

RIFLE: Imputation and Robust Inference from Low Order Marginals

Sina Baharlouei

University of Southern California

baharlou@usc.edu

Kelechi Ogudu

University of Southern California

kogudu@usc.edu

Sze-chuan Suen

University of Southern California

ssuen@usc.edu

Meisam Razaviyayn

University of Southern California

razaviya@usc.edu

Reviewed on OpenReview: <https://openreview.net/forum?id=oud7NyOKQy>

Abstract

The ubiquity of missing values in real-world datasets poses a challenge for statistical inference and can prevent similar datasets from being analyzed in the same study, precluding many existing datasets from being used for new analyses. While an extensive collection of packages and algorithms have been developed for data imputation, the overwhelming majority perform poorly if there are many missing values and low sample sizes, which are unfortunately common characteristics in empirical data. Such low-accuracy estimations adversely affect the performance of downstream statistical models. We develop a statistical inference framework for *regression and classification in the presence of missing data without imputation*. Our framework, RIFLE (Robust Inference via Low-order moment Estimations), estimates low-order moments of the underlying data distribution with corresponding confidence intervals to learn a distributionally robust model. We specialize our framework to linear regression and normal discriminant analysis, and we provide convergence and performance guarantees. This framework can also be adapted to impute missing data. In numerical experiments, we compare RIFLE to several state-of-the-art approaches (including MICE, Amelia, MissForest, KNN-imputer, MIDA, and Mean Imputer) for imputation and inference in the presence of missing values. Our experiments demonstrate that RIFLE outperforms other benchmark algorithms when the percentage of missing values is high and/or when the number of data points is relatively small. RIFLE is publicly available at <https://github.com/optimization-for-data-driven-science/RIFLE>.

1 Introduction

Machine learning algorithms have shown promise when applied to various problems, including healthcare, finance, social data analysis, image processing, and speech recognition. However, this success mainly relied on the availability of large-scale, high-quality datasets, which may be scarce in many practical problems, especially in medical and health applications (Pedersen et al., 2017; Sterne et al., 2009; Beaulieu-Jones et al., 2018). Moreover, many experiments and datasets suffer from the small sample size in such applications. Despite the availability of a small number of data points in each study, an increasingly large number of datasets are publicly available. To fully and effectively utilize information on related research questions from diverse datasets, information across various datasets (e.g., different questionnaires from multiple hospitals with overlapping questions) must be combined in a reliable fashion.



Figure 1: Consider the problem of predicting the trait y from feature vector $(\mathbf{x}_1, \dots, \mathbf{x}_{100})$. Suppose that we have access to three data sets: The first dataset includes the measurements of $(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{40}, y)$ for n_1 individuals. The second dataset collects data from another n_2 individuals by measuring $(\mathbf{x}_{30}, \dots, \mathbf{x}_{80})$ with no measurements of the target variable y in it; and the third dataset contains the measurements from the variables $(\mathbf{x}_{70}, \dots, \mathbf{x}_{100}, y)$ for n_3 number of individuals. How one should learn the predictor $\hat{y} = h(\mathbf{x}_1, \dots, \mathbf{x}_{100})$ from these three datasets?

After integrating data from different studies, the obtained dataset can contain large blocks of missing values, as they may not share the same features (Figure 1).

There are three general approaches for handling missing values in statistical inference (classification and regression) tasks. A Naïve method is to remove the rows containing missing entries. However, such an approach is not an option when the percentage of missingness in a dataset is high. For instance, as demonstrated in Figure 1, the entire dataset will be discarded if we eliminate the rows with at least one missing entry.

The most common methodology for handling missing values in a learning task is to impute them in a pre-processing stage. The general idea behind data imputation is that the missing values can be predicted using the available data entries and correlated features. Imputation algorithms cover a wide range of methods, including imputing missing entries with the columns means Little & Rubin (2019, Chapter 3) (or median), least-square and linear regression-based methods (Raghunathan et al., 2001; Kim et al., 2005; Zhang et al., 2008; Cai et al., 2006; Buuren & Groothuis-Oudshoorn, 2010), matrix completion and expectation maximization approaches Dempster et al. (1977); Ghahramani & Jordan (1994); Honaker et al. (2011), KNN based (Troyanskaya et al., 2001), Tree based methods (Stekhoven & Bühlmann, 2012; Xia et al., 2017), and methods using different neural network structures. Appendix A presents a comprehensive review of these methods.

The imputation of data allows practitioners to run standard statistical algorithms requiring complete data. However, the prediction model’s performance can be highly reliant on the accuracy of the imputer. High error rates in the prediction of missing values by the imputer can lead to the catastrophic performance of the downstream statistical methods executed on the imputed data.

Another class of methods for inference in the presence of missing values relies on robust optimization over the uncertainty sets on missing entries. Shivaswamy et al. (2006) and Xu et al. (2009) adopt robust optimization to learn the parameters of a support vector machine model. They consider uncertainty sets for the missing entries in the dataset and solve a min-max problem over those sets. The obtained classifiers are robust to the uncertainty of missing entries within the uncertainty regions. In contrast to the imputation-based approaches, the robust classification formulation does not carry the imputation error to the classification phase. However, finding appropriate intervals for each missing entry is challenging, and it is unclear how to determine the uncertainty range in many real datasets. Moreover, their proposed algorithms are limited to the SVM classifier.




Figure 2: Prediction of the target variable without imputation. RIFLE estimates confidence intervals for low-order (first and second-order) marginals from the input data containing missing values. Then, it solves a distributionally robust problem over the set of all distributions whose low-order marginals are within the estimated confidence intervals.

In this paper, we propose RIFLE (Robust Inference via Low-order moment Estimations) for the direct inference of a target variable based on a set of features containing missing values. The proposed framework does not require the data to be imputed in a pre-processing stage. However, it can also be used as a pre-processing tool for imputing data. The main idea of the proposed framework is to estimate the first and second-order moments of the data and their confidence intervals by bootstrapping on the available data matrix entries. Then, RIFLE finds the optimal parameters of the statistical model for the worst-case distribution with the low-order moments (mean and variance) within the estimated confidence intervals (See Figure 2). Compared to Shivaswamy et al. (2006); Xu et al. (2009), we estimate uncertainty regions for the low-order marginals using the Bootstrap technique. Furthermore, our framework is not restricted to any particular machine learning model, such as support vector machines (Xu et al., 2009).

Contributions: Our main contributions are as follows:

1. We present a distributionally robust optimization framework over the low-order marginals of the training data distribution for inference in the presence of missing values. The proposed framework does not require data imputation as a pre-processing stage. In Section 3 and Section 4, we specialize the framework to ridge regression and classification models as two case studies respectively. The proposed framework provides a novel strategy for inference in the presence of missing data, especially for datasets with large proportions of missing values.
2. We provide theoretical convergence guarantees and the iteration complexity analysis of the presented algorithms for robust formulations of ridge linear regression and normal discriminant analysis. Moreover, we show the consistency of the prediction under mild assumptions and analyze the asymptotic statistical properties of the solutions found by the algorithms.
3. While the robust inference framework is primarily designed for direct statistical inference in the presence of missing values without performing data imputation, it can also be adopted as an imputation tool. To demonstrate the quality of the proposed imputer, we compare its performance with several widely-used imputation packages such as MICE (Buuren & Groothuis-Oudshoorn, 2010), Amelia (Honaker et al., 2011), MissForest (Stekhoven & Bühlmann, 2012), KNN-Imputer (Trojanovskaya et al., 2001), MIDA (Gondara & Wang, 2018), GAIN (Yoon et al., 2018) on real and synthetic datasets. Generally speaking, our method outperforms all of the mentioned packages when the number of missing entries is large.

2 Robust Inference via Estimating Low-order Moments

RIFLE is based on a distributionally robust optimization (DRO) framework over low-order marginals. Assume that $(\mathbf{x}, y) \in \mathbb{R}^d \times \mathbb{R}$ follows a joint probability distribution P^* . A standard approach for predicting the target variable y given the input vector \mathbf{x} is to find the parameter θ that minimizes the population risk with respect to a given loss function ℓ :

$$\min_{\theta} \mathbb{E}_{(\mathbf{x}, y) \sim P^*} [\ell(\mathbf{x}, y; \theta)]. \quad (1)$$

Since the underlying distribution of data is rarely available in practice, the above problem cannot be directly solved. The most common approach for approximating (1) is to minimize the empirical risk with respect to

n given i.i.d samples $(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)$ drawn from the joint distribution P^* :

$$\min_{\boldsymbol{\theta}} \frac{1}{n} \sum_{i=1}^n \ell(\mathbf{x}_i, y_i; \boldsymbol{\theta}).$$

The above empirical risk formulation assumes that all entries of \mathbf{x}_i and y_i are available. Thus, to utilize the empirical risk minimization (ERM) framework in the presence of missing values, one can either remove or impute the missing data points in a pre-processing stage. Training via robust optimization is a natural alternative in the presence of missing data. Shivaswamy et al. (2006); Xu et al. (2009) suggest the following optimization problem that minimizes the loss function for the worst-case scenario over the defined uncertainty sets per data points:

$$\min_{\boldsymbol{\theta}} \max_{\{\boldsymbol{\delta}_i \in \mathcal{N}_i\}_{i=1}^n} \frac{1}{n} \sum_{i=1}^n \ell(\mathbf{x}_i - \boldsymbol{\delta}_i, y_i; \boldsymbol{\theta}), \quad (2)$$

where \mathcal{N}_i represents the uncertainty region of data point i . Shivaswamy et al. (2006) obtains the uncertainty sets by assuming a known distribution on the missing entries of datasets. The main issue in their approach is that the constraints defined on data points are totally uncorrelated. Xu et al. (2009) on the other hand defines \mathcal{N}_i as a ‘‘box’’ constraint around the data point i such that they can be linearly correlated. For this specific case, they show that solving the corresponding robust optimization problem is equivalent to minimizing a regularized reformulation of the original loss function. Such an approach has several limitations: First, it can only handle a few special cases (SVM loss with linearly correlated perturbations on data points). Furthermore, Xu et al. (2009) is primarily designed for handling outliers and contaminated data. Thus, they do not offer any mechanism for the initial estimation of \mathbf{x}_i when several vector entries are missing. In this work, we instead take a *distributionally robust* approach by considering uncertainty on the data distribution instead of defining an uncertainty set for each data point. In particular, we aim to fit the best parameters of a statistical learning model for the worst distribution in a given uncertainty set by solving the following:

$$\min_{\boldsymbol{\theta}} \max_{P \in \mathcal{P}} \mathbb{E}_{(\mathbf{x}, y) \sim P} [\ell(\mathbf{x}, y; \boldsymbol{\theta})], \quad (3)$$

where \mathcal{P} is an uncertainty set over the underlying distribution of data. A key observation is that defining the uncertainty set \mathcal{P} in (3) is easier and computationally more efficient than defining the uncertainty sets $\{\mathcal{N}_i\}_{i=1}^n$ in (2). In particular, the uncertainty set \mathcal{P} can be obtained naturally by estimating low-order moments of data distribution using only available entries. To explain this idea and to simplify the notations, let $\mathbf{z} = (\mathbf{x}, y)$, $\bar{\boldsymbol{\mu}}^{\mathbf{z}} \triangleq \mathbb{E}[\mathbf{z}]$, and $\bar{\mathbf{C}}^{\mathbf{z}} \triangleq \mathbb{E}[\mathbf{z}\mathbf{z}^T]$. While $\bar{\boldsymbol{\mu}}^{\mathbf{z}}$ and $\bar{\mathbf{C}}^{\mathbf{z}}$ are typically not known exactly, one can estimate them (within certain confidence intervals) from the available data by simply ignoring missing entries (assuming the missing value pattern is completely at random, e.g., MCAR). Moreover, we can estimate the confidence intervals via bootstrapping. Particularly, we can estimate $\boldsymbol{\mu}_{\min}^{\mathbf{z}}, \boldsymbol{\mu}_{\max}^{\mathbf{z}}, \mathbf{C}_{\min}^{\mathbf{z}},$ and $\mathbf{C}_{\max}^{\mathbf{z}}$ from data such that $\boldsymbol{\mu}_{\min}^{\mathbf{z}} \leq \bar{\boldsymbol{\mu}}^{\mathbf{z}} \leq \boldsymbol{\mu}_{\max}^{\mathbf{z}}$ and $\mathbf{C}_{\min}^{\mathbf{z}} \leq \bar{\mathbf{C}}^{\mathbf{z}} \leq \mathbf{C}_{\max}^{\mathbf{z}}$ with high probability (where the inequalities for matrices and vectors denote component-wise relations). In Appendix B, we show how a bootstrapping strategy can be used to obtain the confidence intervals described above. Given these estimated confidence intervals from data, (3) can be reformulated as

$$\begin{aligned} \min_{\boldsymbol{\theta}} \quad & \max_P \mathbb{E}_P[\ell(\mathbf{z}; \boldsymbol{\theta})] \\ \text{s.t.} \quad & \boldsymbol{\mu}_{\min}^{\mathbf{z}} \leq \mathbb{E}_P[\mathbf{z}] \leq \boldsymbol{\mu}_{\max}^{\mathbf{z}}, \\ & \mathbf{C}_{\min}^{\mathbf{z}} \leq \mathbb{E}_P[\mathbf{z}\mathbf{z}^T] \leq \mathbf{C}_{\max}^{\mathbf{z}}. \end{aligned} \quad (4)$$

Gao & Kleywegt (2017) utilize the distributionally robust optimization as (3) over the set of positive semi-definite (PSD) cones for robust inference under uncertainty. While their formulation considers ℓ_2 balls for the constraints on low order moments of the data, we use ℓ_∞ constraints that are computationally more natural in the presence of missing entries when combined with bootstrapping. Furthermore, while it can be applied to general convex losses, their method relies on the ellipsoid and the existence of oracles for performing the steps of the ellipsoid method, which is not applicable in modern high-dimensional problems. Moreover, they assume concavity in data (the existence of some oracle to return the worst-case data points) that is

practically unavailable even in convex loss functions (including linear regression and normal discriminant analysis studied in our work).

In Section 3, we study the proposed distributionally robust framework described in (4) for the ridge linear regression. We design efficient first-order convergent algorithms to solve the problem and show how we can use the algorithms for both inference and imputation in the presence of missing values. Further, in Appendix F, we study the proposed distributionally robust framework for the classification problems under the normality assumption of features. In particular, we show how Framework (4) can be specialized to the robust normal discriminant analysis in the presence of missing values.

3 Robust Linear Regression in the Presence of Missing Values

Let us specialize our framework to the ridge linear regression model. In the absence of missing data, ridge regression finds optimal regressor parameter $\boldsymbol{\theta}$ by solving

$$\min_{\boldsymbol{\theta}} \|\mathbf{X}\boldsymbol{\theta} - \mathbf{y}\|_2^2 + \lambda\|\boldsymbol{\theta}\|_2^2,$$

or equivalently by solving:

$$\min_{\boldsymbol{\theta}} \boldsymbol{\theta}^T \mathbf{X}^T \mathbf{X} \boldsymbol{\theta} - 2\boldsymbol{\theta}^T \mathbf{X}^T \mathbf{y} + \lambda\|\boldsymbol{\theta}\|_2^2. \quad (5)$$

Thus, having the second-order moments of the data $\mathbf{C} = \mathbf{X}^T \mathbf{X}$ and $\mathbf{b} = \mathbf{X}^T \mathbf{y}$ is sufficient for finding the optimal solution. In other words, it suffices to compute the inner product of any two column vectors $\mathbf{a}_i, \mathbf{a}_j$ of \mathbf{X} , and the inner product of any column \mathbf{a}_i of \mathbf{X} with vector \mathbf{y} . Since the matrix \mathbf{X} and vector \mathbf{y} are not fully observed due to the existence of missing values, one can use the available data (see (24) for details) to compute the point estimators \mathbf{C}_0 and \mathbf{b}_0 . These point estimators can be highly inaccurate, especially when the number of non-missing rows for two given columns is small. In addition, if the pattern of missing entries does not follow the MCAR assumption, the point estimators are not unbiased estimators of \mathbf{C} and \mathbf{b} .

3.1 A Distributionally Robust Formulation of Linear Regression

As we mentioned above, to solve the linear regression problem, we only need to estimate the second-order moments of the data ($\mathbf{X}^T \mathbf{X}$ and $\mathbf{X}^T \mathbf{y}$). Thus, the distributionally robust formulation described in (4) is equivalent to the following optimization problem for the linear regression model:

$$\begin{aligned} \min_{\boldsymbol{\theta}} \quad & \max_{\mathbf{C}, \mathbf{b}} \quad \boldsymbol{\theta}^T \mathbf{C} \boldsymbol{\theta} - 2\mathbf{b}^T \boldsymbol{\theta} + \lambda\|\boldsymbol{\theta}\|_2^2 \\ \text{s.t.} \quad & \mathbf{C}_0 - c\boldsymbol{\Delta} \leq \mathbf{C} \leq \mathbf{C}_0 + c\boldsymbol{\Delta}, \\ & \mathbf{b}_0 - c\boldsymbol{\delta} \leq \mathbf{b} \leq \mathbf{b}_0 + c\boldsymbol{\delta}, \\ & \mathbf{C} \succeq 0, \end{aligned} \quad (6)$$

where the last constraint guarantees that the covariance matrix is positive and semi-definite. We discuss the procedure of estimating the confidence intervals ($\mathbf{b}_0, \mathbf{C}_0, \boldsymbol{\delta}$, and $\boldsymbol{\Delta}$) in Appendix B.

3.2 RIFLE for Ridge Linear Regression

Since the objective function in (6) is convex in $\boldsymbol{\theta}$ (ridge regression) and concave in \mathbf{b} and \mathbf{C} (linear), the minimization and maximization sub-problems are interchangeable (Sion et al., 1958). Thus, we can equivalently rewrite Problem (6) as:

$$\begin{aligned} \max_{\mathbf{C}, \mathbf{b}} \quad & g(\mathbf{C}, \mathbf{b}) \\ \text{s.t.} \quad & \mathbf{C}_0 - c\boldsymbol{\Delta} \leq \mathbf{C} \leq \mathbf{C}_0 + c\boldsymbol{\Delta}, \\ & \mathbf{b}_0 - c\boldsymbol{\delta} \leq \mathbf{b} \leq \mathbf{b}_0 + c\boldsymbol{\delta}, \\ & \mathbf{C} \succeq 0, \end{aligned} \quad (7)$$

where $g(\mathbf{b}, \mathbf{C}) = \min_{\boldsymbol{\theta}} \boldsymbol{\theta}^T \mathbf{C} \boldsymbol{\theta} - 2\mathbf{b}^T \boldsymbol{\theta} + \lambda \|\boldsymbol{\theta}\|^2$. Function g can be computed in closed-form given any pair of (\mathbf{C}, \mathbf{b}) by setting $\boldsymbol{\theta} = (\mathbf{C} + \lambda \mathbf{I})^{-1} \mathbf{b}$. Thus, using Danskin's Theorem (Danskin, 2012), we can apply projected gradient ascent to function g to find an optimal solution of (7) as described in Algorithm 1. At each iteration of the algorithm, we first perform one step of projected gradient ascent on matrix \mathbf{C} and vector \mathbf{b} ; then we update $\boldsymbol{\theta}$ in closed-form for the obtained \mathbf{C} and \mathbf{b} . We initialize \mathbf{C} and \mathbf{b} using entrwise point estimation on the available rows (see Equation (24) in Appendix B). The projection of \mathbf{b} to the box

Algorithm 1 RIFLE for Ridge Linear Regression in the Presence of Missing Values

- 1: **Input:** $\mathbf{C}_0, \mathbf{b}_0, \Delta, \delta, T$
 - 2: **Initialize:** $\mathbf{C} = \mathbf{C}_0, \mathbf{b} = \mathbf{b}_0$.
 - 3: **for** $i = 1, \dots, T$ **do**
 - 4: Update $\mathbf{C} = \Pi_{\Delta+} [\mathbf{C} + \alpha \boldsymbol{\theta} \boldsymbol{\theta}^T]$
 - 5: Update $\mathbf{b} = \Pi_{\delta} (\mathbf{b} - 2\alpha \boldsymbol{\theta})$
 - 6: Set $\boldsymbol{\theta} = (\mathbf{C} + \lambda \mathbf{I})^{-1} \mathbf{b}$
-

constraint $\mathbf{b}_0 - c\boldsymbol{\delta} \leq \mathbf{b} \leq \mathbf{b}_0 + c\boldsymbol{\delta}$ can be done entrwise and has the following closed-form

$$\Pi_{\delta}(\mathbf{b}_i) = \begin{cases} \mathbf{b}_i & \text{if } \mathbf{b}_{0i} - c\boldsymbol{\delta}_i \leq \mathbf{b}_i \leq \mathbf{b}_{0i} + c\boldsymbol{\delta}_i, \\ \mathbf{b}_{0i} - c\boldsymbol{\delta}_i & \text{if } \mathbf{b}_i < \mathbf{b}_{0i} - c\boldsymbol{\delta}_i, \\ \mathbf{b}_{0i} + c\boldsymbol{\delta}_i & \text{if } \mathbf{b}_{0i} + c\boldsymbol{\delta}_i < \mathbf{b}_i. \end{cases}$$

Theorem 1. Let $(\tilde{\boldsymbol{\theta}}, \tilde{\mathbf{C}}, \tilde{\mathbf{b}})$ be the optimal solution of (6), $\boldsymbol{\theta}^*(\mathbf{b}, \mathbf{C}) = \arg \min_{\boldsymbol{\theta}} \boldsymbol{\theta}^T \mathbf{C} \boldsymbol{\theta} - 2\mathbf{b}^T \boldsymbol{\theta} + \lambda \|\boldsymbol{\theta}\|^2$, and $D = \|\mathbf{C}_0 - \tilde{\mathbf{C}}\|_F^2 + \|\mathbf{b}_0 - \tilde{\mathbf{b}}\|_2^2$. Assume that for any given \mathbf{b} and \mathbf{C} , within the uncertainty (constraint) sets described in (6), $\|\boldsymbol{\theta}^*(\mathbf{b}, \mathbf{C})\| \leq \tau$. Then Algorithm 1 computes an ϵ -optimal solution of the objective function in (7) in $\mathcal{O}\left(\frac{D(\tau+1)^2}{\lambda\epsilon}\right)$ iterations.

