i^*} \mathbf{p}_i^*, & \hat{\lambda}_i < \lambda_i^*; \\ \mathbf{p}_i^*, & \text{otherwise} \end{cases} \\ \hat{\mathbf{q}}_i &:= \begin{cases} \frac{\hat{\mu}_i}{\mu_i^*} \mathbf{q}_i^*, & \hat{\mu}_i < \mu_i^* \\ \mathbf{q}_i^*, & \text{otherwise.} \end{cases} \end{aligned}$$

Lemma A.4. $\|\boldsymbol{\lambda}^* - \hat{\boldsymbol{\lambda}}\|_1 \leq |1 - \mathbf{e}^T \boldsymbol{\lambda}^*| = |s^*|$ and $\|\boldsymbol{\mu}^* - \hat{\boldsymbol{\mu}}\|_1 \leq |1 - \mathbf{e}^T \boldsymbol{\mu}^*| = |t^*|$.

Proof. We will prove $\|\boldsymbol{\lambda}^* - \hat{\boldsymbol{\lambda}}\|_1 \leq |1 - \mathbf{e}^T \boldsymbol{\lambda}^*| = |s^*|$ only, since the proof of the second expression is analogous.

The equality $s^* = 1 - \mathbf{e}^T \boldsymbol{\lambda}^*$ comes from the explicit definition of the constraint $\mathcal{A}\mathbf{x}^* = \mathbf{b}$ in (4) and the assumption that $(s^*, t^*, \mathbf{w}^*) = \mathbf{v}_D = \mathbf{b} - \mathcal{A}\mathbf{x}^*$.

If $\boldsymbol{\lambda}^* = \mathbf{0}$, since $\hat{\boldsymbol{\lambda}} \in \Delta^j$,

$$\|\boldsymbol{\lambda}^* - \hat{\boldsymbol{\lambda}}\|_1 = \|\hat{\boldsymbol{\lambda}}\|_1 = 1 = |1 - \mathbf{e}^T \boldsymbol{\lambda}^*|.$$

Suppose $\boldsymbol{\lambda}^* \geq \mathbf{0}$ and let $r := \|\boldsymbol{\lambda}^*\|_1 = \sum_{i=1}^j \lambda_i^* = \mathbf{e}^T \boldsymbol{\lambda}^*$. Defining $\bar{\boldsymbol{\lambda}} := \frac{1}{r} \boldsymbol{\lambda}^*$, we have that $\bar{\boldsymbol{\lambda}} \in \Delta^j$, so

$$\begin{aligned} \|\boldsymbol{\lambda}^* - \hat{\boldsymbol{\lambda}}\|_1 &\leq \|\boldsymbol{\lambda}^* - \bar{\boldsymbol{\lambda}}\|_1 \\ &= |1 - q| \frac{\|\boldsymbol{\lambda}^*\|}{q} \\ &= |1 - \mathbf{e}^T \boldsymbol{\lambda}^*|. \end{aligned}$$

□

Lemma A.5. $\|\hat{\mathbf{p}}_i - \mathbf{p}_i^*\| \leq |\hat{\lambda}_i - \lambda_i^*|$ and $\|\hat{\mathbf{q}}_i - \mathbf{q}_i^*\| \leq |\hat{\mu}_i - \mu_i^*|$.

Proof. If $\hat{\lambda}_i \geq \lambda_i^*$, then the inequality is clear. If $\hat{\lambda}_i < \lambda_i^*$, then $\hat{\mathbf{p}}_i = \hat{\lambda}_i \mathbf{p}_i^* / \lambda_i^*$, so $\|\hat{\mathbf{p}}_i - \mathbf{p}_i^*\| = \|(\hat{\lambda}_i - \lambda_i^*) \mathbf{p}_i^* / \lambda_i^*\| \leq |\hat{\lambda}_i - \lambda_i^*|$, since $\|\mathbf{p}_i^*\| \leq \lambda_i^*$. The proof of $\|\hat{\mathbf{q}}_i - \mathbf{q}_i^*\| \leq |\hat{\mu}_i - \mu_i^*|$ is analogous. □

