

DOUBLY ROBUST DATA-DRIVEN DISTRIBUTIONALLY ROBUST OPTIMIZATION

BLANCHET, J., KANG, Y., ZHANG, F., HE, F., AND HU, Z.

ABSTRACT. Data-driven Distributionally Robust Optimization (DD-DRO) via optimal transport has been shown to encompass a wide range of popular machine learning algorithms. The distributional uncertainty size is often shown to correspond to the regularization parameter. The type of regularization (e.g. the norm used to regularize) corresponds to the shape of the distributional uncertainty. We propose a data-driven robust optimization methodology to inform the transportation cost underlying the definition of the distributional uncertainty. We show empirically that this additional layer of robustification, which produces a method we called doubly robust data-driven distributionally robust optimization (DD-R-DRO), allows to enhance the generalization properties of regularized estimators while reducing testing error relative to state-of-the-art classifiers in a wide range of data sets.

1. INTRODUCTION

A wide class of popular machine learning estimators have been recently shown to be particular cases of data-driven Distributionally Robust Optimization (DD-DRO) formulations with a distributional uncertainty set centered around the empirical distribution.

For example, regularized logistic regression, support vector machines and sqrt-Lasso, among many other machine learning formulations can be exactly represented as DD-DRO problems involving an uncertainty set comprised of probability distributions which are within a distance δ from the empirical distribution. The distance is measured in terms of a class of suitably defined Wasserstein distances or, more generally, optimal transport distances between distributions.

Our contribution in this paper is to build an additional robustification layer on top of the DD-DRO formulation which encompasses the machine learning algorithms mentioned earlier. Because of the second layer of robustification, we call our approach DD-R-DRO.

More specifically, we consider a parametric family of optimal transport distances and formulate a data-driven Robust Optimization (RO) problem for the selection of such a distance, which in turn is used to inform the distributional uncertainty region in the type of DD-DRO mentioned in the previous paragraph. In addition, we provide an iterative algorithm for solving such RO problem.

In order to explain DD-R-DRO more precisely, let us discuss different layers of robustness that are added in our various optimization formulations and how these layers translate in terms of machine learning properties.

A DD-DRO problem takes the general form

$$(1) \quad \min_{\beta} \max_{P \in \mathcal{U}_{\delta}(P_n)} \mathbb{E}_P[l(X, Y, \beta)],$$

where β is a decision variable, (X, Y) is a random element, and $l(X, Y, \beta)$ is a loss incurred if the decision β is taken and (X, Y) is realized. The expectation $\mathbb{E}_P[\cdot]$ is taken under the probability model P . The set $\mathcal{U}_{\delta}(P_n)$ is called the distributional uncertainty set; it is centered around the empirical distribution P_n of the data, and it is indexed by the parameter $\delta > 0$, which measures

the size of the distributional uncertainty.

The min-max problem in (1) can be interpreted as a game. We (the outer player) wish to learn a task using a class of machines indexed by β . An adversary (the inner player) is introduced to enhance out-of-sample performance. The adversary has a budget δ and can perturb the data, represented by P_n , in a certain way – this is important and we will return to this point. By introducing the artificial adversary and the distributional uncertainty, the DD-DRO formulation provides a direct mechanism to control the generalization properties of the learning procedure.

To further connect the DD-DRO representation (1) with more mainstream machine learning mechanisms for the control of out-of-sample performance (such as regularization), we recall one of the explicit representations given in Blanchet et al. (2016).

In the context of generalized logistic regression (i.e. if the $l(x, y, \beta) = \log(1 + \exp(-y\beta^T x))$), given an empirical sample $\mathcal{D}_n = \{(X_i, Y_i)\}_{i=1}^n$ with $Y_i \in \{-1, 1\}$ and a judicious choice of the distributional uncertainty $\mathcal{U}_\delta(P_n)$, Blanchet et al. (2016) shows that

$$(2) \quad \min_{\beta} \max_{P \in \mathcal{U}_\delta(P_n)} \mathbb{E}_P[l(X, Y, \beta)] = \min_{\beta} \left(\mathbb{E}_{P_n}[l(X, Y, \beta)] + \delta \|\beta\|_p \right),$$

where $\|\cdot\|_p$ is the l_p norm in \mathbb{R}^d for $p \in [1, \infty)$ and $\mathbb{E}_{P_n}[l(X, Y, \beta)] = n^{-1} \sum_{i=1}^n l(X_i, Y_i, \beta)$.

The definition of $\mathcal{U}_\delta(P_n)$ turns out to be informed by the dual norm $\|\cdot\|_q$ with $1/p + 1/q = 1$. In simple words, the *shape* of the distributional uncertainty $\mathcal{U}_\delta(P_n)$ directly implies the *type of regularization*; and the *size* of the distributional uncertainty, δ , dictates the regularization parameter.

The story behind the connection to sqrt-Lasso, support vector machines and other estimators is completely analogous to that given for (2). A key point in most of the known representations, such as (2), is that they are only partially informed by data. Only the center, P_n , and the size, δ (via cross validation) are informed by data, but not the shape.

In recent work, Blanchet et al. (2017) proposes using metric learning procedures to inform the shape of the distributional uncertainty. But the procedure proposed in Blanchet et al. (2017), though data-driven, is not robustified.

