



Robust linear classification from limited training data

Deepayan Chakrabarti¹

Received: 3 January 2021 / Revised: 2 September 2021 / Accepted: 4 October 2021 /
Published online: 18 November 2021

© The Author(s), under exclusive licence to Springer Science+Business Media LLC, part of Springer Nature 2021

Abstract

We consider the problem of linear classification under general loss functions in the limited-data setting. Overfitting is a common problem here. The standard approaches to prevent overfitting are dimensionality reduction and regularization. But dimensionality reduction loses information, while regularization requires the user to choose a norm, or a prior, or a distance metric. We propose an algorithm called RoLIN that needs no user choice and applies to a large class of loss functions. RoLIN combines “reliable” information from the top principal components with a robust optimization to extract any useful information from “unreliable” subspaces. It also includes a new robust cross-validation that is better than existing cross-validation methods in the limited-data setting. Experiments on 25 real-world datasets and three standard loss functions show that RoLIN broadly outperforms both dimensionality reduction and regularization. Dimensionality reduction has 14%–40% worse test loss on average as compared to RoLIN. Against L_1 and L_2 regularization, RoLIN can be up to 3x better for logistic loss and 12x better for squared hinge loss. The differences are greatest for small sample sizes, where RoLIN achieves the best loss on 2x to 3x more datasets than any competing method. For some datasets, RoLIN with 15 training samples is better than the best norm-based regularization with 1500 samples.

Keywords Classification · Principal components · Maximum entropy · Robust optimization

1 Introduction

In many machine learning applications, the size of the training data is small relative to the number of features, and acquiring more data may be too costly or time-consuming. For example, a standard dataset to predict breast-cancer from gene-expression data has only 99 positive examples for 7, 650 features (Sotiriou et al. 2003), and datasets with smaller training sizes are also of interest (Blagus and Lusa 2013). In computational advertising, we must learn to predict whether an ad is relevant to a person after seeing only limited data for that ad. Since there are many available ads, waiting for more

Editor: Willem Waegeman.

✉ Deepayan Chakrabarti
deepay@utexas.edu

¹ McCombs School of Business, University of Texas, Austin, TX, USA

training data can reduce ad revenue. Limited training data also leads to the “cold-start” problem in recommendation systems, where we must quickly tune our recommendations for new users, or find relevant matches for new items. Thus, the limited-data setting is widely applicable.

We consider the problem of learning a linear classifier from limited training data. In such cases, overfitting is common. In other words, the feature weight vector that minimizes the training loss often has a test loss that is much worse than the training loss. The usual solution for such overfitting is to do dimensionality reduction or regularization. Dimensionality reduction reduces the number of features, while regularization keeps all features but penalizes large weights. For example, for the least-squares loss, Principal Components Regression (PCR) does dimensionality reduction, while Ridge and LASSO do regularization based on the L_2 and L_1 norms respectively.

However, both dimensionality reduction and regularization have weaknesses. Dimensionality reduction ignores information. For example, PCR uses only the top few principal components of the feature matrix. But, if the top principal components are uncorrelated with the response variable, the PCR solution may perform poorly (Jolliffe 1982). For regularization, it is not easy to choose the best L_p norm. It depends on the dataset, the training size, and the loss function. Different norms can lead to significantly different test losses. Existing explanations of regularization rely on priors, or a distance between probability distributions, or a distance metric in feature space. It is not clear why such inputs are needed and how we should choose them in practice. This motivates the following problem:

How can we build a linear classifier that (a) outperforms both dimensionality reduction and norm-based regularization, (b) works for a wide range of loss functions, and (c) needs no user input such as a norm or a prior?

Our proposed method, called RoLIN (RObust LINear classification), aims to achieve this by combining dimensionality reduction with robust optimization. Dimensionality reduction methods such as PCR rely on the idea that solutions constructed from the top singular vectors (principal components) are less prone to overfitting. Like PCR, RoLIN first constructs such a classifier. But, unlike PCR, we do not ignore the bottom singular vectors. Training and test loss can indeed be very different for data projected on to the bottom singular vectors. But even from this “unreliable” projected data, we may be able to estimate some low-order moments, such as the mean and some aspects of the covariance. Now, the loss function depends on the entire data distribution, not just the low-order moments. So, RoLIN constructs a “worst-case” distribution that matches the low-order moments. Then it finds a classifier that has the smallest loss under this distribution. Now, we have two classifiers: one from the top principal components, and a robust one from the rest. RoLIN combines them into a single classifier that captures all available information. Thus, RoLIN goes beyond the top few principal components, but still avoids overfitting. Figure 1 shows the intuition behind RoLIN.

We summarize our main contributions below.

A new approach to avoid overfitting Given limited training data, the top principal components are often reliable while the bottom components are noisy. Motivated by this, RoLIN processes the top and bottom principal components differently. Reliable data from the top components is used directly, while noisy data from the bottom components is filtered through a robust optimization. This ensures that all available reliable information is extracted and used in building RoLIN’s classifier. By limiting the robust optimization to the unreliable subspace, RoLIN avoids becoming too conservative.

No user choice needed In contrast to existing regularization methods, RoLIN does not force the user to choose a norm, prior, or distance metric.

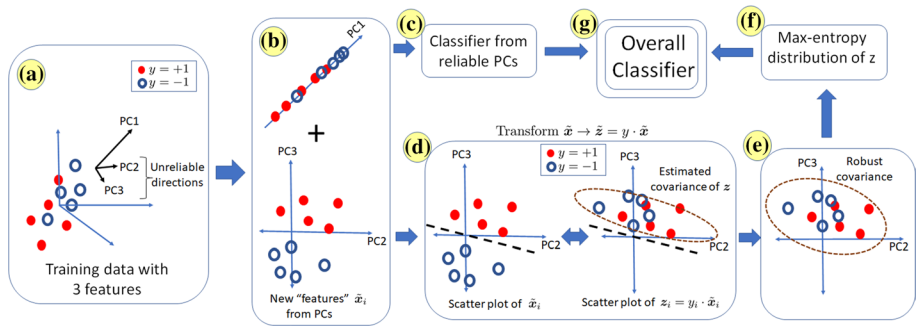


Fig. 1 Overview of RoLIN: **a** We can reliably estimate the top Principal Components (PCs) of the data distribution from limited training data. In this example, only the first PC is reliably estimated. **b** We project the data onto the reliable and unreliable subspaces. **c** We find the classifier with the least training loss on data projected on the reliable subspace. **d** But for the orthogonal subspace, such a classifier may overfit. Linear classifiers such as logistic regression minimize a loss that is a function of $z := y \cdot \tilde{x}$, where \tilde{x} is a projected datapoint. Given limited training data, even the covariance estimate of z can be noisy (Marcenko and Pastur 1967). So, minimizing the loss over the empirical distribution of z can yield classifiers (shown by the dashed line) that have low training loss but much higher test loss. **e** RoLIN builds a robust covariance of z . **f** It uses this to construct a maximum-uncertainty distribution for z . **g** Optimizing on this distribution gives a robust classifier, which RoLIN combines with the reliable classifier from step (c)

Applicable to many loss functions RoLIN works, unchanged, for the logistic, hinge, squared hinge, and modified Huber losses, among others. In particular, we can use RoLIN for both logistic regression and linear SVMs.

Robust cross-validation Existing cross-validation methods identify overfitting classifiers by their poor accuracy on holdout sets. But in limited-data settings, holdout sets are small and holdout accuracy may be too noisy. We develop a new cross-validation method, called ROBUSTCV, that checks for several signs of overfitting that are missed by standard cross-validation.

Empirical results We compare RoLIN against competing methods on three loss functions and 25 real-world datasets, where the number of features ranges from $p = 8$ to $p = 43,680$. We test each dataset under five different training sizes, from $n = 15$ to $n = 200$ samples. RoLIN outperforms dimensionality reduction as well as L_1 and L_2 regularization. Dimensionality reduction has 14% – 40% worse loss on average than RoLIN, under all problem settings. For some datasets, dimensionality reduction can be 4x worse. Under logistic loss, RoLIN can be up to 3x better than the best norm-based regularization. Under squared hinge loss, RoLIN can be up to 12x better. RoLIN performs particularly well for small training sizes, where robustness is most important. When 50 or fewer training samples are available, RoLIN achieves the smallest loss on around 2x to 3x as many datasets as the next best method, depending on the loss function. For some datasets, RoLIN with $n = 15$ samples is better than both L_1 and L_2 regularization with $n = 1500$ samples. Finally, among the competitors of RoLIN, no single method dominates, and it is challenging to choose the best method for a given dataset, loss function, and training size. In contrast, we find that RoLIN works well for all datasets under all problem settings.

The rest of the paper is organized as follows. We present our robust formulation and the main theorems in Sect. 2. Section 3 provides detailed algorithms for RoLIN and ROBUSTCV. Section 4 presents empirical results. We discuss prior work on overfitting in Sect. 5, and we conclude in Sect. 6. All proofs are deferred to Appendix A.

2 Robust minimization of expected loss

We are given n independent training samples from some distribution \mathcal{D} of pairs $(\mathbf{x}, y) \in \mathbb{R}^p \times \{-1, 1\}$, where \mathbf{x} is a feature vector with p features, and y is a binary class label. We want to train a classifier, parameterized by β , to output a positive score $g_\beta(\mathbf{x})$ when it predicts $y = 1$, and a negative score otherwise. The quality of classification is measured by a loss function $\ell(y, g_\beta(\mathbf{x}))$. The best classifier is the one that minimizes the expected loss

$$\min_{\beta} E_{(y,\mathbf{x}) \in \mathcal{D}} \ell(y, g_\beta(\mathbf{x})). \tag{1}$$

We consider linear classifiers where $g_\beta(\mathbf{x}) = \beta_0 + \beta_w^T \mathbf{x}$, where β_0 (intercept) and β_w (feature weights) are the first and the remaining elements of $\beta \in \mathbb{R}^{p+1}$. In this setting, many common losses are functions of $y \cdot g_\beta(\mathbf{x})$, and we denote the loss $\ell(y, g_\beta(\mathbf{x}))$ as $\ell(y \cdot g_\beta(\mathbf{x}))$ henceforth. Common loss functions include

$$\ell(y \cdot g_\beta(\mathbf{x})) = \begin{cases} \log_2(1 + \exp(-y \cdot g_\beta(\mathbf{x}))) & \text{(logistic loss)} \\ \max(0, 1 - y \cdot g_\beta(\mathbf{x})) & \text{(hinge loss)} \\ (\max(0, 1 - y \cdot g_\beta(\mathbf{x})))^2 & \text{(squared hinge loss)} \\ \mathbb{1}_{y \cdot g_\beta(\mathbf{x}) \geq -1} \cdot \max(0, 1 - y \cdot g_\beta(\mathbf{x}))^2 - \mathbb{1}_{y \cdot g_\beta(\mathbf{x}) < -1} \cdot (4y \cdot g_\beta(\mathbf{x})) & \text{(modified Huber loss)} \\ \mathbb{1}_{y \cdot g_\beta(\mathbf{x}) \leq 0} & \text{(zero-one loss).} \end{cases} \tag{2}$$

Well-known classifiers such as logistic regression (logistic loss) and linear SVM (hinge or squared hinge loss) fall under this framework. Such linear classifiers are also the building blocks for popular complex classifiers such as neural networks. Except for zero-one loss, all the other losses are convex in β . In this paper, we seek to minimize the expected loss in Eq. 1 for such convex loss functions.

RoLIN splits this problem into separate problems in different subspaces of the feature space. Next, we discuss the details of subspace separation, our robust optimization, and its solution. But first, we discuss the connection to PCR in more detail, as this helps us explain the unique features of RoLIN.

2.1 Intuition for subspaces via PCR

Consider the problem of least-squares regression:

$$\min_{\beta \in \mathbb{R}^p} E_{\mathbf{x},y} (y - \beta^T \mathbf{x})^2 = \min_{\beta \in \mathbb{R}^p} E[y^2] - 2\beta^T E[y \cdot \mathbf{x}] + \beta^T E[\mathbf{x}\mathbf{x}^T] \beta, \tag{3}$$

where we assume a zero intercept for ease of exposition. Suppose we are given n i.i.d. training samples $(y_i, \mathbf{x}_i) \in \mathbb{R} \times \mathbb{R}^p$. Then we can solve Eq. 3 after replacing the expectation terms with their estimates. But for small n , estimation errors lead to poor out-of-sample performance. Instead, Principal Components Regression (PCR) first projects the features \mathbf{x}_i on to the top few principal components. Then, it solves Eq. 3 only on the projected data. In other words, PCR splits the feature space \mathbb{R}^p into a subspace \mathcal{S}_1 spanned by the top principal components, and the orthogonal subspace \mathcal{S}_2 . It then solves for the best $\beta \in \mathcal{S}_1$ and ignores \mathcal{S}_2 .