Theorem 2.3. Define $\delta_{\mathcal{A}}$ to be the smallest perturbation to the points $\mathbf{c}_i$ and $\mathbf{d}_i$ such that their ellipsoids are no longer separable. Then

$$\delta_{\mathcal{A}} \leq 3\|\mathbf{v}_D\|.$$

Proof. Since $\mathbf{b} - \mathcal{A}\mathbf{x}^* = \mathbf{v}_D = (s^*, t^*, \mathbf{w}^*)$, we have that $\mathbf{b} - \mathcal{A}\hat{\mathbf{x}} = (0, 0, \mathbf{r})$, where

$$\mathbf{r} := \mathbf{w}^* + \sum_{i=1}^j ((\lambda_i^* - \hat{\lambda}_i) \mathbf{c}_i + A_i(\mathbf{p}_i^* - \hat{\mathbf{p}}_i)) + \sum_{i=1}^l (-(\mu_i^* - \hat{\mu}_i) \mathbf{d}_i + B_i(\mathbf{q}_i^* - \hat{\mathbf{q}}_i)).$$

From Lemma A.4,

$$\sum_{i=1}^j |\lambda_i^* - \hat{\lambda}_i| \leq |s^*|, \quad \sum_{i=1}^l |\mu_i^* - \hat{\mu}_i| \leq |t^*|,$$

so from Lemma A.5, we have that

$$\sum_{i=1}^j \|\mathbf{p}_i^* - \hat{\mathbf{p}}_i\| \leq |s^*|, \quad \sum_{i=1}^l \|\mathbf{q}_i^* - \hat{\mathbf{q}}_i\| \leq |t^*|.$$

Also by the normalization assumption (1), $\|\mathbf{c}_i\| \leq 1$, $\|A_i\| \leq 1$, $\|\mathbf{d}_i\| \leq 1$, $\|B_i\| \leq 1$. These inequalities allow us to conclude that

$$\|\mathbf{r}\| \leq \|\mathbf{w}^*\| + 2|s^*| + 2|t^*| \leq 3\|\mathbf{v}_D\|.$$

We have established that

$$\sum_{i=1}^j (\hat{\lambda}_i \mathbf{c}_i + A_i \hat{\mathbf{p}}_i) + \sum_{i=1}^l (-\hat{\mu}_i \mathbf{d}_i + B_i \hat{\mathbf{q}}_i) = \mathbf{r},$$

which is equivalent to

$$\sum_{i=1}^j (\hat{\lambda}_i (\mathbf{c}_i - \mathbf{r}) + A_i \hat{\mathbf{p}}_i) + \sum_{i=1}^l (-\hat{\mu}_i \mathbf{d}_i + B_i \hat{\mathbf{q}}_i) = \mathbf{0}.$$

So we may define $\hat{\mathcal{A}}$ with all the data the same as $\mathcal{A}$ except with centers at $\mathbf{c}_i - \mathbf{r}$ instead of $\mathbf{c}_i$, for each i . After this perturbation, the ESP is no longer separable. It follows that the minimum perturbation is at most $3\|\mathbf{v}_D\|$. $\square$

B Details of Section 3

B.1 An example of a dual formulation of the soft-margin SVM problem with multiple optimizers

Consider the one-dimensional case of (12), where $c_1 = c_2 = c_3 = 1, c_4 = -1$ and $d_1 = d_2 = -1, d_3 = 1$, and $\gamma \geq 1$. The well-classified points are clearly c_1, c_2, c_3 and d_1, d_2 . The KKT conditions for this problem are as follows:

$$\begin{aligned} x &= C^T \mathbf{u} - D^T \mathbf{v} \\ \mathbf{e}^T \mathbf{u} &= \mathbf{e}^T \mathbf{v} \\ \mathbf{z}_1 &= \gamma \mathbf{e} - \mathbf{u} \\ \mathbf{z}_2 &= \gamma \mathbf{e} - \mathbf{v} \\ c_i x + x_i &\geq 1 - s_{1,i}, \quad \forall i \in [j] \\ d_i x + x_i &\leq -1 + s_{2,i}, \quad \forall i \in [l] \\ u_i(1 - s_{1,i} - c_i x - \xi) &= 0, \quad \forall i \in [j] \\ v_i(1 - s_{2,i} + d_i x + \xi) &= 0, \quad \forall i \in [l] \\ \mathbf{z}^T \mathbf{s} &= 0 \\ \mathbf{0} &\leq \mathbf{u} \leq \gamma \mathbf{e} \\ \mathbf{0} &\leq \mathbf{v} \leq \gamma \mathbf{e} \\ \mathbf{s}, \mathbf{z} &\geq \mathbf{0} \end{aligned} \tag{30}$$

