
Efficient Learning of Minimax Risk Classifiers in High Dimensions

Karthek Bondugula¹

Santiago Mazuelas^{1,2}

Aritz Pérez¹

¹Basque Center for Applied Mathematics (BCAM), Bilbao, Spain

²IKERBASQUE-Basque Foundation for Science

Abstract

High-dimensional data is common in multiple areas, such as health care and genomics, where the number of features can be tens of thousands. In such scenarios, the large number of features often leads to inefficient learning. Constraint generation methods have recently enabled efficient learning of L1-regularized support vector machines (SVMs). In this paper, we leverage such methods to obtain an efficient learning algorithm for the recently proposed minimax risk classifiers (MRCs). The proposed iterative algorithm also provides a sequence of worst-case error probabilities and performs feature selection. Experiments on multiple high-dimensional datasets show that the proposed algorithm is efficient in high-dimensional scenarios. In addition, the worst-case error probability provides useful information about the classifier performance, and the features selected by the algorithm are competitive with the state-of-the-art.

1 INTRODUCTION

High-dimensional data is common in multiple areas such as health care and genomics. A typical example of high-dimensional supervised classification problem in genomics is to separate healthy patients from cancer patients based on gene expression data with tens of thousands of features [Guyon et al., 2002]. In addition to high-dimensional raw data, learning methods often perform a high-dimensional representation of the input data vector in order to improve the classification performance [Sonnenburg et al., 2006, Rahimi and Recht, 2008, Liu et al., 2017]. Learning in such high-dimensional settings often leads to highly complex optimization processes because the number of variables involved in the optimization increases with the number of features [Shi et al., 2010, Yuan et al., 2012].

In addition to high-dimensional data, a limited number of samples is common in the above-mentioned applications e.g., tens of thousands of features but less than 100 patients [Guyon and Elisseeff, 2003, Brown et al., 2012]. In such scenarios, the conventional performance assessment based on cross-validation can be unreliable [Varoquaux, 2018, Vabalas et al., 2019]. In addition, these cross-validation estimates also increase the computational cost as they require learning multiple classifiers.

Multiple methods have been proposed to improve the learning efficiency in high dimensions (see e.g., Yuan et al. [2012]). These methods are based on several approaches such as coordinate descent [Hsieh et al., 2008, Yu et al., 2011], interior-point method [Koh et al., 2007], and stochastic subgradient [Shalev-Shwartz et al., 2007]. In addition, multiple classification and regression methods exploit the parameters' sparsity induced by convex penalties and regularization terms in high-dimensional settings [Mignacco et al., 2020, Celentano and Montanari, 2022]. Recently, constraint generation techniques have enabled the efficient learning of L1-regularized support vector machines (SVMs) for cases with a large number of features [Dedieu et al., 2022]. These techniques obtain improved efficiency for the linear program (LP) corresponding with binary L1-SVMs due to the sparsity in the solution.

In this paper, we present a learning algorithm for the recently proposed minimax risk classifiers (MRCs) [Mazuelas et al., 2020, 2023]. The presented algorithm provides efficient learning in high dimensions, obtains worst-case error probabilities that can serve to assess performance, and performs feature selection. Specifically, the main contributions in the paper are as follows.

- We present a learning algorithm for MRCs that utilizes constraint generation methods to significantly improve the efficiency in high dimensions.
- The presented algorithm utilizes a greedy feature selection approach that achieves a fast decrease in worst-case error probability while using a small number of features.

- Our theoretical results show that the proposed algorithm obtains a sequence of classifiers with decreasing worst-case error probabilities that converges to the solution obtained using all the features.
- Using multiple benchmark datasets, we experimentally show that the proposed algorithm performs efficient learning in high-dimensional settings. In addition, the algorithm obtains an error assessment for the classifier and performs effective feature selection.

Notations: For a set \mathcal{S} , we denote its complement as \mathcal{S}^c and its cardinality as $|\mathcal{S}|$; bold lowercase and uppercase letters represent vectors and matrices, respectively; vectors subscripted by a set of indices represent subvectors obtained by the components corresponding to the indices in the set; matrices subscripted by a set of indices represent submatrices obtained by the columns corresponding to the indices in the set; \mathbf{I} denotes the identity matrix; $\mathbb{1}\{\cdot\}$ denotes the indicator function; $\mathbf{1}$ denotes a vector of ones; for a vector \mathbf{v} , $|\mathbf{v}|$ and $(\mathbf{v})_+$, denote its component-wise absolute value and positive part, respectively; $\|\cdot\|_1$ denotes the 1-norm of its argument; \otimes denotes the Kronecker product; \preceq and \succeq denote vector inequalities; $\mathbb{E}_p\{\cdot\}$ denotes the expectation of its argument with respect to distribution p ; and \mathbf{e}_i denotes the i -th vector in a standard basis.

2 PRELIMINARIES

In this section, we describe the setting addressed in the paper and MRC methods that minimize the worst-case error probability.

2.1 PROBLEM FORMULATION

Supervised classification uses instance-label pairs to determine classification rules that assign labels to instances. We denote by \mathcal{X} and \mathcal{Y} the sets of instances and labels, respectively, with \mathcal{Y} represented by $\{1, 2, \dots, |\mathcal{Y}|\}$. We denote by $\mathsf{T}(\mathcal{X}, \mathcal{Y})$ the set of all classification rules (both randomized and deterministic) and we denote by $h(y|x)$ the probability with which rule $h \in \mathsf{T}(\mathcal{X}, \mathcal{Y})$ assigns label $y \in \mathcal{Y}$ to instance $x \in \mathcal{X}$ ($h(y|x) \in \{0, 1\}$ for deterministic classification rules). In addition, we denote by $\Delta(\mathcal{X} \times \mathcal{Y})$ the set of probability distributions on $\mathcal{X} \times \mathcal{Y}$ and by $\ell(h, p)$ the expected 0-1 loss of the classification rule $h \in \mathsf{T}(\mathcal{X}, \mathcal{Y})$ with respect to distribution $p \in \Delta(\mathcal{X} \times \mathcal{Y})$. If $p^* \in \Delta(\mathcal{X} \times \mathcal{Y})$ is the underlying distribution of the instance-label pairs, then $\ell(h, p^*)$ is the classification error probability or classification risk of rule h denoted in the following as $\mathcal{R}(h)$, that is, $\mathcal{R}(h) := \ell(h, p^*)$.