Proof. The proof is relegated to Appendix H. □

In Appendix C, we show how using the acceleration method of Nesterov can improve the convergence rate of Algorithm 1 to $\mathcal{O}\left(\sqrt{\frac{D(\tau+1)^2}{\epsilon\lambda}}\right)$. A technical issue of Algorithm 1 and its accelerated version presented in Appendix C is that projection of \mathbf{C} to the intersection of box constraints and the set of positive semidefinite matrices ($\Pi_{\Delta+}[\mathbf{C}]$) is challenging and cannot be done in closed-form. In the implementation of Algorithm 1, we relax the problem by removing the PSD constraint on \mathbf{C} to avoid this complexity and time-consuming singular value decomposition at each iteration. This relaxation does not drastically change the algorithm's performance, as our experiments show in Section 5. A more systematic approach is to write the dual problem of the maximization problem and handle the resulting constrained minimization problem with the Alternating Direction Method of Multipliers (ADMM). The detailed procedure of such an approach can be found in Appendix D. All these algorithms are provably convergent to the optimal points of Problem (6). In addition to theoretical convergence, we have numerically evaluated the convergence of resulting algorithms in Appendix K. Further, the proposed algorithms are **consistent**, as discussed in Appendix J.

3.3 Performance Guarantees for RIFLE

Thus far, we have discussed how to efficiently solve the robust linear regression problem in the presence of missing values. A natural question in this context is the statistical performance of the obtained optimal solution in the previous section on the unseen test data points. Theorem 2 answers this question from two perspectives: Assuming that the missing values are distributed completely at random, our estimators are consistent. Moreover, for the finite case, Theorem 2 part (b) states that with the proper choice of confidence intervals, with high probability, the test loss of the obtained solution is bounded by the training loss of the estimator. Note that the results regarding the performance of the robust estimator generally hold for MCAR missing pattern. However, we perform several experiments on datasets with MNAR patterns to show how RIFLE works in practice on such datasets in Section 5.

Theorem 2. Assume the data domain is bounded and that the missing pattern of the data follows MCAR. Let $\mathbf{X}^{n \times d}$, \mathbf{y} be the training data drawn i.i.d. from the ground-truth distribution P^* with low-order moments \mathbf{C}^* and \mathbf{b}^* . Further, assume that each entry of \mathbf{X} and \mathbf{y} is missing with probability $p < 1$. Let $(\tilde{\boldsymbol{\theta}}_n, \tilde{\mathbf{C}}_n, \tilde{\mathbf{b}}_n)$ be the solution of Problem (6).

(a) **Consistency of the Covariance Estimator:** As the number of data points goes to infinity, the estimated low-order marginals converge to the ground-truth values, almost surely. More precisely,

$$\lim_{n \rightarrow \infty} \tilde{\mathbf{C}}_n = \mathbb{E}_{P^*}[\mathbf{xx}^T], \quad a.s., \quad (8)$$

$$\lim_{n \rightarrow \infty} \tilde{\mathbf{b}}_n = \mathbb{E}_{P^*}[\mathbf{xy}], \quad a.s. \quad (9)$$

(b) **Defining**

$$\begin{aligned} L_{\text{train}}(\tilde{\boldsymbol{\theta}}_n) &= \tilde{\boldsymbol{\theta}}_n^T \tilde{\mathbf{C}}_n \tilde{\boldsymbol{\theta}}_n - 2\tilde{\mathbf{b}}_n \tilde{\boldsymbol{\theta}}_n + \lambda \|\tilde{\boldsymbol{\theta}}_n\|_2^2 \\ L_{\text{test}}(\tilde{\boldsymbol{\theta}}_n) &= \tilde{\boldsymbol{\theta}}_n^T \mathbf{C}^* \tilde{\boldsymbol{\theta}}_n - 2\mathbf{b}^{*T} \tilde{\boldsymbol{\theta}}_n + \lambda \|\tilde{\boldsymbol{\theta}}_n\|_2^2, \end{aligned}$$

where $\mathbf{C}^* = \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim P^*}[\mathbf{xx}^T]$ and $\mathbf{b}^* = \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim P^*}[\mathbf{xy}]$ are the ground-truth second-order moments. Given $V = \max_{i,j} \text{Var}(X_i X_j)$ (maximum variance of pairwise feature products), with the probability of at least $1 - \frac{d^2 V}{2c^2 \Delta^2 n(1-p)}$, we have:

$$L_{\text{test}}(\tilde{\boldsymbol{\theta}}) \leq L_{\text{train}}(\tilde{\boldsymbol{\theta}}), \quad (10)$$

where $\Delta = \min\{\Delta_{ij}\}$ and c is the hyper-parameter for controlling the size of the confidence intervals as presented in (6)

Proof. The proof is relegated to Appendix H. □

3.4 Imputation of Missing Values and Going Beyond Linear Regression

RIFLE can be used for imputing missing data. To this end, we impute different features of a given dataset independently. More precisely, to impute each feature containing missing values, we consider it as a target variable \mathbf{y} and the rest of the features as the input \mathbf{X} in our methodology. Then, we train a model to predict the feature \mathbf{y} given \mathbf{X} via Algorithm 1 (or its ADMM version, Algorithm 7, in the appendix). Let the obtained optimal solutions be \mathbf{C}^* , \mathbf{b}^* , and $\boldsymbol{\theta}^*$. For a given missing entry, we can use $\boldsymbol{\theta}^*$ only if all other features in the row of that missing entry are available. However, that is not usually the case in practice, as each row can contain more than one missing entry. Therefore, one can learn a separate model for each missing pattern in the dataset. Let us clarify this point through the example in Figure 1. In this example, we have three different missing patterns (one missing pattern for each dataset). For missing entries in Dataset 1, the first forty features are available. Let \mathbf{r}_j denote the vector of the first 40 features in row j . Assume that we aim to impute entry $i \in \{41, \dots, 100\}$ in row j where i denoted by x_{ji} . To this end, we restrict \mathbf{X} to the first 40 features. Moreover, we consider $y = x_i$ as the target variable. Then, we run Algorithm 1 on \mathbf{X} and y to obtain the optimal \mathbf{C}^* , \mathbf{b}_i^* , and $\boldsymbol{\theta}_i^*$. Consequently, we impute x_{ji} as follows:

$$x_{ji} = \mathbf{r}_j^T \boldsymbol{\theta}_i^*$$

We can use the same methodology for imputing missing entries in each feature for missing patterns in Dataset 2 and Dataset 3. While this approach is reasonable for the missing pattern observed in Figure 1, in many practical problems, different rows can have distinct missing patterns. Thus, in the worst case, Algorithm 1 must be executed once for each missing entry. Such an approach is computationally expensive and might be infeasible in large-scale datasets containing large amounts of missing entries. Alternatively, one can perform Algorithm 1 only once to obtain \mathbf{C}^* and \mathbf{b}^* (considered the “worst-case/pessimistic” estimation of the moments). Then to impute each missing entry, \mathbf{C}^* and \mathbf{b}^* are restricted to the features available in that missing entry’s row. Having the restricted \mathbf{C}^* and \mathbf{b}^* , the regressor $\boldsymbol{\theta}^*$ can be obtained in closed-form (line 6 in Algorithm 1). In this approach, we perform algorithm 1 once and find the optimal $\boldsymbol{\theta}^*$ for each missing entry based on the estimated \mathbf{C}^* and \mathbf{b}^* . This approach can lead to sub-optimal solutions compared to the former approach, but it is much faster and more scalable.

Beyond Linear Regression: While the developed methods are primarily designed for ridge linear regression, one can apply non-linear transformations (kernels) to obtain models beyond linear. In Appendix E, we show how to extend the developed algorithms to quadratic models. The RIFLE framework applied to the quadratically transformed data is called **QRIFLE**.

4 Robust Classification Framework

In this section, we study the proposed framework in (4) for the classification tasks in the presence of missing values. Since the target variable $y \in \mathcal{Y} = \{1, \dots, M\}$ takes discrete values in classification tasks, we consider the uncertainty sets over the data's first- and second-order marginals given each target value (label) separately. Therefore, the distributionally robust classification over low-order marginals can be described as:

$$\begin{aligned} \min_{\mathbf{w}} \max_P \mathbb{E}_P[\ell(\mathbf{x}, y, \mathbf{w})] \\ \text{s.t. } \boldsymbol{\mu}_{\min, y} \leq \mathbb{E}_P[\mathbf{x}|y] \leq \boldsymbol{\mu}_{\max, y} \quad \forall y \in \mathcal{Y} \\ \boldsymbol{\Sigma}_{\min, y} \leq \mathbb{E}_P[\mathbf{x}\mathbf{x}^T|y] \leq \boldsymbol{\Sigma}_{\max, y} \quad \forall y \in \mathcal{Y} \end{aligned} \quad (11)$$

where $\boldsymbol{\mu}_{\min}, \boldsymbol{\mu}_{\max}, \boldsymbol{\Sigma}_{\min}$, and $\boldsymbol{\Sigma}_{\max}$ are the estimated confidence intervals for the first and second order of the data distribution. Unlike the robust linear regression task in Section 3, the evaluation of the objective function in (11) might depend on higher-order marginals (beyond second-order) due to the nonlinearity of the loss function. As a result, Problem (11) is a non-convex non-concave intractable min-max optimization problem in general. For the sake of computational traceability, we restrict the distribution in the inner maximization problem to the set of normal distributions. In the following section, we specialize (11) to the quadratic discriminant analysis as a case study. The methodology can be extended to other popular classification algorithms, such as support vector machines and multi-layer neural networks.

4.1 Robust Quadratic Discriminant Analysis

Learning a logistic regression model on datasets containing missing values has been studied extensively in the literature (Fung & Wrobel, 1989; Abonazel & Ibrahim, 2018). Besides deleting missing values and imputation-based approaches, Fung & Wrobel (1989) models the logistic regression task in the presence of missing values as a linear discriminant analysis problem where the underlying assumption is that the predictors follow normal distribution conditional on the labels. Mathematically speaking, they assume that the data points assigned to a specific label follow a Gaussian distribution, i.e., $\mathbf{x}|y = i \sim N(\boldsymbol{\mu}_i, \boldsymbol{\Sigma})$. They use the available data to estimate the parameters of each Gaussian distribution. Therefore, the parameters of the logistic regression model can be assigned based on the estimated parameters of the Gaussian distributions for different classes. Similar to the linear regression case, the estimations of means and covariances are unbiased only when the data satisfies the MCAR condition. Moreover, when the number of data points in the dataset is small, the variance of the estimations can be very high. Thus, to train a logistic regression model that is robust to the percentage and different types of missing values, we specialize the general robust classification framework formulated in Equation (11) to the logistic regression model. Instead of considering a common covariance matrix for the conditional distributions of \mathbf{x} given labels y (linear discriminant analysis), we assume a more general case where each conditional distribution has its own covariance matrix (quadratic discriminant analysis). Assume that $\mathbf{x}|y \sim N(\boldsymbol{\mu}_y, \boldsymbol{\Sigma}_y)$ for $y = 0, 1$. We aim to find the optimal solution to the following problem:

$$\begin{aligned} \min_{\mathbf{w}} \max_{\boldsymbol{\mu}_0, \boldsymbol{\mu}_1, \boldsymbol{\Sigma}_0, \boldsymbol{\Sigma}_1} \mathbb{E}_{\mathbf{x}|y=1 \sim N(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1)} \left[-\log \left(\sigma(\mathbf{w}^T \mathbf{x}) \right) \right] \mathbb{P}(y=1) + \\ \mathbb{E}_{\mathbf{x}|y=0 \sim N(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)} \left[-\log \left(1 - \sigma(\mathbf{w}^T \mathbf{x}) \right) \right] \mathbb{P}(y=0) \\ \text{s.t. } \boldsymbol{\mu}_{\min_0} \leq \boldsymbol{\mu}_0 \leq \boldsymbol{\mu}_{\max_0} \\ \boldsymbol{\mu}_{\min_1} \leq \boldsymbol{\mu}_1 \leq \boldsymbol{\mu}_{\max_1} \\ \boldsymbol{\Sigma}_{\min_0} \leq \boldsymbol{\Sigma}_0 \leq \boldsymbol{\Sigma}_{\max_0} \\ \boldsymbol{\Sigma}_{\min_1} \leq \boldsymbol{\Sigma}_1 \leq \boldsymbol{\Sigma}_{\max_1} \end{aligned} \quad (12)$$

Where $\sigma(\mathbf{x}) = 1 / (1 + \exp(-\mathbf{x}))$ is the sigmoid function.

To solve Problem (12), first, we focus on the scenario when the target variable has no missing values. In this case, each data point contributes to the estimation of either $(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1)$ or $(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$, depending on its label. Similar to the robust linear regression case, we can apply Algorithm 4 to estimate the confidence intervals for $\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i$ using data points whose target variable equals i ($y = i$).

Obviously, the objective function is convex in \mathbf{w} since the logistic regression loss is convex, and the expectation of loss can be seen as a weighted summation, which is convex. Thus, fixing $\boldsymbol{\mu}, \boldsymbol{\Sigma}$ the outer minimization problem can be solved with respect to \mathbf{w} using standard first-order methods such as gradient descent.

Although the robust reformulation of logistic regression stated in (12) is convex in \mathbf{w} and concave in $\boldsymbol{\mu}_0$ and $\boldsymbol{\mu}_1$, the inner maximization problem is intractable with respect to $\boldsymbol{\Sigma}_0$ and $\boldsymbol{\Sigma}_1$. We approximate Problem (12) in the following manner:

$$\begin{aligned} \min_{\mathbf{w}} \max_{\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0, \boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1} & \pi_1 \mathbb{E}_{\mathbf{x}|y=1 \sim N(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1)} \left[-\log \left(\sigma(\mathbf{w}^T \mathbf{x}) \right) \right] + \pi_0 \mathbb{E}_{\mathbf{x}|y=0 \sim N(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)} \left[-\log \left(1 - \sigma(\mathbf{w}^T \mathbf{x}) \right) \right], \\ \text{s.t.} & \boldsymbol{\mu}_{\min_0} \leq \boldsymbol{\mu}_0 \leq \boldsymbol{\mu}_{\max_0} \\ & \boldsymbol{\mu}_{\min_1} \leq \boldsymbol{\mu}_1 \leq \boldsymbol{\mu}_{\max_1} \\ & \boldsymbol{\Sigma}_0 \in \{\boldsymbol{\Sigma}_{01}, \boldsymbol{\Sigma}_{02}, \dots, \boldsymbol{\Sigma}_{0k}\} \\ & \boldsymbol{\Sigma}_1 \in \{\boldsymbol{\Sigma}_{11}, \boldsymbol{\Sigma}_{12}, \dots, \boldsymbol{\Sigma}_{1k}\}, \end{aligned} \quad (13)$$

where $\pi_1 = \mathbb{P}(y = 1)$ and $\pi_0 = \mathbb{P}(y = 0)$. To compute optimal $\boldsymbol{\mu}_0$ and $\boldsymbol{\mu}_1$, we have:

$$\max_{\boldsymbol{\mu}_1} \mathbb{E}_{\mathbf{x} \sim N(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1)} \left[-\log \left(\sigma(\mathbf{w}^T \mathbf{x}) \right) \right] \quad \text{s.t.} \quad \boldsymbol{\mu}_{\min} \leq \boldsymbol{\mu}_1 \leq \boldsymbol{\mu}_{\max} \quad (14)$$

Theorem 3. *Let $\mathbf{a}[i]$ be the i -th element of vector \mathbf{a} . The optimal solution of Problem (14) has the following form:*

$$\boldsymbol{\mu}_1^*[i] = \begin{cases} \boldsymbol{\mu}_{\max}[i], & \text{if } \mathbf{w}[i] \leq 0 \\ \boldsymbol{\mu}_{\min}[i], & \text{if } \mathbf{w}[i] > 0. \end{cases} \quad (15)$$

Note that we relaxed (12) by taking the maximization problem over a finite set of $\boldsymbol{\Sigma}$ estimations. We estimate each $\boldsymbol{\Sigma}$ by bootstrapping on the available data using Algorithm 4. Define $f_i(\mathbf{w})$ as:

$$f_i(\mathbf{w}) = \pi_1 \mathbb{E}_{\mathbf{x} \sim N(\boldsymbol{\mu}_1^*, \boldsymbol{\Sigma}_{i1})} \left[-\log \left(\sigma(\mathbf{w}^T \mathbf{x}) \right) \right] \quad (16)$$

Similarly, we can define:

$$g_i(\mathbf{w}) = \pi_0 \mathbb{E}_{\mathbf{x} \sim N(\boldsymbol{\mu}_0^*, \boldsymbol{\Sigma}_{i0})} \left[-\log \left(1 - \sigma(\mathbf{w}^T \mathbf{x}) \right) \right] \quad (17)$$

Since the maximization problem is over a finite set, we can rewrite Problem (13) as:

$$\begin{aligned} \min_{\mathbf{w}} \max_{i,j \in \{1, \dots, k\}} & f_i(\mathbf{w}) + g_j(\mathbf{w}) = \min_{\mathbf{w}} \max_{p_1, \dots, p_k, q_1, \dots, q_k} \sum_{i=1}^k p_i f_i(\mathbf{w}) + \sum_{j=1}^k q_j g_j(\mathbf{w}) \\ \text{s.t.} & \sum_{i=1}^k p_i = 1, \quad p_i \geq 0 \\ & \sum_{j=1}^k q_j = 1, \quad q_j \geq 0 \end{aligned} \quad (18)$$

Since the maximum of several functions is not necessarily smooth (differentiable), we add a quadratic regularization term to the maximization problem, accelerating the convergence rate (Nouiehed et al., 2019) as follows:

$$\begin{aligned} \min_{\mathbf{w}} \max_{p_1, \dots, p_k, q_1, \dots, q_k} & \sum_{i=1}^k p_i f_i(\mathbf{w}) - \delta \sum_{i=1}^k p_i^2 + \sum_{j=1}^k q_j g_j(\mathbf{w}) - \delta \sum_{j=1}^k q_j^2 \\ \text{s.t.} & \sum_{i=1}^k p_i = 1, \quad p_i \geq 0 \\ & \sum_{j=1}^k q_j = 1, \quad q_j \geq 0 \end{aligned} \quad (19)$$

First, we show how to solve the inner maximization problem. Note that the p_i 's and q_i 's are independent. We show how to find optimal p_i 's. Optimizing with respect to q_i 's is similar. Since the maximization problem is a constrained quadratic program, we can write the Lagrangian function as follows:

$$\begin{aligned} & \max_{p_1, \dots, p_k} \sum_{i=1}^k p_i f_i(\mathbf{w}) - \delta \sum_{i=1}^k p_i^2 - \lambda \left(\sum_{i=1}^k p_i - 1 \right) \\ & \text{s.t. } p_i \geq 0 \end{aligned} \quad (20)$$

Having the optimal λ , the above problem has a closed-form solution with respect to each p_i , which can be written as:

$$p_i^* = \left[\frac{-\lambda + f_i}{2\delta} \right]_+$$

Since p_i^* is a non-increasing function with respect to λ , we can find the optimal value of λ using the following bisection algorithm. Algorithm 2 demonstrates how to find an ϵ -optimal λ and p_i^* 's efficiently using the bisection idea.