One KKT solution of this problem is given by $x = 1$, $\xi = 0$, $\mathbf{z}_1 = \gamma \mathbf{e} - \mathbf{u}$, $\mathbf{z}_2 = \gamma \mathbf{e} - \mathbf{v}$,

$$\mathbf{u} = \begin{pmatrix} \gamma/3 + 1/6 \\ \gamma/3 + 1/6 \\ \gamma/3 + 1/6 \\ \gamma \end{pmatrix}, \quad \mathbf{v} = \begin{pmatrix} \gamma/2 + 1/4 \\ \gamma/2 + 1/4 \\ \gamma/2 + 1/4 \\ \gamma \end{pmatrix}, \quad \mathbf{s}_1 = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 2 \end{pmatrix}, \quad \mathbf{s}_2 = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 2 \end{pmatrix}.$$

Indeed, this optimal solution is desirable, since for the properly classified points, $u_i < \gamma$ and $v_i < \gamma$. However, we can construct another KKT solution from the above by adding $2\gamma/3 - 1/6$ to u_1 and subtracting $\gamma/3 - 1/12$ from both u_2 and u_3 :

$$\mathbf{u} = \begin{pmatrix} \gamma \\ 1/4 \\ 1/4 \\ \gamma \end{pmatrix}, \quad \mathbf{v} = \begin{pmatrix} \gamma/2 + 1/4 \\ \gamma/2 + 1/4 \\ \gamma/2 + 1/4 \\ \gamma \end{pmatrix}.$$

We still set $\mathbf{z}_1 = \gamma \mathbf{e} - \mathbf{u}$ and $\mathbf{z}_2 = \gamma \mathbf{e} - \mathbf{v}$, and the values of all other variables remain the same. This new optimal solution is undesirable in that the well-separated point c_1 has $u_1 = \gamma$, preventing it from being identified by simply verifying the value of u_4 is strictly less than γ . This example can be easily extended to having many more points, in which case even more well-classified points can have corresponding optimal dual equal to γ .

B.2 Solving the perturbed dual

Lemma B.1. *Let $\mathbf{r} := (\mathbf{u}, \mathbf{v}) \in \mathbb{R}^j \times \mathbb{R}^l$ and let*

$$\mathbf{r}^+ := \text{proj}_\Omega \left(\mathbf{r} + \frac{1}{L} \nabla g(\mathbf{r}) \right).$$

Then,

$$\text{dist} \left(\nabla g(\mathbf{r}^+), \mathcal{N}_\Omega(\mathbf{r}^+) \right) \leq 2L \|\mathbf{r} - \mathbf{r}^+\|.$$

Proof. The problem (15) can be rewritten as the unconstrained maximum of the function G , where $G := g(\mathbf{u}, \mathbf{v}) - i_\Omega(\mathbf{u}, \mathbf{v})$, which is equivalent to the convex problem $\min\{-G(\mathbf{r}) : \mathbf{r} \in \Omega\}$. The projection of $\mathbf{r} + \frac{1}{L} \nabla g(\mathbf{r})$ is the minimizer of

$$h(\mathbf{z}) = -g(\mathbf{r}) - \nabla g(\mathbf{r})^T (\mathbf{z} - \mathbf{r}) + \frac{L}{2} \|\mathbf{z} - \mathbf{r}\|^2 + i_\Omega(\mathbf{z}),$$

hence $\mathbf{z} = \mathbf{r}^+$ must satisfy the first-order optimality condition $\mathbf{0} \in \partial h(\mathbf{r}^+)$, or equivalently

$$\mathbf{0} \in -\nabla g(\mathbf{r}) + L(\mathbf{r}^+ - \mathbf{r}) + \mathcal{N}_\Omega(\mathbf{r}^+).$$