Instance-label pairs can be represented as real vectors by using a feature mapping $\Phi : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^m$. The most common way to define such feature mapping is using multiple features over instances together with one-hot encodings of

labels as follows (see e.g., Mohri et al. [2018])

$$\Phi(x, y) = \mathbf{e}_y \otimes \Psi(x) = [\mathbb{1}\{y = 1\}\Psi(x)^T, \quad (1)$$

$$\mathbb{1}\{y = 2\}\Psi(x)^T, \dots, \mathbb{1}\{y = |\mathcal{Y}|\}\Psi(x)^T]^T$$

where the map $\Psi : \mathcal{X} \rightarrow \mathbb{R}^d$ represents instances as real vectors of size d . This map can be just the identity $\Psi(x) = x$ or given by a feature representation such as that provided by random Fourier features (RFF) [Rahimi and Recht, 2008], that is

$$\Psi(x) = [\cos(\mathbf{u}_1^T \mathbf{x}), \cos(\mathbf{u}_2^T \mathbf{x}), \dots, \cos(\mathbf{u}_D^T \mathbf{x}), \quad (2)$$

$$\sin(\mathbf{u}_1^T \mathbf{x}), \sin(\mathbf{u}_2^T \mathbf{x}), \dots, \sin(\mathbf{u}_D^T \mathbf{x})]^T$$

for D random Gaussian vectors $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_D \sim N(0, \gamma \mathbf{I})$ with covariance given by the scaling parameter γ (see e.g., Rahimi and Recht [2008]).

In this paper we consider scenarios in which the features' dimensionality (d) is large, and propose efficient learning methods for the recently presented MRCs, which are briefly described in the following.

2.2 MINIMAX RISK CLASSIFIERS

MRCs are classification rules that minimize the worst-case error probability over distributions in an uncertainty set [Mazuelas et al., 2020, 2022, 2023]. Specifically, such rules are solutions to the minimax risk problem defined as

$$\min_{h \in \mathsf{T}(\mathcal{X}, \mathcal{Y})} \max_{p \in \mathcal{U}} \ell(h, p) \quad (3)$$

where \mathcal{U} is an uncertainty set of distributions determined by expectation estimates as

$$\mathcal{U} = \{p \in \Delta(\mathcal{X} \times \mathcal{Y}) : |\mathbb{E}_p\{\Phi(x, y)\} - \boldsymbol{\tau}| \preceq \boldsymbol{\lambda}\} \quad (4)$$

where $\boldsymbol{\tau}$ denotes the mean vector of expectation estimates corresponding with the feature mapping Φ , and $\boldsymbol{\lambda} \succeq \mathbf{0}$ is a confidence vector that accounts for inaccuracies in the estimate. The mean and confidence vectors can be obtained from the training samples $\{(x_i, y_i)\}_{i=1}^n$ as

$$\boldsymbol{\tau} = \frac{1}{n} \sum_{i=1}^n \Phi(x_i, y_i), \quad \boldsymbol{\lambda} = \frac{\mathbf{s}}{\sqrt{n}} \quad (5)$$

where \mathbf{s} denotes the vector formed by the component-wise sample standard deviations of $\{\Phi(x_i, y_i)\}_{i=1}^n$.

As described in Mazuelas et al. [2020, 2022, 2023] using the expected 0-1 loss ℓ , the MRC rule h solution of (3) is given by a linear combination of the feature mapping with coefficients given by a vector $\boldsymbol{\mu}^* \in \mathbb{R}^m$. Specifically, the MRC rule h given by $\boldsymbol{\mu}^*$ assigns label $\hat{y} \in \mathcal{Y}$ to instance $x \in \mathcal{X}$ with probability

$$h(\hat{y}|x) = \begin{cases} (\Phi(x, \hat{y})^T \boldsymbol{\mu}^* - \varphi(\boldsymbol{\mu}^*))_+ / d_x & \text{if } d_x \neq 0 \\ 1/|\mathcal{Y}| & \text{if } d_x = 0 \end{cases} \quad (6)$$

where $\varphi(\boldsymbol{\mu}) = \max_{x \in \mathcal{X}, \mathcal{C} \subseteq \mathcal{Y}} \left(\frac{\sum_{y \in \mathcal{C}} \Phi(x, y)^T \boldsymbol{\mu} - 1}{|\mathcal{C}|} \right)$
 $d_x = \sum_{y \in \mathcal{Y}} (\Phi(x, y)^T \boldsymbol{\mu}^* - \varphi(\boldsymbol{\mu}^*))_+.$

Note that the corresponding deterministic classification rule h^d assigns a label that maximizes the probability in (6) and is given by $\arg \max_{y \in \mathcal{Y}} h(y|x) = \arg \max_{y \in \mathcal{Y}} \Phi(x, y)^T \boldsymbol{\mu}^*$. In addition, the classification risk of such classification rule h^d is bounded by twice the error probability of the MRC rule h , that is,

$$\mathcal{R}(h^d) \leq 2\mathcal{R}(h) \quad (7)$$

since $1 - h^d(y|x) \leq 2(1 - h(y|x))$ for any $x \in \mathcal{X}, y \in \mathcal{Y}$.

The vector $\boldsymbol{\mu}^*$ that determines the 0-1 MRC rule corresponding to (3) is obtained by solving the convex optimization problem [Mazuelas et al., 2020, 2023]

$$R^* = \min_{\boldsymbol{\mu}} 1 - \boldsymbol{\tau}^T \boldsymbol{\mu} + \varphi(\boldsymbol{\mu}) + \boldsymbol{\lambda}^T |\boldsymbol{\mu}|. \quad (8)$$

The minimum value R^* of (8) equals the minimax risk value of (3). Hence, it is the MRC's worst-case error probability for distributions included in the uncertainty set \mathcal{U} , and is an upper bound on the MRC's classification error when the underlying distribution is included in such uncertainty set.

The convex optimization problem (8) of MRCs is L1-penalized and leads to sparsity in the coefficients $\boldsymbol{\mu}^*$ corresponding to the features in Φ . This sparsity in the coefficients implies that only a subset of features are sufficient to obtain the optimal worst-case error probability for MRCs. Therefore, in practice, efficient learning can be achieved by using the relevant subset of features in Φ to solve the MRC optimization problem. This intuition leads to the efficient learning algorithm presented in this following.

3 EFFICIENT LEARNING OF MRCS IN HIGH DIMENSIONS

This section describes the proposed learning algorithm based on constraint generation over an LP formulation of MRCs learning (8).

3.1 LP FORMULATION FOR 0-1 MRCS

The theorem below presents an LP formulation (primal and dual) for the 0-1 MRC convex optimization problem (8).