Algorithm 2 Finding the optimal λ and p_i 's using the bisection idea

- 1: **Initialize:** $\lambda_{\text{low}} = 0, \lambda_{\text{high}} = \max_i f_i, p_i = 0 \quad \forall i \in \{1, 2, \dots, k\}$.
 - 2: **while** $|\sum_{i=1}^k p_i - 1| > \epsilon$ **do**
 - 3: $\lambda = \frac{\lambda_{\text{low}} + \lambda_{\text{high}}}{2}$
 - 4: Set $p_i = \left[\frac{-\lambda + f_i}{2\delta} \right]_+ \quad \forall i \in \{1, 2, \dots, k\}$
 - 5: **if** $\sum_{i=1}^k p_i < 1$ **then**
 - 6: $\lambda_{\text{high}} = \lambda$
 - 7: **else**
 - 8: $\lambda_{\text{low}} = \lambda$
 - 9: **return** $\lambda, p_1, p_2, \dots, p_k$.
-

Remark 4. An alternative method for finding optimal λ , and p_i 's is to sort f_i values in $\mathcal{O}(k \log k)$ first, and then finding the smallest f_i such that if we set $\lambda = f_i$, the sum of p_i 's is bigger than 1 (let j be the index of that value). Without loss of generality, assume that $f_1 \leq \dots \leq f_k$. Then, $\sum_{i=j}^k \frac{-\lambda + f_i}{2\delta} = 1$, which has a closed-form solution with respect to λ .

To update \mathbf{w} , we need to solve the following optimization problem:

$$\min_{\mathbf{w}} \sum_{i=1}^k p_i^* f_i(\mathbf{w}) + \sum_{j=1}^k q_j^* g_j(\mathbf{w}), \quad (21)$$

Similar to the standard statistical learning framework, we solve the following empirical risk minimization problem by applying the gradient descent to \mathbf{w} on a finite data sample. Define \hat{f}_i as follows:

$$\hat{f}_i(\mathbf{w}) = \pi_1 \sum_{t=1}^n \left[-\log \left(\sigma(\mathbf{w}^T \mathbf{x}_t) \right) \right], \quad (22)$$

where $\mathbf{x}_1, \dots, \mathbf{x}_n$ are generated from the distribution $\mathcal{N}(\boldsymbol{\mu}_1^*, \boldsymbol{\Sigma}_{1i})$. The empirical risk minimization problem can be written as follows:

$$\min_{\mathbf{w}} \sum_{i=1}^k p_i^* \hat{f}_i(\mathbf{w}) + \sum_{j=1}^k q_j^* \hat{g}_j(\mathbf{w}), \quad (23)$$

Algorithm 3 summarizes the robust linear discriminant analysis method for the case where the label of all data points is available. Theorem 5 demonstrates the convergence of gradient descent algorithm applied to (23) in $\mathcal{O}\left(\frac{k}{\epsilon} \log\left(\frac{M}{\epsilon}\right)\right)$ iterations to an ϵ -optimal solution.

Algorithm 3 Robust Quadratic Discriminant Analysis in the Presence of Missing Values

-
- 1: **Input:** $\mathbf{X}_0, \mathbf{X}_1$: matrix of data points with labels 0 and 1 respectively, T : Number of iterations, α : Step-size.
 - 2: Estimate $\boldsymbol{\mu}_{\min_0}$ and $\boldsymbol{\mu}_{\max_0}$ using the available entries of \mathbf{X}_0 .
 - 3: Estimate $\boldsymbol{\mu}_{\min_1}$ and $\boldsymbol{\mu}_{\max_1}$ using the available entries of \mathbf{X}_1 .
 - 4: Estimate $\boldsymbol{\Sigma}_{01}, \dots, \boldsymbol{\Sigma}_{0k}$ using bootstrap estimator on the available data of \mathbf{X}_0 .
 - 5: Estimate $\boldsymbol{\Sigma}_{11}, \dots, \boldsymbol{\Sigma}_{1k}$ using bootstrap estimator on the available data of \mathbf{X}_1 .
 - 6: **for** $i = 1, \dots, T$ **do**
 - 7: Compute $\boldsymbol{\mu}_1^*$ and $\boldsymbol{\mu}_0^*$ by Equation (15).
 - 8: Find optimal p_1, \dots, p_k , and q_1, \dots, q_k using Algorithm 2.
 - 9: $\mathbf{w} = \mathbf{w} - \alpha \left(\sum_{i=1}^k p_i^* \nabla \hat{f}_i(\mathbf{w}) + \sum_{j=1}^k q_j^* \nabla \hat{g}_j(\mathbf{w}) \right)$
-

Theorem 5. Assume that $M = \max_i f_i$. Gradient descent algorithm requires $\mathcal{O}\left(\frac{k}{\epsilon} \log\left(\frac{M}{\epsilon}\right)\right)$ gradient evaluations for converging to an ϵ -optimal saddle point of the optimization problem (23).

In Appendix F, we extend the methodology to the case where y contains missing entries.

5 Experiments

In this section, we evaluate RIFLE’s performance on a diverse set of inference tasks in the presence of missing values. We compare RIFLE’s performance to several state-of-the-art approaches for data imputation on synthetic and real-world datasets. The experiments are designed in a manner that the sensitivity of the model to factors such as the number of samples, data dimension, types, and proportion of missing values can be evaluated. The description of all datasets used in the experiments can be found in Appendix I.

5.1 Evaluation Metrics

We need access to the ground-truth values of the missing entries to evaluate RIFLE and other state-of-the-art imputation approaches. Hence, we artificially mask a proportion of available data entries and predict them with different imputation methods. A method performs better than others if the predicted missing entries are closer to the ground-truth values. To measure the performance of RIFLE and the existing approaches on a regression task for a given test dataset consisting of N data points, we use normalized root mean squared error (NRMSE), defined as:

$$\text{NRMSE} = \frac{\sqrt{\frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2}}{\sqrt{\frac{1}{N} \sum_{i=1}^N (y_i - \bar{y})^2}}$$

where y_i , \hat{y}_i , and \bar{y} represent the true value of the i -th data point, the predicted value of the i -th data point, and the average of true values of data points, respectively. In all experiments, generated missing entries follow either a missing completely at random (MCAR) or a missing not at random (MNAR) pattern. A discussion on the procedure of generating these patterns can be found in Appendix G.

5.2 Tuning Hyper-parameters of RIFLE

The hyper-parameter c in (7) controls the robustness of the model by adjusting the size of confidence intervals. This parameter is tuned by performing a cross-validation procedure over the set $\{0.1, 0.25, 0.5, 1, 2, 5, 10, 20, 50, 100\}$, and the one with the lowest NMRSE is chosen. The default value in the implementation is $c = 1$ since it consistently performs well over different experiments. Furthermore, λ , the hyper-parameter for the ridge regression regularizer, is tuned by choosing 20% of the data as the validation set from the set $\{0.01, 0.1, 0.5, 1, 2, 5, 10, 20, 50\}$. To tune K , the number of bootstrap samples for estimating the confidence intervals, we tried 10, 20, 50, and 100. No significant difference is observed in terms of the test performance for the above values.

Furthermore, we tune the hyper-parameters of the competing packages as follows. For KNN-Imputer (Troyanskaya et al., 2001), we try $\{2, 10, 20, 50\}$ for the number of neighbors (K) and pick the one with the highest performance. For MICE (Buuren & Groothuis-Oudshoorn, 2010) and Amelia (Honaker et al., 2011), we generate 5 different imputed data and pick the one with the highest performance on the test data. MissForest has multiple hyper-parameters. We keep the criterion as ‘‘MSE’’ since our performance evaluation measure is NRMSE. Moreover, we tune the number of iterations and number of estimations (number of trees) by checking values from $\{5, 10, 20\}$ and $\{50, 100, 200\}$, respectively. We do not change the structure of the neural networks for MIDA (Gondara & Wang, 2018) and GAIN (Yoon et al., 2018), and the default versions are performed for imputing datasets.

5.3 RIFLE Consistency

In Theorem 2 Part (a), we demonstrated that RIFLE is consistent. In Figure 3, we investigate the consistency




Figure 3: Comparing the consistency of RIFLE, MissForest, KNN Imputer, MICE, Amelia, and Expectation Maximization methods on a synthetic dataset containing 40% of missing values.

of RIFLE on synthetic datasets with different proportions of missing values. The synthetic data has 50 input features following a jointly normal distribution with the mean whose entries are randomly chosen from the interval $(-100, 100)$. Moreover, the covariance matrix equals $\Sigma = SS^T$ where S elements are randomly picked from $(-1, 1)$. The dimension of S is 50×20 . The target variable is a linear function of input features added to a mean zero normal noise with a standard deviation of 0.01. As depicted in Figure 3, RIFLE requires fewer samples to recover the ground-truth parameters of the model compared to MissForest, KNN Imputer, Expectation Maximization (Dempster et al., 1977), and MICE. Amelia’s performance is significantly good since the predictors have a joint normal distribution and the linear underlying model. Note that by increasing the number of samples, the NRMSE of our framework converges to 0.01, which is the standard deviation of the zero-mean Gaussian noise added to each target value (the dashed line).

5.4 Data Imputation via RIFLE

As explained in Section 3, while the primary goal of RIFLE is to learn a robust regression model in the presence of missing values, it can also be used as an imputation tool. We run RIFLE and several state-of-the-art approaches on five datasets from the UCI repository (Dua & Graff, 2017) (Spam, Housing, Clouds, Breast Cancer, and Parkinson datasets) with different proportions of MCAR missing values (the description of the datasets can be found in Appendix I). Then, we compute the NRMSE of imputed entries. Table 1 shows the performance of RIFLE compared to other approaches for the datasets where the proportion of missing values are relatively high ($\frac{n(1-p)}{d} \approx \mathcal{O}(1)$). RIFLE outperforms these methods in almost all cases and performs slightly better than MissForest, which uses a highly non-linear model (random forest) to impute missing values.

Dataset Name	RIFLE	QRIFLE	MICE	Amelia	GAIN	MissForest	MIDA	EM
Spam (30%)	0.87 \pm 0.009	0.82 \pm 0.009	1.23 \pm 0.012	1.26 \pm 0.007	0.91 \pm 0.005	0.90 \pm 0.013	0.97 \pm 0.008	0.94 \pm 0.004
Spam (50%)	0.90 \pm 0.013	0.86 \pm 0.014	1.29 \pm 0.018	1.33 \pm 0.024	0.93 \pm 0.015	0.92 \pm 0.011	0.99 \pm 0.011	0.97 \pm 0.008
Spam (70%)	0.92 \pm 0.017	0.91 \pm 0.019	1.32 \pm 0.028	1.37 \pm 0.032	0.97 \pm 0.014	0.95 \pm 0.016	0.99 \pm 0.018	0.98 \pm 0.017
Housing (30%)	0.86 \pm 0.015	0.89 \pm 0.018	1.03 \pm 0.024	1.02 \pm 0.016	0.82 \pm 0.015	0.84 \pm 0.018	0.93 \pm 0.025	0.95 \pm 0.011
Housing (50%)	0.88 \pm 0.021	0.90 \pm 0.024	1.14 \pm 0.029	1.09 \pm 0.027	0.88 \pm 0.019	0.88 \pm 0.018	0.98 \pm 0.029	0.96 \pm 0.016
Housing (70%)	0.92 \pm 0.026	0.95 \pm 0.028	1.22 \pm 0.036	1.18 \pm 0.038	0.95 \pm 0.027	0.93 \pm 0.024	1.02 \pm 0.037	0.98 \pm 0.017
Clouds (30%)	0.81 \pm 0.018	0.79 \pm 0.019	0.98 \pm 0.024	1.04 \pm 0.027	<u>0.76</u> \pm 0.021	0.71 \pm 0.011	0.83 \pm 0.022	0.86 \pm 0.013
Clouds (50%)	0.84 \pm 0.026	0.84 \pm 0.028	1.10 \pm 0.041	1.13 \pm 0.046	<u>0.82</u> \pm 0.027	0.75 \pm 0.023	0.88 \pm 0.033	0.89 \pm 0.018
Clouds (70%)	0.87 \pm 0.029	0.90 \pm 0.033	1.16 \pm 0.044	1.19 \pm 0.048	0.89 \pm 0.035	0.81 \pm 0.031	0.93 \pm 0.044	0.92 \pm 0.023
Breast Cancer (30%)	0.52 \pm 0.023	0.54 \pm 0.027	0.74 \pm 0.031	0.81 \pm 0.032	0.58 \pm 0.024	0.55 \pm 0.016	0.70 \pm 0.026	0.67 \pm 0.014
Breast Cancer (50%)	0.56 \pm 0.026	0.59 \pm 0.027	0.79 \pm 0.029	0.85 \pm 0.033	0.64 \pm 0.025	0.59 \pm 0.022	0.76 \pm 0.035	0.69 \pm 0.022
Breast Cancer (70%)	0.59 \pm 0.031	0.65 \pm 0.034	0.86 \pm 0.042	0.92 \pm 0.044	0.70 \pm 0.037	0.63 \pm 0.028	0.82 \pm 0.035	0.67 \pm 0.014
Parkinson (30%)	0.57 \pm 0.016	0.55 \pm 0.016	0.71 \pm 0.019	0.67 \pm 0.021	0.53 \pm 0.015	0.54 \pm 0.010	0.62 \pm 0.017	0.64 \pm 0.011
Parkinson (50%)	0.62 \pm 0.022	0.64 \pm 0.025	0.77 \pm 0.029	0.74 \pm 0.034	0.61 \pm 0.022	0.65 \pm 0.014	0.71 \pm 0.027	0.69 \pm 0.022
Parkinson (70%)	0.67 \pm 0.027	0.74 \pm 0.033	0.85 \pm 0.038	0.82 \pm 0.037	<u>0.69</u> \pm 0.031	0.73 \pm 0.022	0.78 \pm 0.038	0.75 \pm 0.029

Table 1: Performance comparison of RIFLE, QRIFLE (Quadratic RIFLE), and state-of-the-art methods on several UCI datasets. We applied to impute methods on three different missing-value proportions for each dataset. The best imputer is highlighted with bold font, and the second-best imputer is underlined. Each experiment is done 5 times, and the average and the standard deviation of performances are reported.

5.5 Sensitivity of RIFLE to the Number of Samples and Proportion of Missing Values

In this section, we analyze the sensitivity of RIFLE and other state-of-the-art approaches to the number of samples and the proportion of missing values. In the experiment in Figure 4, we create 5 datasets containing

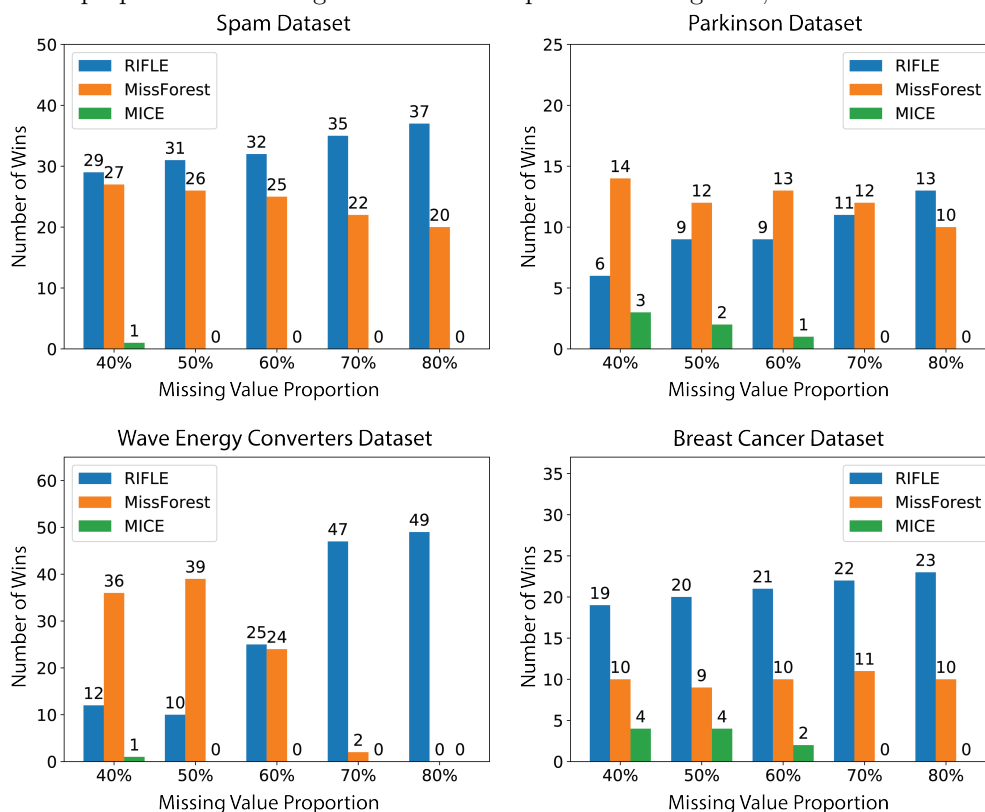


Figure 4: Performance Comparison of RIFLE, MICE, and MissForest on four UCI datasets: Parkinson, Spam, Wave Energy Converter, and Breast Cancer. For each dataset, we count the number of features that each method outperforms the others.

40%, 50%, 60%, 70%, and 80% of MCAR missing values, respectively, for four real datasets (Spam, Parkinson, Wave Energy Converter, and Breast Cancer) from UCI Repository (Dua & Graff, 2017) (the description of the datasets can be found in Appendix I). Given a feature in a dataset containing missing values, we say an imputer wins that feature if the imputation error in terms of NRMSE for that imputer is less than the error

of the other imputers. Figure 4 reports the number of features won by each imputer on the created datasets described above. As we observe, the number of wins for RIFLE increases as we increase the proportion of missing values. This observation shows that the sensitivity of RIFLE as an imputer to the proportion of missing values is less than MissForest and MICE in general.