The reason for the success of PCR is as follows. The principal directions and singular values correspond to the eigenvectors and eigenvalues of the matrix $\hat{M} = \sum_i \mathbf{x}_i \mathbf{x}_i^T / n$. The top eigenvalues and eigenvectors of \hat{M} are often close to those of the expectation matrix

$M = E[\hat{M}] = E[\mathbf{x}\mathbf{x}^T]$, even for small training sizes. This is because the estimation error for an eigenvector depends on the gap between its eigenvalue and all other eigenvalues (Davis and Kahan 1970; Yu et al. 2015). A larger gap implies smaller estimation error. For many datasets, this gap is large for the top eigenvalues. So the top principal directions are well estimated, and the same holds for the singular values too (Zhao et al. 2019). Hence, \hat{M} and M have similar projections on the subspace \mathcal{S}_1 spanned by these well-estimated principal directions. So for any $\beta \in \mathcal{S}_1$, $\beta^T \hat{M}$ is close to $\beta^T M$, and so $\beta^T \hat{M} \beta \approx \beta^T M \beta$. Applying this in Eq. 3, $\min_{\beta \in \mathcal{S}_1} E(y - \beta^T x)^2 \approx \min_{\beta \in \mathcal{S}_1} \sum_i (y_i - \beta^T x_i)^2 / n$, and this becomes PCR's solution. In contrast, the remaining principal components are poorly estimated when n is small. So, the least-squares training loss is not a reliable indicator of the expected loss in the orthogonal subspace \mathcal{S}_2 . Therefore PCR ignores \mathcal{S}_2 .

The loss functions we consider in Eq. 2 are not restricted to just second moments as in Eq. 3. However, the basic ideas underlying PCR are still applicable, as we discuss next.

2.2 Subspace separation

We want to minimize the expected loss under a linear score function. The loss is given by $\ell(y \cdot g_\beta(x)) = \ell(y\beta_0 + \beta_w^T(y \cdot x)) = \ell(y\beta_0 + \beta_w^T z)$, where $\beta = (\beta_0 \beta_w)^T$ with intercept $\beta_0 \in \mathbb{R}$ and feature vector $\beta_w \in \mathbb{R}^p$, and $z = y \cdot x$. Extending the PCR argument, we propose to split the space \mathbb{R}^p into three subspaces \mathcal{S}_0 , \mathcal{S}_1 , and \mathcal{S}_2 . The subspace \mathcal{S}_0 is spanned by the top few eigenvectors of $\hat{M} = \sum_i x_i x_i^T / n$. We expect that the loss function can be reliably estimated in this subspace:

$$E[\ell(y\beta_0 + \beta_w^T z)] \approx \mathbb{P}_n[\ell(y\beta_0 + \beta_w^T z)] \quad \text{for any } \beta_w \in \mathcal{S}_0, \quad (4)$$

where $\mathbb{P}_n[\cdot]$ represents the empirical mean. The next few eigenvectors span the subspace \mathcal{S}_1 . Here, we can reliably estimate only the first and second moments of the distribution of z projected on to \mathcal{S}_1 :

$$\beta_w^T E[zz^T] \beta_w \approx \beta_w^T \mathbb{P}_n[zz^T] \beta_w \quad \text{for any } \beta_w \in \mathcal{S}_1 \quad (5)$$

But this may not be enough to estimate the loss function accurately. We must also consider the subspace \mathcal{S}_2 that is orthogonal to both \mathcal{S}_0 and \mathcal{S}_1 . Here, we only expect first moments to be well-estimated. For the subspace \mathcal{S}_2 , orthogonal to both \mathcal{S}_0 and \mathcal{S}_1 , only first moments are well-estimated.

$$\beta_w^T E[z] \approx \beta_w^T \mathbb{P}_n[z] \quad \text{for any } \beta_w \in \mathcal{S}_1 \cup \mathcal{S}_2 \quad (6)$$

But the second moments under \mathcal{S}_2 are not arbitrary. Note that \mathcal{S}_1 and \mathcal{S}_2 are constructed from separate sets of eigenvectors of the sample covariance matrix $\mathbb{P}_n[zz^T] = \mathbb{P}_n[\mathbf{x}\mathbf{x}^T] = \hat{M}$ (since $z = y \cdot x$ and $y \in \{+1, -1\}$). So they are orthogonal under \hat{M} , that is, $\mathbb{P}_n[(P_{\mathcal{S}_1} z)(P_{\mathcal{S}_2} z)^T] = 0$, where $P_{\mathcal{S}_i} \in \mathbb{R}^{p \times p}$ is a matrix that projects any vector on to \mathcal{S}_i , for $i \in \{0, 1, 2\}$. We expect \mathcal{S}_1 and \mathcal{S}_2 to remain nearly orthogonal under the population covariance M :

$$E[(P_{\mathcal{S}_1} z)(P_{\mathcal{S}_2} z)^T] \approx 0. \quad (7)$$

Finally, we expect the eigenvalues of the second-moment matrix under \mathcal{S}_2 to be smaller than those under \mathcal{S}_1 .

$$\sigma_{max}\left(E\left[(P_{S_2}z)(P_{S_2}z)^T\right]\right) < \sigma_{min}\left(\mathbb{P}_n\left[(P_{S_1}z)(P_{S_1}z)^T\right]\right), \tag{8}$$

where σ_{max} and σ_{min} refer to the maximum and minimum non-zero eigenvalues of a matrix.

2.3 Different optimizations over subspaces

To construct our solution, we first optimize the training loss over S_0 . By Eq. 4, the training loss accurately reflects the expected loss of a solution β with intercept β_0 and $\beta_w \in S_0$. So, we choose an intercept $\beta_0 \in \mathbb{R}$ and $\beta_{S_0} \in S_0$ that solves the following optimization:

$$\begin{aligned} \min_{\beta_0 \in \mathbb{R}, \beta_{S_0} \in S_0} & \frac{1}{n} \sum_{i=1}^n \ell\left(\beta_0 \cdot y_i + \beta_{S_0}^t(P_{S_0}z_i)\right) \\ & = \min_{\beta_0, \beta_{S_0}} \frac{1}{n} \sum_{i=1}^n \ell\left(y_i \cdot \left(\beta_0 + \beta_{S_0}^t(P_{S_0}x_i)\right)\right) \end{aligned} \tag{9}$$

This is the usual parameter-fitting problem in classification but with projected features $P_{S_0}x_i$. We can use off-the-shelf solvers for logistic regression (for logistic loss) or linear SVM (hinge or squared hinge losses). For other convex loss functions, such as the modified Huber loss, we can use standard optimizers such as stochastic gradient descent.

Now, unlike PCR, we do not ignore $S_1 \cup S_2$. Suppose we set $\beta_w = \beta_{S_0} + \beta_{S_1 \cup S_2}$ for some vector $\beta_{S_1 \cup S_2} \in S_1 \cup S_2$. Then, the expected loss $E[\ell(\beta_0 \cdot y + \beta_w^t z)]$ equals $E[\ell(\beta_0 \cdot y + \beta_{S_0}^t(P_{S_0}z) + \beta_{S_1 \cup S_2}^t(P_{S_1 \cup S_2}z))]$. If we change $\beta_{S_1 \cup S_2}$, it affects the third term but not the first two. Setting $\beta_{S_1 \cup S_2} = 0$ corresponds to dimensionality reduction, because we only use the top principal components in S_0 . But a careful choice of $\beta_{S_1 \cup S_2}$ can reduce the loss obtained from β_{S_0} alone. But we cannot just project the data on to $S_1 \cup S_2$ and pick the $\beta_{S_1 \cup S_2}$ that minimizes training loss. This is because we cannot reliably estimate the loss function in this subspace. Instead, we need a $\beta_{S_1 \cup S_2}$ that is robust to estimation errors. To avoid being too conservative, we still need to use all available information about $S_1 \cup S_2$ (Eqs. 5–8). We formulate this as a robust optimization problem, which we discuss next.

2.4 Robust formulation

To select a robust $\beta_{S_1 \cup S_2}$, we need to characterize the distribution of the data projected on to $S_1 \cup S_2$. The empirical distribution is unreliable here. Instead, we will construct distributions that are “worst-case”, in that they have the maximum uncertainty subject to the constraints in Eqs. 5–8. Then, we pick the $\beta_{S_1 \cup S_2}$ with the best worst-case performance. This prevents $\beta_{S_1 \cup S_2}$ from overfitting to incidental aspects of the empirical distribution, while still using all reliable information about moments.

We note that our worse-case distribution depends on the data, but not on the weight vector $\beta_{S_1 \cup S_2}$. An alternative notion of robustness is to let the worst-case distribution depend on $\beta_{S_1 \cup S_2}$ as well. This corresponds to setting, for each possible choice of $\beta_{S_1 \cup S_2}$, the worst possible higher-order moments of the data distribution. Since our only constraints are on the mean and covariance, setting all other moments to their worst-case values is overly conservative. Fixing the data distribution to the maximum-uncertainty distribution helps us achieve robustness in a more practical way.

We will formulate our robust model assuming that Eqs. 5–7 are equalities. The first moments of the distribution of $P_{S_1 \cup S_2} \mathbf{z}$ can be taken to be the first moments of the empirical distribution. For the second moments of this distribution, we have only partial information. Let $V_{S_1 \cup S_2}$ be a matrix whose columns are the eigenvectors that span S_1 and S_2 . Writing the second-moment matrix of $P_{S_1 \cup S_2} \mathbf{z}$ in the basis $V_{S_1 \cup S_2}$, we get a block-wise form:

$$\begin{aligned}
 & V_{S_1 \cup S_2}^T E \left[(P_{S_1 \cup S_2} \mathbf{z})(P_{S_1 \cup S_2} \mathbf{z})^T \right] V_{S_1 \cup S_2} \\
 &= \begin{bmatrix} V_{S_1}^T E \left[(P_{S_1} \mathbf{z})(P_{S_1} \mathbf{z})^T \right] V_{S_1} & V_{S_1}^T E \left[(P_{S_1} \mathbf{z})(P_{S_2} \mathbf{z})^T \right] V_{S_2} \\ V_{S_2}^T E \left[(P_{S_2} \mathbf{z})(P_{S_1} \mathbf{z})^T \right] V_{S_1} & V_{S_2}^T E \left[(P_{S_2} \mathbf{z})(P_{S_2} \mathbf{z})^T \right] V_{S_2} \end{bmatrix} =: \begin{bmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{bmatrix} \quad (10)
 \end{aligned}$$

Now,

$$\begin{aligned}
 B_{11} &= V_{S_1}^T E \left[(P_{S_1} \mathbf{z})(P_{S_1} \mathbf{z})^T \right] V_{S_1} \\
 &= V_{S_1}^T \mathbb{P}_n \left[(P_{S_1} \mathbf{z})(P_{S_1} \mathbf{z})^T \right] V_{S_1} \\
 &= \mathbb{P}_n \left[(V_{S_1}^T \mathbf{z})(V_{S_1}^T \mathbf{z})^T \right],
 \end{aligned}$$

where the second equality follows from applications of Eq. 5, and the third equality follows from $P_{S_1} = V_{S_1} V_{S_1}^T$. So B_{11} is the second-moment matrix of the data projected on to V_{S_1} . Also, by Eq. 7,

$$B_{12} = B_{21} = 0. \quad (11)$$

For B_{22} , we have no estimates but only a bound (Eq. 8). This suggests the following uncertainty set for B_{22} :

$$B_{22} \in \mathcal{U} := \left\{ W \mid W \geq \frac{1}{n} \sum_{i=1}^n (V_{S_2}^T \mathbf{z}_i)(V_{S_2}^T \mathbf{z}_i)^T, \|W\| \leq \sigma_{bound} \right\}, \quad (12)$$

where $\sigma_{bound} = \sigma_{\min} \left(\mathbb{P}_n \left[(P_{S_1} \mathbf{z})(P_{S_1} \mathbf{z})^T \right] \right)$. Note that by construction, this uncertainty set is non-empty. Equations 10–12 thus characterize the second moments of the distribution of $P_{S_1 \cup S_2} \mathbf{z}$.

Now, we construct our worst-case distribution for $P_{S_1 \cup S_2} \mathbf{z}$; call it $q(P_{S_1 \cup S_2} \mathbf{z})$. We choose $q(P_{S_1 \cup S_2} \mathbf{z})$ to be the distribution with the maximum entropy (and hence the most ‘‘uncertainty’’) subject to the first and second moments specified above. It is well known that the maximum entropy is achieved by the exponential family distribution with those moments (Cover and Thomas 2006). Now, we pick $\beta_{S_1 \cup S_2}$ that performs best under $q(P_{S_1 \cup S_2} \mathbf{z})$:

$$\beta_{S_1 \cup S_2} = \arg \min_{b \in S_1 \cup S_2} \max_{B_{22} \in \mathcal{U}} \frac{1}{n} \sum_{i=1}^n E_{r \sim q(\cdot)} \left[\ell \left(\beta_0 \cdot y_i + \beta_{S_0}^T (P_{S_0} \mathbf{z}_i) + \mathbf{b}^T \mathbf{r} \right) \right], \quad (13)$$

where \mathbf{r} is a random variable that represents $P_{S_1 \cup S_2} \mathbf{z}$, and β_0 and β_{S_0} are the solutions of Eq. 9. Note that $q(\cdot)$ depends on B_{22} .