Theorem 1. *Let \mathcal{S} be the set of pairs of instances and labels' subsets, i.e., $\mathcal{S} = \{(x, \mathcal{C}) : x \in \mathcal{X}, \mathcal{C} \subseteq \mathcal{Y}, \mathcal{C} \neq \emptyset\}$, and for each $i = 1, 2, \dots, |\mathcal{S}|$ corresponding to pair (x, \mathcal{C}) let*

$$\mathbf{g}_i = \frac{\sum_{y \in \mathcal{C}} \Phi(x, y)}{|\mathcal{C}|}, \quad b_i = \frac{1}{|\mathcal{C}|} - 1. \quad (9)$$

Then, the 0-1 MRC problem (8) is equivalent to the LP

$$\mathcal{P} : \min_{\boldsymbol{\mu}_1, \boldsymbol{\mu}_2, \nu} -(\boldsymbol{\tau} - \boldsymbol{\lambda})^T \boldsymbol{\mu}_1 + (\boldsymbol{\tau} + \boldsymbol{\lambda})^T \boldsymbol{\mu}_2 + \nu \quad (10)$$

$$\text{s.t.} \quad \mathbf{F}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) - \nu \mathbf{1} \preceq \mathbf{b}$$

$$\boldsymbol{\mu}_1, \boldsymbol{\mu}_2 \succeq 0$$

where the i -th row of matrix $\mathbf{F} \in \mathbb{R}^{|\mathcal{S}| \times m}$ is given by \mathbf{g}_i^T , and the i -th component of vector $\mathbf{b} \in \mathbb{R}^{|\mathcal{S}|}$ is given by b_i . In addition, the Lagrange dual of (10) is

$$\mathcal{D} : \max_{\boldsymbol{\alpha}} -\mathbf{b}^T \boldsymbol{\alpha} \quad (11)$$

$$\text{s.t.} \quad \boldsymbol{\tau} - \boldsymbol{\lambda} \preceq \mathbf{F}^T \boldsymbol{\alpha} \preceq \boldsymbol{\tau} + \boldsymbol{\lambda}$$

$$\mathbf{1}^T \boldsymbol{\alpha} = 1, \boldsymbol{\alpha} \succeq 0$$

and if $\boldsymbol{\mu}_1^*, \boldsymbol{\mu}_2^*, \nu^*$ is a solution of (10), we have that $\boldsymbol{\mu}^* = \boldsymbol{\mu}_1^* - \boldsymbol{\mu}_2^*$ is a solution of (8).

Proof. In the convex optimization problem (8), we introduce the additional variable $\nu \in \mathbb{R}$ given by $\nu = \varphi(\boldsymbol{\mu}) + 1$. Then, using \mathbf{g}_i and \mathbf{b} as defined above, we have

$$\nu = \max_{i \in \{1, 2, \dots, |\mathcal{S}|\}} \mathbf{g}_i^T \boldsymbol{\mu} - b_i. \quad (12)$$

Therefore, the optimization problem (8) becomes

$$\min_{\boldsymbol{\mu}, \nu} -\boldsymbol{\tau}^T \boldsymbol{\mu} + \boldsymbol{\lambda}^T |\boldsymbol{\mu}| + \nu \quad (13)$$

$$\text{s.t.} \quad \mathbf{F} \boldsymbol{\mu} - \nu \mathbf{1} \preceq \mathbf{b}.$$

The optimization in (13) can be reformulated as the LP in (10) using the change of variables $\boldsymbol{\mu}_1 = (\boldsymbol{\mu})_+$ and $\boldsymbol{\mu}_2 = (-\boldsymbol{\mu})_+$ which implies $\boldsymbol{\mu} = \boldsymbol{\mu}_1 - \boldsymbol{\mu}_2$ and $|\boldsymbol{\mu}| = \boldsymbol{\mu}_1 + \boldsymbol{\mu}_2$. Finally, the dual in (11) is directly obtained from the LP (10) (see e.g., Section 4.2 in Bertsimas and Tsitsiklis [1997]). \square

The number of variables of the primal problem \mathcal{P} in (10) (constraints in dual problem \mathcal{D} in (11)) is given by the number of features of Φ , $m = d|\mathcal{Y}|$. Therefore, the complexity of MRC learning given by such LP formulation is $O(m^3)$ which is not affordable in high-dimensional settings. The next section presents an efficient learning algorithm for MRCs in high dimensions based on constraint generation methods.

3.2 ALGORITHM FOR EFFICIENT LEARNING

MRC learning can be efficiently carried out in high-dimensional settings because the solution of (10) is a sparse vector. This sparsity is due to the implicit L1-penalization in (8), $\boldsymbol{\lambda}^T |\boldsymbol{\mu}|$, and the fact that usually only a small subset of features are informative in high-dimensional settings [Ghosh and Cabrera, 2022]. The sparsity of the MRCs' coefficients enables to carry out the learning process considering only a small subset of variables in (10) (respectively constraints in (11)). In the

following, we propose an algorithm based on constraint generation methods that solves the MRC LP problem in (10) by iteratively selecting small subsets of features of Φ .

Algorithm 1 Efficient learning algorithm for 0-1 MRCs

Input: \mathbf{F} , \mathbf{b} , $\boldsymbol{\tau}$, $\boldsymbol{\lambda}$, initial subset of features \mathcal{I} , dual constraints' violation threshold ϵ , maximum number of iterations k_{\max} , and maximum number of features selected in each iteration n_{\max}

Output: selected features $\mathcal{I}^* \subseteq \{1, 2, \dots, m\}$, optimal solution $\boldsymbol{\mu}^* \in \mathbb{R}^m$, and optimal worst-case error probability R^*

- 1: $\boldsymbol{\mu}^1 = [0, 0, \dots, 0] \in \mathbb{R}^m$
- 2: LPSOLVE($\mathcal{P}_{\mathcal{I}}$)
 $\boldsymbol{\mu}_1^1, \boldsymbol{\mu}_2^1, \nu^1 \leftarrow$ Solution of primal
 $\boldsymbol{\alpha}^1 \leftarrow$ Solution of dual
 $R^1 \leftarrow$ Optimal value (worst-case error probability)
- 3: $\boldsymbol{\mu}_{\mathcal{I}}^1 = \boldsymbol{\mu}_1^1 - \boldsymbol{\mu}_2^1$
- 4: $\mathcal{J} \leftarrow$ SELECT(\mathbf{F} , $\boldsymbol{\tau}$, $\boldsymbol{\lambda}$, n_{\max} , \mathcal{I} , ϵ , $\boldsymbol{\alpha}^1$)
- 5: $k = 1$
- 6: **while** $\mathcal{J} \setminus \mathcal{I} \neq \emptyset$ and $k \leq k_{\max}$ **do**
- 7: $k = k + 1$
- 8: $\boldsymbol{\mu}_1^{k-1} = (\boldsymbol{\mu}^{k-1}_{\mathcal{J}})_+$, $\boldsymbol{\mu}_2^{k-1} = (-\boldsymbol{\mu}^{k-1}_{\mathcal{J}})_+$
- 9: LPSOLVE($\mathcal{P}_{\mathcal{J}}$, $\boldsymbol{\mu}_1^{k-1}$, $\boldsymbol{\mu}_2^{k-1}$, ν^{k-1})
 $\boldsymbol{\mu}_1^k, \boldsymbol{\mu}_2^k, \nu^k \leftarrow$ Solution of primal
 $\boldsymbol{\alpha}^k \leftarrow$ Solution of dual
 $R^k \leftarrow$ Optimal value (worst-case error probability)
- 10: $\boldsymbol{\mu}^k = \boldsymbol{\mu}^{k-1}$
- 11: $\boldsymbol{\mu}^k_{\mathcal{J}} = \boldsymbol{\mu}_1^k - \boldsymbol{\mu}_2^k$
- 12: $\mathcal{I} \leftarrow \mathcal{J}$
- 13: $\mathcal{J} \leftarrow$ SELECT(\mathbf{F} , $\boldsymbol{\tau}$, $\boldsymbol{\lambda}$, n_{\max} , \mathcal{I} , ϵ , $\boldsymbol{\alpha}^k$)
- 14: **end while**
- 15: $\mathcal{I}^* = \mathcal{I}$, $\boldsymbol{\mu}^* = \boldsymbol{\mu}^k$, $R^* = R^k$