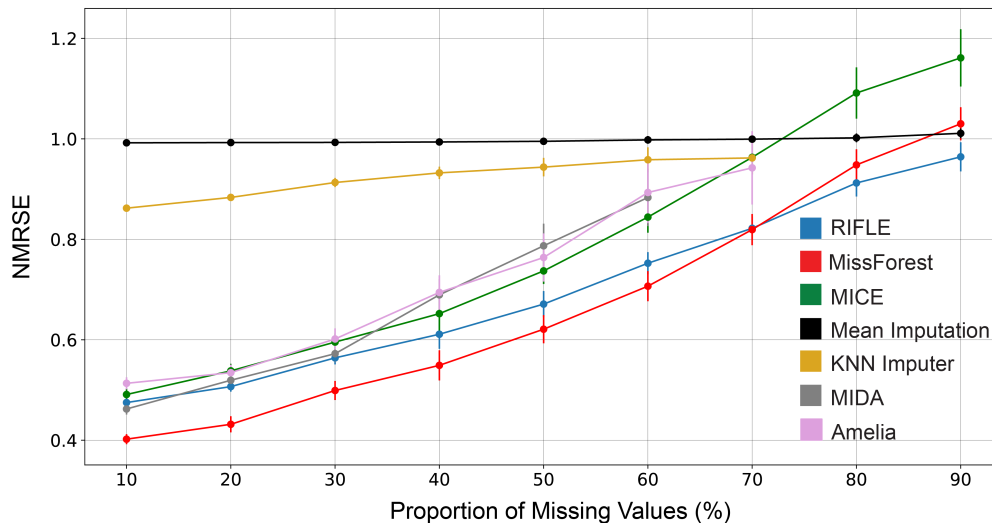


Figure 5: Sensitivity of RIFLE, MissForest, Amelia, KNN Imputer, MIDA, and Mean Imputer to the percentage of missing values on the Drive dataset. Increasing the percentage of missing value entries degrades the benchmarks’ performance compared to RIFLE. KNN-imputer implementation cannot be executed on datasets containing 80% (or more) missing entries. Moreover, Amelia and MIDA do not converge to a solution when the percentage of missing value entries is higher than 70%.

Figure 4 does not show how the increase in the proportion of missing values changes the NRMSE of imputers. Next, we analyze the sensitivity of RIFLE and several imputers to change in missing value proportions. Fixing the proportion of missing values, we generate 10 random datasets containing missing values in random locations on the Drive dataset (the description of datasets is available in Appendix I). We impute the missing values for each dataset with RIFLE, MissForest, Mean Imputation, and MICE. Figure 5 shows the average and the standard deviation of these 4 imputers’ performances for different proportions of missing values (10% to 90%). Figure 5 depicts the sensitivity of MissForest and RIFLE to the proportion of missing values in the Drive dataset. We select 400 data points for each experiment with different proportions of missing values (from 10% to 90%) and report the average NRMSE of imputed entries. Finally, in Figure 6, we have evaluated RIFLE and other methods on the BlogFeedback dataset (see Appendix I) containing 40% missing values. The results show that RIFLE’s performance is less sensitive to decreasing the number of samples.

5.6 Performance Comparison on Real Datasets

In this section, we compare the performance of RIFLE to several state-of-the-art approaches, including MICE (Buuren & Groothuis-Oudshoorn, 2010), Amelia (Honaker et al., 2011), MissForest (Stekhoven & Bühlmann, 2012), KNN Imputer (Raghunathan et al., 2001), and MIDA (Gondara & Wang, 2018). There are two primary ways to do this. One method to predict a continuous target variable in a dataset with many missing values is first to impute the missing data with a state-of-the-art package, then run a linear regression. An alternative approach is to directly learn the target variable, as we discussed in Section 3.

Table 2 compares the performance of mean imputation, MICE, MIDA, MissForest, and KNN to that of RIFLE on three datasets: NHANES, Blog Feedback, and superconductivity. Both Blog Feedback and Superconductivity datasets contain 30% of MNAR missing values generated by Algorithm 9, with 10000 and 20000 training samples, respectively. The description of the NHANES data and its distribution of missing values can be found in Appendix I.




Figure 6: Sensitivity of RIFLE, MissForest, MICE, Amelia, Mean Imputer, KNN Imputer, and MIDA to the number of samples for the imputations of Blog Feedback dataset containing 40% of MCAR missing values. When the number of samples is limited, RIFLE outperforms other methods, and its performance is very close to the non-linear imputer MissForest for larger samples.

Efficiency of RIFLE: We perform RIFLE for 1000 iterations and the step size of 0.01 in the above experiments. At each iteration, the main operation is to find the optimal θ for any given \mathbf{b} and \mathbf{C} . The average time of each method on each dataset is reported in Table 5 in Appendix L. The main reason for the time efficiency of RIFLE compared to MICE, MissForest, MIDA, and KNN Imputer is that it directly predicts the target variable without imputation of all missing entries.

Methods	Datasets		
	Super Conductivity	Blog Feedback	NHANES
Regression on Complete Data	0.4601	0.7432	0.6287
RIFLE	0.4873 ± 0.0036	0.8326 ± 0.0085	0.6304 ± 0.0027
Mean Imputer + Regression	0.6114 ± 0.0006	0.9235 ± 0.0003	0.6329 ± 0.0008
MICE + Regression	0.5078 ± 0.0124	0.8507 ± 0.0325	0.6612 ± 0.0282
EM + Regression	0.5172 ± 0.0162	0.8631 ± 0.0117	0.6392 ± 0.0122
MIDA Imputer + Regression	0.5213 ± 0.0274	0.8394 ± 0.0342	0.6542 ± 0.0164
MissForest	0.4925 ± 0.0073	0.8191 ± 0.0083	0.6365 ± 0.0094
KNN Imputer	0.5438 ± 0.0193	0.8828 ± 0.0124	0.6427 ± 0.0135

Table 2: Normalized RMSE of RIFLE and several state-of-the-art Methods on Superconductivity, blog feedback, and NHANES datasets. The first two datasets contain 30% Missing Not At Random (MNAR) missing values in the training phase generated by Algorithm 9. Each method applied 5 times to each dataset, and the result is reported as the average performance \pm standard deviation of experiments in terms of NRMSE.

Since MICE and MIDA cannot predict values during the test phase without data imputation, we use them in a pre-processing stage to impute the data. Then we apply the linear regression to the imputed dataset. On the other hand, RIFLE, KNN imputer, and MissForest can predict the target variable without imputing the training dataset. Table 2 shows that RIFLE outperforms all other state-of-the-art approaches executed on the three mentioned datasets. In particular, RIFLE outperforms MissForest, while the underlying model RIFLE uses is simpler (linear) compared to the nonlinear random forest model utilized by Missforest.

Number of Training Data Points	Method		
	LDA	Robust LDA	Robust QDA
50	52.38% \pm 3.91%	62.14% \pm 1.78%	61.36% \pm 1.62%
100	61.24% \pm 1.89%	68.46% \pm 1.04%	70.07% \pm 0.95%
200	73.49% \pm 0.97%	73.35% \pm 0.67%	73.51% \pm 0.52%

Table 3: Sensitivity of Linear Discriminant Analysis, Robust LDA (Common Covariance Matrices), and Robust QDA (Different Covariance matrices for two groups) to the number of training samples.

5.6.1 Performance of RIFLE on Classification Tasks

In Section 4, we discussed how to specialize RIFLE to robust normal discriminant analysis in the presence of missing values. Since the maximization problem over the second moments of the data (Σ) is intractable, we solved the maximization problem over a set of k covariance matrices estimated by bootstrap sampling. To investigate the effect of choosing k on the performance of the robust classifier, we train robust normal discriminant analysis models for different values of k on two training datasets (Avila and Magic) containing 40% MCAR missing values. The description of the datasets can be found in Appendix I. For $k = 1$, there is no maximization problem, and thus, it is equivalent to the classifier proposed in Fung & Wrobel (1989). As shown in Figure 7, increasing the number of covariance estimations generally enhances the accuracy of the classifier in the test phase. However, as shown in Theorem 5, the required time for completing the training phase grows linearly regarding the number of covariance estimations.




Figure 7: Effect of the number of covariance estimations on the performance (left) and run time (right) of robust LDA on Avila and Magic datasets. Increasing the number of covariance estimations (k) improves the model’s accuracy on the test data. However, it takes longer training time.

5.6.2 Comparison of Robust Linear Regression and Robust QDA

An alternative approach to the robust QDA presented in Section 4 is to apply the robust linear regression algorithm (Section 3) and mapping the solutions to each one of the classes by thresholding (positive value maps to Label 1 and negative values to label -1).

Table 4 compares the performance of two classifiers on three different datasets. As demonstrated in the table, when all features are continuous, quadratic discriminant analysis has a better performance. It shows the QDA model relies highly on the normality assumption, while robust linear regression handles the categorical features better than robust QDA.

Limitations and Future Directions: The proposed framework for robust regression in the presence of missing values is limited to linear models. While in Appendix E, we use polynomial kernels to apply non-linear transformations on the data, such an approach can potentially increase the number of missing values in the

Dataset	Feature Type	Accuracy of Methods					
		RIFLE	Robust QDA	MissForest	MICE	KNN Imputer	EM
Glass Identification	Continuous	67.12% \pm 1.84%	69.54% \pm 1.97%	65.76% \pm 1.49%	62.48% \pm 2.45%	60.37% \pm 1.12%	68.21% \pm 0.94%
Annealing	Mixed	63.41% \pm 2.44%	59.51% \pm 2.21%	64.91% \pm 1.35%	60.66% \pm 1.59%	57.44% \pm 1.44%	59.43% \pm 1.29%
Abalone	Mixed	68.41% \pm 0.74%	63.27% \pm 0.76%	69.40% \pm 0.42%	63.12% \pm 0.98%	62.43% \pm 0.38%	62.91% \pm 0.37%
Lymphography	Discrete	66.32% \pm 1.05%	58.15% \pm 1.21%	66.11% \pm 0.94%	55.73% \pm 1.24	57.39% \pm 0.88%	59.55% \pm 0.68%
Adult	Discrete	72.42% \pm 0.06%	60.36% \pm 0.08	70.34% \pm 0.03%	63.30% \pm 0.14%	60.14% \pm 0.00	60.69% \pm 0.01%

Table 4: Accuracy of RIFLE, MICE, KNN-Imputer, Expectation Maximization (EM), and Robust QDA on different discrete, mixed, and continuous datasets. Robust QDA can perform better than other methods when the input features are continuous, and the target variable is discrete. However, RIFLE results in higher accuracy in mixed and discrete settings.

kernel space generated by the composition of the original features. A future direction is to develop efficient algorithms for non-linear regression models such as multi-layer neural networks, decision tree regressors, gradient boosting regressors, and support vector regression models. In the case of robust classification, the methodology is extendable to any loss beyond quadratic discriminant analysis. Unlike the regression case, a limitation of the proposed method for robust classification is its reliance on the Gaussianity assumption of data distribution (conditioned on each data label). A natural extension is to assume the underlying data distribution follows a mixture of Gaussian distributions.

Conclusion: In this paper, we proposed a distributionally robust optimization framework over the distributions with the low-order marginals within the estimated confidence intervals for inference and imputation of datasets in the presence of missing values. We developed algorithms for regression and classification with convergence guarantees. The method’s performance is evaluated on synthetic and real datasets with different numbers of samples, dimensions, missing value proportions, and types of missing values. In most experiments, RIFLE consistently outperforms other existing methods.

Acknowledgments

This work was supported by the NIH/NSF Grant 1R01LM013315-01, the NSF CAREER Award CCF-2144985, and the AFOSR Young Investigator Program Award FA9550-22-1-0192.

References

- Mohamed Reda Abonazel and Mohamed Gamal Ibrahim. On estimation methods for binary logistic regression model with missing values. *International Journal of Mathematics and Computational Science*, 4(3): 79–85, 2018.
- Jakob Assländer, Martijn A Cloos, Florian Knoll, Daniel K Sodickson, Jürgen Hennig, and Riccardo Lattanzi. Low rank alternating direction method of multipliers reconstruction for mr fingerprinting. *Magnetic resonance in medicine*, 79(1):83–96, 2018.
- Babak Barazandeh and Meisam Razaviyayn. Solving non-convex non-differentiable min-max games using proximal gradient method. *arXiv preprint arXiv:2003.08093*, 2020.
- Brett K Beaulieu-Jones, Daniel R Lavage, John W Snyder, Jason H Moore, Sarah A Pendergrass, and Christopher R Bauer. Characterizing and managing missing structured data in electronic health records: data analysis. *JMIR medical informatics*, 6(1):e11, 2018.
- Dimitris Bertsimas, Colin Pawlowski, and Ying Daisy Zhuo. From predictive methods to missing data imputation: an optimization approach. *The Journal of Machine Learning Research*, 18(1):7133–7171, 2017.
- Sébastien Bubeck. Convex optimization: Algorithms and complexity. *arXiv preprint arXiv:1405.4980*, 2014.
- S van Buuren and Karin Groothuis-Oudshoorn. mice: Multivariate imputation by chained equations in r. *Journal of statistical software*, pp. 1–68, 2010.
- Zhipeng Cai, Maysam Heydari, and Guohui Lin. Iterated local least squares microarray missing value imputation. *Journal of bioinformatics and computational biology*, 4(05):935–957, 2006.
- John M Danskin. *The theory of max-min and its application to weapons allocation problems*, volume 5. Springer Science & Business Media, 2012.
- Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1):1–22, 1977.
- Dheeru Dua and Casey Graff. UCI machine learning repository, 2017. URL <http://archive.ics.uci.edu/ml>.
- Karen Yuen Fung and Barbara A Wrobel. The treatment of missing values in logistic regression. *Biometrical Journal*, 31(1):35–47, 1989.
- Daniel Gabay and Bertrand Mercier. A dual algorithm for the solution of nonlinear variational problems via finite element approximation. *Computers & mathematics with applications*, 2(1):17–40, 1976.
- Rui Gao and Anton J Kleywegt. Distributionally robust stochastic optimization with dependence structure. *arXiv preprint arXiv:1701.04200*, 2017.
- Zoubin Ghahramani and Michael I Jordan. Supervised learning from incomplete data via an em approach. In *Advances in neural information processing systems*, pp. 120–127, 1994.
- Lovedeep Gondara and Ke Wang. MIDA: Multiple imputation using denoising autoencoders. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pp. 260–272. Springer, 2018.
- Bingsheng He and Xiaoming Yuan. On non-ergodic convergence rate of douglas–rachford alternating direction method of multipliers. *Numerische Mathematik*, 130(3):567–577, 2015.
- James Honaker, Gary King, and Matthew Blackwell. Amelia ii: A program for missing data. *Journal of statistical software*, 45(7):1–47, 2011.
- Mingyi Hong, Zhi-Quan Luo, and Meisam Razaviyayn. Convergence analysis of alternating direction method of multipliers for a family of nonconvex problems. *SIAM Journal on Optimization*, 26(1):337–364, 2016.

- Hyunsoo Kim, Gene H Golub, and Haesun Park. Missing value estimation for dna microarray gene expression data: local least squares imputation. *Bioinformatics*, 21(2):187–198, 2005.
- Roderick JA Little and Donald B Rubin. *Statistical analysis with missing data*, volume 793. John Wiley & Sons, 2019.
- Yurii E Nesterov. A method for solving the convex programming problem with convergence rate $O(1/k^2)$. In *Dokl. akad. nauk Sssr*, volume 269, pp. 543–547, 1983.
- Maher Nouiehed, Maziar Sanjabi, Tianjian Huang, Jason D Lee, and Meisam Razaviyayn. Solving a class of non-convex min-max games using iterative first order methods. *Advances in Neural Information Processing Systems*, 32:14934–14942, 2019.
- AB Pedersen, EM Mikkelsen, D Cronin-Fenton, NR Kristensen, TM Pham, L Pedersen, et al. Missing data and multiple imputation in clinical epidemiological research. *clin epidemiol.* 2017; 9: 157-66, 2017.
- Trivellore E Raghunathan, James M Lepkowski, John Van Hoewyk, Peter Solenberger, et al. A multivariate technique for multiply imputing missing values using a sequence of regression models. *Survey methodology*, 27(1):85–96, 2001.
- Joseph L Schafer. *Analysis of incomplete multivariate data*. Chapman and Hall/CRC, 1997.
- Pannagadatta K Shivaswamy, Chiranjib Bhattacharyya, and Alexander J Smola. Second order cone programming approaches for handling missing and uncertain data. *Journal of Machine Learning Research*, 7 (Jul):1283–1314, 2006.
- Maurice Sion et al. On general minimax theorems. *Pacific Journal of mathematics*, 8(1):171–176, 1958.
- Daniel J Stekhoven and Peter Bühlmann. Missforest—non-parametric missing value imputation for mixed-type data. *Bioinformatics*, 28(1):112–118, 2012.
- Jonathan AC Sterne, Ian R White, John B Carlin, Michael Spratt, Patrick Royston, Michael G Kenward, Angela M Wood, and James R Carpenter. Multiple imputation for missing data in epidemiological and clinical research: potential and pitfalls. *Bmj*, 338:b2393, 2009.
- Olga Troyanskaya, Michael Cantor, Gavin Sherlock, Pat Brown, Trevor Hastie, Robert Tibshirani, David Botstein, and Russ B Altman. Missing value estimation methods for dna microarrays. *Bioinformatics*, 17 (6):520–525, 2001.
- Jing Xia, Shengyu Zhang, Guolong Cai, Li Li, Qing Pan, Jing Yan, and Gangmin Ning. Adjusted weight voting algorithm for random forests in handling missing values. *Pattern Recognition*, 69:52–60, 2017.
- Huan Xu, Constantine Caramanis, and Shie Mannor. Robustness and regularization of support vector machines. *Journal of machine learning research*, 10(7), 2009.
- Jinsung Yoon, James Jordon, and Mihaela Van Der Schaar. Gain: Missing data imputation using generative adversarial nets. *arXiv preprint arXiv:1806.02920*, 2018.
- Tianyun Zhang, Shaokai Ye, Kaiqi Zhang, Jian Tang, Wujie Wen, Makan Fardad, and Yanzhi Wang. A systematic dnn weight pruning framework using alternating direction method of multipliers. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 184–199, 2018.
- Xiaobai Zhang, Xiaofeng Song, Huinan Wang, and Huanping Zhang. Sequential local least squares imputation estimating missing value of microarray data. *Computers in biology and medicine*, 38(10):1112–1120, 2008.

A A Review of Missing Value Imputation Methods in the Literature

The fundamental idea behind many data imputation approaches is that the missing values can be predicted based on the available data of other data points and correlated features. One of the most straightforward imputation techniques is to replace missing values by the mean (or median) of that feature calculated from what data is available see Little & Rubin (2019, Chapter 3). However, this naïve approach ignores the correlation between features and does not preserve the variance of features. Another class of imputers has been developed based on the least-square methods (Raghunathan et al., 2001; Kim et al., 2005; Zhang et al., 2008; Cai et al., 2006). Raghunathan et al. (2001) learns a linear model with multivariate Gaussian noise for the feature with the least missing entries. It repeats the same procedure on the updated data to impute the next feature with the least missing entries until all features are completely imputed. One drawback of this approach is that the error from the imputation of previous features can be propagated to subsequent features. To impute entries of a given feature in a dataset, Kim et al. (2005) learns several univariate regression models that consider that feature as the response. Then it takes the average of these predictions as the final value of imputation. This approach fails to learn the correlations involving more than two features.