2.5 The solution of the robust objective

The solution $\beta_{S_1 \cup S_2}$ of Eq. 13 depends not only on the distribution $q(P_{S_1 \cup S_2} \mathbf{z})$ but also on β_0 , β_{S_0} , and $P_{S_0} \mathbf{z}_i$. This suggests that solving the robust optimization might be difficult. However, we show a surprising result. While the scale of $\beta_{S_1 \cup S_2}$ indeed depends on all the above factors, the direction of $\beta_{S_1 \cup S_2}$ does not. In fact, in many cases, **the direction does not even depend on the specific loss function.**

Theorem 1 (Direction of the Robust solution) *Suppose the loss function $\ell(\cdot)$ is non-negative, monotonically non-increasing, convex, differentiable, and the absolute value of its first derivative $|\ell'(\cdot)|$ has finite non-zero expectation under the standard Normal distribution. Then, any solution of Eq. 13 satisfies*

$$\begin{aligned} \beta_{S_1 \cup S_2} &= c \cdot V_{S_1 \cup S_2} \Sigma^{-1} \mu, \\ \text{where } \Sigma &= \begin{bmatrix} \frac{1}{n} \sum_{i=1}^n (V_{S_1}^T \mathbf{z}_i)(V_{S_1}^T \mathbf{z}_i)^T & 0 \\ 0 & \sigma_{bound} \cdot I \end{bmatrix}, \\ \mu &= \frac{1}{n} \sum_{i=1}^n V_{S_1 \cup S_2}^T \mathbf{z}_i, \end{aligned} \tag{14}$$

for some scalar c . Further, if there is a sequence of loss functions $\ell^{(m)}(\cdot)$ satisfying the properties mentioned above such that $\lim_{m \rightarrow \infty} \sup_{x \in \mathbb{R}} |\ell^{(m)}(x) - \ell(x)| = 0$, then there is a solution of the form of Eq. 14 that is arbitrarily close to the optimal.

Corollary 1 (Wide applicability) *The minimizer $\beta_{S_1 \cup S_2}$ of Eq. 13 has the form of Eq. 14 for logistic, hinge, squared hinge, and modified Huber losses.*

These results are significant from both a theoretical and practical standpoint. It is challenging to formulate tractable robust optimizations. Uncertainty sets are often chosen for their ease of analysis. So, it is encouraging to see a simple closed-form structure emerge from a well-motivated formulation. Further, we do not need separate analyses for each loss function. Armed with Theorem 1, we only need to pick a single scalar, which is the magnitude $\|\beta_{S_1 \cup S_2}\|$. We will choose this by cross-validation.

Computing the direction of $\beta_{S_1 \cup S_2}$ is also easy because Σ is a diagonal matrix. To see this, let X be the matrix with \mathbf{x}_i as its i^{th} row, and let $X = UDV^T$ be its singular value decomposition (SVD). The SVD of X is related to the eigenvectors and eigenvalues of \hat{M} by the formula $n \cdot \hat{M} = \sum_i \mathbf{z}_i \mathbf{z}_i^T = \sum_i \mathbf{x}_i \mathbf{x}_i^T = VD^2V^T$. For $i \in \{0, 1, 2\}$, let D_{S_i} be the diagonal matrix of singular values corresponding to the eigenvectors in V_{S_i} . Then, the top-left block of Σ equals $\sum_i (V_{S_1}^T \mathbf{z}_i)(V_{S_1}^T \mathbf{z}_i)^T / n = D_{S_1}^2 / n$. So, Σ is a diagonal matrix with entries $D_{S_1}^2 / n$ and σ_{bound} . Since $\sigma_{bound} = \min(D_{S_1}^2 / n) > \max(D_{S_2}^2 / n)$ from Eq. 8, we may write $\Sigma = \text{diag}(\max(D_{S_1 \cup S_2}^2 / n, \sigma_{bound}))$. We propose using a smooth upper-bound of this: $\Sigma_{smooth} = \text{diag}(D_{S_1 \cup S_2}^2 / n + \sigma_{bound})$. By varying σ_{bound} , we get smooth transitions between different choices for S_1 and S_2 . Using Σ_{smooth} also reveals a curious connection between our robust solution and ridge regression.

Theorem 2 (Connection to ridge regression) *The robust solution $\beta_{S_1 \cup S_2}$ using Σ_{smooth} is also the solution, up to a scaling factor, for regressing y_i on $P_{S_1 \cup S_2} \mathbf{x}_i$ with a ridge penalty:*

$$\beta_{S_1 \cup S_2} \propto \arg \min_b \sum_{i=1}^n (y_i - \mathbf{b}^T P_{S_1 \cup S_2} \mathbf{x}_i)^2 + n\sigma_{bound} \cdot \|\mathbf{b}\|^2.$$

Thus, we can view RoLIN as a mix of standard classification over the well-estimated top principal components and ridge regression over the poorly-estimated orthogonal subspace.

3 Algorithm and Robust cross-validation

Algorithm 1 Calculate β for RoLIN.

```

1: function CALCBETA( $\{\mathbf{x}_i \in \mathbb{R}^p, y_i \in \{1, -1\} \mid i = 1, \dots, n\}, k, \sigma_{ratio}, b_{max}$ )
2:    $\mathbf{z}_i \leftarrow y_i \cdot \mathbf{x}_i$ 
3:    $Z \leftarrow n \times p$  matrix whose  $i^{th}$  row is  $\mathbf{z}_i$ 
4:    $U, D, V \leftarrow \text{SVD}(Z)$   $\triangleright Z = UDV^T, D_{i,i}$  in descending order
5:    $V_{S_0} \leftarrow$  first  $k$  columns of  $V$   $\triangleright$  Basis vectors for subspace  $S_0$ 
6:    $V_{S_1 \cup S_2} \leftarrow$  last  $p - k$  columns of  $V$ 
7:    $D_{S_1 \cup S_2} \leftarrow$  diagonal matrix with entries  $D_{k+1, k+1}, D_{k+2, k+2}, \dots, D_{p, p}$ 
8:    $\beta_0, \gamma \leftarrow \arg \min_{\beta_0 \in \mathbb{R}, \gamma \in \mathbb{R}^k} \sum_{i=1}^n \ell(\beta_0 \cdot y_i + \gamma^T V_{S_0}^T \mathbf{z}_i)$ 
9:    $\beta_{S_0} \leftarrow V_{S_0} \gamma$ 
10:   $\sigma_{bound} \leftarrow \sigma_{ratio} \cdot \max(D_{S_1 \cup S_2}^2)$ 
11:   $\nu \leftarrow V_{S_1 \cup S_2} (D_{S_1 \cup S_2}^2 + \sigma_{bound} \cdot I)^{-1} V_{S_1 \cup S_2}^T Z^T \mathbf{1}$ 
12:  if  $\|\nu\| \neq 0$  then
13:     $\eta_{S_1 \cup S_2} \leftarrow V_{S_1 \cup S_2} \nu / \|\nu\|$ 
14:     $\|\beta_{S_1 \cup S_2}\| \leftarrow \arg \min_{c \in [0, b_{max}]} \sum_{i=1}^n \ell(\beta_0 \cdot y_i + (\beta_{S_0} + c \cdot \eta_{S_1 \cup S_2})^T \mathbf{z}_i)$ 
15:     $\beta_{S_1 \cup S_2} \leftarrow \|\beta_{S_1 \cup S_2}\| \cdot \eta_{S_1 \cup S_2}$ 
16:  else
17:     $\beta_{S_1 \cup S_2} \leftarrow 0$ 
18:  end if
19:   $\beta \leftarrow$  intercept  $\beta_0$  and feature weights  $\beta_{S_0} + \beta_{S_1 \cup S_2}$ 
20:  return  $\beta$ 
21: end function

```

RoLIN combines two algorithms: the CALCBETA algorithm to calculate the solution β , and the ROBUSTCV algorithm to robustly select model parameters for CALCBETA. We now provide details for both these algorithms.

Calculation of the solution vector β Algorithm 1 shows the steps in calculating β . Apart from the data itself, it requires three inputs. This first input is the number k of top principal components that comprise the subspace S_0 . The second input is a parameter σ_{ratio} from which we construct σ_{bound} . Tuning σ_{ratio} allows for smooth transitions between S_1 and S_2 . Setting $\sigma_{ratio} = 0$ corresponds to setting $S_2 = \emptyset$, while a large σ_{ratio} corresponds to $S_1 = \emptyset$. Third, we need an upper bound b_{max} on the magnitude $\|\beta_{S_1 \cup S_2}\|$.

We first construct the matrix Z with rows $\mathbf{z}_i = y_i \cdot \mathbf{x}_i$. The singular value decomposition of Z gives the diagonal matrix D of singular values and the matrix V of singular vectors (step 4)¹. We form V_{S_0} from the first k singular vectors and $V_{S_1 \cup S_2}$ from the remaining singular vectors (steps 5, 6). V_{S_0} and $V_{S_1 \cup S_2}$ span the subspaces S_0 and $S_1 \cup S_2$, respectively. For

¹ The matrix Z has the same singular values/vectors as the matrix X with rows \mathbf{x}_i .

\mathcal{S}_0 , we compute the optimal intercept β_0 and weight vector $\beta_{\mathcal{S}_0} \in \mathcal{S}_0$ via Eq. 9 (steps 8, 9). As discussed in Sect. 2.3, this step can use any convex minimizer. Then, for $\mathcal{S}_1 \cup \mathcal{S}_2$, we compute the direction vector $\eta_{\mathcal{S}_1 \cup \mathcal{S}_2}$ using Theorem 1 (steps 11–13). Here, we use Σ_{smooth} with $\sigma_{bound} = \sigma_{ratio} * \max(D_{\mathcal{S}_1 \cup \mathcal{S}_2}^2)$, where $D_{\mathcal{S}_1 \cup \mathcal{S}_2}$ contains singular values corresponding to $V_{\mathcal{S}_1 \cup \mathcal{S}_2}$. Finally, we choose the best magnitude of $\beta_{\mathcal{S}_1 \cup \mathcal{S}_2}$ over the training samples, but under the bound $\|\beta_{\mathcal{S}_1 \cup \mathcal{S}_2}\| \leq b_{max}$ (step 14). Bounded norm solutions have small generalization error (Sect. 5), so this is appropriate for the poorly-estimated subspace $\mathcal{S}_1 \cup \mathcal{S}_2$. Note that the question of the “right” norm does not arise. We must bound the L_2 -norm $\|\beta_{\mathcal{S}_1 \cup \mathcal{S}_2}\|$ since we already know the direction of the vector $\beta_{\mathcal{S}_1 \cup \mathcal{S}_2}$.