The proposed learning algorithm obtains the optimal solution $\boldsymbol{\mu}^*$ to the original problem by iteratively solving a sequence of subproblems (see details in Algorithm 1). These subproblems correspond with small subsets of features selected by a constraint generation method (SELECT function detailed in Algorithm 2).

The subproblem of (10) corresponding to the subset of features $\mathcal{J} \subseteq \{1, 2, \dots, m\}$ is defined as

$$\begin{aligned} \mathcal{P}_{\mathcal{J}} : \min_{\boldsymbol{\mu}_1, \boldsymbol{\mu}_2, \nu} & -(\boldsymbol{\tau}_{\mathcal{J}} - \boldsymbol{\lambda}_{\mathcal{J}})^T \boldsymbol{\mu}_1 + (\boldsymbol{\tau}_{\mathcal{J}} + \boldsymbol{\lambda}_{\mathcal{J}})^T \boldsymbol{\mu}_2 + \nu \\ \text{s.t.} & \mathbf{F}_{\mathcal{J}}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) + \nu \mathbf{1} \preceq \mathbf{b} \\ & \boldsymbol{\mu}_1, \boldsymbol{\mu}_2 \succeq 0 \end{aligned} \quad (14)$$

where $\boldsymbol{\tau}_{\mathcal{J}}$ and $\boldsymbol{\lambda}_{\mathcal{J}}$ denote the subvectors of $\boldsymbol{\tau}$ and $\boldsymbol{\lambda}$ corre-

sponding to the \mathcal{J} components, and $\mathbf{F}_{\mathcal{J}}$ denotes the submatrix of \mathbf{F} corresponding to the \mathcal{J} columns. In addition, the dual of (14) is

$$\begin{aligned} \mathcal{D}_{\mathcal{J}} : \max_{\boldsymbol{\alpha}} & -\mathbf{b}^T \boldsymbol{\alpha} \\ \text{s.t.} & \boldsymbol{\tau}_{\mathcal{J}} - \boldsymbol{\lambda}_{\mathcal{J}} \preceq \mathbf{F}_{\mathcal{J}}^T \boldsymbol{\alpha} \preceq \boldsymbol{\tau}_{\mathcal{J}} + \boldsymbol{\lambda}_{\mathcal{J}} \\ & \mathbf{1}^T \boldsymbol{\alpha} = 1, \boldsymbol{\alpha} \succeq 0. \end{aligned} \quad (15)$$

At each iteration k of Algorithm 1 (Line 9), the LPSOLVE function solves such subproblems and obtains the primal solution $\boldsymbol{\mu}_1^k, \boldsymbol{\mu}_2^k, \nu^k$ and dual solution $\boldsymbol{\alpha}^k$ along with the worst-case error probability R^k given by the optimal value. The dual solution is used by the SELECT function (Line 4 and 13) to obtain a subsequent subset of features. The primal solution can be used by the LPSOLVE function (Line 9) to warm-start the optimization in the next iteration. This iterative process ends when the SELECT function returns a subset of features that does not contain any new feature.

Notice that at each iteration k , the algorithm obtains coefficients $\boldsymbol{\mu}^k$, features subset \mathcal{J} , and worst-case error probability R^k . In Section 3.5, we show that such coefficients $\boldsymbol{\mu}^k$ provide an MRC with worst-case error probability R^k corresponding with the uncertainty set given by features \mathcal{J} . In addition, we show that such worst-case error probability monotonically decreases with the number of iterations and converges to the optimal corresponding with all the features.

3.3 GREEDY FEATURE SELECTION

The SELECT function determines the subset of features for the next iteration based on the previous subset and the current solution. Specifically, this greedy selection process selects features corresponding to the most violated constraints in the dual and removes features corresponding to the constraints satisfied with a positive slack (overly-satisfied constraints). Such a process aims to achieve the fastest decrease in worst-case error probability while using the smallest number of features. As described in the literature for constraint generation methods [Bertsimas and Tsitsiklis, 1997, Desrosiers and Lübbecke, 2005], this type of selection process can be implemented in several alternative ways such as adding any subset of violated constraints or not removing constraints. In practice, we observe that adding a subset of violated constraints and removing all the overly-satisfied constraints achieves the fastest convergence.

The implementation details for function SELECT are provided in Algorithm 2. The vector \mathbf{v} quantifies the violations in the dual constraints. In particular, the features corresponding to the negative values of vector \mathbf{v} (oversatisfied constraints) are removed, and the features corresponding to the n_{\max} largest positive values of vector \mathbf{v} are added. Moreover, the selection of features is also restricted by hyperparameter $\epsilon \geq 0$ that represents the minimum vio-

Algorithm 2 SELECT (greedy feature selection)

Input: \mathbf{F} , $\boldsymbol{\tau}$, $\boldsymbol{\lambda}$, n_{\max} , current subset of features \mathcal{I} , dual constraints' violation threshold ϵ , and dual solution $\boldsymbol{\alpha}$

Output: selected features $\mathcal{J} \subseteq \{1, \dots, m\}$

- 1: $\mathbf{v} = |\mathbf{F}^T \boldsymbol{\alpha} - \boldsymbol{\tau}| - \boldsymbol{\lambda}$
 - 2: $\mathcal{J} = \{i : i \in \mathcal{I}, v_i = 0\}$
 - 3: $\mathcal{A} \leftarrow \text{ARGNMAX}(\mathbf{v}, n_{\max})$
 - 4: $\mathcal{A} = \{i : i \in \mathcal{A}, v_i > \epsilon\}$
 - 5: $\mathcal{J} \leftarrow \mathcal{J} \cup \mathcal{A}$
-

lation for dual constraints.