Hence, the optimality conditions are $L(\mathbf{r} - \mathbf{r}^+) + \nabla g(\mathbf{r}) \in \mathcal{N}_\Omega(\mathbf{r}^+)$. If we add $-\nabla g(\mathbf{r}^+)$ to both sides we obtain

$$L(\mathbf{r} - \mathbf{r}^+) + \nabla g(\mathbf{r}) - \nabla g(\mathbf{r}^+) \in -\nabla g(\mathbf{r}^+) + \mathcal{N}_\Omega(\mathbf{r}^+) = \partial(-G(\mathbf{r}^+)),$$

where the final equality uses that the subdifferential sum rule holds since g is proper and convex with full domain, and Ω has nonempty interior. Since $L(\mathbf{r} - \mathbf{r}^+) + \nabla g(\mathbf{r}) - \nabla g(\mathbf{r}^+)$ is contained in $\partial(-G(\mathbf{r}^+))$, it gives an upper bound on $\text{dist}(\mathbf{0}, \partial(-G(\mathbf{r}^+)))$:

$$\begin{aligned} \text{dist}(\mathbf{0}, \partial(-G(\mathbf{r}^+))) &\leq \|L(\mathbf{r} - \mathbf{r}^+) + \nabla g(\mathbf{r}) - \nabla g(\mathbf{r}^+)\| \\ &\leq L \|\mathbf{r} - \mathbf{r}^+\| + \|\nabla g(\mathbf{r}) - \nabla g(\mathbf{r}^+)\| \\ &\leq 2L \|\mathbf{r} - \mathbf{r}^+\|, \end{aligned}$$

where we used the triangle inequality and L -smoothness of g . $\square$

B.3 Obtaining a hyperplane that separates the well-classified points

Lemma B.2. $\psi(\theta)$ is Lipschitz continuous with constant γ .

Proof. The derivative of $\psi(\theta)$ defined in (21) is continuous, and on each interval we have the upper bound $|\psi'(\theta)| \leq \gamma$. $\square$

Theorem B.1. Let $\bar{K}$ and K be given according to (25). Suppose $\|(\mathbf{u}_k, \mathbf{v}_k) - (\mathbf{u}^*, \mathbf{v}^*)\| \leq \Delta$, $\mathbf{w}_k := C^T \mathbf{u}_k - D^T \mathbf{v}_k$, and t_k is the solution to (27). Then $(\mathbf{w}_k, t_k)$ encode a separating hyperplane for the well-classified points, in particular

$$\begin{aligned} \mathbf{c}_i^T \mathbf{w}_k + t_k &\geq 1 - \bar{K}, \quad \forall i \in J, \\ \mathbf{d}_i^T \mathbf{w}_k + t_k &\leq -1 + \bar{K}, \quad \forall i \in L. \end{aligned}$$

Proof. Fix an $i_c \in J$. Then

$$\begin{aligned} \sum_{i=1}^j \psi(\mathbf{c}_i^T \mathbf{w}_k + t_k) &= j_0 \psi(\mathbf{c}_{i_c}^T \mathbf{w}_k + t_k) \\ &\quad + \sum_{i \in J \setminus \{i_c\}} (\psi(\mathbf{c}_i^T \mathbf{w}_k + t_k) - \psi(\mathbf{c}_{i_c}^T \mathbf{w}_k + t_k)) \\ &\quad + \sum_{i \notin J} \psi(\mathbf{c}_i^T \mathbf{w}_k + t_k). \end{aligned} \tag{31}$$

Let us focus on bounding $j_0 \psi(\mathbf{c}_{i_c}^T \mathbf{w}_k + t_k)$. Suppose $(\mathbf{w}^*, t^*)$ are optimal solutions to (20). By definition of t_k and since minimizing the objective in (27) is equivalent to minimizing

$\sum_{i=1}^j \psi(\mathbf{c}_i^T \mathbf{w}_k + t) + \sum_{i=1}^l \psi(-\mathbf{d}_i^T \mathbf{w}_k - t)$, we may write

$$\begin{aligned}
\sum_{i=1}^j \psi(\mathbf{c}_i^T \mathbf{w}_k + t_k) + \sum_{i=1}^l \psi(-\mathbf{d}_i^T \mathbf{w}_k - t_k) &\leq \sum_{i=1}^j \psi(\mathbf{c}_i^T \mathbf{w}_k + t^*) + \sum_{i=1}^l \psi(-\mathbf{d}_i^T \mathbf{w}_k - t^*) \\
&= \sum_{i=1}^j \psi(\mathbf{c}_i^T \mathbf{w}^* + t^*) + \sum_{i=1}^l \psi(-\mathbf{d}_i^T \mathbf{w}^* - t^*) \\
&\quad + \sum_{i=1}^j (\psi(\mathbf{c}_i^T \mathbf{w}_k + t^*) - \psi(\mathbf{c}_i^T \mathbf{w}^* + t^*)) \\
&\quad + \sum_{i=1}^l (\psi(-\mathbf{d}_i^T \mathbf{w}_k - t^*) - \psi(-\mathbf{d}_i^T \mathbf{w}^* - t^*)).
\end{aligned}$$