Algorithm 2 Robust Cross-validation for ROLIN.

```

1: function ROBUSTCV(dataset  $\mathcal{D}$ ,  $\Theta_{ratio}$ ,  $\Theta_{slack}$ ,  $\Theta_{gain}$ )
2:   Split  $\mathcal{D}$  into multiple train/holdout splits  $\{(\mathcal{D}_j^{tr}, \mathcal{D}_j^{ho}) \mid j = 1, \dots, J\}$ 
3:    $C_\sigma, C_b \leftarrow$  fixed set of choices for  $\sigma_{ratio}$  and  $b_{max}$  ▷ Initialization
4:    $k_{max} \leftarrow$  MAXRELIABLEPCS( $\Theta_{ratio}$ ) ▷ Find the maximum number of reliable PCs
5:    $\Psi_{\mathcal{S}_0} \leftarrow \{\psi = (k, 0, 0) \mid k \in \{1, \dots, k_{max}\}\}$ 
6:    $\psi_{\mathcal{S}_0}^{rob} \leftarrow$  ROBUSTPARAMS( $\Psi_{\mathcal{S}_0}$ ,  $\Theta_{slack}$ ) ▷ Robust solution using only top PCs
7:    $\Psi \leftarrow \{\psi = (k, \sigma_{ratio}, b_{max}) \mid k \in \{1, \dots, k_{max}\}, \sigma_{ratio} \in C_\sigma, b_{max} \in C_b\}$ 
8:    $\psi^{rob} \leftarrow$  ROBUSTPARAMS( $\Psi$ ,  $\Theta_{slack}$ ) ▷ Robust solution for general setting
9:    $\psi^{best} \leftarrow \psi_{\mathcal{S}_0}^{rob}$  if  $\text{cost}(\psi^{rob}) \geq (1 - \Theta_{gain}) \text{cost}(\psi_{\mathcal{S}_0}^{rob})$  else  $\psi^{rob}$ 
10:  return  $\psi^{best}$ 
11: end function

12: function MAXRELIABLEPCS( $\Theta_{ratio}$ )
13:  for all  $k \in \{1, \dots, p\}$  do
14:     $\psi \leftarrow (k, \sigma_{ratio} = 0, b_{max} = 0)$ 
15:     $\text{loss}_{ratio}(\psi), \text{cost}(\psi) \leftarrow$  CALCCOST( $\{(\mathcal{D}_j^{tr}, \mathcal{D}_j^{ho})\}, \psi)$ 
16:  end for
17:   $k_{max} \leftarrow \max(\{k \mid \text{loss}_{ratio}(k, 0, 0) \leq \Theta_{ratio} \quad \forall m \leq k\})$ 
18:  return  $k_{max}$ 
19: end function

20: function ROBUSTPARAMS( $\Psi$ ,  $\Theta_{slack}$ )
21:  for all  $\psi \in \Psi$  do
22:     $\text{loss}_{ratio}(\psi), \text{cost}(\psi) \leftarrow$  CALCCOST( $\{(\mathcal{D}_j^{tr}, \mathcal{D}_j^{ho})\}, \psi)$ 
23:  end for
24:   $\psi^* \leftarrow \arg \min_{\psi \in \Psi} \text{cost}(\psi)$ 
25:   $\Psi_{slack} \leftarrow \{\psi \in \Psi \mid \text{cost}(\psi) \leq (1 + \Theta_{slack}) \text{cost}(\psi^*)\}$  ▷ Nearly-min-cost candidates
26:   $\psi^{rob} \leftarrow \arg \min_{\psi \in \Psi_{slack}} (\text{cost}(\psi) + \text{loss}_{max}(\psi))$  ▷ Most robust candidate
27:  return  $\psi^{rob}$ 
28: end function

29: function CALCCOST(data splits  $\{(\mathcal{D}_j^{tr}, \mathcal{D}_j^{ho}) \mid j = 1, \dots, J\}, \psi = (k, \sigma_{ratio}, b_{max})$ )
30:  for all  $(\mathcal{D}_j^{tr}, \mathcal{D}_j^{ho})$  do
31:     $\beta \leftarrow$  CalcBeta( $\mathcal{D}_j^{(tr)}$ ,  $\psi$ )
32:     $\text{loss}_j^{tr}, \text{loss}_j^{ho} \leftarrow$  training and holdout loss using  $\beta$ 
33:  end for
34:   $\text{loss}_{avg} \leftarrow \frac{1}{J} \sum_{j=1}^J \text{loss}_j^{ho}$ 
35:   $\text{loss}_{max} \leftarrow \max_j \text{loss}_j^{ho}$ 
36:   $\text{loss}_{ratio} \leftarrow \frac{1}{J} \sum_{j=1}^J \frac{\text{loss}_j^{ho}}{\text{loss}_j^{tr}}$ 
37:   $\text{cost} \leftarrow \text{loss}_{avg}$  if  $\text{loss}_{ratio} \leq \Theta_{ratio}$  else  $\text{loss}_{max}$  ▷ Robust cost
38:  return  $\text{loss}_{ratio}, \text{cost}$ 
39: end function

```

Robust cross-validation for choosing model parameters. Now we need to select the input parameters $\psi = (k, \sigma_{ratio}, b_{max})$ for CALC BETA (Algorithm 1). A poor ψ leads to an overconfident classifier. But, there may only be a few holdout samples where an overconfident classifier incurs significant losses. The averaging step of cross-validation can hide these few large losses. To counter this, we develop a new robust cross-validation method called ROBUSTCV (Algorithm 2).

ROBUSTCV guards against overconfidence by using three signals. The first signal is the loss ratio, which we define as the ratio of the holdout loss to training loss, averaged over all cross-validation splits (step 36). A ψ with a high loss ratio indicates that overfitting is likely. To implement this idea, we use a loss ratio threshold Θ_{ratio} . Below the threshold, we use the average holdout loss as a measure of the cost of ψ , like standard cross-validation. But above the threshold, the cost of ψ is set to the maximum holdout loss (step 37). We also use the loss ratio to find an upper bound k_{max} for the number of principal components that are well-estimated (step 4). Throughout our algorithm, we restrict the parameter k in ψ to $k \leq k_{max}$.

The second warning sign of overconfidence is a significant difference between the average holdout loss and the maximum holdout loss. Standard cross-validation picks the ψ^* with the smallest cost. But, we find a robust parameter setting ψ^{rob} whose cost is within a factor $(1 + \Theta_{stack})$ of ψ^* , but whose worst-case holdout loss is better (steps 24–26).

Third, realizing that the distribution of $P_{S_1 \cup S_2} z$ is poorly estimated, we check if a solution constructed from S_0 alone is good enough. That is, we select a $\psi_{S_0}^{rob}$ that ignores the subspace $S_1 \cup S_2$ unless the cost improves by a factor of Θ_{gain} using ψ^{rob} (step 9). Together, these steps ensure that ROBUSTCV selects a ψ^{best} that is robust but not too conservative. RoLIN then runs CALC BETA under ψ^{best} over the entire training sample. The output of this is RoLIN's solution.

4 Experiments

We will first compare RoLIN against competing methods for 25 real-world datasets and three loss functions. Then, we will contrast RoLIN run with ROBUSTCV versus alternative cross-validation schemes. Finally, we will present a sensitivity analysis for RoLIN's parameters.

Datasets We use 25 benchmark real-world datasets from the UCI repository². These span many domains, and the number of features range for $p = 8$ to $p = 43,680$. We convert categorical variables into binary “dummy” variables, and count each dummy variable as a separate feature.

Evaluation methodology We run experiments with logistic loss, squared hinge loss, and modified Huber loss (Eq. 2). All three are standard loss functions, and the first two are widely used in logistic regression and linear support vector machines. Our focus is on the limited-data setting because this is where estimation errors are significant, and finding a good solution is difficult. So we vary the number of training samples from $n = 15$ to $n = 200$ for each dataset. For each experiment, we randomly choose n points as the training samples and the remainder as the test samples. Then, we compute the average test loss for RoLIN and all competing methods. We repeat this process 50 times. For each n , we report

² <https://archive.ics.uci.edu/ml/datasets>.

the trimmed mean of the losses, because it is robust to the occasional outlier. That is, we drop the five best and five worst test losses from the 50 repetitions, and calculate the average of the remaining test losses. We note that the mean losses have the same pattern as the trimmed means, with RoLin outperforming other methods by an even wider margin.

Competing methods The closest competitors to RoLin are norm-based regularization (using L_1 and L_2 norms), and dimensionality reduction using the top few principal components (Top PCs). For norm-based regularization, we use Python's `scikit-learn` implementations for all losses and regularizations. We select the regularization parameter via standard cross-validation. Note that we calculate cross-validation loss using the actual loss function which we want to minimize, and not zero-one loss as in common in practice. Using zero-one loss gives sub-optimal results (see Appendix B). For Top PCs, the number of principal components is chosen by cross-validation.

We also contrast RobustCV with two cross-validation methods. One is the standard CV, which picks the parameter with the best average holdout loss. The second method (CV-1-SD) considers all parameters whose loss is within one standard deviation of the best loss, and picks the parameter that achieves the most regularization (Hastie et al. 2009).

Implementation details. For RoLin³, we run RobustCV with $\Theta_{ratio} = 5$, $\Theta_{slack} = 0.1$, and $\Theta_{gain} = 0.05$. We split the training data using five instances of 5-fold cross-validation (De Brabanter et al. 2002). We vary b_{max} from 0.01 to 0.1 for $n = 15$. For larger n , we scale the upper range with \sqrt{n} . This allows more weight to be placed on the robust solution when more data is available. We vary σ_{ratio} from 1 to 10. We also use RobustCV to choose whether to normalize the features. We use the above settings for all experiments except for the sensitivity analysis in Sect. 4.3.

4.1 Accuracy of RoLin

The detailed plots of the performance of every method for each dataset and each loss function are shown in Fig. 10 in the appendix. In the following, we summarize our results and present our main observations.

RoLin outperforms the competing methods. Figure 2 shows the aggregate statistics comparing RoLin against competing methods over all 25 datasets. The top panel of Fig. 2 counts the number of datasets on which any given method achieves the best loss. We see that RoLin is the best in all settings except for logistic loss with 200 training samples. RoLin is particularly dominant for small training sizes, since this is when robustness to estimation error is most needed. For $n = 15$ training samples, RoLin is the best performer on at least 15 datasets, irrespective of the loss function. When 50 or fewer training samples are available, RoLin achieves the smallest loss on around 2x to 3x as many datasets as the next best method. RoLin also works very well for modified Huber Loss; it is best for 14 or more datasets for any training size.

We observe the same pattern when we compare the actual value of the loss. The bottom panel of Fig. 2 shows the loss incurred by each method compared against that of RoLin, averaged over all datasets. RoLin always has a better loss on average, for all loss functions and training sizes. The greatest difference is for the smallest training size $n = 15$, where the next best method is on average 14% – 40% worse than RoLin, depending on the loss

³ The code is available from <https://faculty.mcombs.utexas.edu/deepayan.chakrabarti/mywww/software/ROLIN-2021.tgz>.

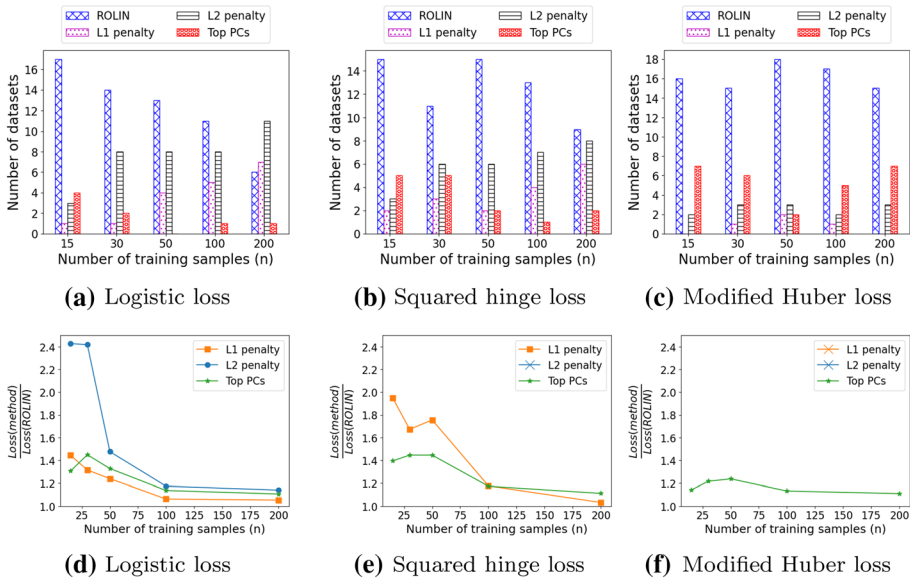


Fig. 2 Comparison of RoLIN against competing methods: The top row shows the number of datasets on which each method is best. RoLIN performs best for most losses and training sizes. The bottom row shows the ratio of the trimmed means of losses of each method against that of RoLIN, averaged over all 25 datasets. Again, RoLIN works best for all loss functions and training sizes. We do not show L_2 regularization for the squared hinge and modified Huber losses, and L_1 regularization for the modified Huber loss, because their average loss is too large

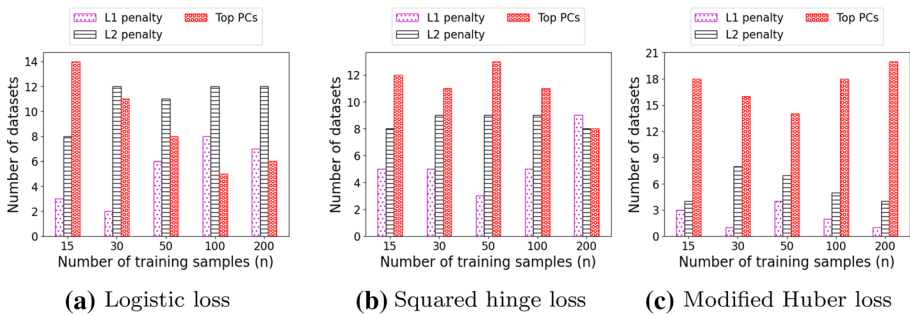


Fig. 3 Among the competitors of RoLIN, no one method is best: We compare the three competitors of RoLIN against each other (ignoring RoLIN). Top PCs works well for Modified Huber loss, but for other losses, there is no one method that works best

function. But even with $n = 200$ training samples, every method is worse on average than RoLIN. For the modified Huber loss, the average improvement of RoLIN over L_1 and L_2 regularization is too large to fit on the plot.