3.4 COMPUTATIONAL COMPLEXITY

The computational complexity of Algorithm 1 is given by the number of iterations and the complexity per iteration that depend on the maximum number of features selected in each iteration n_{\max} and the hyperparameter ϵ . Decreasing n_{\max} results in a reduced complexity per iteration at the expense of an increased number of iterations. Increasing ϵ can decrease the number of iterations and complexity per iteration at the expense of achieving approximate solutions. For instance, if $n_{\max} = m$ and $\epsilon = 0$ the algorithm finds the exact solution in only one iteration but the complexity of such iteration is large. In Section 4, we show that in practice $\epsilon = 0.0001$ and $n_{\max} \in [50, 500]$ obtains efficient and accurate learning of MRCs.

The number of iterations can be further reduced by using an adequate choice for the initial subset of features. This initial subset of features \mathcal{I} can be obtained efficiently using a simple approach for coarse feature selection such as correlation screening [Tibshirani et al., 2012] or few iterations of a first order optimization method [Mazuelas et al., 2023]. Similarly, the complexity per iteration is reduced by using a warm-start in LPSOLVE. In the proposed algorithm, a warm-start for iteration k is obtained directly from the previous solution $\boldsymbol{\mu}^{k-1}$ of the LP at iteration $k - 1$. Note that this warm-start is a basic feasible solution for the LP at iteration k since it is obtained by removing and adding features for which the corresponding coefficients at $k - 1$ are zero.

3.5 THEORETICAL ANALYSIS OF THE ALGORITHM

The following shows the theoretical properties of the proposed learning algorithm. In particular, the algorithm's iterations provide a sequence of MRCs with decreasing worst-case error probabilities that converges to the MRC corresponding with all the features.

The following theorem shows that the algorithm provides a decreasing sequence of worst-case error probabilities that can provide an upper bound to the classification error of the corresponding MRCs.

Theorem 2. *Let $\boldsymbol{\mu}^k$ and R^k be the coefficients and worst-case error probability obtained by the proposed algorithm at iteration k . If h^k is the MRC given by $\boldsymbol{\mu}^k$, then we have that*

$$\mathcal{R}(h^k) \leq R^k + (|\mathbb{E}_{p^*}\{\Phi(x, y)_{\mathcal{J}}\} - \boldsymbol{\tau}_{\mathcal{J}}| - \boldsymbol{\lambda}_{\mathcal{J}})^T |\boldsymbol{\mu}_{\mathcal{J}}^k| \quad (16)$$

where \mathcal{J} is the subset of features used at iteration k . In addition, the algorithm provides MRCs with decreasing worst-case error probabilities, that is,

$$R^{k+1} \leq R^k, \quad k = 1, 2, \dots \quad (17)$$

Proof. At each intermediate step, $\boldsymbol{\mu}^k$ and R^k are the coefficients and worst-case error probability of the MRC corresponding with uncertainty set

$$\mathcal{U}^k = \{p \in \Delta(\mathcal{X} \times \mathcal{Y}) : |\mathbb{E}_{p^*}\{\Phi(x, y)_{\mathcal{J}}\} - \boldsymbol{\tau}_{\mathcal{J}}| \preceq \boldsymbol{\lambda}_{\mathcal{J}}\} \quad (18)$$

using Theorem 1 for the subproblem (14). Therefore, inequality (16) is obtained using the bounds for error probabilities for MRCs in Mazuelas et al. [2023]. The second result is obtained since the warm-start at iteration $k + 1$ is a feasible solution with an objective value that equals the worst-case error probability R^k at iteration k . The warm-start is feasible at $k + 1$ and has value R^k because it is obtained from the solution at iteration k by removing and adding features corresponding to zero coefficients. \square

Inequality (16) bounds the classification error of the MRC classifier at any intermediate iteration. Specifically, the worst-case error probability R^k at any iteration k corresponding to a feature subset \mathcal{J} directly provides such a bound if the error in the mean vector estimate is not underestimated, i.e., $\boldsymbol{\lambda}_{\mathcal{J}} \succeq |\mathbb{E}_{p^*}\{\Phi(x, y)_{\mathcal{J}}\} - \boldsymbol{\tau}_{\mathcal{J}}|$. In particular, if $\boldsymbol{\lambda}_{\mathcal{J}}$ is a confidence vector with coverage probability $1 - \delta$, that is, $\mathbb{P}\{|\mathbb{E}_{p^*}\{\Phi(x, y)_{\mathcal{J}}\} - \boldsymbol{\tau}_{\mathcal{J}}| \preceq \boldsymbol{\lambda}_{\mathcal{J}}\} \geq 1 - \delta$, then

$$\mathcal{R}(h^k) \leq R^k \quad (19)$$

with probability at least $1 - \delta$. In other cases, R^k still provides approximate bounds as long as the underestimation $|\mathbb{E}_{p^*}\{\Phi(x, y)_{\mathcal{J}}\} - \boldsymbol{\tau}_{\mathcal{J}}| - \boldsymbol{\lambda}_{\mathcal{J}}$ is not substantial. Similar bounds also hold for the sequence of deterministic classification rules due to inequality in (7).

The next theorem shows the convergence of the sequence of MRCs obtained by the proposed algorithm to the MRC corresponding with all the features.

Theorem 3. *Let $\boldsymbol{\mu}^*$ and R^* be the MRC coefficients and worst-case error probability obtained by solving (10) using*

all the features. If R^k , $k = 1, 2, \dots$, is the sequence of worst-case error probabilities obtained by the proposed algorithm using dual constraints' violation threshold ϵ . Then, there exists $k_0 \geq 1$ such that

$$R^* \leq R^k \leq R^* + \epsilon \|\boldsymbol{\mu}^*\|_1 \quad (20)$$

for any $k \geq k_0$.

Proof. The first inequality follows by noting that the feasible set of (11) corresponding with all the features is contained in the feasible set of any subproblem given a subset of features. In the following, we prove the second inequality.