Since ψ is always nonnegative,

$$\sum_{i=1}^j \psi(\mathbf{c}_i^T \mathbf{w}_k + t_k) \leq \sum_{i=1}^j \psi(\mathbf{c}_i^T \mathbf{w}_k + t_k) + \sum_{i=1}^l \psi(-\mathbf{d}_i^T \mathbf{w}_k - t_k),$$

so using (31), we have an upper bound on $j_0 \psi(\mathbf{c}_{i_c}^T \mathbf{w}_k + t_k)$, namely

$$\begin{aligned}
j_0 \psi(\mathbf{c}_{i_c}^T \mathbf{w}_k + t_k) &\leq \sum_{i=1}^j \psi(\mathbf{c}_i^T \mathbf{w}^* + t^*) + \sum_{i=1}^l \psi(-\mathbf{d}_i^T \mathbf{w}^* - t^*) \\
&\quad + \sum_{i=1}^j (\psi(\mathbf{c}_i^T \mathbf{w}_k + t^*) - \psi(\mathbf{c}_i^T \mathbf{w}^* + t^*)) \\
&\quad + \sum_{i=1}^l (\psi(-\mathbf{d}_i^T \mathbf{w}_k - t^*) - \psi(-\mathbf{d}_i^T \mathbf{w}^* - t^*)) \\
&\quad - \sum_{i \in J \setminus \{i_c\}} (\psi(\mathbf{c}_i^T \mathbf{w}_k + t_k) - \psi(\mathbf{c}_{i_c}^T \mathbf{w}_k + t_k)) \\
&\quad - \sum_{i \notin J} \psi(\mathbf{c}_i^T \mathbf{w}_k + t_k).
\end{aligned}$$

Using Lemma B.2, we can upper bound the following

$$\begin{aligned}
\sum_{i=1}^j (\psi(\mathbf{c}_i^T \mathbf{w}_k + t^*) - \psi(\mathbf{c}_i^T \mathbf{w}^* + t^*)) &\leq \sum_{i=1}^j \gamma |\mathbf{c}_i^T \mathbf{w}_k - \mathbf{c}_i^T \mathbf{w}^*| \leq j \gamma R \|\mathbf{w}_k - \mathbf{w}^*\|, \\
\sum_{i=1}^l (\psi(-\mathbf{d}_i^T \mathbf{w}_k - t^*) - \psi(-\mathbf{d}_i^T \mathbf{w}^* - t^*)) &\leq \sum_{i=1}^l \gamma |\mathbf{d}_i^T \mathbf{w}_k - \mathbf{d}_i^T \mathbf{w}^*| \leq l \gamma R \|\mathbf{w}_k - \mathbf{w}^*\|.
\end{aligned}$$

To bound $\|\mathbf{w}_k - \mathbf{w}^*\|$, we use that $\|(\mathbf{u}_k, \mathbf{v}_k) - (\mathbf{u}^*, \mathbf{v}^*)\| \leq \Delta$ along with the KKT conditions

(17), which tell us that $\mathbf{w}^* = C^T \mathbf{u}^* - D^T \mathbf{v}^*$:

$$\begin{aligned} \|\mathbf{w}_k - \mathbf{w}^*\| &= \|(C^T \mathbf{u}_k - C^T \mathbf{u}^*) + (D^T \mathbf{v}^* - D^T \mathbf{v}_k)\| \\ &\leq \|C^T\|_2 \|\mathbf{u}_k - \mathbf{u}^*\| + \|D^T\|_2 \|\mathbf{v}_k - \mathbf{v}^*\| \\ &\leq \sqrt{j}R \|\mathbf{u}_k - \mathbf{u}^*\| + \sqrt{l}R \|\mathbf{v}_k - \mathbf{v}^*\| \\ &\leq (\sqrt{j} + \sqrt{l})R\Delta. \end{aligned}$$