One of the driving points of using robust optimization techniques in machine learning is that the introduction of an adversary can be seen as a tool to control the testing error. While the data-driven procedure in Blanchet et al. (2017) is rich in the use of information, and hence it is able to improve the generalization performance, the lack of robustification exposes the testing error to potentially high variability. So, our contribution in this paper is to design an RO procedure for choosing the shape of $\mathcal{U}_\delta(P_n)$ using a suitable parametric family. In the context of logistic regression, for example, the parametric family that we consider includes the type of choice leading to (2) as a particular case. In turn, the choice of $\mathcal{U}_\delta(P_n)$ is applied to formulation (1) in order to obtain a *doubly-robustified* estimator.

Figure 1 shows the various combinations of information and robustness which have been studied in the literature so far. The figure shows four diagrams. Diagram (A) represents standard empirical risk minimization (ERM); which fully uses the information but often leads to high variability in testing error and, therefore, poor out-of-sample performance. Diagram (B) represents DD-DRO where only the center, P_n , and the size of the uncertainty, δ , are data driven; this choice controls out-of-sample performance but does not use data to shape the type of perturbation, thus potentially resulting in testing error bounds which might be pessimistic. Diagram (C) represents DD-DRO with data-driven shape information for perturbation type using metric learning techniques; this construction can reduce the testing error bounds at the expense of increase in the variability of the testing error estimates. Diagram (D) represents DD-R-DRO, the shape of the perturbation allowed

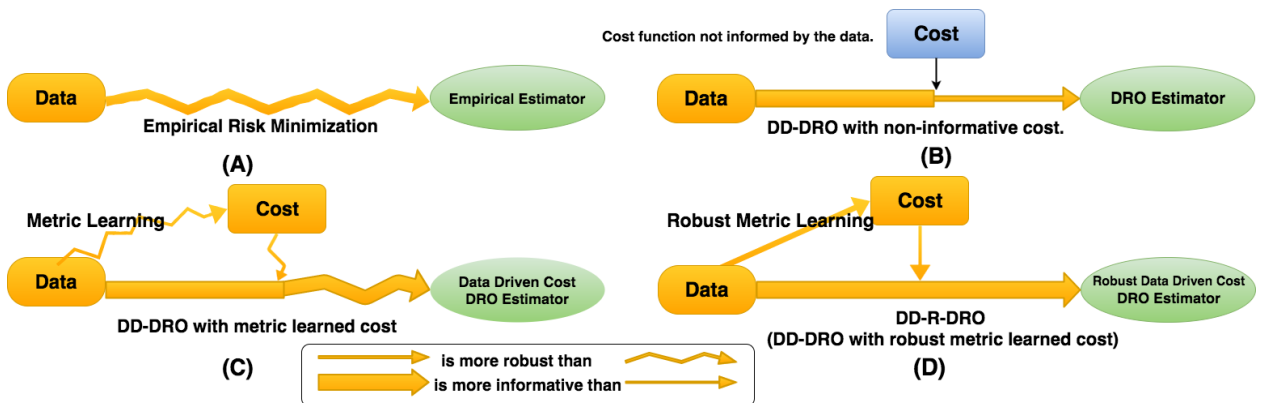


FIGURE 1. Four diagrams illustrating information on robustness.

for the adversary player is estimated using an RO procedure; this double robustification, as we shall show in the numerical experiments is able to control the variability present in the third diagram.

In the diagrams, the straight arrows represent the use of a robustification procedure. A wide arrow represents the use of high degree of information. A wiggly arrow indicates potentially noisy testing error estimates.

The contributions of this paper can be stated, in order of importance, as follows:

- 1) The fourth diagram, DD-R-DRO, illustrates the main contribution of this paper, namely, a double robustification approach which reduces the generalization error, utilizes information efficiently, and controls variability.
- 2) An explicit RO formulation for metric learning tasks.
- 3) Iterative procedures for the solution of these RO problems.

2. DD-DRO, OPTIMAL TRANSPORT, AND MACHINE LEARNING

Let us consider a supervised machine learning classification problem, where we have a response $Y \in \{-1, 1\}$ and predictors $X \in \mathbb{R}^d$. Underlying there is a general loss function $l(x, y, \beta)$ and a class of classifiers indexed by the parameter β . The distributional uncertainty region in (1) takes the form

$$(3) \quad \mathcal{U}_\delta(P_n) = \{P : D_c(P, P_n) \leq \delta\},$$

where $D_c(P, P_n)$ is a suitably defined notion of discrepancy between P and P_n so that $D_c(P, P_n) = 0$ implies that $P = P_n$.

Other notions of discrepancy have been considered in the DRO literature for example the Kullback-Leibler divergence (or another divergence notion which depends on the likelihood ratio) is utilized [Hu and Hong \(2013\)](#). Unfortunately, divergence criteria which relies on the existence of the likelihood ratio between P and P_n ultimately forces P to share the same support as P_n , therefore potentially inducing undesirable out-of-sample performance.

Instead, we follow the approach in [Esfahani and Kuhn \(2015\)](#), [Shafieezadeh-Abadeh et al. \(2015\)](#), and [Blanchet et al. \(2016\)](#), and define $D_c(P, P_n)$ as the optimal transport discrepancy between P and P_n .

2.1. Optimal Transport Distances and Discrepancies. Assume that the cost function $c : \mathbb{R}^{d+1} \times \mathbb{R}^{d+1} \rightarrow [0, \infty]$ is lower semicontinuous. We also assume that $c(u, v) = 0$ if and only if $u = v$.