There is no clear second-best method among the competitors of RoLIN. In practice, we must choose a single method to apply to a dataset. Figure 2 shows that while RoLIN is best, there is no clear second-best method. In terms of the loss, Top PCs works well everywhere. However, Fig. 2d shows that L_1 regularization is better for most training sizes for




Fig. 4 RoLIN versus Top PCs: RoLIN is consistently better

logistic loss. Further, if we consider the instances where some method outperforms RoLIN, that method is often L_2 regularization (Figs. 2a and 2b. But the average loss for L_2 regularization can be much worse than the other methods (Fig. 2e, f). Figure 3 compares only the competitors of RoLIN to each other. Top PCs works well for the modified Huber loss, but for other losses, L_2 regularization is comparable or sometimes better. Thus, among the competitors of RoLIN, no single method dominates.

Robust optimization contributes significantly to RoLIN’s performance. The difference between Top PCs and RoLIN is that Top PCs ignores the bottom principal components, while RoLIN uses a robust optimization for them. Hence, the importance of robust optimization can be gauged from the difference between these two methods⁴.

The previous results show that Top PCs works reasonably well on all loss functions and training sizes. But the consistency of Top PCs comes at a cost: it rarely outperforms RoLIN (top panel of Fig. 2). Figure 4 shows the ratio of the loss of Top PCs against RoLIN on a log-scale. For every loss function, and for any training size, RoLIN is better than Top PCs on at least 75% of the datasets (shown by the bottom of the boxes being around one). Further, Top PCs can be up to 4x worse than RoLIN. Even with $n = 200$ training samples, Top PCs can still be 2x worse. This clearly demonstrates the need for the robust optimization step in RoLIN.

Norm-based regularization can occasionally have very large losses. For particular datasets and settings, RoLIN can be much better than both L_1 and L_2 norm-based regularizations. It can be up to 3x better under logistic loss, and 12x better under squared hinge loss. For modified Huber loss, no norm-based regularization yields a reasonable classifier for the Credit and Gas sensor datasets (Fig. 10(12) and 10(18)). Indeed, for several datasets, the classifiers obtained from norm-based regularization have such poor test loss that they do not appear in the plots for Fig. 10.

To illustrate this, Fig. 5 compares RoLIN against norm-based regularization for three specific datasets. Note that the y-axis is on a log scale, and we report trimmed means which remove outliers. Plot (a) shows an instance where RoLIN with $n = 15$ training samples is better than both norm-based methods with $n = 1500$ samples. Plot (c) shows a similar situation. In plot (b), the losses for norm-based regularization become much worse when training size is reduced. This cannot be due to occasional outliers, because the trimmed mean ignores the worst five test losses. Further, L_1 -regularization in plot (c) is not close to convergence even with $n = 1500$. These examples highlight the perils of choosing the “wrong” norm. RoLIN sidesteps this issue entirely.

⁴ Top PCs uses standard CV while RoLIN uses ROBUSTCV. Changing Top PCs to use ROBUSTCV only increases the gap in performance between Top PCs and RoLIN.

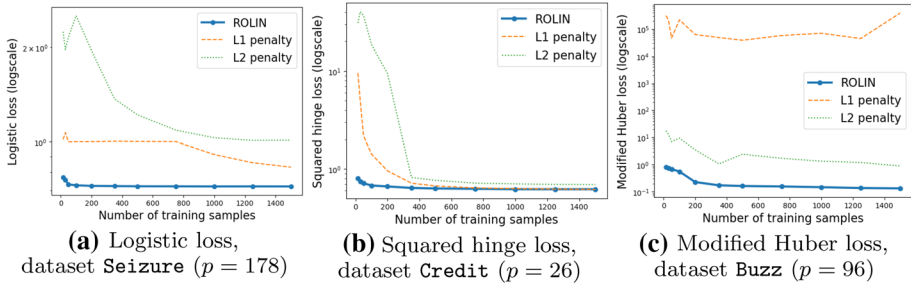


Fig. 5 Test loss comparison on three example datasets: RoLIN is compared against L_1 and L_2 regularizations for three loss functions, over a wide range of training sizes n . For each n , we report the trimmed mean of the test losses over 50 repeated experiments. Note that the y-axis is plotted on a log-scale




Fig. 6 Importance of ROBUSTCV: We plot the loss when RoLIN is run with CV and CV-1-SD, versus ROBUSTCV. ROBUSTCV is much better for small training sizes

Finally, we note that RoLIN performs as well or better than competing methods for zero-one loss (or, misclassification rate). Since our focus is on convex losses, we defer these results to Appendix B.

4.2 Importance of ROBUSTCV

To examine the influence of ROBUSTCV, we run RoLIN with standard cross-validation (CV) and another common variant (CV-1-SD). For each dataset and loss function, we calculate the trimmed mean of test loss of RoLIN with CV and CV-1-SD, for $n = 15$ to $n = 200$. We compare these against the trimmed means using ROBUSTCV. Figure 6 show that CV is better than CV-1-SD for all losses. Between CV and ROBUSTCV, ROBUSTCV outperforms for small training sizes. For $n \leq 30$, CV is 5% – 25% worse on average than ROBUSTCV, depending on the loss function. The differences mostly disappear when more training samples are available.

These results show the usefulness of robustness in cross-validation for small sample sizes. In such scenarios, an overconfident classifier may correctly classify all but a few points, and only these few points provide any warning about the unsuitability of the classifier. ROBUSTCV is designed to look for these warning signals and hence can avoid overconfident classifiers. Standard CV averages over all holdout sets, and this attenuates or even hides the warning signs.




Fig. 7 Sensitivity to the parameters of ROBUStCV: We plot the relative difference in trimmed means for logistic loss when the parameters $(\Theta_{ratio}, \Theta_{slack}, \Theta_{gain})$ are varied from their default values of (5, 0.1, 0.05). Positive values imply larger losses. RO_{LIN} is seen to be robust to a wide range of parameter choices

4.3 Sensitivity analysis

Recall that ROBUStCV requires three parameters. The first is Θ_{ratio} , which is the threshold ratio of holdout to training loss above which we distrust the average holdout loss. The second is Θ_{slack} , which is the importance we assign to the maximum holdout loss versus the average holdout loss. The third is Θ_{gain} , which characterizes our preference for solutions constructed only from the top principal components. Now, we vary these parameters one at a time from their default values and report results on ten datasets.

Figure 7 shows the relative increase in the trimmed means under logistic loss for different values of these parameters. Plot (a) shows that for Θ_{ratio} , any value in the range $\Theta_{ratio} \in [2.5, 100]$ works well (the default is 5). Larger values of Θ_{ratio} mean that we ignore instances where training loss is much smaller than holdout loss, which is a clear sign of overfitting. Smaller values mean that we always use the maximum holdout loss instead of the average holdout loss. Always focusing on maximum loss is too conservative, so it performs poorly for our expected test loss objective.

Plot (b) shows that any choice of $\Theta_{slack} \leq 0.15$ yields similar results (the default is 0.1). Losses become worse for larger values of Θ_{slack} . A large Θ_{slack} means that we downplay the average holdout loss and focus on the maximum holdout loss. Like a small Θ_{ratio} , this is too conservative and does not work for the same reason.

Plot (c) shows that any $\Theta_{gain} \leq 0.05$ yields good results (the default is 0.05). Higher values imply a preference for solutions based on only the top few principal components, ignoring the robust solution from the remaining principal components. When $\Theta_{gain} \rightarrow \infty$, we get Top PCs. We see that high Θ_{gain} leads to a significant increase in the test loss, showing the importance of the robust component of RO_{LIN}.

Extreme values for any of these parameters correspond to either standard cross-validation or very conservative choices. The former is bad for small n , while the latter performs poorly for large n . But for a broad range of parameters, ROBUStCV achieves good results.

5 Prior work

A common approach to deal with limited data is regularization. Here, we add to the desired objective an extra term that penalizes large feature weights. This term is typically some L_q norm of the feature weight vector, with $q = 1$ and $q = 2$ being common choices:

$$\min_{\beta} \mathbb{P}_n \ell(y \cdot g_{\beta}(\mathbf{x})) + \lambda \cdot \|\beta_w\|_q. \quad (15)$$

There are several competing justifications of the regularization term in Eq. 15. Regularization can emerge from a prior, or as the solution of a robust optimization, or as a way to bound estimation errors. Next, we discuss these, and contrast them with RoLIN.

Prior. We can cast regularization as a prior on the parameter vector β . Then, the solution of Eq. 15 is the *maximum a posteriori* (MAP) estimate of β . For example, a zero-mean spherical Gaussian prior for β_w gives L_2 regularization, while a Laplace prior yields L_1 regularization. But one may construct a prior for any L_q -norm, or any Mahalanobis distance metric. Choosing the best prior for a dataset is difficult, but it matters a lot, as we showed in Sect. 4. RoLIN does not assume a prior, so it sidesteps this difficulty entirely.

Priors are also useful for dealing with corrupted data (Kordzakhia et al. 2001; Feng et al. 2014; Tibshirani and Manning 2014). Further, L_1 priors induce sparsity in the solution, which makes the model easier to interpret (Tibshirani 1996). We do not consider data corruption or interpretability in this paper.

Robust optimization. Many optimization problems have parameters or constraints that must be learned from data. Robust optimization methods protect against corrupted data, outliers, and incorrect assumptions (Ben-Tal et al. 2009). These methods first construct uncertainty sets that reflect the ambiguity in the data. Then, they optimize a worst-case objective over the uncertainty set. For some uncertainty sets, this worst-case objective matches norm-based regularization.

Robust optimization methods typically fall into two groups. Methods in the first group assume that the training samples are perturbed. The perturbation could be because of uncertain or missing data (Trafalis and Gilbert 2006; Gao Huang et al. 2012; Wang and Pardalos 2014; Tzelepis et al. 2018), adversarial opponents (Globerson and Roweis 2006), or different training and test distributions (Bi and Zhang 2004). Robustness to perturbations is also equivalent to robustness under chance constraints (Bhattacharyya 2004; Shivaswamy et al. 2006). To achieve robustness, we assume that the “true” data fall inside an uncertainty set constructed from the “perturbed” data. Standard uncertainty sets impose a bound on some norm of the perturbation. Choosing a particular norm gives a corresponding norm-based regularization (El Ghaoui and Le Bret 1997; Xu et al. 2009a, b).

The second group of robust optimization methods constructs uncertainty sets of probability distributions. They assume that the true distribution of (y, \mathbf{x}) lies in this uncertainty set and optimize for the worst-case distribution in this set. Delage and Ye (2010); Goh and Sim (2010); Wiesemann et al. (2014) consider distributions with appropriately bounded moments. Others choose distributions within a bounded distance from the empirical distribution. The distance can be the Prohorov metric (Erdoğ an and Iyengar 2006), KL-divergence (Jiang and Guan 2016), or Wasserstein distance (Wozabal 2012; Shafieezadeh-Abadeh et al. 2015, 2017; Mohajerin Esfahani and Kuhn 2018). For Wasserstein distance, the user must also choose a distance metric in feature space. Choosing a distance metric based on some norm yields a regularization using that norm.

Thus, both types of robust optimization approaches rely on the user to choose a distance metric or a norm. This choice determines the form of the regularization term. The “best” choice for a dataset is unclear. Further, robust optimization emphasizes worst-case performance. This can make robust algorithms too conservative for our average-loss objective. In contrast, RoLIN does not require any user inputs. Also, RoLIN restricts robust optimization to just the bottom principal components, which are much noisier than the top components. This protects RoLIN from becoming too conservative.

Estimation error bounds. Regularization also ensures that the training loss is close to the expected loss. Let $\Delta_\beta := E\ell(y \cdot g_\beta(\mathbf{x})) - \mathbb{P}_n\ell(y \cdot g_\beta(\mathbf{x}))$ be the difference between the expected and training losses. For large training sizes n , this is small for any β . But, for small n , Δ_β can be much greater than zero for some values of β . However, if β has a small norm, Δ_β can be upper-bounded. For example, under zero-one loss, if $\|\beta\|_1 \leq 1/\rho$ (bounded L_1 norm), then Δ_β decays with ρ and the square-root of n (see Mohri et al. 2018). Similar results hold when β has a small L_2 norm. Now, regularization biases the objective of Eq. 15 towards a β with a small norm. This ensures that the expected loss of the solution is comparable to its training loss. Hence, regularization avoids overfitting.

Still, bounding Δ_β is not enough. Our aim is a β with small expected loss, not a small Δ_β . Further, such bounds may hold for many norms or Mahalanobis distances. Choosing the best norm for a dataset is difficult. For small n , the bounds on Δ_β can be loose⁵. So, we cannot just pick the norm with the best bound. Finally, while a β with a small norm may have a small Δ_β , the converse need not be true. There may be other solutions that have a small Δ_β and also a low loss.