Many more complex algorithms have been developed for imputation, although many are sensitive to initial assumptions and may not converge. For instance, KNN-Imputer imputes a missing feature of a data point by taking the mean value of the K closest complete data points (Troyanskaya et al., 2001). MissForest, on the other hand, imputes the missing values of each feature by learning a random forest classifier using other training data features (Stekhoven & Bühlmann, 2012). MissForest does not need to assume that all features are continuous (Honaker et al., 2011) or categorical (Schafer, 1997). However, both KNN-imputer and MissForest do not guarantee statistical or computational convergence for their algorithms. Moreover, when the proportion of missing values is high, both are likely to have a severe drop in performance, as demonstrated in Section 5. The Expectation Maximization (EM) algorithm is another popular approach that learns the parameters of a prior distribution on the data using available values based on the EM algorithm of Dempster et al. (1977); see also Ghahramani & Jordan (1994) and Honaker et al. (2011). The EM algorithm is also used in Amelia, which fits a jointly normal distribution to the data using EM and the bootstrap technique (Honaker et al., 2011). While Amelia demonstrates a superior performance on datasets following a normal distribution, it is highly sensitive to the violation of the normality assumption (as discussed in Bertsimas et al. (2017)). Ghahramani & Jordan (1994) adopt the EM algorithm to learn a joint Bernoulli distribution for the categorical data and a joint Gaussian distribution for the continuous variables independently. While those algorithms can be viewed as inference methods based on low-order estimates of moments, they do not consider *uncertainty* in such low-order moments estimates. By contrast, our framework utilizes *robust optimization* to consider the uncertainty around the estimated moments. Moreover, our optimization procedure for imputation and prediction is guaranteed to converge despite some of the algorithms mentioned above.

Another popular method for data imputation is multiple imputations by chained equations (MICE). MICE learns a parametric distribution for each feature conditional on the remaining features. For instance, it assumes that the current target variable is a linear function of other features with a zero-mean Gaussian noise. Each feature can have its distinct distribution and parameters (e.g., Poisson regression, logistic regression). Based on the learned parameters of conditional distributions, MICE can generate one or more imputed datasets (Buuren & Groothuis-Oudshoorn, 2010). More recently, several neural network-based imputers have been proposed. GAIN (Generative Adversarial Imputation Network) learns a generative adversarial network based on the available data and then imputes the missing values using the trained generator (Yoon et al., 2018). One advantage of GAIN over other existing GAN imputers is that it does not need a complete dataset during the training phase. MIDA (Multiple Imputation using Denoising Autoencoders) is an auto-encoder-based approach that trains a denoising auto-encoder on the available data considering the missing entries as noise. Similar to other neural network-based methods, these algorithms suffer from their black-box nature. They are challenging to interpret/explain, making them unpopular in mission-critical healthcare approaches. In addition, no statistical or computational guarantees are provided for these algorithms.

Bertsimas et al. (2017) formulates the imputation task as a constrained optimization problem where the constraints are determined by the underlying classification model such as KNN (k -nearest neighbors), SVM

(Support Vector Machine), and Decision Trees. Their general framework is non-convex, and the authors relax the optimization for each choice of the cost function using first-order methods. The block coordinate descent algorithm then optimizes the relaxed problem. They show the convergence and accuracy of their proposed algorithm numerically, while a theoretical analysis that guarantees the algorithm’s convergence is absent in their work.

B Estimating Confidence Intervals of Low-order Moments

In this section, we explain the methodology of estimating confidence intervals for $\mathbb{E}[\mathbf{z}_i]$ and $\mathbb{E}[\mathbf{z}_i\mathbf{z}_j]$. Let $\mathbf{X}^{n \times d}$ and \mathbf{y} be the data matrix and target variables for n given data points respectively whose entries are in $\tilde{\mathbb{R}} = \mathbb{R} \cup \{*\}$, where $*$ symbol represents a missing entry. Moreover, assume that \mathbf{a}_i represents the i -th column (feature) of matrix \mathbf{X} . We define:

$$\tilde{\mathbf{a}}_i(k) = \begin{cases} \mathbf{a}_i(k) & \text{if } \mathbf{a}_i(k) \neq * \\ 0 & \text{if } \mathbf{a}_i(k) = * \end{cases}$$

Thus, $\tilde{\mathbf{a}}$ is obtained by replacing the missing values with 0. We estimate the confidence intervals for the mean and covariance of features using multiple bootstrap samples on the available data. Let $\mathbf{C}_0[i][j]$ and $\Delta_0[i][j]$ be the center and the radius of the confidence interval for $\mathbf{C}[i][j]$, respectively. We compute the center of the confidence interval for $\mathbf{C}[i][j]$ as follows:

$$\mathbf{C}_0[i][j] = \frac{1}{m_{ij}} \tilde{\mathbf{a}}_i^T \tilde{\mathbf{a}}_j \quad (24)$$

where $m_i = |\{k : \mathbf{a}_i(k) \neq *\}|$ and $m_{ij} = |\{k : \mathbf{a}_i(k) \neq *, \mathbf{a}_j(k) \neq *\}|$. This estimator is obtained from the rows where both features are available. More precisely, let \mathbf{M} be the mask of the input data matrix \mathbf{X} defined as:

$$\mathbf{M}_{ij} = \begin{cases} 0, & \text{if } \mathbf{X}_{ij} \text{ is missing,} \\ 1, & \text{otherwise.} \end{cases}$$

Assume that $m_{ij} = (\mathbf{M}^T \mathbf{M})_{ij}$, which is the number of rows in the dataset where both features i and j are available. To estimate the confidence intervals for \mathbf{C}_{ij} , we use Algorithm 4. First, we select multiple (K) samples of size $N = m_{ij}$ from the rows where both features are available. Each one of these samples with size m_{ij} is obtained by applying a bootstrap sampler (sampling with replacement) on the m_{ij} rows where both features are available. Then, we compute the second-order moment of two features for each sample.

To find the radius of confidence intervals for each given pair (i, j) of features, we choose k different bootstrap samples with length n on the rows where both features i and j are available. Then, we compute $\mathbf{C}_0[i][j]$ of two features in each bootstrap sample. The standard deviation of these estimations determines the radius of the corresponding confidence interval. Algorithm 4 summarizes the required steps for computing the confidence interval radius for the ij -th entry of covariance matrix Δ . Note that the confidence intervals for $\boldsymbol{\mu}$ can be computed similarly. Having \mathbf{C}_0 and Δ , the confidence interval for the matrix \mathbf{C} is computed as

Algorithm 4 Estimating Confidence Interval Length Δ_{ij} for Feature i and Feature j .

- 1: **Input:** K : Number of bootstrap estimations
 - 2: **for** $t = 1, \dots, K$ **do**
 - 3: Pick n samples with replacement from the rows where both i -th and j -th are available.
 - 4: Let $(\hat{X}_{i1}, \hat{X}_{j1}), \dots, (\hat{X}_{in}, \hat{X}_{jn})$ be the i -th and j -th features of the selected samples
 - 5: $C_t = \frac{1}{n} \sum_{r=1}^n \hat{X}_{ir} \hat{X}_{jr}$
 - 6: $\Delta_{ij} = \text{std}(C_1, C_2, \dots, C_K)$
-

follows:

$$\begin{aligned} \mathbf{C}_{\min} &= \mathbf{C}_0 - c\Delta \\ \mathbf{C}_{\max} &= \mathbf{C}_0 + c\Delta, \end{aligned}$$

Computing \mathbf{b}_{\min} and \mathbf{b}_{\max} can be done in the same manner. The hyper-parameter c is defined to control the robustness of the model by tuning the length of confidence intervals. A larger c corresponds to bigger confidence intervals and, thus, a more robust estimator. On the other hand, large values for c lead to very large confidence intervals that can adversely affect the performance of the trained model.

Remark 6. *Since the computation of confidence intervals for different entries of the covariance matrix are independent of each other, they can be computed in parallel. In particular, if γ cores are available, $\lceil d/\gamma \rceil$ features (columns of the covariance matrix) can be assigned to each one of the available cores.*

C Solving Robust Ridge Regression with the Optimal Convergence Rate

The convergence rate of Algorithm 1 to the optimal solution of Problem (6) can be slow in practice since the algorithm requires to do a matrix inversion for updating $\boldsymbol{\theta}$ and applying the box constraint to \mathbf{C} and \mathbf{b} at each iteration. While we update the minimization problem in closed-form with respect to $\boldsymbol{\theta}$, we can speed up the convergence rate of the maximization problem by applying Nesterov’s acceleration method to function $g(\mathbf{b}, \mathbf{C})$ in (7). Since function g is the minimum of convex functions, its gradient with respect to \mathbf{C} and \mathbf{b} can be computed using Danskin’s theorem. Algorithm 5 describes the steps to optimize Problem (7) using Nesterov’s acceleration method.

Algorithm 5 Applying the Nesterov’s Acceleration Method to Robust Linear Regression

- 1: $\mathbf{C}_0, \mathbf{b}_0, \Delta, \delta, T$
 - 2: **Initialize:** $\mathbf{C}_1 = \mathbf{C}_0, \mathbf{b}_1 = \mathbf{b}_0, \gamma_0 = 0, \gamma_1 = 1.$
 - 3: **for** $i = 1, \dots, T$ **do**
 - 4: $\gamma_{i+1} = \frac{1 + \sqrt{1 + 4\gamma_i^2}}{2}$
 - 5: $Y_{\mathbf{C}_i} = \mathbf{C}_i + \frac{\gamma_i - 1}{\gamma_{i+1}}(\mathbf{C}_i - \mathbf{C}_{i-1})$
 - 6: $\mathbf{C}_{i+1} = \Pi_{\Delta+} \left(Y_{\mathbf{C}_i} + \frac{1}{L} \boldsymbol{\theta} \boldsymbol{\theta}^T \right)$
 - 7: $Y_{\mathbf{b}_i} = \mathbf{b}_i + \frac{\gamma_i - 1}{\gamma_{i+1}}(\mathbf{b}_i - \mathbf{b}_{i-1})$
 - 8: $\mathbf{b}_{i+1} = \Pi_{\delta} \left(Y_{\mathbf{b}_i} - \frac{2\boldsymbol{\theta}}{L} \right)$
 - 9: Set $\boldsymbol{\theta} = (\mathbf{C}_{i+1} + \lambda I)^{-1} \mathbf{b}_{i+1}$
-

Theorem 7. *Let $(\tilde{\boldsymbol{\theta}}, \tilde{\mathbf{C}}, \tilde{\mathbf{b}})$ be the optimal solution of (6) and $D = \|\mathbf{C}_0 - \tilde{\mathbf{C}}\|_F^2 + \|\mathbf{b}_0 - \tilde{\mathbf{b}}\|_2^2$. Assume that for any given \mathbf{b} and \mathbf{C} , within the uncertainty sets described in (6), $\|\boldsymbol{\theta}^*(\mathbf{b}, \mathbf{C})\| \leq \tau$. Then, Algorithm 1 computes an ϵ -optimal solution of the objective function in $\mathcal{O}\left(\sqrt{\frac{D(\tau+1)^2}{\lambda\epsilon}}\right)$ iterations.*

Proof. The proof is relegated to Appendix H. □

D Solving the Dual Problem of the Robust Ridge Linear Regression via ADMM

The Alternating Direction Method of Multipliers (ADMM) is a popular algorithm for efficiently solving linearly constrained optimization problems (Gabay & Mercier, 1976; Hong et al., 2016). It has been extensively applied to large-scale optimization problems in machine learning and statistical inference in recent years (Assländer et al., 2018; Zhang et al., 2018). Consider the following optimization problem consisting of two blocks of variables \mathbf{x} and \mathbf{y} that are linearly coupled:

$$\begin{aligned} \min_{\mathbf{w}, \mathbf{z}} \quad & f(\mathbf{w}) + g(\mathbf{z}) \\ \text{s.t.} \quad & \mathbf{A}\mathbf{w} + \mathbf{B}\mathbf{z} = \mathbf{c}, \end{aligned} \tag{25}$$

The augmented Lagrangian of the above problem can be written as:

$$\min_{\mathbf{w}, \mathbf{z}} f(\mathbf{w}) + g(\mathbf{z}) + \langle \mathbf{A}\mathbf{w} + \mathbf{B}\mathbf{z} - \mathbf{c}, \boldsymbol{\lambda} \rangle + \frac{\rho}{2} \|\mathbf{A}\mathbf{w} + \mathbf{B}\mathbf{z} - \mathbf{c}\|^2 \tag{26}$$

ADMM schema updates the primal and dual variables iteratively as presented in Algorithm 6.

Algorithm 6 General ADMM Algorithm

- 1: **for** $t = 1, \dots, T$ **do**
 - 2: $\mathbf{w}^{t+1} = \arg \min_{\mathbf{w}} f(\mathbf{w}) + \langle \mathbf{A}\mathbf{w} + \mathbf{B}\mathbf{z}^t - \mathbf{c}, \boldsymbol{\lambda} \rangle + \frac{\rho}{2} \|\mathbf{A}\mathbf{w} + \mathbf{B}\mathbf{z}^t - \mathbf{c}\|^2$
 - 3: $\mathbf{z}^{t+1} = \arg \min_{\mathbf{z}} f(\mathbf{w}^{t+1}) + \langle \mathbf{A}\mathbf{w}^{t+1} + \mathbf{B}\mathbf{z} - \mathbf{c}, \boldsymbol{\lambda} \rangle + \frac{\rho}{2} \|\mathbf{A}\mathbf{w}^{t+1} + \mathbf{B}\mathbf{z} - \mathbf{c}\|^2$
 - 4: $\boldsymbol{\lambda}^{t+1} = \boldsymbol{\lambda}^t + \rho(\mathbf{A}\mathbf{w}^{t+1} + \mathbf{B}\mathbf{z}^{t+1} - \mathbf{c})$
-

As we mentioned earlier, simultaneous projection of \mathbf{C} to the set of positive semi-definite matrices and the box constraint $\mathbf{C}_{\min} \leq \mathbf{C} \leq \mathbf{C}_{\max}$ in Algorithm 1 and Algorithm 5 is computationally expensive. Moreover, careful step-size tuning is necessary to avoid inconsistency and guarantee convergence in that algorithm.

An alternative approach for solving Problem (6) that avoids removing the PSD constraint in the implementation of Algorithm 1 and Algorithm 5 is to solve the dual of the inner maximization problem. Since the maximization problem is concave with respect to \mathbf{C} and \mathbf{b} , and the relative interior of the feasible set of constraints is non-empty, the duality gap is zero. Hence, instead of solving the inner maximization problem, we can solve its dual which is a minimization problem. Theorem 8 describes the dual problem of the inner maximization problem in (6). Thus, Problem (6) can be alternatively formulated as a minimization problem rather than a min-max problem. We can solve such a constrained minimization problem efficiently via the ADMM algorithm. As we will show, the ADMM algorithm applied to the dual problem does not need tuning of step-size or applying simultaneous projections to the box constraints and positive semi-definite (PSD) constraints.

Theorem 8. (*Dual Problem*) *The inner maximization problem described in (6) can be equivalently formulated as:*

$$\begin{aligned}
 & \min_{\mathbf{A}, \mathbf{B}, \mathbf{d}, \mathbf{e}, \mathbf{H}} -\langle \mathbf{b}_{\min}, \mathbf{d} \rangle + \langle \mathbf{b}_{\max}, \mathbf{e} \rangle - \langle \mathbf{C}_{\min}, \mathbf{A} \rangle + \langle \mathbf{C}_{\max}, \mathbf{B} \rangle + \lambda \|\boldsymbol{\theta}\|^2 \\
 \text{s.t.} \quad & -\boldsymbol{\theta}\boldsymbol{\theta}^T - \mathbf{A} + \mathbf{B} - \mathbf{H} = 0, \\
 & 2\boldsymbol{\theta} - \mathbf{d} + \mathbf{e} = 0, \\
 & \mathbf{A}, \mathbf{B}, \mathbf{d}, \mathbf{e} \geq 0, \\
 & \mathbf{H} \succeq 0.
 \end{aligned}$$

Therefore, Problem (6) can be alternatively written as:

$$\begin{aligned}
 & \min_{\boldsymbol{\theta}, \mathbf{A}, \mathbf{B}, \mathbf{d}, \mathbf{e}, \mathbf{H}} -\langle \mathbf{b}_{\min}, \mathbf{d} \rangle + \langle \mathbf{b}_{\max}, \mathbf{e} \rangle - \langle \mathbf{C}_{\min}, \mathbf{A} \rangle + \langle \mathbf{C}_{\max}, \mathbf{B} \rangle + \lambda \|\boldsymbol{\theta}\|^2 \\
 \text{s.t.} \quad & -\boldsymbol{\theta}\boldsymbol{\theta}^T - \mathbf{A} + \mathbf{B} - \mathbf{H} = 0, \\
 & 2\boldsymbol{\theta} - \mathbf{d} + \mathbf{e} = 0, \\
 & \mathbf{A}, \mathbf{B}, \mathbf{d}, \mathbf{e} \geq 0, \\
 & \mathbf{H} \succeq 0.
 \end{aligned} \tag{27}$$

Proof. The proof is relegated to Appendix H. □

To apply the ADMM method to the dual problem, we require to divide the optimization variables into two blocks as in (25) such that both sub-problems in Algorithm 6 can be efficiently solved. To do so, first, we introduce the auxiliary variables \mathbf{d}' , \mathbf{e}' , $\boldsymbol{\theta}'$, \mathbf{A}' and \mathbf{B}' to the dual problem. Also, let $\mathbf{G} = \mathbf{H} + \boldsymbol{\theta}'\boldsymbol{\theta}'^T$.

Therefore, Problem (27) is equivalent to:

$$\begin{aligned}
& \min_{\boldsymbol{\theta}, \mathbf{A}, \mathbf{B}, \mathbf{d}, \mathbf{e}, \mathbf{H}} && -\langle \mathbf{b}_{\min}, \mathbf{d} \rangle + \langle \mathbf{b}_{\max}, \mathbf{e} \rangle - \langle \mathbf{C}_{\min}, \mathbf{A} \rangle + \langle \mathbf{C}_{\max}, \mathbf{B} \rangle + \lambda \|\boldsymbol{\theta}\|^2 \\
& \text{s.t.} && \mathbf{B} - \mathbf{A} = \mathbf{G}, \\
& && 2\boldsymbol{\theta} - \mathbf{d} + \mathbf{e} = 0, \\
& && \mathbf{A} = \mathbf{A}', \mathbf{B} = \mathbf{B}', \\
& && \mathbf{d} = \mathbf{d}', \mathbf{e} = \mathbf{e}', \boldsymbol{\theta} = \boldsymbol{\theta}', \\
& && \mathbf{A}', \mathbf{B}', \mathbf{d}', \mathbf{e}' \geq 0, \\
& && \mathbf{G} \succeq \boldsymbol{\theta}' \boldsymbol{\theta}'^T.
\end{aligned} \tag{28}$$

Since handling both constraints on $\boldsymbol{\theta}$ in Problem (27) is difficult, we interchange $\boldsymbol{\theta}$ with $\boldsymbol{\theta}'$ in the first constraint. Moreover, the non-negativity constraints on $\mathbf{A}, \mathbf{B}, \mathbf{d}$ and \mathbf{e} are exchanged with non-negativity constraints on $\mathbf{A}', \mathbf{B}', \mathbf{d}'$ and \mathbf{e}' . For the simplicity of presentation, assume that $\mathbf{c}_1^t = \mathbf{b}_{\min} - \boldsymbol{\mu}_d^t + \rho \mathbf{d}'^t + \boldsymbol{\eta}^t$, $\mathbf{c}_2^t = -\mathbf{b}_{\max} - \boldsymbol{\mu}_e^t + \rho \mathbf{e}'^t - \boldsymbol{\eta}^t$, $\mathbf{c}_3^t = -\boldsymbol{\mu}_{\boldsymbol{\theta}} + \rho \boldsymbol{\theta}'^t - 2\boldsymbol{\eta}^t$, $\mathbf{D}_1^t = \rho \mathbf{A}'^t - \rho \mathbf{G}^t + \boldsymbol{\Gamma}^t - \mathbf{M}_A^t + \mathbf{C}_{\min}$, and $\mathbf{D}_2^t = \rho \mathbf{B}'^t + \rho \mathbf{G}^t - \boldsymbol{\Gamma}^t - \mathbf{M}_B^t - \mathbf{C}_{\max}$. Algorithm 7 describes the ADMM algorithm applied to Problem (28).