Given two distributions P and Q , with supports \mathcal{S}_P and \mathcal{S}_Q , respectively, we define the optimal transport discrepancy, D_c , via

$$(4) \quad D_c(P, Q) = \inf\{E_\pi[c(U, V)] : \pi \in \mathcal{P}(\mathcal{S}_P \times \mathcal{S}_Q), \pi_U = P, \pi_V = Q\},$$

where $\mathcal{P}(\mathcal{S}_P \times \mathcal{S}_Q)$ is the set of probability distributions π supported on $\mathcal{S}_P \times \mathcal{S}_Q$, and π_U and π_V denote the marginals of U and V under π , respectively. Because $c(\cdot)$ is non-negative we have that $D_c(P, Q) \geq 0$. Moreover, requiring that $c(u, v) = 0$ if and only if $u = v$ guarantees that $D_c(P, Q) = 0$ if and only if $P = Q$.

If, in addition, $c(\cdot)$ is symmetric (i.e. $c(u, v) = c(v, u)$), and there exists $\varrho \geq 1$ such that $c^{1/\varrho}(u, w) \leq c^{1/\varrho}(u, v) + c^{1/\varrho}(v, w)$ (i.e. $c^{1/\varrho}(\cdot)$ satisfies the triangle inequality) then it can be easily verified (see Villani (2008)) that $D_c^{1/\varrho}(P, Q)$ is a metric. For example, if $c(u, v) = \|u - v\|_q^\varrho$ for $q \geq 1$ (where $\|u - v\|_q$ denotes the l_q norm in \mathbb{R}^{d+1}) then $D_c(\cdot)$ is known as the Wasserstein distance of order ϱ .

Observe that (4) is obtained by solving a linear programming problem. For example, suppose that $Q = P_n$, so $Q \in \mathcal{P}(\mathcal{D}_n)$ and assume that the support \mathcal{S}_P of P is finite. Then, using $U = (X, Y)$, we have that $D_c(P, P_n)$ is obtained by computing

$$(5) \quad \begin{aligned} & \min_{\pi} \sum_{u \in \mathcal{S}_P} \sum_{v \in \mathcal{D}_n} c(u, v) \pi(u, v) : \\ \text{s.t. } & \sum_{u \in \mathcal{S}_P} \pi(u, v) = \frac{1}{n} \quad \forall v \in \mathcal{D}_n \\ & \sum_{v \in \mathcal{D}_n} \pi(u, v) = P(\{u\}) \quad \forall u \in \mathcal{X}_N, \\ & \pi(u, v) \geq 0 \quad \forall (u, v) \in \mathcal{S}_P \times \mathcal{D}_n \end{aligned}$$

A completely analogous linear program (LP), albeit an infinite dimensional one, can be defined if \mathcal{S}_P has infinitely many elements. This LP has been extensively studied in great generality in the context of Optimal Transport under the name of Kantorovich's problem (see Villani (2008)). Requiring $c(\cdot)$ to be lower semicontinuous guarantees the existence of an optimal solution to Kantorovich's problem.

Note that $D_c(P, P_n)$ can be interpreted as the minimal cost of rearranging (i.e. transporting the mass of) the distribution P_n into the distribution P . The rearrangement mechanism has a transportation cost $c(u, w) \geq 0$ for moving a unit of mass from location u in the support of P_n to location w in the support of P . For instance, in the setting of (2) we have that

$$(6) \quad c((x, y), (x', y')) = \|x - x'\|_q^2 I(y = y') + \infty \cdot I(y \neq y').$$

The infinite contribution in the definition of c (i.e. $\infty \cdot I(y \neq y')$) indicates that the adversary player in the DRO formulation is not allowed to perturb the response variable.

3. DATA-DRIVEN SELECTION OF OPTIMAL TRANSPORT COST FUNCTION

By suitably choosing $c(\cdot)$ we might further improve the generalization properties of the DD-DRO estimator based on (1). To fix ideas, consider a suitably parameterized family of transportation costs

as follows. Let Λ be a positive semidefinite matrix (denoted as $\Lambda \succeq 0$) and define $\|x\|_{\Lambda}^2 = x^T \Lambda x$. Inspired by (6), consider the cost function

$$(7) \quad c_{\Lambda}((x, y), (x', y')) = d_{\Lambda}^2(x, x') I(y = y') + \infty I(y \neq y'),$$

where $d_{\Lambda}^2(x, x') = \|x - x'\|_{\Lambda}^2$. Then, Blanchet et al. (2017) shows that in the generalized logistic regression setting (i.e. $l(x, y, \beta) = \log(1 + \exp(-y\beta^T x))$), if Λ is positive definite, we obtain

$$(8) \quad \min_{\beta} \max_{P: D_{c_{\Lambda}}(P, P_n) \leq \delta} \mathbb{E}[l(X, Y, \beta)] = \min_{\beta} \mathbb{E}_{P_n}[l(X, Y, \beta)] + \delta \|\beta\|_{\Lambda^{-1}}.$$

If the choice of Λ is data driven in order to impose a penalty on transportation costs whose outcomes that are highly impactful in terms of risk, then we would be able to control the risk bound induced by the DD-DRO formulation. This is the strategy studied in Blanchet et al. (2017) in which metric learning procedures have been implemented precisely to achieve such control. Our contribution, as we shall explain in the next section is the use of a robust optimization formulation to calibrate $c_{\Lambda}(\cdot)$. We emphasize that once $c_{\Lambda}(\cdot)$ is calibrated, it can be used to multiple learning tasks and arbitrary loss functions (not only logistic regression).