All the above justifications for regularization need the user to choose a norm, or a prior, or a distance metric. The right choice depends on the dataset, the training size, and the loss function. This choice is challenging but also crucial because the wrong choice can significantly hurt performance. In contrast, RoLin needs no user input and does not force the solution to have a small norm. This suggests that explicit norm-based regularization of the form of Eq. 15 is unnecessary.

6 Conclusions

Our goal is to build a linear classifier with two properties. First, it should optimize for general loss functions, instead of the usual zero-one loss. This can be interpreted as accurately predicting class probabilities and not just the binary class labels. Second, its accuracy should gracefully degrade with smaller training sample sizes. The usual approach is to do dimensionality reduction via principal components, or to add to the loss function a regularization term based on a norm chosen by the user. But dimensionality reduction loses data, while regularization is sensitive to the choice of norm. Our proposed method, called RoLin, overcomes these flaws. Unlike dimensionality reduction, it does not ignore the bottom principal components. Unlike regularization, RoLin is entirely automatic and needs no user input. Further, it works well with many loss functions.

RoLin first projects the data on to its top principal components and minimizes training loss on the projected data. The resulting classifier does not overfit because the top principal components are stable. But this classifier ignores the subspace orthogonal to the top principal components. We cannot minimize training loss in this subspace, because estimates of loss are unreliable. So RoLin constructs a robust classifier here. Finally, RoLin combines the two classifiers to get the benefits of both.

To select the parameters of RoLin, we develop a new robust cross-validation algorithm called RobustCV. This checks for several warning signs of overfitting missed by standard cross-validation. RobustCV helps RoLin work well even with small training sizes.

⁵ For instance, if the bound needs to hold with probability greater than 0.95, then the error term in the bound is at least 0.13 for $n \leq 100$ (see Mohri et al. 2018, Thms. 13.3 and 13.4). For comparison, the zero-one loss of even a baseline classifier is at most 0.5.

Experiments on 25 real-world datasets and three loss functions show that RoLIN outperforms existing state of the art methods. RoLIN does particularly well for small training sizes. For $n = 15$ training samples, RoLIN has 14% – 40% lower loss on average than the next-best competitor, under all problem settings. When 50 or fewer training samples are available, RoLIN achieves the smallest loss on around 2x to 3x as many datasets as the next best method. For the modified Huber loss, RoLIN dominates other methods for all training sizes. Further, among the competitors of RoLIN, no single method is best. Norm-based regularization is close to RoLIN for logistic regression, but is not comparable for modified Huber loss. On some datasets, RoLIN achieves with $n = 15$ samples an accuracy that regularization fails to reach with $n = 1500$ samples. Dimensionality reduction via the top principal components rarely outperforms RoLIN, especially for logistic loss. Finally, the best norm for regularization depends on the dataset, training size, and loss function. So, for a new problem setting, picking the right norm is difficult. In contrast, RoLIN works well for all datasets and settings.

There are several ways to extend RoLIN. We can try to use RoLIN for non-linear classification via the kernel trick. Here, each test point \mathbf{x} is classified based on a linear combination of $K(\mathbf{x}, \mathbf{x}_i)$ where \mathbf{x}_i is a training point and $K(\cdot, \cdot)$ is a kernel function. This suggests that we can use RoLIN on the kernel matrix instead of the features matrix. We can also use RoLIN for multiclass classification via one-versus-the-rest binary classification. Finally, we note that RoLIN does not handle outliers or different training and test distributions. The top principal components of the training and test distributions may not be similar in this setting. One possibility is to project RoLIN's solution on to the set of small-norm solutions, which may be more robust under outliers.

Appendices

A Proofs

Proof (Theorem 1) Consider the case of a loss function $\ell(\cdot)$ that satisfies the properties stated in the theorem. Choose any feasible $\mathbf{b} \in \mathcal{S}_1 \cup \mathcal{S}_2$, and let $\tilde{\mathbf{b}} = V_{\mathcal{S}_1 \cup \mathcal{S}_2}^T \mathbf{b}$. Choose any $B_{22} \in \mathcal{U}$. This fixes the second moment of $P_{\mathcal{S}_1 \cup \mathcal{S}_2} \mathbf{z}$ as $V_{\mathcal{S}_1 \cup \mathcal{S}_2} \Sigma_{B_{22}} V_{\mathcal{S}_1 \cup \mathcal{S}_2}^T$, where

$$\Sigma_{B_{22}} =: \begin{bmatrix} B_{11} & 0 \\ 0 & B_{22} \end{bmatrix},$$

with $B_{11} = \mathbb{P}_n \left[\left(V_{\mathcal{S}_1}^T \mathbf{z} \right) \left(V_{\mathcal{S}_1}^T \mathbf{z} \right)^T \right]$. Also, from Eq. 6, the first moment is fixed at $(1/n) \sum_i P_{\mathcal{S}_1 \cup \mathcal{S}_2} \mathbf{z}_i = V_{\mathcal{S}_1 \cup \mathcal{S}_2} \boldsymbol{\mu}$. Now, the maximum entropy distribution with given first and second moments is the Gaussian distribution with those moments:

$$q(\cdot) = \mathcal{N} \left(V_{\mathcal{S}_1 \cup \mathcal{S}_2} \boldsymbol{\mu}, V_{\mathcal{S}_1 \cup \mathcal{S}_2} \left(\Sigma_{B_{22}} - \boldsymbol{\mu} \boldsymbol{\mu}^T \right) V_{\mathcal{S}_1 \cup \mathcal{S}_2}^T \right).$$

Since $\mathbf{r} \sim q(\cdot)$, for any training sample i , we have

$$\beta_0 \cdot y_i + \beta_{S_0}^T (P_{S_0} z_i) + \mathbf{b}^T \mathbf{r} \sim \mathcal{N}(m_i, s^2), \tag{16}$$

$$\text{where } m_i = \beta_0 \cdot y_i + \beta_{S_0}^T (P_{S_0} z_i) + \tilde{\mathbf{b}}^T \boldsymbol{\mu} \tag{17}$$

$$s = \tilde{\mathbf{b}}^T (\Sigma_{B_{22}} - \boldsymbol{\mu} \boldsymbol{\mu}^T) \tilde{\mathbf{b}}. \tag{18}$$

Hence, the expected loss under $q(\cdot)$ is

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n E_{r \sim q(\cdot)} \left[\ell \left(\beta_0 \cdot y_i + \beta_{S_0}^T (P_{S_0} z_i) + \mathbf{b}^T \mathbf{r} \right) \right] &= \frac{1}{n} \sum_{i=1}^n E_{w \sim \mathcal{N}(m_i, s^2)} [\ell(w)] \\ &= \frac{1}{n} \sum_{i=1}^n E_{w \sim \mathcal{N}(0,1)} [\ell(m_i + w \cdot s)]. \end{aligned} \tag{19}$$

Taking the partial derivative with respect to s , we have

$$\begin{aligned} \frac{\partial}{\partial s} E_{w \sim \mathcal{N}(0,1)} [\ell(m_i + w \cdot s)] &= E[w \cdot \ell'(m_i + w \cdot s)] \\ &= \int_0^\infty w [\ell'(m_i + w \cdot s) - \ell'(m_i - w \cdot s)] \cdot \phi(w) dw \\ &> 0, \end{aligned}$$

where the interchange of differentiation and integration in the first equality follows from the dominated convergence theorem since $|\ell'(\cdot)|$ has finite expectation, and the last inequality is because $\ell'(\cdot)$ is monotonically non-decreasing (due to convexity) and not zero everywhere (because $|\ell'(\cdot)|$ has non-zero expectation). Thus, for any \mathbf{b} , the worst case expected loss is achieved when $s = \tilde{\mathbf{b}}^T (\Sigma_{B_{22}} - \boldsymbol{\mu} \boldsymbol{\mu}^T) \tilde{\mathbf{b}}$ is maximized. From the uncertainty set of Eq. 12, the maximum is achieved at $B_{22} = \sigma_{bound} \cdot I$. Hence, the worst-case is achieved with $\Sigma_{B_{22}} = \Sigma$, and $s = \tilde{\mathbf{b}}^T (\Sigma - \boldsymbol{\mu} \boldsymbol{\mu}^T) \tilde{\mathbf{b}}$, where Σ is defined in the theorem statement.

Next, taking the partial derivative with respect to m_i , we have

$$\frac{\partial}{\partial m_i} E_{w \sim \mathcal{N}(0,1)} [\ell(m_i + w \cdot s)] = E_{w \sim \mathcal{N}(0,1)} [\ell'(m_i + w \cdot s)] < 0,$$

where $\ell'(\cdot) \leq 0$ because $\ell(\cdot)$ is monotonically non-increasing, and $\ell'(x) < 0$ for some x because $|\ell'(\cdot)|$ has non-zero expectation. Now, observe that each m_i (for $i = 1 \dots n$) increases with $\tilde{\mathbf{b}}^T \boldsymbol{\mu}$. Hence, if Eq. 13 achieves its optimal at $s = s^*$, then the optimal solution $\beta_{S_1 \cup S_2}$ is of the form $\beta_{S_1 \cup S_2} = V_{S_1 \cup S_2} \tilde{\mathbf{b}}$, where $\tilde{\mathbf{b}}$ solves

$$\begin{aligned} &\text{maximize } \tilde{\mathbf{b}}^T \boldsymbol{\mu} \\ &\text{subject to } \tilde{\mathbf{b}}^T (\Sigma - \boldsymbol{\mu} \boldsymbol{\mu}^T) \tilde{\mathbf{b}} = s^*. \end{aligned}$$

It is easily shown that the solution must be of the form

$$\beta_{S_1 \cup S_2} = c \cdot V_{S_1 \cup S_2} (\Sigma - \boldsymbol{\mu} \boldsymbol{\mu}^T)^{-1} \boldsymbol{\mu}, \tag{20}$$

for some scalar c that depends on s^* . Finally, by the Sherman-Morrison formula,

$$(\Sigma - \mu\mu^T)^{-1} \mu = \frac{1}{1 - \mu^T \Sigma^{-1} \mu} \cdot \Sigma^{-1} \mu.$$

Combining this with Eq. 20 gives the statement of the theorem.

Now, consider a loss function $\ell(\cdot)$ that is the limit of a sequence of loss functions $\ell^{(m)}(\cdot)$ such that $\lim_{m \rightarrow \infty} \sup_{x \in \mathbb{R}} |\ell^{(m)}(x) - \ell(x)| = 0$. Then, for any $\epsilon > 0$, there exists an N such that for all $m > N$, $\sup_{x \in \mathbb{R}} |\ell^{(m)}(x) - \ell(x)| < \epsilon$. Define $h(\mathbf{b}) = 1/n \cdot \sum_i E[\ell(m_i + w \cdot s)]$, where m_i and s are functions of \mathbf{b} ; define $h^{(m)}(\mathbf{b})$ similarly. Then, for any \mathbf{b} and for large enough m ,

$$|h^{(m)}(\mathbf{b}) - h(\mathbf{b})| < \epsilon. \tag{21}$$

There is a sequence of minimizers $\mathbf{b}^{(m)}$ under $\ell^{(m)}$ such that $\mathbf{b}^{(m)} = c^{(m)} \cdot \Sigma^{-1} \mu$. Hence,

$$h^{(m)}(\mathbf{b}^{(m)}) \leq h^{(m)}(\beta_{S_1 \cup S_2}) \leq h^{(m)}(\mathbf{b}^{(m)}) + 2\epsilon,$$

where the first statement follows from the optimality of $\mathbf{b}^{(m)}$ under $\ell^{(m)}$, and the second from two applications of Eq. 21. So, the minimizers under $\ell^{(m)}$ provide arbitrarily good solutions under ℓ , for m large enough. \square

Proof (Corollary 1) The logistic loss $\ell(x) = \log(1 + \exp(-x))$ is clearly non-negative, monotonically decreasing, and convex. Also, $|\ell'(x)| = \exp(-x)/(1 + \exp(-x))$, so $0 < |\ell'(x)| < 1$, so it has finite non-zero expectation under the Gaussian. Hence, the conditions of Theorem 1 are satisfied.

The squared hinge loss $\ell(x) = (\max(0, 1 - x))^2$ is non-negative, monotonically decreasing, and convex. Also, $|\ell'(x)| = 2 \cdot \max(0, 1 - x) \leq 2(1 + |x|)$, so $0 < E|\ell'(x)| \leq 2 + 2\sqrt{2/\pi}$. Hence, Theorem 1 applies.

The hinge loss $\ell(x) = \max(0, 1 - x)$ is the limit of a sequence of functions $\ell^{(m)}(\cdot)$ indexed by a monotonically decreasing sequence $\alpha_m \rightarrow 0$:

$$\ell^{(m)}(x) = \begin{cases} \max(0, 1 - x) & \text{if } x \leq 1 - \alpha_m \\ \frac{(x - (1 + \alpha_m))^2}{4\alpha_m} & 1 - \alpha_m < x \leq 1 + \alpha_m \\ 0 & x > 1 + \alpha_m \end{cases}$$

Each $\ell^{(m)}(\cdot)$ is non-negative, monotonically decreasing, and convex. Also, $|\ell^{(m)'}(x)| \leq 1$ with equality for $x \leq 1 - \alpha_m$, so it has finite non-zero expectation. Again, Theorem 1 applies.