So we get the bound

$$\begin{aligned} \sum_{i=1}^j (\psi(\mathbf{c}_i^T \mathbf{w}_k + t^*) - \psi(\mathbf{c}_i^T \mathbf{w}^* + t^*)) &\leq j (\sqrt{j} + \sqrt{l}) \gamma R^2 \Delta, \\ \sum_{i=1}^l (\psi(-\mathbf{d}_i^T \mathbf{w}_k - t^*) - \psi(-\mathbf{d}_i^T \mathbf{w}^* - t^*)) &\leq l (\sqrt{j} + \sqrt{l}) \gamma R^2 \Delta. \end{aligned} \tag{32}$$

We may again use that ψ is Lipschitz to bound $\psi(\mathbf{c}_{i_c}^T \mathbf{w}_k + t_k) - \psi(\mathbf{c}_i^T \mathbf{w}_k + t_k)$:

$$\begin{aligned} - \sum_{i \in J \setminus \{i_c\}} (\psi(\mathbf{c}_i^T \mathbf{w}_k + t_k) - \psi(\mathbf{c}_{i_c}^T \mathbf{w}_k + t_k)) &= \sum_{i \in J \setminus \{i_c\}} (\psi(\mathbf{c}_{i_c}^T \mathbf{w}_k + t_k) - \psi(\mathbf{c}_i^T \mathbf{w}_k + t_k)) \\ &\leq \sum_{i \in J \setminus \{i_c\}} \gamma |\mathbf{c}_{i_c}^T \mathbf{w}_k - \mathbf{c}_i^T \mathbf{w}_k| \\ &\leq \sum_{i \in J \setminus \{i_c\}} \gamma \|\mathbf{c}_{i_c} - \mathbf{c}_i\| \|\mathbf{w}_k\| \\ &\leq (j_0 - 1) \gamma 2\rho \left((1 + \xi) + (\sqrt{j} + \sqrt{l}) R\Delta \right), \end{aligned} \tag{33}$$

where the last line uses $\|\mathbf{c}_{i_c} - \mathbf{c}_i\| \leq 2\rho$, and $\|\mathbf{w}_k\| \leq \|\mathbf{w}^*\| + \|\mathbf{w}_k - \mathbf{w}^*\| \leq (1 + \xi) + (\sqrt{j} + \sqrt{l})R\Delta$. Since $\mathbf{c}_i^T \mathbf{w}^* + t^* \geq 1 - K$ and $-\mathbf{d}_i^T \mathbf{w}^* - t^* \geq 1 - K$ for all well-classified points, we have that for well-classified points, $\psi(\mathbf{c}_i^T \mathbf{w}^* + t^*)$ and $\psi(-\mathbf{d}_i^T \mathbf{w}^* - t^*)$ are bounded above by $\gamma K^2 = 64K^2/n$. It follows that

$$\sum_{i \in J} \psi(\mathbf{c}_i^T \mathbf{w}^* + t^*) + \sum_{i \in L} \psi(-\mathbf{d}_i^T \mathbf{w}^* - t^*) \leq n_0 \gamma K^2.$$

For points that are not well-classified, we use the bound on t^* from the statement of Corollary 3.1:

$$(1 - K) - \mathbf{c}_{j_c}^T \mathbf{w}^* \leq t^* \leq (-1 + K) - \mathbf{d}_{j_d}^T \mathbf{w}^*,$$

where $j_c := \operatorname{argmin}_{j \in J} \mathbf{c}_j^T \mathbf{w}^*$ and $j_d := \operatorname{argmin}_{j \in L} -\mathbf{d}_j^T \mathbf{w}^*$. Then for all $i \notin J$, we have

$$\mathbf{c}_i^T \mathbf{w}^* + t^* \geq (1 - K) + (\mathbf{c}_i - \mathbf{c}_{j_c})^T \mathbf{w}^* \geq (1 - K) + \|\mathbf{c}_i - \mathbf{c}_{j_c}\| \|\mathbf{w}^*\|.$$

Under the data assumptions, $\|\mathbf{c}_i - \mathbf{c}_{j_c}\| \geq -(\sigma_1 + \sigma_2) - 2\rho$, so

$$\mathbf{c}_i^T \mathbf{w}^* + t^* \geq (1 - K) - (\sigma_1 + \sigma_2 + 2\rho) (1 + \xi).$$