Let $k_0 \geq 1$ be an iteration in which SELECT function in Algorithm 1 does not add any new feature. Such a case occurs after a finite number of iterations due to the properties of constraint generation methods [Bertsimas and Tsitsiklis, 1997]. If $\boldsymbol{\alpha}^{k_0}$ is the dual solution obtained at iteration k_0 with subset of features \mathcal{J} , then we have that

$$\boldsymbol{\tau}_{\mathcal{J}} - \boldsymbol{\lambda}_{\mathcal{J}} \preceq \mathbf{F}_{\mathcal{J}} \boldsymbol{\alpha}^{k_0} \preceq \boldsymbol{\tau}_{\mathcal{J}} + \boldsymbol{\lambda}_{\mathcal{J}} \quad (21)$$

$$\boldsymbol{\tau}_{\mathcal{J}^c} - \boldsymbol{\lambda}_{\mathcal{J}^c} - \epsilon \mathbf{1} \preceq \mathbf{F}_{\mathcal{J}^c} \boldsymbol{\alpha}^{k_0} \preceq \boldsymbol{\tau}_{\mathcal{J}^c} + \boldsymbol{\lambda}_{\mathcal{J}^c} + \epsilon \mathbf{1}. \quad (22)$$

On combining (21) and (22), we have

$$\boldsymbol{\tau} - \boldsymbol{\lambda} - \epsilon \mathbf{1} \preceq \mathbf{F} \boldsymbol{\alpha}^{k_0} \preceq \boldsymbol{\tau} + \boldsymbol{\lambda} + \epsilon \mathbf{1}. \quad (23)$$

Now, consider the following primal problem corresponding to the dual (11) with constraints as in (23).

$$\begin{aligned} \min_{\boldsymbol{\mu}_1, \boldsymbol{\mu}_2, \nu} & -(\boldsymbol{\tau} - \boldsymbol{\lambda} - \epsilon \mathbf{1})^T \boldsymbol{\mu}_1 + (\boldsymbol{\tau} + \boldsymbol{\lambda} + \epsilon \mathbf{1})^T \boldsymbol{\mu}_2 + \nu \\ \text{s.t.} & \mathbf{F}^T (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) - \nu \mathbf{1} \preceq \mathbf{c} \\ & \boldsymbol{\mu}_1, \boldsymbol{\mu}_2 \succeq 0. \end{aligned} \quad (24)$$

If $\boldsymbol{\mu}_1^*$, $\boldsymbol{\mu}_2^*$, and ν^* correspond to the MRC coefficients obtained by solving (10) using all the features, then it is also a feasible solution to the problem (24) since they have the same feasible set. Therefore, using weak duality (Theorem 4.3 in Bertsimas and Tsitsiklis [1997]), we have that

$$-\mathbf{b}^T \boldsymbol{\alpha}^{k_0} \leq -(\boldsymbol{\tau} - \boldsymbol{\lambda} - \epsilon \mathbf{1})^T \boldsymbol{\mu}_1^* + (\boldsymbol{\tau} + \boldsymbol{\lambda} + \epsilon \mathbf{1})^T \boldsymbol{\mu}_2^* + \nu^* \quad (25)$$

since $\boldsymbol{\alpha}^{k_0}$ is a feasible solution for the dual of (24). Therefore,

$$R^{k_0} \leq R^* + \epsilon \mathbf{1} (\boldsymbol{\mu}_1^* + \boldsymbol{\mu}_2^*) = R^* + \epsilon \|\boldsymbol{\mu}^*\|_1 \quad (26)$$

since $\boldsymbol{\mu}_1^* = (\boldsymbol{\mu}^*)_+$ and $\boldsymbol{\mu}_2^* = (-\boldsymbol{\mu}^*)_+$. Hence, the second inequality holds for any $k \geq k_0$ since $R^k \leq R^{k_0}$ due to the monotonic decrease of the worst-case error probability shown in Theorem 2. \square

Table 1. High-dimensional data sets.

Number	Data set	Variables	Samples	Classes
1	Arcene	10000	200	2
2	Colon	2000	62	2
3	CLL_SUB_111	11340	111	3
4	Dorothea	100000	1150	2
5	GLI_85	22283	85	2
6	GLIOMA	4434	50	4
7	Leukemia	7129	72	2
8	Lung	12600	203	5
9	MLL	12582	72	3
10	Ovarian	15154	253	2
11	Prostate_GE	5966	102	2
12	SMK_CAN_187	19993	187	2
13	TOX_171	5748	171	4

Inequality (20) shows that if $\epsilon = 0$ the algorithm finds the MRC corresponding to all the features. In other cases, it finds an approximate solution whose accuracy depends on the hyperparameter ϵ . Therefore, the hyperparameter ϵ can serve to reduce the complexity of the algorithm while obtaining near optimal solutions. In the next section, we further analyze the effect of hyperparameter ϵ through numerical experiments using multiple real datasets.

4 RESULTS

In this section, we present a set of experimental results that analyze the effect of hyperparameters and quantify the training time improvements in practice. In addition, we describe the benefits of having the worst-case error probabilities and show that the presented algorithm performs efficient feature selection. The proposed Algorithm 1 (MRC-CG) is compared with the LP formulation for MRCs in 10 using all the features (MRC-LP), and the constraint generation method for L1-SVMs (SVM-CG) [Dedieu et al., 2022]. In addition, the quality of the feature selection is compared with SVM-CG, recursive feature elimination (RFE) for SVMs [Guyon et al., 2002], minimum redundancy maximum relevancy (MRMR), and analysis of variance (ANOVA) [Ding and Peng, 2005, Peng et al., 2005]. The experimental results are obtained using 13 real-world high-dimensional datasets as shown in Table 1. The datasets 4 and 10 can be obtained from the UCI [Dua and Graff, 2017] repository and the remaining from <https://jundongl.github.io/scikit-feature/datasets.html>. The implementation of the algorithm proposed is available in the library MRCpy [Bondugula et al., 2023] and the experimental setup in <https://github.com/MachineLearningBCAM/Constraint-Generation-for-MRCs>.

The hyperparameter ϵ : We present results that show the influence of ϵ , and illustrate in practice the theoretical prop-




Figure 1. Monotonic decrease of the worst-case error probability R^k increasing the number of iterations k .




Figure 2: Improved efficiency of MRC-CG over MRC-LP for different values of maximum number of features selected n_{\max}

erties of MRC-CG using datasets "Ovarian" and "Arcene". The worst-case error probability R^* is obtained by solving (10) using all the features and the sequence of worst-case error probabilities R^k is obtained by MRC-CG for increasing number of iterations k . In Figure 1, we show the convergence of R^k to R^* using $\epsilon = \{0.01, 0.001, 0.0001\}$ and $n_{\max} = 100$. In practice, we observe that R^k converges to R^* with differences atmost in the order of 10^{-3} even for $\epsilon = 0.01$. In addition, the results show that R^k is monotonically decreasing (as shown in Theorem 2) and achieves significant convergence in few iterations. We also observe that the smaller values of ϵ lead to more accurate results, as shown in Theorem 3. On average, highly accurate results are obtained in 20 iterations using $\epsilon = 0.0001$.