3.1. Data-Driven Cost Functions via Metric Learning Procedures. We quickly review the elements of standard metric learning procedures. Our data is of the form $\mathcal{D}_n = \{(X_i, Y_i)\}_{i=1}^n$ and $Y_i \in \{-1, +1\}$. The prediction variables are assumed to be standardized.

Motivated by applications such as social networks, in which there is a natural graph which can be used to connect instances in the data, we assume that one is given sets \mathcal{M} and \mathcal{N} , where \mathcal{M} is the set of the pairs that should be close (so that we can connect them) to each other, and \mathcal{N} , on contrary, is characterizing the relations that the pairs should be far away (not connected), we define them as

$$\begin{aligned} \mathcal{M} &= \{(X_i, X_j) \mid X_i \text{ and } X_j \text{ must connect}\}, \\ \text{and } \mathcal{N} &= \{(X_i, X_j) \mid X_i \text{ and } X_j \text{ should not connect}\}. \end{aligned}$$

While it is typically assumed that \mathcal{M} and \mathcal{N} are given, one may always resort to k -Nearest-Neighbor (k -NN) method for the generation of these sets. This is the approach that we follow in our numerical experiments. But we emphasize that choosing any criterion for the definition of \mathcal{M} and \mathcal{N} should be influenced by the learning task in order to retain both interpretability and performance. In our experiments we let (X_i, X_j) belong to \mathcal{M} if, in addition to being sufficiently close (i.e. in the k -NN criterion), $Y_i = Y_j$. If $Y_i \neq Y_j$, then we have that $(X_i, X_j) \in \mathcal{N}$.

In addition, we consider the relative constraint set \mathcal{R} containing data triplets with relative relation defined as

$$\mathcal{R} = \{(i, j, k) \mid d_{\Lambda}(X_i, X_j) \text{ should be smaller than } d_{\Lambda}(X_i, X_k)\}.$$

Let us consider the following two formulations of metric learning, the so-called Absolute Metric Learning formulation

$$(9) \quad \min_{\Lambda \succeq 0} \sum_{(i,j) \in \mathcal{M}} d_{\Lambda}^2(X_i, X_j) \quad \text{s.t.} \quad \sum_{(i,j) \in \mathcal{N}} d_{\Lambda}^2(X_i, X_j) \geq 1,$$

and the Relative Metric Learning formulation,

$$(10) \quad \min_{\Lambda \succeq 0} \sum_{(i,j,k) \in \mathcal{R}} (d_{\Lambda}^2(X_i, X_j) - d_{\Lambda}^2(X_i, X_k) + 1)_+.$$

Both formulations have their merits, (9) exploits both the constraint sets \mathcal{M} and \mathcal{N} , while (10) is only based on information in \mathcal{R} . Further intuition or motivation of those two formulations can be found in Xing et al. (2002) and Weinberger and Saul (2009), respectively. We will show how to formulate and solve the robust counterpart of those two representative examples by robustifying a

single constraint set or two sets simultaneously. For simplicity we only discuss these two formulations, but many metric learning algorithms are based on natural generalizations of those two forms, as mentioned in the survey [Bellet et al. \(2013\)](#).

Once formulation (9) or (10) are considered and the matrix Λ has been calibrated, one may then consider the cost function $c_\Lambda(\cdot)$ in (7) and solve the problem (8). This is the benchmark that we will consider in our numerical experiments. And we will contrast this approach versus a method which chooses Λ using a robust optimization version of (9) or (10) as we shall explain next.

4. ROBUST OPTIMIZATION FOR METRIC LEARNING

In this section, we review a robust optimization method to metric learning optimization problem to learn a robust data-driven cost function. RO is a family of optimization techniques that deals with uncertainty or misspecification in the objective function and constraints. RO was first proposed in [Ben-Tal et al. \(2009\)](#) and has attracted increasing attentions in the recent decades [El Ghaoui and Lebret \(1997\)](#) and [Bertsimas et al. \(2011\)](#). RO has been applied in machine learning to regularize statistical learning procedures, for example, in [Xu et al. \(2009a\)](#) and [Xu et al. \(2009b\)](#) robust optimization was employed for SR-Lasso and support vector machines. We apply RO, as we shall demonstrate, to reduce the variability in testing error when implementing DD-DRO.

4.1. Robust Optimization for Relative Metric Learning. The RO formulation that we shall use for (10) is based on the work of [Huang et al. \(2012\)](#). In order to motivate this formulation, suppose that we know that only α level, e.g. $\alpha = 90\%$, of the constraints are satisfied, but we do not have information on exactly which of them are ultimately satisfied. The value of α may be inferred using cross validation.

Instead of optimizing over all subsets of constraints, we try to minimize the worst case loss function over all possible $\alpha|\mathcal{R}|$ constraints (where $|\cdot|$ is cardinality of a set) and obtain the following min-max formulation

$$(11) \quad \min_{\Lambda \succeq 0} \max_{\tilde{q} \in \mathcal{T}(\alpha)} \sum_{(i,j,k) \in \mathcal{R}} q_{i,j,k} (d_\Lambda^2(X_i, X_j) - d_\Lambda^2(X_i, X_k) + 1)_+,$$

where $\mathcal{T}(\alpha)$ is a robust uncertainty set of the form

$$\mathcal{T}(\alpha) = \left\{ \tilde{q} = \{q_{i,j,k} | (i, j, k) \in \mathcal{R}\} \mid 0 \leq q_{i,j,k} \leq 1, \sum_{(i,j,k) \in \mathcal{R}} q_{i,j,k} \leq \alpha \times |\mathcal{R}| \right\},$$

which is a convex and compact set.