Corollary 9. *If the feasible set of Problem (6) has non-empty interior, then Algorithm 7 converges to an ϵ -optimal solution of Problem (28) in $\mathcal{O}(\frac{1}{\epsilon})$ iterations.*

Proof. Since the inner maximization problem, in (6) is convex, and its feasible interior set is not empty, the duality gap is zero by Slater's condition. Thus, according to Theorem 6.1 in He & Yuan (2015), Algorithm 7 converges to an optimal solution of the primal-dual problem with a linear rate. Moreover, the sequence of constraint residuals converges to zero with a linear rate as well. \square

Remark 10. *The optimal solution obtained from the ADMM algorithm can be different from the one given by Algorithm 1 because we remove the positive semi-definite constraint on \mathbf{C} in the latter. We investigate the difference between solutions of two algorithms in three cases: First, we generate a small positive semi-definite matrix \mathbf{C}^* and the matrix of confidence intervals ($\boldsymbol{\Delta}$) as follows:*

$$\mathbf{C}^* = \begin{bmatrix} 97 & 40 & 92 \\ 40 & 17 & 38 \\ 92 & 38 & 88 \end{bmatrix}, \quad \boldsymbol{\Delta} = \begin{bmatrix} 0.2 & 0.3 & 0.2 \\ 0.3 & 0.1 & 0.2 \\ 0.1 & 0.3 & 0.1 \end{bmatrix}.$$

Moreover, let \mathbf{b}^* and $\boldsymbol{\delta}$ are generated as follows:

$$\mathbf{b}^* = \begin{bmatrix} 6.65 \\ 8.97 \\ 5.40 \end{bmatrix}, \quad \boldsymbol{\delta} = \begin{bmatrix} 0.1 \\ 0.2 \\ 0.2 \end{bmatrix}.$$

Initializing both algorithms with a random matrix within $\mathbf{C}_{\min} = \mathbf{C}^* - \boldsymbol{\Delta}$ and $\mathbf{C}_{\max} = \mathbf{C}^* + \boldsymbol{\Delta}$, and a random vector within $\mathbf{b}_{\min} = \mathbf{b}^* - \boldsymbol{\delta}$ and $\mathbf{b}_{\max} = \mathbf{b}^* + \boldsymbol{\delta}$, ADMM algorithm returns a different solution from Algorithm 1. Besides, the difference in the performance of algorithms during the test phase can be observed in the experiments on synthetic datasets depicted in Figure 3 as well, especially when the number of samples is smaller.

Algorithm 7 Applying ADMM to the Dual Reformulation of Robust Linear Regression

-
- 1: **Given:** $\mathbf{b}_{\min}, \mathbf{b}_{\max}, \mathbf{C}_{\min}, \mathbf{C}_{\max}, \lambda, \rho$
 - 2: **Initialize:** $\mathbf{C}_1 = \mathbf{C}_0, \mathbf{b}_1 = \mathbf{b}_0, \gamma_0 = 0, \gamma_1 = 1.$
 - 3: **for** $t = 0, \dots, T$ **do**
 - 4: $\boldsymbol{\theta}^{t+1} = \frac{1}{6\lambda+7\rho}(2\mathbf{c}_1^t - 2\mathbf{c}_2^t - 3\mathbf{c}_3^t)$
 - 5: $\mathbf{d}^{t+1} = \frac{1}{6\lambda+7\rho}(\frac{6\rho+4\lambda}{\rho}\mathbf{c}_1^t + \frac{4\rho+4\lambda}{\rho}\mathbf{c}_2^t + 2\mathbf{c}_3^t)$
 - 6: $\mathbf{e}^{t+1} = \frac{2}{6\lambda+7\rho}(\frac{\rho+2\lambda}{\rho}\mathbf{c}_1^t + \frac{3\rho+2\lambda}{\rho}\mathbf{c}_2^t - \mathbf{c}_3^t)$
 - 7: $\mathbf{A}'^{t+1} = \max(\mathbf{A}^t + \frac{\mathbf{M}_A^t}{\rho}, 0)$
 - 8: $\mathbf{B}'^{t+1} = \max(\mathbf{B}^t + \frac{\mathbf{M}_B^t}{\rho}, 0)$
 - 9: $\mathbf{G}^{t+1} = [\mathbf{B}^t - \mathbf{A}^t + \frac{\boldsymbol{\Gamma}^t}{\rho} - \boldsymbol{\theta}'^t \boldsymbol{\theta}'^{tT}]_+ + \boldsymbol{\theta}'^t \boldsymbol{\theta}'^{tT}$
 - 10: $\mathbf{d}'^{t+1} = \max(\mathbf{d}^t + \frac{\boldsymbol{\mu}_d^t}{\rho}, 0)$
 - 11: $\mathbf{e}'^{t+1} = \max(\mathbf{e}^t + \frac{\boldsymbol{\mu}_e^t}{\rho}, 0)$
 - 12: $\boldsymbol{\theta}'^{t+1} = \arg \min_{\boldsymbol{\theta}'^t} \|\boldsymbol{\theta}^{t+1} - \boldsymbol{\theta}'^t\|^2 + \langle \boldsymbol{\mu}_\theta^t, \boldsymbol{\theta}^{t+1} - \boldsymbol{\theta}'^t \rangle \quad \text{s.t. } \mathbf{G}^{t+1} \succeq \boldsymbol{\theta}'^{tT} \boldsymbol{\theta}'^t$
 - 13: $\mathbf{A}^{t+1} = \frac{1}{3\rho}(2\mathbf{D}_1^t + \mathbf{D}_2^t)$
 - 14: $\mathbf{B}^{t+1} = \frac{1}{3\rho}(\mathbf{D}_1^t + 2\mathbf{D}_2^t)$
 - 15: $\mathbf{M}_A^{t+1} = \mathbf{M}_A^t + \rho(\mathbf{A}^{t+1} - \mathbf{A}'^{t+1})$
 - 16: $\mathbf{M}_B^{t+1} = \mathbf{M}_B^t + \rho(\mathbf{B}^{t+1} - \mathbf{B}'^{t+1})$
 - 17: $\boldsymbol{\mu}_d^{t+1} = \boldsymbol{\mu}_d^t + \rho(\mathbf{d}^{t+1} - \mathbf{d}'^{t+1})$
 - 18: $\boldsymbol{\mu}_e^{t+1} = \boldsymbol{\mu}_e^t + \rho(\mathbf{e}^{t+1} - \mathbf{e}'^{t+1})$
 - 19: $\boldsymbol{\mu}_\theta^{t+1} = \boldsymbol{\mu}_\theta^t + \rho(\boldsymbol{\theta}^{t+1} - \boldsymbol{\theta}'^{t+1})$
 - 20: $\boldsymbol{\eta}^{t+1} = \boldsymbol{\eta}^t + \rho(2\boldsymbol{\theta}^{t+1} - \mathbf{d}^{t+1} + \mathbf{e}^{t+1})$
 - 21: $\boldsymbol{\Gamma}^{t+1} = \boldsymbol{\Gamma}^t + \rho(\mathbf{B}^{t+1} - \mathbf{A}^{t+1} - \mathbf{G}^{t+1})$
-

Now, we show how to apply ADMM schema to Problem (28) to obtain Algorithm 7. As we discussed earlier, we consider two separate blocks of variables $\mathbf{w} = (\boldsymbol{\theta}, \mathbf{d}, \mathbf{e}, \mathbf{G}, \mathbf{B}', \mathbf{A}')$ and $\mathbf{z} = (\mathbf{d}', \mathbf{e}', \boldsymbol{\theta}', \mathbf{B}, \mathbf{A})$. Assigning $\boldsymbol{\Gamma}, \boldsymbol{\eta}, \mathbf{M}_A, \mathbf{M}_B, \boldsymbol{\mu}_d, \boldsymbol{\mu}_e$, and $\boldsymbol{\mu}_\theta$ to the constraints of Problem (28) in order, we can write the corresponding augmented Lagrangian function as:

$$\begin{aligned}
& \min_{\boldsymbol{\theta}, \boldsymbol{\theta}', \mathbf{A}, \mathbf{A}', \mathbf{B}, \mathbf{B}', \mathbf{d}, \mathbf{d}', \mathbf{e}, \mathbf{e}', \mathbf{G}} \quad -\langle \mathbf{b}_{\min}, \mathbf{d} \rangle + \langle \mathbf{b}_{\max}, \mathbf{e} \rangle - \langle \mathbf{C}_{\min}, \mathbf{A} \rangle + \langle \mathbf{C}_{\max}, \mathbf{B} \rangle + \lambda \|\boldsymbol{\theta}\|^2 \\
& \quad + \langle \mathbf{A} - \mathbf{A}', \mathbf{M}_A \rangle + \frac{\rho}{2} \|\mathbf{A} - \mathbf{A}'\|_F^2 \\
& \quad + \langle \mathbf{B} - \mathbf{B}', \mathbf{M}_B \rangle + \frac{\rho}{2} \|\mathbf{B} - \mathbf{B}'\|_F^2 \\
& \quad + \langle \mathbf{d} - \mathbf{d}', \boldsymbol{\mu}_d \rangle + \frac{\rho}{2} \|\mathbf{d} - \mathbf{d}'\|^2 \\
& \quad + \langle \mathbf{e} - \mathbf{e}', \boldsymbol{\mu}_e \rangle + \frac{\rho}{2} \|\mathbf{e} - \mathbf{e}'\|^2 \\
& \quad + \langle \boldsymbol{\theta} - \boldsymbol{\theta}', \boldsymbol{\mu}_\theta \rangle + \frac{\rho}{2} \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|^2 \\
& \quad + \langle 2\boldsymbol{\theta} - \mathbf{d} + \mathbf{e}, \boldsymbol{\eta} \rangle + \frac{\rho}{2} \|2\boldsymbol{\theta} - \mathbf{d} + \mathbf{e}\|^2 \\
& \quad + \langle \mathbf{B} - \mathbf{A} - \mathbf{G}, \boldsymbol{\Gamma} \rangle + \frac{\rho}{2} \|\mathbf{B} - \mathbf{A} - \mathbf{G}\|_F^2 \\
& \quad \text{s.t.} \quad \mathbf{A}', \mathbf{B}', \mathbf{d}', \mathbf{e}' \geq 0, \\
& \quad \mathbf{G} \succeq \boldsymbol{\theta}' \boldsymbol{\theta}'^T,
\end{aligned} \tag{29}$$

At each iteration of the ADMM algorithm, the parameters of one block are fixed, and the optimization problem is solved with respect to the parameters of the other block. For the simplicity of presentation, let $\mathbf{c}_1^t = \rho\boldsymbol{\theta}^t - \boldsymbol{\mu}_\theta^t - 2\boldsymbol{\eta}^t$, $\mathbf{c}_2^t = \rho\mathbf{d}'^t - \boldsymbol{\mu}_d^t - \mathbf{b}_{\min} + \boldsymbol{\eta}^t$, $\mathbf{c}_3^t = \rho\mathbf{e}'^t - \boldsymbol{\mu}_e^t + \mathbf{b}_{\max} - \boldsymbol{\eta}^t$, $\mathbf{D}_1^t = \rho\mathbf{A}'^t - \rho\mathbf{G}^t + \boldsymbol{\Gamma}^t - \mathbf{M}_A^t + \mathbf{C}_{\min}$, and $\mathbf{D}_2^t = \rho\mathbf{B}'^t + \rho\mathbf{G}^t - \boldsymbol{\Gamma}^t - \mathbf{M}_B^t - \mathbf{C}_{\max}$.

We have two non-trivial problems containing positive semi-definite constraints. The sub-problem with respect to \mathbf{G} can be written as:

$$\begin{aligned} \min_{\mathbf{G}} \langle \mathbf{B}^t - \mathbf{A}^t - \mathbf{G}, \mathbf{\Gamma}^t \rangle + \frac{\rho}{2} \|\mathbf{B}^t - \mathbf{A}^t - \mathbf{G}\|_F^2 \\ \text{s.t. } \mathbf{G} \succeq \boldsymbol{\theta}'^t \boldsymbol{\theta}'^{tT}, \end{aligned} \quad (30)$$

By completing the square, and changing the variable $\mathbf{G}' = \mathbf{G} - \boldsymbol{\theta}'^t \boldsymbol{\theta}'^{tT}$, equivalently we require to solve the following problem:

$$\begin{aligned} \min_{\mathbf{G}'} \frac{\rho}{2} \|\mathbf{G}' - (\mathbf{B}^t - \mathbf{A}^t - \boldsymbol{\theta}'^t \boldsymbol{\theta}'^{tT} + \frac{\mathbf{\Gamma}^t}{\rho})\|_F^2 \\ \text{s.t. } \mathbf{G}' \succeq 0, \end{aligned} \quad (31)$$

Thus, $\mathbf{G}'^* = [\mathbf{B}^t - \mathbf{A}^t + \frac{\mathbf{\Gamma}^t}{\rho} - \boldsymbol{\theta}'^t \boldsymbol{\theta}'^{tT}]_+$, where $[\mathbf{A}]_+$ is the projection to the set of PSD matrices, which can be done by setting the negative eigenvalues of \mathbf{A} in its singular value decomposition to zero.

The other non-trivial sub-problem in Algorithm (7) is the minimization with respect to $\boldsymbol{\theta}'$ (Line 10). By completing the square, it can be equivalently formulated as:

$$\begin{aligned} \min_{\boldsymbol{\theta}'} \|\boldsymbol{\theta}' - (\boldsymbol{\theta}^{t+1} + \frac{\boldsymbol{\mu}^t \boldsymbol{\theta}}{\rho})\|_2^2 \\ \text{s.t. } \mathbf{G}^{t+1} \succeq \boldsymbol{\theta}' \boldsymbol{\theta}'^T, \end{aligned} \quad (32)$$

Let $\mathbf{G} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^T$ be the singular value decomposition of the matrix \mathbf{G} where $\mathbf{\Lambda}$ is a diagonal matrix containing the eigenvalues of the matrix \mathbf{G} . Set $\boldsymbol{\alpha} = \boldsymbol{\theta}^{t+1} + \frac{\boldsymbol{\mu}^t \boldsymbol{\theta}}{2}$. Since $\mathbf{U}^T \mathbf{U} = \mathbf{I}$, we have:

$$\|\mathbf{U}^T \boldsymbol{\theta} - \mathbf{U}^T \boldsymbol{\alpha}\|_2^2 = \|\boldsymbol{\theta} - \boldsymbol{\alpha}\|_2^2$$

Set $\boldsymbol{\beta} = \mathbf{U}^T \boldsymbol{\theta}'$, then Problem (32) can be reformulated as:

$$\begin{aligned} \min_{\boldsymbol{\beta}} \|\boldsymbol{\beta} - \mathbf{U}^T \boldsymbol{\alpha}\|_2^2 \\ \text{s.t. } \boldsymbol{\beta} \boldsymbol{\beta}^T \preceq \mathbf{\Lambda}. \end{aligned} \quad (33)$$

Note that the constraint of the above optimization problem is equivalent to the following:

$$\boldsymbol{\beta} \boldsymbol{\beta}^T \preceq \mathbf{\Lambda} \Leftrightarrow \begin{bmatrix} 1 & \boldsymbol{\beta}^T \\ \boldsymbol{\beta} & \mathbf{\Lambda} \end{bmatrix} \succeq 0 \Leftrightarrow \boldsymbol{\beta}^T \mathbf{\Lambda}^{-1} \boldsymbol{\beta} \leq 1 \Leftrightarrow \sum_{i=1}^n \frac{\beta_i^2}{\lambda_i} \leq 1,$$

where $\lambda_i = \mathbf{\Lambda}_{ii}$. Since the block matrix is symmetric, using Schur Complement, it is positive semi-definite if and only if $\mathbf{\Lambda}$ is positive semi-definite and $1 - \boldsymbol{\beta}^T \mathbf{\Lambda}^{-1} \boldsymbol{\beta} \geq 0$ (the third inequality above).

Set $\boldsymbol{\gamma} = \mathbf{U}^T \boldsymbol{\alpha}$, then we can write Problem (33) as:

$$\begin{aligned} \min_{\boldsymbol{\beta}} \|\boldsymbol{\beta} - \boldsymbol{\gamma}\|_2^2 \\ \text{s.t. } \sum_{i=1}^n \frac{\beta_i^2}{\lambda_i} \leq 1, \end{aligned} \quad (34)$$

It can be easily shown that the optimal solution has the form $\beta_i^* = \frac{\gamma_i}{1 + \frac{\mu^*}{\lambda_i}}$, where μ^* is the optimal Lagrangian multiplier corresponding to the constraint of Problem (34). The optimal Lagrangian multiplier can be obtained by the bisection algorithm similar to Algorithm 2. Having β^* , the optimal $\boldsymbol{\theta}$ can be computed by solving the linear equation $\mathbf{U}^T \boldsymbol{\theta}^* = \boldsymbol{\beta}^*$.

E Quadratic RIFLE: Using Kernels to Go Beyond Linearity

A natural extension of RILFE to non-linear models is to transform the original data via multiple Kernels and then apply RIFLE to the obtained data. To this end, we applied Polynomial Kernels to the original data that

considers the polynomial transformations of features and their interactions. A drawback of this approach is that if the original data contains d features, and the order of the polynomial Kernel is t , the number of features in the transformed data will be $\mathcal{O}(d^t)$ that increases the runtime of the prediction/imputation drastically. Thus, we only consider $t = 2$, which leads to a dataset containing the interaction of different features of the original data. We call the RIFLE algorithm applied on the data transformed by Quadratic Kernel Quadratic RIFLE (QRIFLE). Table 1 demonstrates the performance of QRIFLE alongside RIFLE and other state-of-the-art approaches. Moreover, we applied QRIFLE on a regression task where the correlation between predictors and the target variable is quadratic (Figure 8). We can observe that QRIFLE works better than RIFLE when the percentage of missing values is not high.

E.1 Performance of RIFLE and QRIFLE on Synthetic Non-linear Data

A natural question is how RIFLE performs when the underlying model is non-linear. To evaluate RIFLE and other methods, we have generated jointly normal data similar to the experiment in Figure 9. Here, we have 5000 data points, and the data dimension is $d = 5$. The target variable has the following quadratic relationship with the input features:

$$y = x_1^2 + 3x_3^2 - 6x_5^2 - 0.9x_1x_4 + 9x_2x_3 + 3.2x_4x_5 - 1.7x_2x_5 - 5x_1 - 2x_3 + 7x_4 + 4.6$$

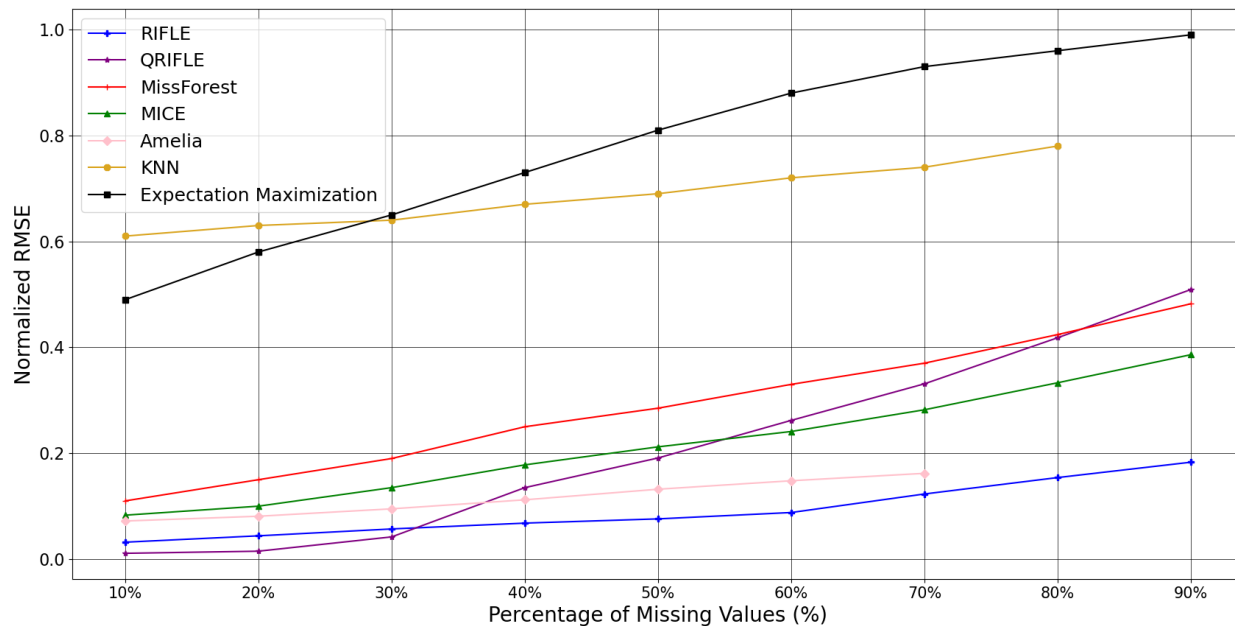


Figure 8: Performance of RIFLE, QRIFLE, MissForest, Amelia, KNN Imputer, MICE, Expectation Maximization to the percentage of missing values on Quadratic artificial datasets with different percentages of missing values.