The modified Huber loss is similar to the functions $\ell^{(m)}(\cdot)$ above, for which the conditions of the Theorem are satisfied. \square

Proof (Theorem 2) Let \tilde{Z} be a matrix with $\tilde{z}_i := P_{S_1 \cup S_2} z_i$ in row i . Then, using Σ_{smooth} in Theorem 1,




Fig. 8 *Zero-one loss*: We compare the misclassification loss of RoLIN against other methods. Positive values imply that RoLIN is more accurate than the competing method. We pick the feature weights that minimize logistic loss but choose the hyperparameters by cross-validation to minimize zero-one loss. RoLIN outperforms L_1 regularization and Top-PCs, and is comparable to L_2 regularization

$$\begin{aligned}
 \beta_{S_1 \cup S_2} &\propto V_{S_1 \cup S_2} \Sigma_{smooth}^{-1} V_{S_1 \cup S_2}^T \left(\sum_i z_i / n \right) \\
 &= \left(V_{S_1 \cup S_2} \Sigma_{smooth} V_{S_1 \cup S_2}^T \right)^+ \left(\sum_i P_{S_1 \cup S_2} z_i / n \right) \\
 &= \left(V_{S_1 \cup S_2} D_{S_1 \cup S_2}^2 V_{S_1 \cup S_2}^T + n \sigma_{bound} \cdot V_{S_1 \cup S_2} V_{S_1 \cup S_2}^T \right)^+ (\tilde{Z}^T \mathbf{1}) \\
 &= \left(\tilde{Z}^T \tilde{Z} + n \sigma_{bound} \cdot V_{S_1 \cup S_2} V_{S_1 \cup S_2}^T \right)^+ (\tilde{Z}^T \mathbf{1}) \\
 &= (\tilde{Z}^T \tilde{Z} + n \sigma_{bound} \cdot I)^{-1} (\tilde{Z}^T \mathbf{1}),
 \end{aligned}
 \tag{22}$$

where M^+ is the Moore-Penrose pseudoinverse of matrix M , and $\mathbf{1}$ is a vector of all ones. The first equality follows from the observation that Σ_{smooth} is diagonal, and the fact that $P_{S_0} z$ is in the kernel of the pseudoinverse. The change from $V_{S_1 \cup S_2} V_{S_1 \cup S_2}^T$ to the identity matrix in the last step makes the matrix invertible but does not change the solution, since the rows of \tilde{Z} lie in the subspace $S_1 \cup S_2$. Equation 22 is the solution of the following ridge regression problem:

$$\begin{aligned}
 \min_b \sum_{i=1}^n (1 - \mathbf{b}^T \tilde{z}_i)^2 + n \sigma_{bound} \cdot \|\mathbf{b}\|^2 \\
 = \min_b \sum_{i=1}^n (y_i - \mathbf{b}^T \tilde{\mathbf{x}}_i)^2 + n \sigma_{bound} \cdot \|\mathbf{b}\|^2,
 \end{aligned}$$

where $\tilde{\mathbf{x}}_i = P_{S_1 \cup S_2} \mathbf{x}_i$, and we use the fact that $y_i \in \{-1, 1\}$. □

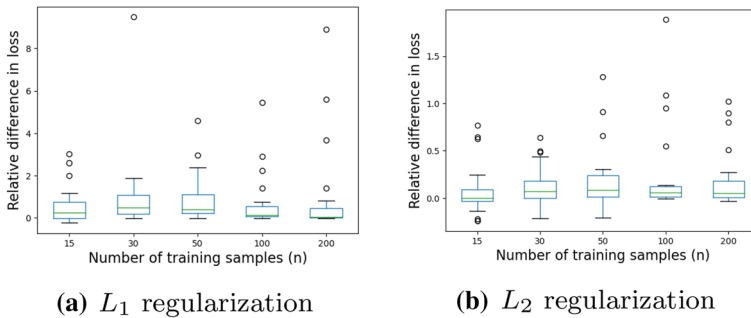


Fig. 9 Relative difference in loss if cross-validation is run with zero-one loss instead of logistic loss: Optimizing for zero-one loss in cross-validation leads to worse logistic loss on the test set. In other words, the predicted class probabilities are inaccurate under cross-validation with zero-one loss

B Results for Zero-One Loss

In many applications, we only care about zero-one (misclassification) loss. For such tasks, we pick the hyperparameters (e.g., the regularization λ in Eq. 15) that minimize zero-one loss in cross-validation. After fixing these hyperparameters, we optimize the feature weights using a convex loss such as logistic loss. Thus, the parameter-fitting step uses two losses. In contrast, our previous experiments used a single convex loss everywhere. We now explore the effects of this “double-loss” optimization.

Figure 8 shows the misclassification loss of various methods after “double-loss” parameter-fitting. RoLIN outperforms L_1 regularization and Top-PCs, and is similar to L_2 regularization. Thus, even though RoLIN is aimed at convex losses, it is useful even for zero-one loss.

However, the double-loss optimization improves zero-one loss only at the cost of worse values for the logistic loss. Figure 9 shows the relative increase in the logistic loss on test samples when we move from the single-loss to the double-loss parameter-fitting. The test logistic loss increases for both L_1 and L_2 regularization for all training sizes. The test logistic loss can be up to 10x larger for L_1 and up to 2x larger for L_2 regularization. Hence, the confidence scores (or equivalently, the class probabilities) learned by the logistic regression classifier can be inaccurate if we use zero-one loss to choose the hyperparameters.

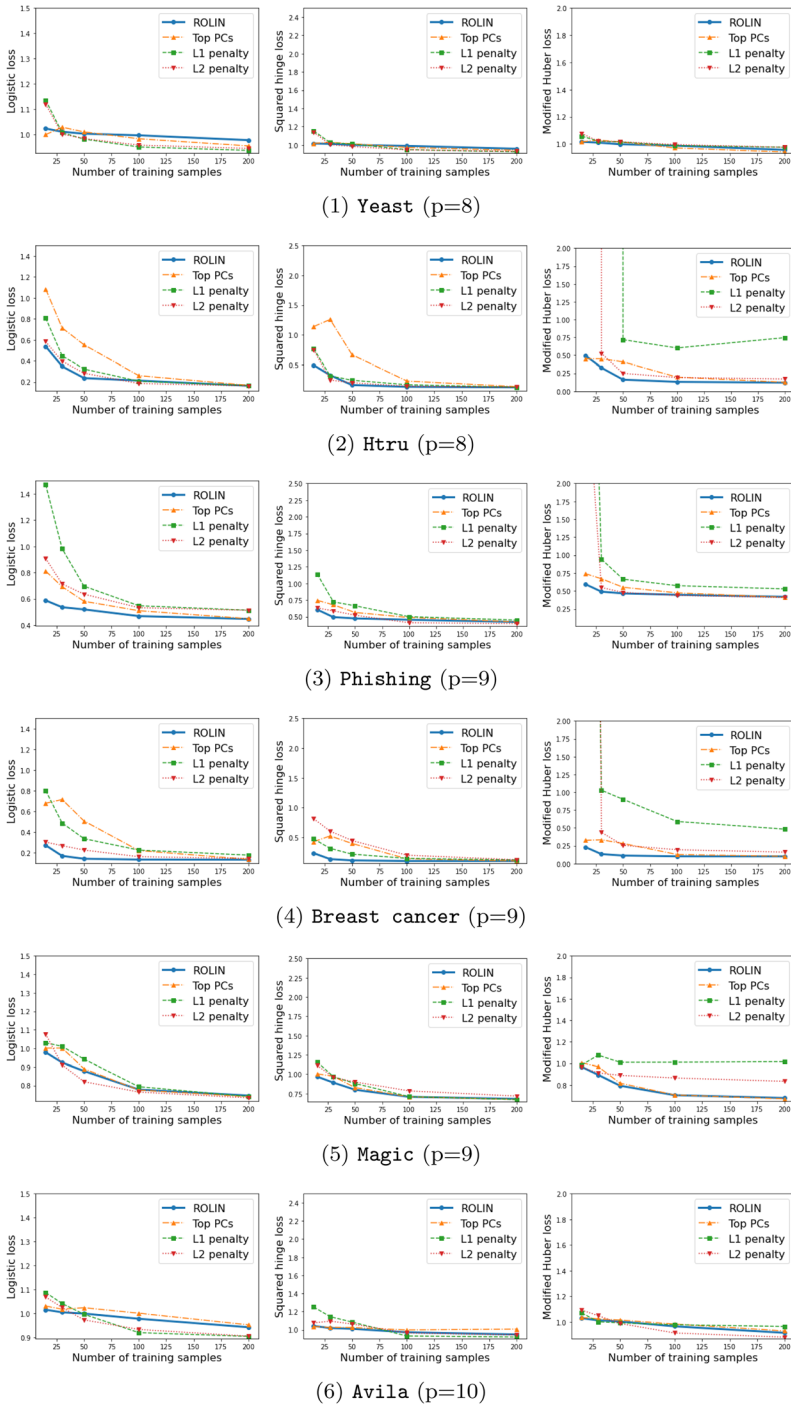


Fig. 10 Logistic, squared hinge, and modified Huber losses for 25 datasets under varying training sizes. Each plot compares the trimmed mean loss of RoLIN against L_1 and L_2 regularization. Regularized losses are too large to fit in some plots, such as plots (16) and (20)

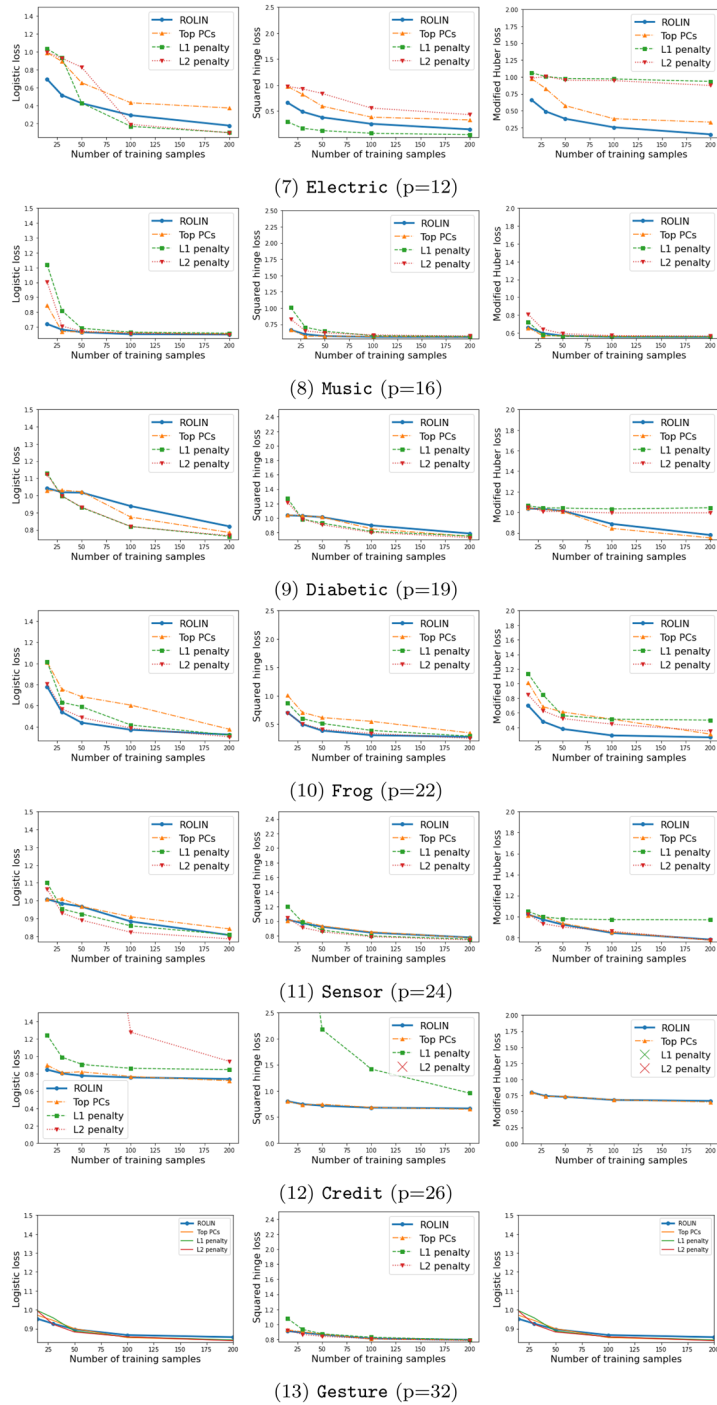


Fig. 10 (continued)