In addition, the objective function in (10) is convex in Λ and concave (linear) in \tilde{q} , so we can switch the order of min-max by resorting to Sion's min-max theorem ([Terkelsen \(1973\)](#)). This important observation suggests an iterative algorithm. For a fixed $\Lambda \succeq 0$, the inner maximization is linear in \tilde{q} , and the optimal \tilde{q} satisfy $\tilde{q}_{i,j,k} = 1$ whenever $(d_\Lambda(X_i, X_j) - d_\Lambda(X_i, X_k) + 1)_+$ ranks in the top $\alpha|\mathcal{R}|$ largest values and equals $\tilde{q}_{i,j,k}$ otherwise.

Let us use $\mathcal{R}_\alpha(\Lambda)$ to denote the subset of constraints satisfying that the corresponding loss function $(d_\Lambda(X_i, X_j) - d_\Lambda(X_i, X_k) + 1)_+$ ranks at the top $\alpha|\mathcal{R}|$ largest values among the corresponding loss function values of the triplets in \mathcal{R} .

For fixed \tilde{q} , the optimization problem is convex in Λ , we can solve this problem using sub-gradient or smoothing approximation algorithms ([Nesterov \(2005\)](#)). Particularly, as we discussed above, if \tilde{q} is the solution for fixed Λ , we know, solving Λ is equivalent to solving its non-robust counterpart (10), replacing \mathcal{R} by $\mathcal{R}_\alpha(\Lambda)$, where $\mathcal{R}_\alpha(\Lambda)$ is a subset of \mathcal{R} that contains the constraints have top $\alpha|\mathcal{R}|$ violation, i.e.

$$\mathcal{R}_\alpha(\Lambda) = \left\{ (i, j, k) \in \mathcal{R} \mid (d_\Lambda^2(X_i, X_j) - d_\Lambda^2(X_i, X_k) + 1)_+ \text{ ranks top } \alpha \text{ within } \mathcal{R} \right\}.$$

We summarize the sub-gradient based sequentially update algorithm as in Algorithm 1.

Algorithm 1 Sequential Coordinate-wise Metric Learning Using Relative Relations

- 1: **Initialize** $\Lambda = I_d$, learning rate $\alpha = 0.01$ tracking error $Error = 1000$ as a large number. Then randomly sample α proportion of elements from \mathcal{R} to construct $\mathcal{R}_\alpha(\Lambda)$.
- 2: **while** $Error > 10^{-3}$ **do**
- 3: Update Λ using projected (projected to positive semidefinite matrix cone) subgradient descent technique.

$$\Lambda = \pi_{\mathbb{S}_+} \left(\Lambda - \alpha \sum_{(i,j,k) \in \mathcal{R}_\alpha(\Lambda)} \nabla_{\Lambda} (d_{\Lambda}^2(X_i, X_j) - d_{\Lambda}^2(X_i, X_k) + 1)_+ \right)$$

- 4: Update tracking error $Error$ as the norm of difference between latest matrix Λ and average of last 50 iterations.
 - 5: Every few steps (5 or 10 iterations), compute $(d_{\Lambda}^2(X_i, X_j) - d_{\Lambda}^2(X_i, X_k) + 1)_+$ for all $(i, j, k) \in \mathcal{R}$, then update $\mathcal{R}_\alpha(\Lambda)$.
 - 6: **Output** Λ .
-

As a remark, we would like to highlight the following observations. While we focus on metric learning simply as a loss minimization procedure as in (10) and (11) for simplicity, in practice people usually add a regularization term (such as $\|\Lambda\|_F$) to the loss minimization, as is common in metric learning literature (see Bellet et al. (2013)). It is easy to observe our discussion above regarding the min-max exchange uses Sion's min-max theorem and everything else remains largely intact if we consider regularization. Likewise, one can use a more general loss functions than the hinge loss used in (10) and (11).

4.2. Robust Optimization for Absolute Metric Learning. The RO formulation that we present here for (9) appears to be novel in the literature. Note that (9) can be written into the Lagrangian form,

$$\min_{\Lambda \succeq 0} \max_{\lambda \geq 0} \sum_{(i,j) \in \mathcal{M}} d_{\Lambda}^2(X_i, X_j) + \lambda \left(1 - \sum_{(i,j) \in \mathcal{N}} d_{\Lambda}^2(X_i, X_j) \right).$$

Following similar discussion for \mathcal{R} , let us assume that the sets \mathcal{M} and \mathcal{N} are noisy or inaccurate at level α (i.e. $\alpha \cdot 100\%$ of their elements are incorrectly assigned). We can construct robust uncertainty sets $\mathcal{W}(\alpha)$ and $\mathcal{V}(\alpha)$ from the constraints in \mathcal{M} and \mathcal{N} as follows,

$$\begin{aligned} \mathcal{W}(\alpha) &= \{ \tilde{\eta} = \{ \eta_{ij} : (i, j) \in \mathcal{M} \} \mid 0 \leq \eta_{ij} \leq 1, \sum_{(i,j) \in \mathcal{M}} \eta_{ij} \leq \alpha \times |\mathcal{M}| \}, \\ \mathcal{V}(\alpha) &= \{ \tilde{\xi} = \{ \xi_{ij} : (i, j) \in \mathcal{N} \} \mid 0 \leq \xi_{ij} \leq 1, \sum_{(i,j) \in \mathcal{N}} \xi_{ij} \geq \alpha \times |\mathcal{N}| \}. \end{aligned}$$