We evaluated the performance of KNN-Imputer (Troyanskaya et al., 2001), MICE (Buuren & Groothuis-Oudshoorn, 2010), Amelia (Honaker et al., 2011), MissForest (Stekhoven & Bühlmann, 2012), and Expectation Maximization (Dempster et al., 1977), alongside the RIFLE and QRIFLE. QRIFLE is the RIFLE application on the original data transformed by a polynomial kernel with the degree of 2. Although QRIFLE can learn the quadratic models, the number of missing values in the new features (interaction terms) will be higher than the original data. For instance, if, on average, 50% of entries are missing in the original features, there will be 75% of missing entries in the interaction terms. Moreover, the computation complexity will be increased since we have d^2 features instead of d if we use QRIFLE. Figure 8 demonstrates the performance of the aforementioned methods on the artificial data with 5000 samples containing different percentages of missing values. We generated 5 artificial datasets for each missing value percentage, and each method is performed 5 times on the datasets. We reported the average performances for each method in Figure 8. For

small percentages of missing values, QRIFLE performs better than other approaches. However, by increasing the percentage of missing values, QRIFLE performance drops, and RIFLE works much better than RIFLE.

F Robust Quadratic Discriminant Analysis (Presence of Missing Values in the Target Feature)

In Section 4 we formalized robust quadratic discriminant analysis assuming the target variable is fully available. In this appendix, we study Problem (12) when the target variable contains missing values.

If the target feature contains missing values, the proposed algorithm for solving the optimization problem (13) does not exploit the data points whose target feature is unavailable. However, such points can contain valuable statistical information about the underlying data distribution. Thus, we apply an Expectation Maximization (EM) procedure on the dataset as follows:

Assume that a dataset consisting of $n + m$ samples. Let $(\mathbf{X}_1, y_1), \dots, (\mathbf{X}_n, y_n)$ be n samples whose target variable is available and $(\mathbf{X}_{n+1}, z_1), \dots, (\mathbf{X}_{n+m}, z_m)$ are m samples where their corresponding labels are missing. Similar to the previous case, we assume:

$$\mathbf{X}_i | z_i = j \sim \mathcal{N}(\boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j), \quad j = 0, 1.$$

Thus, the probability of observing a data point x_i can be written as:

$$\begin{aligned} P(\mathbf{X}_i = \mathbf{x}_i) &= \pi_0 P(\mathbf{X}_i = \mathbf{x}_i | z_i = 0) + \pi_1 P(\mathbf{X}_i = \mathbf{x}_i | z_i = 1) \\ &= \pi_0 \mathcal{N}(\mathbf{x}_i; \boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0) + \pi_1 \mathcal{N}(\mathbf{x}_i; \boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1) \end{aligned}$$

The log of likelihood function can be formulated as follows:

$$\ell(\boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0, \boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1) = \sum_{i=1}^{n+m} \log \left(\pi_0 \mathcal{N}(\mathbf{x}_i; \boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0) + \pi_1 \mathcal{N}(\mathbf{x}_i; \boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1) \right)$$

We apply Expectation Maximization procedure to jointly update $\boldsymbol{\Sigma}_0, \boldsymbol{\Sigma}_1, \boldsymbol{\mu}_0, \boldsymbol{\mu}_1$ and z_i 's. Note that the posterior distribution of z_i can be written as:

$$P(Z_i = t | \mathbf{X}_i = \mathbf{x}_i) = \frac{P(\mathbf{X}_i = \mathbf{x}_i | Z_i = t) P(Z_i = t)}{P(\mathbf{X}_i = \mathbf{x}_i)} = \frac{\pi_t \mathcal{N}(\mathbf{x}_i; \boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t)}{P(\mathbf{X}_i = \mathbf{x}_i)}$$

We update z_i values in the E-step by comparing the posterior probabilities for two possible labels. Precisely, we assign label 1 to Z_i if and only if:

$$\pi_1 \mathcal{N}(\mathbf{x}_i; \boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1) > \pi_0 \mathcal{N}(\mathbf{x}_i; \boldsymbol{\mu}_0, \boldsymbol{\Sigma}_0)$$

In M-step, we estimate $\boldsymbol{\Sigma}_0, \boldsymbol{\Sigma}_1, \boldsymbol{\mu}_0, \boldsymbol{\mu}_1, \pi_0$ and π_1 by fixing the z_i values. Since in M-step, all labels (both already available y_i 's and estimated z_i 's in E-step) are assigned, updating the aforementioned parameters

can be done as follows:

$$\boldsymbol{\mu}_1[j] = \frac{1}{|\mathcal{S}_1 \cap \mathcal{T}_j|} \sum_{i \in \mathcal{S}_1 \cap \mathcal{T}_j} \mathbf{x}_i[j] \quad (35)$$

$$\boldsymbol{\mu}_0[j] = \frac{1}{|\mathcal{S}_0 \cap \mathcal{T}_j|} \sum_{i \in \mathcal{S}_0 \cap \mathcal{T}_j} \mathbf{x}_i[j] \quad (36)$$

$$\boldsymbol{\Sigma}_1[i][j] = \frac{1}{|\mathcal{S}_1 \cap \mathcal{T}_i \cap \mathcal{T}_j|} \sum_{t \in \mathcal{S}_1 \cap \mathcal{T}_i \cap \mathcal{T}_j} \mathbf{x}_t[i] \mathbf{x}_t[j] \quad (37)$$

$$\boldsymbol{\Sigma}_0[i][j] = \frac{1}{|\mathcal{S}_0 \cap \mathcal{T}_i \cap \mathcal{T}_j|} \sum_{t \in \mathcal{S}_0 \cap \mathcal{T}_i \cap \mathcal{T}_j} \mathbf{x}_t[i] \mathbf{x}_t[j] \quad (38)$$

$$\pi_1 = \frac{|\mathcal{S}_1|}{|\mathcal{S}_0 \cup \mathcal{S}_1|} \quad (39)$$

$$\pi_0 = \frac{|\mathcal{S}_0|}{|\mathcal{S}_0 \cup \mathcal{S}_1|} \quad (40)$$

We apply the M-step and E-step iteratively to obtain $\boldsymbol{\Sigma}_1$ and $\boldsymbol{\Sigma}_0$. Based on the random initialization of z_i 's we can obtain different values for $\boldsymbol{\mu}_0, \boldsymbol{\mu}_1, \boldsymbol{\Sigma}_0$ and $\boldsymbol{\Sigma}_1$. Having these estimations, we apply Algorithm 3 to solve the robust normal discriminant analysis formulated in (13).

Algorithm 8 Expectation Maximization Procedure for Learning a Robust Normal Discriminant Analysis

- 1: **Input:** T : Number of EM iterations, k : Number of covariance estimations at each iteration.
 - 2: **Initialize:** Set each missing labels randomly to 0 or 1.
 - 3: **for** $i = 1, \dots, T$ **do**
 - 4: Estimate k covariance matrices by sampling with replacement from the available entries
 - 5: Find an optimal \mathbf{w} for Problem (18)
 - 6: Update the missing labels using the new \mathbf{w} obtained above.
-

G Generating Missing Values Patterns in Numerical Experiments

In this appendix, we define MCAR and MNAR patterns and discuss how to generate them in a given dataset. Formally, the distribution of missing values in a dataset follows a missing completely at random (MCAR) pattern if the probability of having a missing value for a given entry is constant, independent of other available and missing entries. On the other hand, a dataset follows a Missing At Random (MAR) pattern if the missingness of each entry only depends on the available data of other features. Finally, if the distribution of missing values does not follow an MCAR or MAR pattern, we call it missing not at random (MNAR).

To generate the MCAR pattern on a given dataset, we fix a constant probability $0 < p < 1$ and make each data entry unavailable with the probability of p . On the other hand, the generation of the MNAR pattern is based on the idea that if the value of an entry is farther from the mean of its corresponding feature, then the probability of missingness for that entry is larger.

The generation of the MNAR pattern is based on the idea that if the value of an entry is farther from the mean of its corresponding feature, then the probability of missingness for that entry is larger. Algorithm 9 describes the procedure of generating MNAR missing values for a given column of a dataset:

Algorithm 9 Generating MNAR Pattern for a Given Column of a Dataset

-
- 1: **Input:** x_1, x_2, \dots, x_n : The entries of the current column in the dataset, a, b : Hyper-parameters controlling the percentage of missing values
 - 2: **Initialize:** Set $\mu = \frac{1}{n} \sum_{i=1}^n x_i$ and $\sigma^2 = \frac{1}{n} \sum_{i=1}^n x_i^2 - \mu^2$.
 - 3: **for** $i = 1, \dots, n$ **do**
 - 4: $x'_i = \frac{x_i - \mu}{\sigma}$
 - 5: $p_i = F(a|x'_i| + b)$
 - 6: Set $x_i = *$ with probability of p_i
-

Note that F in the above algorithm is the cumulative distribution function of a standard Gaussian random variable. a and b control the percentage of missing values in the given column. As a and b increase, the probability of having more missing values is higher. Since the availability of each data entry depends on its value, the generated missing pattern is missing not at random (MNAR).

H Proof of Lemmas and Theorems

In this appendix, we prove all lemmas and theorems presented in the article. First, we prove the following lemma that is useful in several convergence proofs:

Lemma 11. Let $\theta^*(\mathbf{b}, \mathbf{C}) = \arg \min_{\theta} \theta^T \mathbf{C} \theta - 2\mathbf{b}^T \theta + \lambda \|\theta\|^2$. Assume that for any given \mathbf{b} and \mathbf{C} , $\|\theta^*(\mathbf{b}, \mathbf{C})\| \leq \tau$. Then, the Lipschitz constant of the gradient of the function $g(\mathbf{b}, \mathbf{C}) = \min_{\theta} \theta^T \mathbf{C} \theta - 2\mathbf{b}^T \theta + \lambda \|\theta\|^2$ used in Problem (7) is equal to $L = \frac{2(\tau+1)^2}{\lambda}$.

Proof. Since the problem is convex in θ and concave in \mathbf{C} and \mathbf{b} , we have:

$$\min_{\theta} \max_{\mathbf{C}, \mathbf{b}} \theta^T \mathbf{C} \theta - 2\mathbf{b}^T \theta + \lambda \|\theta\|^2 = - \min_{\mathbf{C}, \mathbf{b}} \max_{\theta} -\theta^T \mathbf{C} \theta + 2\mathbf{b}^T \theta - \lambda \|\theta\|^2$$

Assume that $h(\theta, \mathbf{C}, \mathbf{b}) \triangleq -\theta^T \mathbf{C} \theta + 2\mathbf{b}^T \theta - \lambda \|\theta\|^2$. Define L_{11}, L_{12} as follows:

$$\begin{aligned} \|\nabla_{\mathbf{b}, \mathbf{C}} h(\theta, \mathbf{b}_1, \mathbf{C}_1) - \nabla_{\mathbf{b}, \mathbf{C}} h(\theta, \mathbf{b}_2, \mathbf{C}_2)\| &\leq L_{11} \|(\mathbf{C}_1, \mathbf{b}_1) - (\mathbf{C}_2, \mathbf{b}_2)\| \\ \|\nabla_{\theta} h(\theta, \mathbf{b}_1, \mathbf{C}_1) - \nabla_{\theta} h(\theta, \mathbf{b}_2, \mathbf{C}_2)\| &\leq L_{12} \|(\mathbf{C}_1, \mathbf{b}_1) - (\mathbf{C}_2, \mathbf{b}_2)\| \end{aligned}$$

$h(\theta, \mathbf{b}, \mathbf{C})$ is convex in \mathbf{C} and \mathbf{b} and strongly concave with respect to θ . According to Lemma 1 in Barzandeh & Razaviyayn (2020), $g' = -g = \max_{\theta} h(\theta, \mathbf{b}, \mathbf{C})$ is Lipschitz continuous with the Lipschitz constant equal to:

$$L_g = L_{g'} = L_{11} + \frac{L_{12}^2}{\sigma},$$

where $\sigma = 2\lambda$ is the strong-concavity modulus of $-\theta^T \mathbf{C} \theta + 2\mathbf{b}^T \theta - \lambda \|\theta\|^2$. Note that

$$\begin{aligned} \nabla_{\mathbf{b}, \mathbf{C}} h(\theta, \mathbf{b}, \mathbf{C}) &= -\theta \theta^T + 2\theta \rightarrow \\ \nabla_{\mathbf{b}, \mathbf{C}} h(\theta, \mathbf{b}_1, \mathbf{C}_1) - \nabla_{\mathbf{b}, \mathbf{C}} h(\theta, \mathbf{b}_2, \mathbf{C}_2) &= 0 \end{aligned}$$

Thus, $L_{11} = 0$. On the other hand,

$$\begin{aligned} \nabla_{\theta} h(\theta, \mathbf{b}, \mathbf{C}) &= -2\theta \mathbf{C} + 2\mathbf{b} - 2\lambda \theta \rightarrow \\ \nabla_{\theta} h(\theta, \mathbf{b}_1, \mathbf{C}_1) - \nabla_{\theta} h(\theta, \mathbf{b}_2, \mathbf{C}_2) &= -2(\mathbf{C}_1 - \mathbf{C}_2)\theta + 2(\mathbf{b}_1 - \mathbf{b}_2) \leq 2\|\mathbf{C}_1 - \mathbf{C}_2\|_2 \|\theta\|_2 + 2\|\mathbf{b}_1 - \mathbf{b}_2\|_2 \\ &\leq 2\|(\mathbf{C}_1, \mathbf{b}_1) - (\mathbf{C}_2, \mathbf{b}_2)\|_2 \|\theta\|_2 + 2\|(\mathbf{C}_1, \mathbf{b}_1) - (\mathbf{C}_2, \mathbf{b}_2)\|_2 \leq (2\|\theta\|_2 + 2)\|(\mathbf{C}_1, \mathbf{b}_1) - (\mathbf{C}_2, \mathbf{b}_2)\|_2 \end{aligned}$$

Therefore, $L_{12} = 2 \max \|\theta\|_2 + 2$, which means $L_g = \frac{2(\max \|\theta\|_2 + 1)^2}{\lambda}$. Note that θ is computed exactly in Algorithm 1 and Algorithm 5 at each iterations. Thus, during the optimization procedure the norm of θ is bounded by the maximum norm of θ for any given \mathbf{b} and \mathbf{C} :

$$\max \|\theta\|_2 \leq \max_{\mathbf{b}, \mathbf{C}} \theta^*(\mathbf{b}, \mathbf{C}) \leq \tau$$

As a result, $L_g = \frac{2(\tau+1)^2}{\lambda}$. □

Proof of Theorem 1: Since the set of feasible solutions for \mathbf{b} and \mathbf{C} defines a compact set, and function g is a concave function with respect to \mathbf{b} and \mathbf{C} , the projected gradient ascent algorithm converges to the global maximizer of g in $T = \mathcal{O}(\frac{LD}{\epsilon})$ iterations (Bubeck, 2014, Theorem 3.3), where $D = \|\mathbf{C}_0 - \mathbf{C}^*\|_F^2 + \|\mathbf{b}_0 - \mathbf{b}^*\|_2^2$ and L is the Lipschitz constant of function g , which is equal to $\frac{2(\tau+1)^2}{\lambda}$ according to Lemma 11.

Proof of Theorem 7 Algorithm 5 applies the projected Nesterov acceleration method on the concave function g . As proved in Nesterov (1983), the rate of convergence of this method conforms to the lower bound of first-order oracles for the general convex minimization (concave maximization) problems, which is $\mathcal{O}(\sqrt{\frac{LD^2}{\epsilon}})$. We compute the Lipschitz constant L that appeared in the iteration complexity bound by Lemma 11.

Proof of Theorem 8: First, note that if we multiply the objective function by -1 , Problem (6) can be equivalently formulated as:

$$\begin{aligned} \max_{\boldsymbol{\theta}} \quad & \min_{\mathbf{C}, \mathbf{b}} -\boldsymbol{\theta}^T \mathbf{C} \boldsymbol{\theta} + 2\mathbf{b}^T \boldsymbol{\theta} - \lambda \|\boldsymbol{\theta}\|_2^2 \\ \text{s.t.} \quad & -\mathbf{C} + \mathbf{C}_{\min} \leq 0, \\ & \mathbf{C} - \mathbf{C}_{\max} \leq 0, \\ & -\mathbf{b} + \mathbf{b}_{\min} \leq 0, \\ & \mathbf{b} - \mathbf{b}_{\max} \leq 0, \\ & -\mathbf{C} \preceq 0 \end{aligned} \tag{41}$$

If we assign $\mathbf{A}, \mathbf{B}, \mathbf{d}, \mathbf{e}, \mathbf{H}$ to the constraints respectively, then the Lagrangian function can be written as:

$$\begin{aligned} L(\mathbf{C}, \mathbf{b}, \mathbf{A}, \mathbf{B}, \mathbf{d}, \mathbf{e}, \mathbf{H}) = & -\boldsymbol{\theta}^T \mathbf{C} \boldsymbol{\theta} + 2\mathbf{b}^T \boldsymbol{\theta} + \langle \mathbf{A}, -\mathbf{C} + \mathbf{C}_{\min} \rangle \\ & + \langle \mathbf{B}, \mathbf{C} - \mathbf{C}_{\max} \rangle + \langle \mathbf{d}, -\mathbf{b} + \mathbf{b}_{\min} \rangle \\ & + \langle \mathbf{e}, \mathbf{b} - \mathbf{b}_{\max} \rangle - \langle \mathbf{C}, \mathbf{H} \rangle - \lambda \|\boldsymbol{\theta}\|_2^2, \end{aligned} \tag{42}$$

The dual problem is defined as:

$$\max_{\mathbf{A}, \mathbf{B}, \mathbf{d}, \mathbf{e}, \mathbf{H}} \min_{\mathbf{C}, \mathbf{b}} L(\mathbf{C}, \mathbf{b}, \mathbf{A}, \mathbf{B}, \mathbf{d}, \mathbf{e}, \mathbf{H}) \tag{43}$$

The minimization of L takes the following form:

$$\begin{aligned} \min_{\mathbf{C}, \mathbf{b}} \quad & \langle \mathbf{C}, -\boldsymbol{\theta} \boldsymbol{\theta}^T - \mathbf{A} + \mathbf{B} - \mathbf{H} \rangle + \langle \mathbf{b}, 2\boldsymbol{\theta} - \mathbf{d} + \mathbf{e} \rangle - \lambda \|\boldsymbol{\theta}\|_2^2 \\ & - \langle \mathbf{B}, \mathbf{C}_{\max} \rangle + \langle \mathbf{A}, \mathbf{C}_{\min} \rangle - \langle \mathbf{e}, \mathbf{b}_{\max} \rangle + \langle \mathbf{d}, \mathbf{b}_{\min} \rangle, \end{aligned} \tag{44}$$

To avoid $-\infty$ value for the above minimization problem, it is required to set $-\boldsymbol{\theta} \boldsymbol{\theta}^T - \mathbf{A} + \mathbf{B} - \mathbf{H}$ and $2\boldsymbol{\theta} - \mathbf{d} + \mathbf{e}$ to zero. Thus the dual problem of (41) is formulated as:

$$\begin{aligned} \max_{\mathbf{A}, \mathbf{B}, \mathbf{d}, \mathbf{e}, \mathbf{H}} \quad & \mathbf{b}_{\min}^T \mathbf{d} - \mathbf{b}_{\max}^T \mathbf{e} + \langle \mathbf{C}_{\min}, \mathbf{A} \rangle - \langle \mathbf{C}_{\max}, \mathbf{B} \rangle - \lambda \|\boldsymbol{\theta}\|_2^2 \\ \text{s.t.} \quad & -\boldsymbol{\theta} \boldsymbol{\theta}^T - \mathbf{A} + \mathbf{B} - \mathbf{H} = 0, \\ & 2\boldsymbol{\theta} - \mathbf{d} + \mathbf{e} = 0, \\ & \mathbf{A}, \mathbf{B}, \mathbf{d}, \mathbf{e} \geq 0, \\ & \mathbf{H} \succeq 0 \end{aligned} \tag{45}$$

Since the duality gap is zero, Problem (6) can be equivalently formulated as:

$$\begin{aligned}
& \max_{\boldsymbol{\theta}, \mathbf{A}, \mathbf{B}, \mathbf{d}, \mathbf{e}, \mathbf{H}} \mathbf{b}_{\min}^T \mathbf{d} - \mathbf{b}_{\max}^T \mathbf{e} + \langle \mathbf{C}_{\min}, \mathbf{A} \rangle - \langle \mathbf{C}_{\max}, \mathbf{B} \rangle - \lambda \|\boldsymbol{\theta}\|^2 \\
& \text{s.t.} \quad -\boldsymbol{\theta}\boldsymbol{\theta}^T - \mathbf{A} + \mathbf{B} - \mathbf{H} = 0, \\
& \quad \quad 2\boldsymbol{\theta} - \mathbf{d} + \mathbf{e} = 0, \\
& \quad \quad \mathbf{A}, \mathbf{B}, \mathbf{d}, \mathbf{e} \geq 0, \\
& \quad \quad \mathbf{H} \succeq 0.
\end{aligned} \tag{46}$$

We can multiply the objective function by -1 and change the maximization to minimization, which gives the dual problem described in (27).