Feeding this into ψ and bounding σ_1, σ_2 by 1, we get that

$$\psi(\mathbf{c}_i^T \mathbf{w}^* + t^*) \leq \left(1 - \frac{1}{4}\right) \gamma + (1 + 2\rho + K + \xi(2 + 2\rho)) \gamma.$$

Similar work can be done to show the same bound for $\psi(-\mathbf{d}_i^T \mathbf{w}^* + t^*)$. In all, we obtain the below bound, where $n_0 = j_0 + l_0$.

$$\sum_{i=1}^j \psi(\mathbf{c}_i^T \mathbf{w}^* + t^*) + \sum_{i=1}^l \psi(-\mathbf{d}_i^T \mathbf{w}^* - t^*) \leq n_0 \gamma K^2 + n_{\text{noise}} \left(2 - \frac{1}{4} + 2\rho + K + \xi(2 + 2\rho)\right) \gamma. \quad (34)$$

Putting together (32), (33), (34), and that $-\psi(\mathbf{c}_i^T \mathbf{w}_k + t_k) \leq 0$, we obtain the bound

$$\begin{aligned} j_0 \psi(\mathbf{c}_{i_c}^T \mathbf{w}_k + t_k) &\leq n_0 \gamma K^2 + n_{\text{noise}} \left(\frac{7}{4} + 2\rho + K + \xi(2 + 2\rho)\right) \gamma \\ &\quad + n(\sqrt{j} + \sqrt{l}) \gamma R^2 \Delta \\ &\quad + (j_0 - 1) \gamma 2\rho \left((1 + \xi) + (\sqrt{j} + \sqrt{l}) R \Delta\right) \\ &\leq n_0 \gamma K^2 + \nu n \left(\frac{7}{4} + 2\rho + K + \xi(2 + 2\rho)\right) \gamma \\ &\quad + n \sqrt{2n} \gamma R^2 \Delta \\ &\quad + j_0 \gamma 2\rho \left((1 + \xi) + \sqrt{2n} R \Delta\right). \end{aligned}$$

Grouping terms with Δ we have

$$\begin{aligned} j_0 \psi(\mathbf{c}_{i_c}^T \mathbf{w}_k + t_k) &\leq n_0 \gamma K^2 + \nu n \left(\frac{7}{4} + 2\rho + K + \xi(2 + 2\rho)\right) \gamma \\ &\quad + j_0 \gamma 2\rho (1 + \xi) + \left(n \sqrt{2n} R^2 + j_0 2\rho \sqrt{2n} R\right) \gamma \Delta. \end{aligned}$$

Since $l_0 \geq n/3$, and $n_{\text{noise}} > 0$, $j_0 = n - l_0 - n_{\text{noise}} \leq 2n/3$, thus

$$\begin{aligned} n \sqrt{2n} R^2 + j_0 2\rho \sqrt{2n} R &\leq n \sqrt{2n} R^2 + \frac{4n}{3} \rho \sqrt{2n} R \\ &= 4\sqrt{2n}^{3/2} \left(1 + \frac{2}{3} \rho\right), \end{aligned}$$

where on the last line we use that $R \leq 2$. The bound on $j_0 \psi(\mathbf{c}_{i_c}^T \mathbf{w}_k + t_k)$ now becomes

$$\begin{aligned} j_0 \psi(\mathbf{c}_{i_c}^T \mathbf{w}_k + t_k) &\leq n_0 \gamma K^2 + \nu n \left(\frac{7}{4} + 2\rho + K + \xi(2 + 2\rho)\right) \gamma \\ &\quad + j_0 \gamma 2\rho (1 + \xi) + 4\sqrt{2n}^{3/2} \left(1 + \frac{2}{3} \rho\right) \gamma \Delta, \end{aligned}$$

and if we substitute Δ from (26), this becomes

$$\begin{aligned} j_0 \psi(\mathbf{c}_{i_c}^T \mathbf{w}_k + t_k) &\leq n_0 \gamma K^2 + \nu n \left(\frac{7}{4} + 2\rho + K + \xi(2 + 2\rho) \right) \gamma + j_0 \gamma 2\rho(1 + \xi) + \frac{n}{3} (\bar{K} - 3K^2) \gamma \\ &\quad - n\rho(1 + \xi) \left(\frac{4}{3} + 2\nu \right) \gamma - \nu n \left(\frac{9}{4} + 2\xi \right) \gamma - \frac{n}{6} \gamma. \end{aligned}$$