Then we can write the RO counterpart for the loss minimization problem of metric learning as

$$(12) \quad \min_{\Lambda \succeq 0} \max_{\lambda \geq 0} \max_{\tilde{\eta} \in \mathcal{W}(\alpha), \tilde{\xi} \in \mathcal{V}(\alpha)} \sum_{(i,j) \in \mathcal{M}} \eta_{i,j} d_{\Lambda}^2(X_i, X_j) + \lambda \left(1 - \sum_{(i,j) \in \mathcal{N}} \xi_{i,j} d_{\Lambda}^2(X_i, X_j) \right)$$

Note that the Cartesian product $\mathcal{W}(\alpha) \times \mathcal{V}(\alpha)$ is a compact set, and the objective function is convex in Λ and concave (linear) in pair $(\tilde{\eta}, \tilde{\xi})$, so we can apply Sion's min-max Theorem again (see in Terkelsen (1973)) to switch the order of \min_{Λ} - $\max_{(\tilde{\eta}, \tilde{\xi})}$ (after switching \max_{λ} and $\max_{(\tilde{\eta}, \tilde{\xi})}$, which can be done in general). Then we can develop a sequential iterative algorithm to solve this problem as we describe next.

At the k -th step, given fixed $\Lambda_{k-1} \succeq 0$ and $\lambda_{k-1} > 0$ (it is easy to observe that optimal solution λ is positive, i.e. the constraint is active so we may safely assume $\lambda_{k-1} > 0$), the inner maximization problem, becomes,

$$\max_{\tilde{\eta} \in \mathcal{W}(\alpha)} \sum_{(i,j) \in \mathcal{M}} \eta_{i,j} d_{\Lambda_{k-1}}^2(X_i, X_j) + \lambda \left(1 - \min_{\tilde{\xi} \in \mathcal{V}(\alpha)} \sum_{(i,j) \in \mathcal{N}} \xi_{i,j} d_{\Lambda_{k-1}}^2(X_i, X_j) \right).$$

As we discussed for relative constraints case, the optimal solution for $\tilde{\eta}$ and $\tilde{\xi}$ is, $\tilde{\eta}_{i,j}$ is 1, if $d_{\Lambda_{k-1}}^2(X_i, X_j)$ ranks top α within \mathcal{M} and equals 0 otherwise; while, on the contrary, $\tilde{\xi}_{i,j} = 1$ if $d_{\Lambda_{k-1}}^2(X_i, X_j)$ ranks bottom α within \mathcal{N} and equals 0 otherwise.

Similar as $\mathcal{R}_\alpha(\Lambda)$, we can define $\mathcal{M}_\alpha(\Lambda_{k-1})$ ($\mathcal{N}_\alpha(\Lambda_{k-1})$) as subset of \mathcal{M} (\mathcal{N}), which contains the constraints with largest α percent of $d_{\Lambda_{k-1}}(\cdot)$ within in \mathcal{M} ; and $\mathcal{N}_\alpha(\Lambda_{k-1})$ as subset of \mathcal{N} , which contains the constraints with smallest α percent of $d_{\Lambda_{k-1}}(\cdot)$ within in \mathcal{N} . As we observe that the optimal $\tilde{\eta}_{i,j} = 1$ if $(i, j) \in \mathcal{M}_\alpha(\Lambda_{k-1})$ and $\tilde{\xi}_{i,j} = 1$ if $(i, j) \in \mathcal{N}_\alpha(\Lambda_{k-1})$, thus for fixed $\tilde{\eta}$ and $\tilde{\xi}$, we can write the optimization problem over Λ in the constrained case as

$$\min_{\Lambda \succeq 0} \sum_{(i,j) \in \mathcal{M}_\alpha(\Lambda_{k-1})} d_\Lambda^2(X_i, X_j) \quad \text{s.t.} \quad \sum_{(i,j) \in \mathcal{N}_\alpha(\Lambda_{k-1})} d_\Lambda^2(X_i, X_j) \geq 1.$$

This formulation falls within the setting of the problem stated in (9) and thus it can be solved by using techniques discussed in Xing et al. (2002). We summarize the details in Algorithm 2.

Algorithm 2 Sequential Coordinate-wise Metric Learning Using Absolute Constraints

- 1: **Initialize** $A = I_d$, tracking error $Error = 1000$ as a large number. Then randomly sample α proportion of elements from \mathcal{M} (resp. \mathcal{N}) to construct $\mathcal{M}_\alpha(A)$ (resp. $\mathcal{N}_\alpha(A)$).
 - 2: **while** $Error > 10^{-3}$ **do**
 - 3: Update A using procedure provided in Xing et al. (2002).
 - 4: Update tracking error $Error$ as the norm of difference between latest matrix A and average of last 50 iterations.
 - 5: Every few steps (5 or 10 iterations), compute $d_A(W_i, W_j)$ for all $(i, j) \in \mathcal{M} \cup \mathcal{N}$, then update $\mathcal{M}_\alpha(A)$ and $\mathcal{N}_\alpha(A)$.
 - 6: **Output** A .
-

Other robust methods have also been considered in the metric learning literature, see Zha et al. (2009) and Lim et al. (2013) although the connections to RO are not fully exposed.

5. NUMERICAL EXPERIMENTS

We proceed to numerical experiments to verify the performance of our DD-R-DRO method empirically using six binary classification real data sets from UCI machine learning data base Lichman (2013).