The hyperparameter n_{\max} : We present results that show the influence of the hyperparameter n_{\max} on the

training time of MRC-CG in comparison with MRC-LP using datasets "Prostate_GE", "Arcene", "Ovarian" and "SMK_CAN_187". The parameters of MRC-CG are taken as $\epsilon = 0.0001$ and $k_{\max} = 20$, and the training times are averaged over 50 random repetitions using 90% of the data. Figure 2 presents the relative time computed as the time taken by MRC-CG over the time taken by MRC-LP for values of n_{\max} in the range of 1 to 1,500. The figure shows the trade-off due to the choice of hyperparameter n_{\max} . Increasing the value of n_{\max} decreases the number of iterations for convergence at the expense of increasing the complexity per iteration. In practice, with n_{\max} in the interval [50, 500], we obtain a good compromise between the number of iterations required for convergence and the complexity per iteration. On average, for $n_{\max} = 100$, MRC-CG is 10 times faster than MRC-LP. Based on these results, in the remainder of this section we take $n_{\max} = 100$, $\epsilon = 0.0001$, and $k_{\max} = 10$.

Scalability with increasing number of features: We present results to compare the scalability of MRC-CG with MRC-LP, and SVM-CG for increasing number of features on the datasets "Ovarian" and "Arcene". In particular, the results show the effect of the number of features on the training time of MRC-CG and compare it with MRC-LP and SVM-CG. The number of features range from 100 to 25,000 obtained using RFFs. Figure 3 presents the average training times for MRC-CG, MRC-LP, and SVM-CG over 20 random repetitions. We observe that MRC-CG is faster than MRC-LP for all cases and improves the scaling with the number of features significantly better than MRC-LP. In addition, the training times of MRC-CG are competitive with training times of SVM-CG, especially for large number of features.




Figure 3. Comparison of training times (in secs) of MRC-CG, MRC-LP, and SVM-CG for increasing number of RFF.

Table 2. Comparison of training times (in secs) of MRC-CG and SVM-CG along with the worst-case error probabilities R^* and cross-validated error estimate for MRC-CG using multiple high dimensional datasets.

Dataset	MRC-CG			SVM-CG
	R^*	Error	Time (in secs)	Time (in secs)
1	0.20	0.30 ± 0.10	3.56 ± 0.04	8.25 ± 0.54
2	0.24	0.13 ± 0.12	0.20 ± 0.02	0.58 ± 0.07
3	0.16	0.15 ± 0.09	6.84 ± 0.04	3.55 ± 0.35
4	0.07	0.07 ± 0.01	31.4 ± 0.64	220 ± 10.5
5	0.13	0.18 ± 0.17	1.50 ± 0.07	1.36 ± 0.18
6	0.18	0.38 ± 0.20	4.55 ± 0.02	2.58 ± 0.24
7	0.11	0.02 ± 0.05	0.70 ± 0.05	0.37 ± 0.08
8	0.15	0.05 ± 0.05	44.9 ± 0.96	16.6 ± 0.76
9	0.11	0.04 ± 0.06	5.06 ± 0.03	1.45 ± 0.20
10	0.10	0.00 ± 0.00	3.86 ± 0.26	3.35 ± 0.32
11	0.14	0.08 ± 0.07	0.79 ± 0.09	1.00 ± 0.10
12	0.23	0.25 ± 0.11	2.65 ± 0.13	8.9 ± 0.63
13	0.18	0.03 ± 0.03	17.7 ± 0.21	10.5 ± 0.22

Comparison using real-world datasets: Table 2 presents results to compare the average training times for MRC-CG and SVM-CG using the 13 high-dimensional datasets. We observe that MRC-CG is competitive with the SVM-CG in terms of training times. In addition, Table 2 presents worst-case error probability R^* given by MRC-CG along with the 10-fold cross-validated error estimate. We observe that the worst-case error probability R^* obtained using all the data at training can provide an error assessment for MRC-CG without requiring cross-validation. For instance, the R^* is in the confidence interval of the estimated error in 8 out of 13 datasets, and is an upper bound to the confidence interval of the estimated error in the remaining datasets. This worst-case error probability is particularly useful in the addressed setting since the scarcity of training samples results in a high variability for the error estimates based on cross-validation, as shown in the table.

Feature selection: Table 3 presents results to compare the MRC-CG as a feature selection approach with SVM-CG, RFE, MRMR, and ANOVA. The table shows the 10-fold cross-validated errors for logistic regression (LR) and decision tree (DT) using the features selected by the different methods. We assess the quality of features selected based on the errors and the number of features selected. Note that the number of features selected by RFE, MRMR, and ANOVA were set to the number of features selected by SVM-CG. The errors show that the presented MRC-CG provides state-of-the-art results for feature selection and can effectively select the most relevant features.

5 CONCLUSION

In this paper, we presented a learning algorithm for the recently proposed minimax risk classifiers (MRCs) that is efficient in high-dimensional settings. The algorithm utilizes

Table 3. Comparison of number of features selected and the error estimate obtained for logistic regression (LR) and decision tree (DT) using MRC-CG in comparison with multiple methods for feature selection.

Dataset	MRC-CG			SVM-CG			RFE		MRMR		ANOVA	
	Error		No. of features	Error		No. of features	Error		Error		Error	
	LR	DT		LR	DT		LR	DT	LR	DT	LR	DT
1	.29 ± .10	.33 ± .06	172 ± 3	.27 ± .11	.41 ± .09	150 ± 3	.28 ± .05	.31 ± .12	.31 ± .10	.39 ± .12	.27 ± .07	.31 ± .09
2	.20 ± .10	.19 ± .10	33 ± 2	.19 ± .10	.22 ± .13	32 ± 1	.24 ± .11	.20 ± .10	.22 ± .10	.29 ± .17	.22 ± .08	.26 ± .16
5	.23 ± .13	.17 ± .12	76 ± 0	.14 ± .07	.29 ± .11	51 ± 2	.20 ± .11	.31 ± .13	.14 ± .12	.20 ± .11	.13 ± .10	.20 ± .09
7	.04 ± .06	.12 ± .09	63 ± 1	.02 ± .05	.15 ± .12	37 ± 2	.05 ± .06	.16 ± .18	.04 ± .06	.15 ± .16	.04 ± .06	.16 ± .12
10	.00 ± .00	.02 ± .02	118 ± 6	.00 ± .00	.01 ± .02	33 ± 2	.00 ± .00	.03 ± .02	.00 ± .00	.02 ± .04	.00 ± .00	.01 ± .02
11	.08 ± .06	.16 ± .08	77 ± 2	.10 ± .08	.13 ± .10	45 ± 1	.10 ± .10	.18 ± .12	.08 ± .07	.14 ± .12	.08 ± .06	.16 ± .14
12	.29 ± .12	.41 ± .11	135 ± 3	.27 ± .09	.37 ± .09	127 ± 4	.33 ± .08	.41 ± .11	.32 ± .12	.39 ± .09	.30 ± .08	.37 ± .09