Proof of Theorem 2:

(a) Let $\boldsymbol{\Delta}_n$ be the estimated confidence matrix obtained from n samples. The first part of the theorem is true, if $\boldsymbol{\Delta}_n$ converges to 0 as n , the number of samples goes to infinity (the same argument works for b and δ). Assume that $\{(\mathbf{x}_{i1}, \mathbf{x}_{i2})\}_{i=1}^n$ is an i.i.d bootstrap sample over data points that both features \mathbf{X}_1 and \mathbf{X}_2 are available. Since the distribution of missing values is completely at random (MCAR), we have $\mathbb{E}[\mathbf{x}_{i1}\mathbf{x}_{i2}] = \mathbb{E}[\mathbf{X}_1\mathbf{X}_2]$. Therefore, $\mathbb{E}[\frac{1}{n}\sum_{i=1}^n \mathbf{x}_{i1}\mathbf{x}_{i2}] = \mathbb{E}[\mathbf{X}_1\mathbf{X}_2]$. Moreover, since the samples are drawn independently, $\text{Var}[\frac{1}{n}\sum_{i=1}^n \mathbf{x}_{i1}\mathbf{x}_{i2}] = \frac{1}{n^2}\sum_{i=1}^n \text{Var}[\mathbf{x}_{i1}\mathbf{x}_{i2}] = \frac{n}{n^2}\text{Var}[\mathbf{X}_1\mathbf{X}_2] = \frac{1}{n}\text{Var}[\mathbf{X}_1\mathbf{X}_2]$. Since the variance of the product of every two features is bounded, according to the weak law of large numbers:

$$\lim_{n \rightarrow \infty} \Pr\left(\left|\frac{1}{n}\sum_{i=1}^n \mathbf{x}_{i1}\mathbf{x}_{i2} - \mathbb{E}[X_1X_2]\right| \geq \epsilon\right) = 0$$

Therefore, for any given bootstrap sample of features X_1 and X_2 , the estimation converges in probability to the ground-truth value. This means the size of the confidence interval Δ_{12} converges in probability to 0. Therefore, the estimation of $\mathbb{E}[X_1X_2]$ is consistent by the definition of consistency. With the same argument, we can prove the consistency of the estimator for any given features X_i and X_j .

(b) Fix two features i and j . Let $(\hat{X}_{i1}, \hat{X}_{j1}), \dots, (\hat{X}_{im}, \hat{X}_{jm})$ be $m = \lceil n(1-p) \rceil$ i.i.d pairs sampled via bootstrap from the entries where both features i and j are available. Define $Z_t = \hat{X}_{it}\hat{X}_{jt}$ (for simplicity we do not consider the dependence of Z_t to i and j in the notation). Assume that we initialize $\mathbf{C}_{ij} = \frac{1}{m}\sum_{t=1}^m Z_t$. Note that, $\mathbb{E}[Z_t] = \mathbb{E}[\hat{X}_{it}\hat{X}_{jt}] = \mathbf{C}_{ij}^*$. According to Chebyshev's inequality, we have:

$$\Pr\left[\left|\frac{1}{m}\sum_{t=1}^m (Z_t - \mathbb{E}[Z_t])\right| \geq \Delta_{ij}\right] \leq \frac{\text{Var}(\frac{1}{m}\sum_{t=1}^m Z_t)}{c^2\Delta_{ij}^2}$$

Note that Z_t 's are iid samples, thus:

$$\text{Var}(\frac{1}{m}\sum_{t=1}^m Z_t) = \frac{1}{m}\text{Var}(Z_t) \leq \frac{1}{m}\max_{i,j} \text{Var}(\hat{X}_i\hat{X}_j) = \frac{V}{m} = \frac{V}{n(1-p)}$$

Let $\boldsymbol{\Delta} = \min\{\Delta_{ij}\}$. Then, based on the two above inequalities, we have:

$$\Pr\left[\left|\mathbf{C}_0[i][j] - \mathbf{C}^*[i][j]\right| \geq \boldsymbol{\Delta}\right] \leq \frac{V}{c^2\boldsymbol{\Delta}^2n(1-p)}$$

Using a union bound argument, with the probability of at least $1 - \frac{Vd^2}{2c^2\boldsymbol{\Delta}^2n(1-p)}$, we have: $\mathbf{C}_0 - c\boldsymbol{\Delta} \leq \mathbf{C}^* \leq \mathbf{C}_0 + c\boldsymbol{\Delta}$, which means the actual covariance matrix is within the confidence intervals we have considered. In that case, for $(\tilde{\boldsymbol{\theta}}, \tilde{\mathbf{b}}, \tilde{\mathbf{C}})$, we have:

$$\tilde{\boldsymbol{\theta}}^T \tilde{\mathbf{C}} \tilde{\boldsymbol{\theta}} - 2\tilde{\mathbf{b}}^T \tilde{\boldsymbol{\theta}} = \max_{\mathbf{C}, \mathbf{b}} \tilde{\boldsymbol{\theta}}^T \mathbf{C} \tilde{\boldsymbol{\theta}} - 2\tilde{\mathbf{b}}^T \tilde{\boldsymbol{\theta}} \geq \tilde{\boldsymbol{\theta}}^T \mathbf{C}^* \tilde{\boldsymbol{\theta}} - 2\tilde{\mathbf{b}}^T \tilde{\boldsymbol{\theta}},$$

which completes the proof.

Proof of Theorem 3: Since the objective function is **convex** with respect to μ_1 , and the constraint on μ_1 is closed and bounded (compact), an optimal solution exists to the problem on the boundaries (note that the problem is convex **maximization**.) Therefore, for any entry of the μ_1 , it should either take $\mu_{\min}[i]$ or $\mu_{\max}[i]$, which gives the provided solution in the theorem.

I Dataset Descriptions

In this section, we introduce the datasets used in Section 5 to evaluate the performance of RIFLE on regression and classification tasks. Except for the NHANES, all other datasets contain no missing values. For those datasets, we generate MCAR and MNAR missing values artificially (for MNAR patterns, we apply Algorithm 9 to the datasets).

Datasets for Evaluating RIFLE on Regression and Imputation Tasks

- **NHANES:** The percentage of missing values varies for different features of the NHANES dataset. There are two sources of missing values in NHANES data: Missing entries during data collection and missing entries resulting from merging different datasets in the NHANES collection. On average, approximately 20% of data is missing.
- **Super Conductivity**¹: Super Conductivity datasets contains 21263 samples describing superconductors and their relevant features (81 attributes). All features are continuous, and the assigned task is to predict the critical temperature based on the given 81 features. We have used this dataset in experiments summarized in Figure 10, Figure 11, and Table 2.
- **BlogFeedback**²: BlogFeedback data is a collection of 280 features extracted from HTML-documents of the blog posts. The assigned task is to predict the number of comments in the upcoming 24 hours based on the features of more than 60K data training data points. The test dataset is fixed and is originally separated from the training data. The dataset is used in experiments described in Table 2.
- **Breast Cancer(Prognostic)**³: The dataset consists of 34 features and 198 instances. Each record represents follow-up data for one breast cancer case collected in 1984. We have done several experiments to impute the MCAR missing values generated artificially with different proportions. The results are depicted in Table 1 and Figure 4.
- **Parkinson**⁴: The dataset describes a range of biomedical voice recording from 31 people, 23 with Parkinson’s disease (PD). The assigned task is to discriminate healthy people from those with PD. There are 193 records and 23 features in the dataset. The dataset is processed similarly to the Breast Cancer dataset and used in the same experiments.
- **Spam Base**⁵: The dataset consists of 4601 instances and 57 attributes. The assigned classification task is to predict whether the email is spam. To evaluate different imputation methods, we randomly mask a proportion of data entries and impute them with different approaches. The results are depicted in Table 1 and Figure 4.
- **Boston Housing**⁶: Boston Housing dataset contains 506 instances and 14 columns. We generate random missing entries with different proportions and impute them with RIFLE and several state-of-the-art approaches. The results are demonstrated in Table 1 and Figure 4.
- **Cloud**⁷: The dataset has 1024 instances and 10 features extracted from clouds images. We use this dataset in experiments depicted in Table 1 with 70% artificial MCAR missing values.

¹<https://archive.ics.uci.edu/ml/datasets/Superconductivity+Data>

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

³[https://archive.ics.uci.edu/ml/datasets/Breast+Cancer+Wisconsin+\(Prognostic\)](https://archive.ics.uci.edu/ml/datasets/Breast+Cancer+Wisconsin+(Prognostic))

⁴<https://archive.ics.uci.edu/ml/datasets/parkinsons>

⁵<https://archive.ics.uci.edu/ml/datasets/spambase>

⁶<https://www.kaggle.com/c/boston-housing>

⁷<https://archive.ics.uci.edu/ml/datasets/Cloud>

- **Wave Energy Converters**⁸: We sample a subset of 3000 instances with 49 features from the original Wave Energy Converter dataset. We have executed several imputation methods on the dataset, and the results are shown in Figure 4.
- **Sensorless Drive Diagnosis**⁹: The 49 continuous features in this dataset are extracted from electric current drive signals, and the associated classification task is to determine the condition of device's motor. We choose different random samples with size 400 to run experiments (imputation) in Figure 5.

Datasets for Evaluating Robust QDA on Classification Tasks

- **Avila**¹⁰: The Avila dataset consists of 10 attributes extracting from 800 images of "Avila Bible". The associated classification task is to match each pattern (an instance of the dataset) to a copyist. We put 40% of MCAR missing values (both input features and the target variable) on 10 different random samples of the dataset with size 1000. The average accuracy of the robust LDA method on the 10 datasets is demonstrated in Figure 7 for each value of k (the number of covariance estimations).
- **Magic Gamma Telescope**¹¹: The dataset consists of 11 continuous MC-generated features from contributing to the prediction of the type of event (signal or background). We used the same procedure as the above dataset for the results depicted in Figure 7 (random sampling subsets of 1000 data points out of more than 19000).
- **Glass Identification**¹²: This dataset is composed of 10 continuous features and 214 instances. The assigned classification task is to predict the type of glasses based on the materials used for making them. We have assigned 40% of MCAR missing values to the dataset for the experiments reported in Table 4.
- **Annealing**¹³: This dataset is a mix of categorical and numerical features (37 in total), and the associated task is to predict the class (5 classes) of instances (metals). The number of instances in this dataset is 798. We use 500 data points as the training data and the rest as the test. We apply 40% of MCAR missing values to both input features and the target variable. The accuracy of different models is reported in Table 4.
- **Abalone**¹⁴: This dataset consists of 4177 instances and 8 categorical and continuous features. The goal is to predict the age of abalone based on physical measurements. The first 1000 samples are used as the training data and the rest as the test data. We applied the same pre-processing procedure as the above dataset to generate missing values on the training data. The accuracy of different models is reported in Table 4.
- **Lymphography**¹⁵: Lymphography is a categorical dataset containing 18 features and 148 data points obtained from the University Medical Centre, Institute of Oncology, Ljubljana, Yugoslavia. 100 data points are used as the training data; the rest are test data points (with no missing values). We applied the same pre-processing described for the above dataset to generate MCAR missing values.
- **Adult**¹⁶: The adult dataset contains census information of individuals, including education, gender, and capital gain. The assigned classification task is to predict whether a person earns over 50k annually. The train and test sets are two separate files consisting of 32,000 and 16,000 samples, respectively. We consider gender and race as the sensitive attributes (For the experiments involving

⁸<https://archive.ics.uci.edu/ml/datasets/Wave+Energy+Converters>

⁹<https://archive.ics.uci.edu/ml/datasets/dataset+for+sensorless+drive+diagnosis>

¹⁰<https://archive.ics.uci.edu/ml/datasets/Avila>

¹¹<https://archive.ics.uci.edu/ml/datasets/MAGIC+Gamma+Telescope>

¹²<https://archive.ics.uci.edu/ml/datasets/glass+identification>

¹³<https://archive.ics.uci.edu/ml/datasets/Annealing>

¹⁴<https://archive.ics.uci.edu/ml/datasets/abalone>

¹⁵<https://archive.ics.uci.edu/ml/datasets/Lymphography>

¹⁶<https://archive.ics.uci.edu/ml/datasets/adult>.

one sensitive attribute, we have chosen gender). Learning a logistic regression model on the training dataset (without imposing fairness) shows that only 3 features out of 14 have larger weights than the gender attribute.

J Further Discussion on the Consistency of RIFLE

The three developed algorithms in Section 3 for solving robust ridge regression are all consistent. To show this, we have generated a synthetic dataset with 50 input features following a jointly normal distribution. As observed in Figure 9, by increasing the number of samples, the NRMSE of all three algorithms converges to 0.01, which is the standard deviation of the zero-mean Gaussian noise added to each target value (the dashed line). The pattern can be observed for different percentages of missing values.

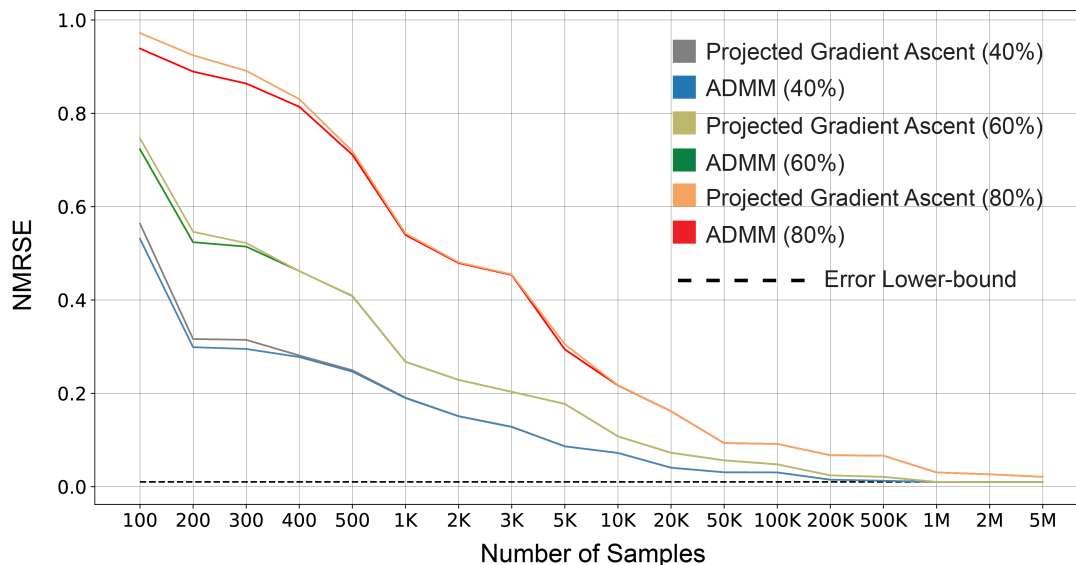


Figure 9: Consistency of ADMM (Algorithm 7) and Projected Gradient Ascent on function g (Algorithm 1) on the synthetic datasets with 40%, 60% and 80% missing values.

K Numerical Experiments for Convergence of RIFLE Algorithms

We presented three algorithms for solving the robust linear regression problem formulated in (6): Projected gradient ascent (Algorithm 1, Nesterov acceleration method (Algorithm 5), and Alternating Direction Method of Multipliers (ADMM) (Algorithm 7) applied on the dual problem.




Figure 10: Convergence of ADMM algorithm to the optimal solution of Problem (27) for different values of ρ . The left plot measures the objective function of Problem (27) per iteration (without considering the constraints), while the right plot demonstrates the constraint violation of the algorithm per iteration. The constraint violation can be measured by adding all regularization terms in the augmented Lagrangian function formulated in Problem (29).

We established the convergence rate of the gradient ascent and Nesterov acceleration methods in Theorem 1 and Theorem 7, respectively. To investigate the convergence of the ADMM algorithm and its dependence on ρ , we perform Algorithm 7 on the Super Conductivity dataset (Description in Appendix I) with 30% MCAR missing values. Figure 10 demonstrates the convergence of the ADMM algorithm for multiple values of ρ applied to the Super Conductivity dataset as described above. As can be observed, decreasing the value of ρ accelerates the ADMM convergence to the optimal value. Note that for $\rho = 0.2$, the objective function is smaller than the final value in the first few iterations. The reason is that for those iterations, the solution is not feasible (as observed in the right figure). The final solution is the optimal **feasible** solution.

In the next experiment, we compare the three proposed algorithms regarding the number of iterations required to reach a certain level of test accuracy on the Super Conductivity dataset. The number of training samples is 1000, containing 40% of MCAR missing values on both input features and the target variable. The test dataset contains 2000 samples. As depicted in Figure 11, ADMM and Nesterov’s algorithms require less number of iterations to reach the ϵ -optimal solution compared to Algorithm 1. However, the cost per iteration of the ADMM algorithm (Algorithm 7) is higher than the Nesterov acceleration and Algorithm 1.




Figure 11: The performance of the Nesterov acceleration method, projected gradient ascent, and ADMM on the Super Conductivity dataset vs. the number of iterations.

L Execution Time Comparison of RIFLE and Other State-of-the-art Approaches

This section reports the average execution time of the RIFLE and other approaches presented in Table 2.

Methods	Datasets		
	Super Conductivity	Blog Feedback	NHANES
Regression on Complete Data	0.3 sec	0.7 sec	0.4 sec
RIFLE	87 sec	471 sec	125 sec
Mean Imputer + Regression	0.4 sec	0.9 sec	0.5 sec
MICE + Regression	112 sec	573 sec	294 sec
EM + Regression	171 sec	612 sec	351 sec
MIDA Imputer + Regression	245 sec	726 sec	599 sec
MissForest	94 sec	321 sec	132 sec
KNN Imputer	66 sec	292 sec	144 sec

Table 5: Execution time of RIFLE and other SOTA methods on three datasets.