We can bound $\frac{n}{3}(\bar{K} - 3K^2)$:

$$\frac{n}{3}(\bar{K} - 3K^2) \leq \frac{n}{3}(\bar{K} - K^2) - \frac{2n}{3}K^2 \leq j_0(\bar{K} - K^2) - l_0 K^2 = j_0 \bar{K} - n_0 K^2,$$

and also

$$-\frac{n}{6} = -\frac{2n/3}{4n} n \leq -\frac{j_0}{4n} n = -\frac{j_0}{4}, \quad \text{and} \quad 2j_0 \rho(1 + \xi) \leq \frac{4n\rho}{3}(1 + \xi).$$

So

$$\begin{aligned} j_0 \psi(\mathbf{c}_{i_c}^T \mathbf{w}_k + t_k) &\leq j_0 \gamma \bar{K} + \nu n \left(\frac{7}{4} + 2\rho + K + \xi(2 + 2\rho) \right) \gamma + \gamma \frac{4n\rho}{3}(1 + \xi) \\ &\quad - n\rho(1 + \xi) \left(\frac{4}{3} + 2\nu \right) \gamma - \nu n \left(\frac{9}{4} + 2\xi \right) \gamma - \frac{j_0}{4} \gamma \\ &\leq j_0 \gamma \bar{K} + \nu n \left(\frac{9}{4} + 2\rho(1 + \xi) + 2\xi \right) \gamma + \gamma \frac{4n\rho}{3}(1 + \xi) \\ &\quad - n\rho(1 + \xi) \left(\frac{4}{3} + 2\nu \right) \gamma - \nu n \left(\frac{9}{4} + 2\xi \right) \gamma - \frac{j_0}{4} \gamma \\ &= j_0 \left(\bar{K} - \frac{1}{4} \right) \gamma, \end{aligned}$$

where in the second inequality we used the assumption that $K < 1/2$. We have shown that $\psi(\mathbf{c}_{i_c}^T \mathbf{w}_k + t_k) \leq (\bar{K} - \frac{1}{4}) \gamma$. Since, i_c was chosen arbitrarily, it follows that for any well-classified $\mathbf{c}_i$

$$\psi(\mathbf{c}_i^T \mathbf{w}_k + t_k) \leq \gamma \left(\bar{K} - \frac{1}{4} \right) = \gamma \left(\frac{3}{4} - (1 - \bar{K}) \right).$$

By our choice of γ and μ , for $\theta \leq \frac{1}{2}$

$$\psi(\theta) = \gamma \left(\frac{3}{4} - \theta \right).$$

Moreover, since $\bar{K} \in [1/2, 1)$, we have $1 - \bar{K} \in (0, 1/2]$, and

$$\psi(1 - \bar{K}) = \gamma \left(\frac{3}{4} - (1 - \bar{K}) \right) \geq \psi(\mathbf{c}_i^T \mathbf{w}_k + t_k).$$

We know ψ is decreasing, so the only way this can happen is if $\mathbf{c}_i^T \mathbf{w}_k + t_k \geq 1 - \bar{K}$. Identical steps show that for any well-classified $\mathbf{d}_i$, we have that $-\mathbf{d}_i^T \mathbf{w}_k - t_k \geq 1 - \bar{K}$. $\square$

C Numerical experiment on unbalanced data for the ESP

Table 3: Comparing the number of iterations and time taken before finding a separating hyperplane in unbalanced data between early-stopping and tolerance-based stopping ($\text{tol} = 10^{-6}$).

0 Ellipsoids	1 Ellipsoids	Columns in $\mathcal{A}$	FISTA (early-stopping)		FISTA (tolerance-based)	
			iterations	time(s)	iterations	time(s)
5	25	23550	5	0.268	1270	15.088
25	5	23550	6	0.285	1166	13.612
10	100	86350	10	0.734	1773	97.589
100	10	86350	11	0.722	1468	79.105
1	125	98910	24	1.476	1516	97.442
125	1	98910	34	1.933	1906	101.261