We consider logistic regression (LR), regularized logistic regression (LRL1), DD-DRO with cost function learned using absolute constraints (DD-DRO (absolute)) and its $\alpha = 50\%, 90\%$ level of doubly robust DRO (DD-R-DRO (absolute)); DD-DRO with cost function learned using relative constraints (DD-DRO (relative)) and its $\alpha = 50\%, 90\%$ level of doubly robust DRO (DD-R-DRO (relative)). For each data and each experiment, we randomly split the data into training and testing and fit models on training set and evaluate on testing set.

We report the mean and standard deviation of training error, testing error, and testing accuracy via 200 independent experiments for each data sets, and summarize the detailed results and data

set information (including split setting) in Table 1.

For solving the DD-DRO and DD-R-DRO problem, we apply the smoothing approximation algorithm introduced in Blanchet et al. (2017) to solve the DRO problem directly, where the size of uncertainty δ is chosen via 5-fold cross-validation.

TABLE 1. Numerical results for real data sets.

		BC	BN	QSAR	Magic	MB	SB
LR	Train	0 ± 0	$.008 \pm .003$	$.026 \pm .008$	$.213 \pm .153$	0 ± 0	0 ± 0
	Test	8.75 ± 4.75	2.80 ± 1.44	35.5 ± 12.8	17.8 ± 6.77	18.2 ± 10.0	14.5 ± 9.04
	Accur	$.762 \pm .061$	$.926 \pm .048$	$.701 \pm .040$	$.668 \pm .042$	$.678 \pm .059$	$.789 \pm .035$
LRL1	Train	$.185 \pm .123$	$.080 \pm .030$	$.614 \pm .038$	$.548 \pm .087$	$.401 \pm .167$	$.470 \pm .040$
	Test	$.428 \pm .338$	$.340 \pm .228$	$.755 \pm .019$	$.610 \pm .050$	$.910 \pm .131$	$.588 \pm .140$
	Accur	$.929 \pm .023$	$.930 \pm .042$	$.646 \pm .036$	$.665 \pm .045$	$.717 \pm .041$	$.811 \pm .034$
DD-DRO (absolute)	Train	$.022 \pm .019$	$.197 \pm .112$	$.402 \pm .039$	$.469 \pm .064$	$.294 \pm .046$	$.166 \pm .031$
	Test	$.126 \pm .034$	$.275 \pm .093$	$.557 \pm .023$	$.571 \pm .043$	$.613 \pm .053$	$.333 \pm .023$
	Accur	$.954 \pm .015$	$.919 \pm .050$	$.733 \pm .026$	$.727 \pm .039$	$.714 \pm .032$	$.887 \pm .011$
DD-R-DRO (absolute) $\alpha = 90\%$	Train	$.029 \pm .013$	$.078 \pm .031$	$.397 \pm .036$	$.420 \pm .063$	$.249 \pm .055$	$.194 \pm .031$
	Test	$.126 \pm .023$	$.259 \pm .086$	$.554 \pm .019$	$.561 \pm .035$	$.609 \pm .044$	$.331 \pm .018$
	Accur	$.954 \pm .012$	$.910 \pm .042$	$.736 \pm .025$	$.729 \pm .032$	$.709 \pm .025$	$.890 \pm .008$
DD-R-DRO (absolute) $\alpha = 50\%$	Train	$.040 \pm .055$	$.137 \pm .030$	$.448 \pm .032$	$.504 \pm .041$	$.351 \pm .048$	$.166 \pm .030$
	Test	$.132 \pm .015$	$.288 \pm .059$	$.579 \pm .017$	$.590 \pm .029$	$.623 \pm .029$	$.337 \pm .013$
	Accur	$.952 \pm .012$	$.918 \pm .037$	$.733 \pm .025$	$.710 \pm .033$	$.715 \pm .021$	$.888 \pm .008$
DD-DRO (relative)	Train	$.086 \pm .038$	$.436 \pm .138$	$.392 \pm .040$	$.457 \pm .071$	$.322 \pm .061$	$.181 \pm .036$
	Test	$.153 \pm .060$	$.329 \pm .124$	$.559 \pm .025$	$.582 \pm .033$	$.613 \pm .031$	$.332 \pm .016$
	Accur	$.946 \pm .018$	$.916 \pm .075$	$.714 \pm .029$	$.710 \pm .027$	$.704 \pm .021$	$.890 \pm .008$
DD-R-DRO (relative) $\alpha = 90\%$	Train	$.030 \pm .014$	$.244 \pm .121$	$.375 \pm .038$	$.452 \pm .067$	$.402 \pm .058$	$.234 \pm .032$
	Test	$.141 \pm .054$	$.300 \pm .108$	$.556 \pm .022$	$.577 \pm .032$	$.610 \pm .024$	$.332 \pm .011$
	Accur	$.949 \pm .019$	$.921 \pm .070$	$.729 \pm .023$	$.717 \pm .025$	$.710 \pm .020$	$.892 \pm .007$
DD-R-DRO (relative) $\alpha = 50\%$	Train	$.031 \pm .016$	$.232 \pm .094$	$.445 \pm .032$	$.544 \pm .057$	$.365 \pm .054$	$.288 \pm .029$
	Test	$.154 \pm .049$	$.319 \pm .078$	$.570 \pm .019$	$.594 \pm .018$	$.624 \pm .018$	$.357 \pm .008$
	Accur	$.948 \pm .019$	$.918 \pm .081$	$.705 \pm .023$	$.699 \pm .028$	$.698 \pm .018$	$.881 \pm .005$
Num Predictors		30	4	30	10	20	56
Train Size		40	20	80	30	30	150
Test Size		329	752	475	9990	125034	2951

We observe that the doubly robust DRO framework, in general, get robust improvement comparing to its non-robust counterpart with $\alpha = 90\%$. More importantly, the robust methods tend to enjoy the variance reduction property due to RO. Also, as the robust level increases, i.e. $\alpha = 50\%$, where we believe in higher noise in cost function learning, we can observe, the doubly robust based approach seems to shrink towards to LRL1, and benefits less from the data-driven cost structure.