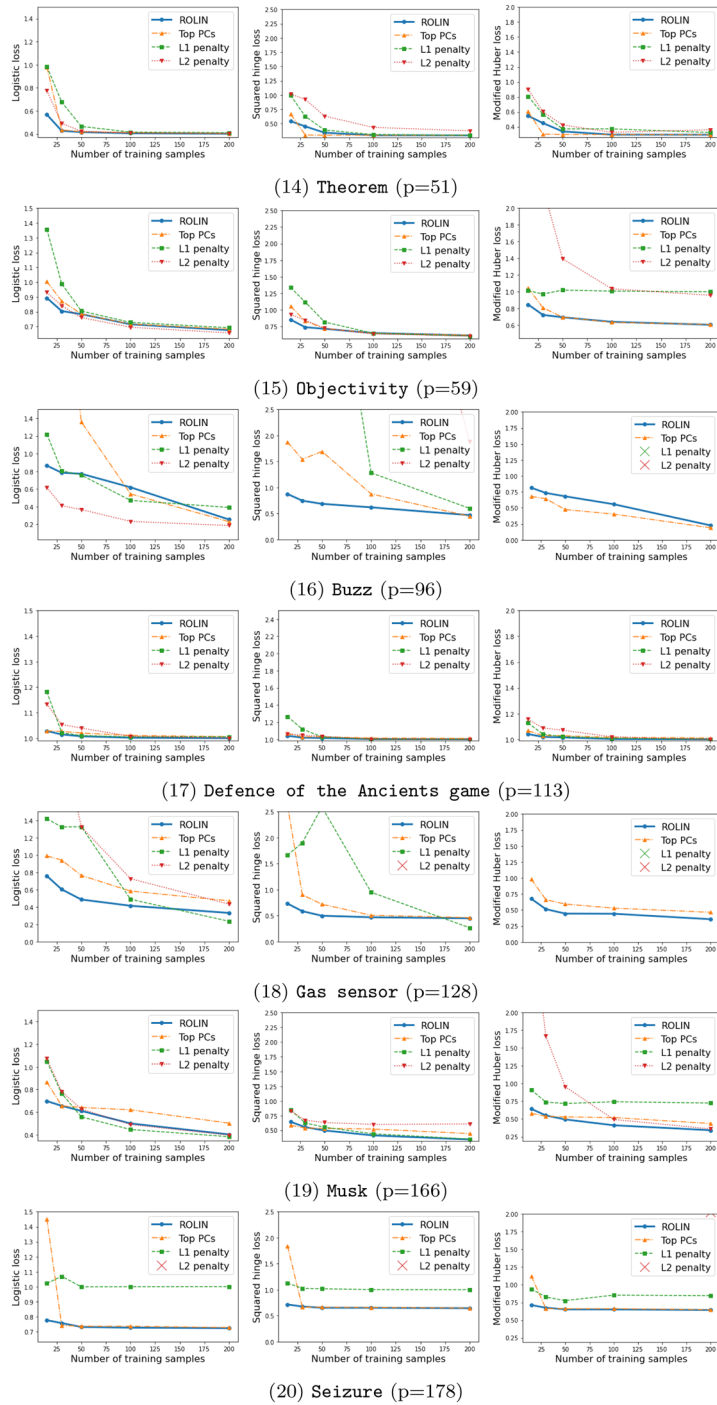


Fig. 10 (continued)

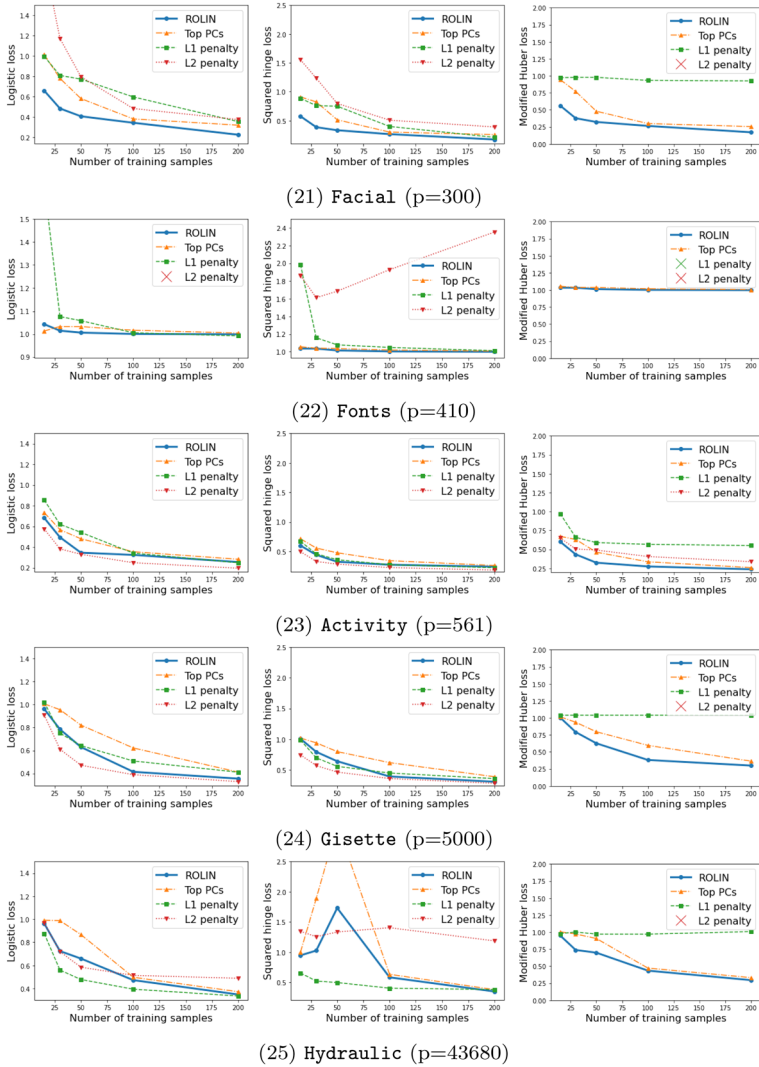


Fig. 10 (continued)

Author Contributions This is the sole work of D. Chakrabarti.

Funding The author thanks Facebook’s Research and Academic Relations Program for a faculty award.

Data availability The datasets used in the paper are publicly available at <https://archive.ics.uci.edu/ml/datasets>.

Code availability The code is available from <https://faculty.mcombs.utexas.edu/deepayan.chakrabarti/mywww/software/ROLIN-2021.tgz>.

Declaration

Conflict of interest The author declares that he has no conflict of interest.

References

- Ben-Tal A, El Ghaoui L, Nemirovski A (2009) Robust optimization
- Bhattacharyya C (2004) Robust classification of noisy data using second order cone programming approach. In: Proceedings of the International Conference on Intelligent Sensing and Information Processing, pp 433–438.
- Bi J, Zhang T (2004) Support Vector Classification with Input Data Uncertainty. *Neural Information Processing Systems* pp 161–168.
- Blagus, R., & Lusa, L. (2013). SMOTE for high-dimensional class-imbalanced data. *BMC Bioinformatics*, 14(1), 106.
- Cover, T. M., & Thomas, J. A. (2006). *Elements of Information Theory (Wiley Series in Telecommunications and Signal Processing)*. USA: Wiley-Interscience.
- Davis, C., & Kahan, W. M. (1970). The rotation of eigenvectors by a perturbation. III. *SIAM Journal on Numerical Analysis*, 7(1), 1–46.
- De Brabanter, J., Pelckmans, K., Suykens, J. A. K., & Vandewalle, J. (2002). Robust Cross-Validation Score Function for Non-linear Function Estimation. *Artificial Neural Networks - ICANN*, 2415, 713–719.
- Delage, E., & Ye, Y. (2010). Distributionally Robust Optimization Under Moment Uncertainty with Application to Data-Driven Problems. *Operations Research*, 58(3), 595–612.
- El Ghaoui, L., & Le Bret, H. (1997). Robust solutions to least-squares problems with uncertain data. *SIAM Journal on matrix analysis and applications*, 18(4), 1035–1064.
- Erdoğan, E., & Iyengar, G. (2006). Ambiguous chance constrained problems and robust optimization. *Mathematical Programming*, 107(1–2), 37–61.
- Feng J, Xu H, Mannor S, Yan S (2014) Robust Logistic Regression and Classification. In: Neural Information Processing Systems, pp 253–261
- Globerson A, Roweis S (2006) Nightmare at test time: Robust learning by feature deletion. In: Proceedings of the 23rd International Conference on Machine Learning, ICML '06, pp 353–360
- Goh J, Sim M (2010) Distributionally Robust Optimization and Its Tractable Approximations. *Operations Research* 58(4-part-1):902–917
- Hastie T, Tibshirani R, Friedman J (2009) The Elements of Statistical Learning, 2nd edn. Springer Series in Statistics, Springer New York Inc.
- Huang, Gao, Song, Shiji, Cheng, Wu., & You, Keyou. (2012). Robust Support Vector Regression for Uncertain Input and Output Data. *IEEE Transactions on Neural Networks and Learning Systems*, 23(11), 1690–1700.
- Jiang, R., & Guan, Y. (2016). Data-driven chance constrained stochastic program. *Mathematical Programming*, 158(1–2), 291–327.
- Jolliffe IT (1982) A note on the use of principal components in regression. *Applied Statistics* pp 300–303
- Kordzakhia, N., Mishra, G. D., & Reiersølmoen, L. (2001). Robust estimation in the logistic regression model. *Journal of Statistical Planning and Inference*, 98(1), 211–223.
- Marcenko, V. A., & Pastur, L. A. (1967). Distribution of Eigenvalues for Some Sets of Random Matrices. *Mathematics of the USSR-Sbornik*, 1(4), 457–483.
- Mohajerin Esfahani, P., & Kuhn, D. (2018). Data-driven distributionally robust optimization using the Wasserstein metric: performance guarantees and tractable reformulations. *Mathematical Programming*, 171(1–2), 115–166.
- Mohri M, Rostamizadeh A, Talwalkar A (2018) Foundations of Machine Learning, 2nd edn
- Shafieezadeh-Abadeh S, Esfahani PM, Kuhn D (2015) Distributionally robust logistic regression. In: Neural Information Processing Systems, NIPS'15, pp 1576–1584
- Shafieezadeh-Abadeh S, Kuhn D, Esfahani PM (2017) Regularization via Mass Transportation. [arXiv: 1710.10016](https://arxiv.org/abs/1710.10016) [cs, math, stat] [ArXiv: 1710.10016](https://arxiv.org/abs/1710.10016)
- Shivaswamy, P. K., Bhattacharyya, C., & Smola, A. J. (2006). Second Order Cone Programming Approaches for Handling Missing and Uncertain Data. *Journal of Machine Learning Research*, 7, 1283–1314.
- Sotiriou, C., Neo, S. Y., McShane, L. M., Korn, E. L., Long, P. M., Jazaeri, A., et al. (2003). Breast cancer classification and prognosis based on gene expression profiles from a population-based study. *Proceedings of the National Academy of Sciences*, 100(18), 10393–10398.

- Tibshirani, R. (1996). Regression Shrinkage and Selection via the Lasso. *Journal of the Royal Statistical Society Series B (Methodological)*, 58(1), 267–288.
- Tibshirani J, Manning CD (2014) Robust Logistic Regression using Shift Parameters. In: Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), pp 124–129
- Trafalis, T. B., & Gilbert, R. C. (2006). Robust classification and regression using support vector machines. *European Journal of Operational Research*, 173(3), 893–909.
- Tzelepis, C., Mezaris, V., & Patras, I. (2018). Linear Maximum Margin Classifier for Learning from Uncertain Data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(12), 2948–2962.
- Wang, X., & Pardalos, P. M. (2014). A Survey of Support Vector Machines with Uncertainties. *Annals of Data Science*, 1(3–4), 293–309.
- Wiesemann, W., Kuhn, D., & Sim, M. (2014). Distributionally Robust Convex Optimization. *Operations Research*, 62(6), 1358–1376.
- Wozabal, D. (2012). A framework for optimization under ambiguity. *Annals of Operations Research*, 193(1), 21–47.
- Xu H, Caramanis C, Mannor S (2009a) Robust regression and lasso. In: Neural Information Processing Systems, pp 1801–1808
- Xu, H., Caramanis, C., & Mannor, S. (2009b). Robustness and Regularization of Support Vector Machines. *Journal of Machine Learning Research*, 10, 1485–1510.
- Yu, Y., Wang, T., & Samworth, R. J. (2015). A useful variant of the Davis-Kahan theorem for statisticians. *Biometrika*, 102(2), 315–323.
- Zhao, L., Chakrabarti, D., & Muthuraman, K. (2019). Portfolio construction by mitigating error amplification: The bounded-noise portfolio. *Operations Research*, 67(4), 965–983.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.