a greedy feature selection approach that iteratively removes and selects features achieving a fast decrease in worst-case error probability while using a small number of features. We prove theoretically that the proposed iterative algorithm obtains a sequence of MRCs with decreasing worst-case error probabilities that converge to the solution obtained using all the features. The numerical results assess the efficiency of the presented algorithm and compare it with the state-of-the-art using 13 high-dimensional datasets with a large number of features. The results show that the presented algorithm converges to the solution obtained using all the features in a few iterations, and provides a significant efficiency increase, especially in cases with a large number of features. In addition, the algorithm provides the worst-case probability error that is particularly useful in high-dimensional scenarios with limited number of samples, which can suffer from a high variability of assessments based on cross-validation. The results also show that the algorithm can provide state-of-the-art results for feature selection and can effectively select the most relevant features.

Acknowledgements

Funding in direct support of this work has been provided by projects PID2019-105058GA-I00, CNS2022-135203, and CEX2021-001142-S funded by MCIN/AEI/10.13039/501100011033 and the European Union “NextGenerationEU”/PRTR, and by programmes ELKARTEK and BERC-2022-2025 funded by the Basque Government.

References

- Dimitris Bertsimas and John N. Tsitsiklis. *Introduction to Linear Optimization*, volume 6. Athena Scientific, Belmont, MA, 1997.
- Kartheek Bondugula, Verónica Álvarez, Jose I. Segovia-Martín, Aritz Pérez, and Santiago Mazuelas. MRCpy: A library for minimax risk classifiers. *arXiv preprint*, arXiv:2108.01952, 2023.
- Gavin Brown, Adam Pocock, Ming-Jie Zhao, and Mikel Luján. Conditional likelihood maximisation: a unifying framework for information theoretic feature selection. *Journal of Machine Learning Research*, 13:27–66, 2012.
- Michael Celentano and Andrea Montanari. Fundamental barriers to high-dimensional regression with convex penalties. *The Annals of Statistics*, 50(1):170–196, 2022.
- Antoine Dedieu, Rahul Mazumder, and Haoyue Wang. Solving L1-regularized SVMs and related linear programs: Revisiting the effectiveness of column and constraint generation. *Journal of Machine Learning Research*, 23:1–41, 2022.
- Jacques Desrosiers and Marco E. Lübbecke. *A primer in column generation*. Springer, 2005.
- Chris Ding and Hanchuan Peng. Minimum redundancy feature selection from microarray gene expression data. *Journal of Bioinformatics and Computational Biology*, 3: 185–205, 2005.
- Dheeru Dua and Casey Graff. UCI Machine Learning Repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- Debopriya Ghosh and Javier Cabrera. Enriched random forest for high dimensional genomic data. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 19:2817–2828, 2022.
- Isabelle Guyon and André Elisseeff. An introduction to variable and feature selection. *Journal of Machine Learning Research*, 3:1157–1182, 2003.
- Isabelle Guyon, Jason Weston, Stephen Barnhill, and Vladimir Vapnik. Gene selection for cancer classification using support vector machines. *Machine learning*, 46: 389–422, 2002.
- Cho-Jui Hsieh, Kai-Wei Chang, Chih-Jen Lin, S Sathiya Keerthi, and Sellamanickam Sundararajan. A dual coordinate descent method for large-scale linear SVM. In *Proceedings of the 25th International Conference on Machine Learning*, pages 408–415, 2008.

- Kwangmoo Koh, Seung-Jean Kim, and Stephen Boyd. An interior-point method for large-scale L1-regularized logistic regression. *Journal of Machine Learning Research*, 8:1519–1555, 2007.
- Yong Liu, Shizhong Liao, Hailun Lin, Yinliang Yue, and Weiping Wang. Infinite kernel learning: generalization bounds and algorithms. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, 2017.
- Santiago Mazuelas, Andrea Zanoni, and Aritz Perez. Minimax classification with 0-1 loss and performance guarantees. In *Advances in Neural Information Processing Systems*, volume 33, pages 302–312, 2020.
- Santiago Mazuelas, Yuan Shen, and Aritz Pérez. Generalized maximum entropy for supervised classification. *IEEE Transactions on Information Theory*, 68:2530–2550, 2022.
- Santiago Mazuelas, Mauricio Romero, and Peter Grünwald. Minimax risk classifiers with 0-1 loss. *arXiv preprint*, arXiv:2201.06487, 2023.
- Francesca Mignacco, Florent Krzakala, Yue Lu, Pierfrancesco Urbani, and Lenka Zdeborova. The role of regularization in classification of high-dimensional noisy gaussian mixture. In *Proceedings of the 37th International Conference on Machine Learning*, pages 6874–6883, 2020.
- Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of machine learning*. MIT press, Cambridge, MA, second edition, 2018.
- Hanchuan Peng, Fuhui Long, and Chris Ding. Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27:1226–1238, 2005.
- Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Advances in Neural Information Processing Systems*, volume 20, pages 1177–1184, 2008.
- 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*, pages 807–814, 2007.
- Jianing Shi, Wotao Yin, Stanley Osher, and Paul Sajda. A fast hybrid algorithm for large-scale L1-regularized logistic regression. *Journal of Machine Learning Research*, 11:713–741, 2010.
- Sören Sonnenburg, Gunnar Rätsch, Christin Schäfer, and Bernhard Schölkopf. Large scale multiple kernel learning. *Journal of Machine Learning Research*, 7:1531–1565, 2006.
- Robert Tibshirani, Jacob Bien, Jerome Friedman, Trevor Hastie, Noah Simon, Jonathan Taylor, and Ryan J Tibshirani. Strong rules for discarding predictors in lasso-type problems. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 74:245–266, 2012.
- Andrius Vabalas, Emma Gowen, Ellen Poliakoff, and Alexander J Casson. Machine learning algorithm validation with a limited sample size. *PloS One*, 14, 2019.
- Gaël Varoquaux. Cross-validation failure: Small sample sizes lead to large error bars. *Neuroimage*, 180:68–77, 2018.
- Hsiang-Fu Yu, Fang-Lan Huang, and Chih-Jen Lin. Dual coordinate descent methods for logistic regression and maximum entropy models. *Machine Learning*, 85:41–75, 2011.
- Guo-Xun Yuan, Chia-Hua Ho, and Chih-Jen Lin. Recent advances of large-scale linear classification. *Proceedings of the IEEE*, 100:2584–2603, 2012.