6. DISCUSSION AND CONCLUSION

We have proposed a novel methodology, DD-R-DRO, which calibrates a transportation cost function by using a data-driven approach based on RO. In turn, DD-R-DRO uses this cost function in the description of a DRO formulation based on optimal transport uncertainty region. The overall methodology is doubly robust. On one hand, DD-DRO, which fully uses the training data to estimate the underlying transportation cost enhances out-of-sample performance by allowing an adversary to perturb the data (represented by the empirical distribution) in order to obtain bounds

on the testing risk which are tight. On the other hand, the tightness of bounds might come at the cost of potentially introducing noise in the testing error performance. The second layer of robustification, as shown in the numerical examples, mitigates precisely the presence of this noise.

REFERENCES

- [1] Bellet, A., Habrard, A., and Sebban, M. (2013). A survey on metric learning for feature vectors and structured data. *arXiv preprint arXiv:1306.6709*.
- [2] Ben-Tal, A., El Ghaoui, L., and Nemirovski, A. (2009). *Robust optimization*. Princeton University Press.
- [3] Bertsimas, D., Brown, D. B., and Caramanis, C. (2011). Theory and applications of robust optimization. *SIAM review*, 53(3):464–501.
- [4] Blanchet, J., Kang, Y., and Murthy, K. (2016). Robust wasserstein profile inference and applications to machine learning. *arXiv preprint arXiv:1610.05627*.
- [5] Blanchet, J., Kang, Y., Zhang, F., and Murthy, K. (2017). Data-driven optimal transport cost selection for distributionally robust optimization. *arXiv preprint arXiv:*.
- [6] El Ghaoui, L. and Lebret, H. (1997). Robust solutions to least-squares problems with uncertain data. *SIAM Journal on Matrix Analysis and Applications*, 18(4):1035–1064.
- [7] Esfahani, P. M. and Kuhn, D. (2015). Data-driven distributionally robust optimization using the wasserstein metric: Performance guarantees and tractable reformulations. *arXiv preprint arXiv:1505.05116*.
- [8] Hu, Z. and Hong, L. J. (2013). Kullback-leibler divergence constrained distributionally robust optimization. *Available at Optimization Online*.
- [9] Huang, K., Jin, R., Xu, Z., and Liu, C.-L. (2012). Robust metric learning by smooth optimization. *arXiv preprint arXiv:1203.3461*.
- [10] Lichman, M. (2013). UCI machine learning repository.
- [11] Lim, D., McFee, B., and Lanckriet, G. R. (2013). Robust structural metric learning. In *ICML-13*, pages 615–623.
- [12] Nesterov, Y. (2005). Smooth minimization of non-smooth functions. *Mathematical programming*, 103(1):127–152.
- [13] Shafieezadeh-Abadeh, S., Esfahani, P. M., and Kuhn, D. (2015). Distributionally robust logistic regression. In *Advances in Neural Information Processing Systems*, pages 1576–1584.
- [14] Terkelsen, F. (1973). Some minimax theorems. *Mathematica Scandinavica*, 31(2):405–413.
- [15] Villani, C. (2008). *Optimal transport: old and new*, volume 338. Springer Science & Business Media.
- [16] Weinberger, K. Q. and Saul, L. K. (2009). Distance metric learning for large margin nearest neighbor classification. *JMLR*, 10(Feb):207–244.
- [17] Xing, E. P., Ng, A. Y., Jordan, M. I., and Russell, S. (2002). Distance metric learning with application to clustering with side-information. In *NIPS-2002*, volume 15, page 12.
- [18] Xu, H., Caramanis, C., and Mannor, S. (2009a). Robust regression and lasso. In *NIPS-2009*, pages 1801–1808.
- [19] Xu, H., Caramanis, C., and Mannor, S. (2009b). Robustness and regularization of support vector machines. *JMLR*, 10(Jul):1485–1510.
- [20] Zha, Z.-J., Mei, T., Wang, M., Wang, Z., and Hua, X.-S. (2009). Robust distance metric learning with auxiliary knowledge. In *IJCAI*, pages 1327–1332.

COLUMBIA UNIVERSITY, DEPARTMENT OF STATISTICS AND DEPARTMENT OF INDUSTRIAL ENGINEERING & OPERATIONS RESEARCH, NEW YORK, NY 10027, UNITED STATES.

E-mail address: jose.blanchet@columbia.edu

COLUMBIA UNIVERSITY, DEPARTMENT OF STATISTICS. NEW YORK, NY 10027, UNITED STATES.

E-mail address: yangkang@stat.columbia.edu

COLUMBIA UNIVERSITY, DEPARTMENT OF STATISTICS AND DEPARTMENT OF INDUSTRIAL ENGINEERING & OPERATIONS RESEARCH. NEW YORK, NY 10027, UNITED STATES.

E-mail address: fz2222@columbia.edu

COLUMBIA UNIVERSITY, DEPARTMENT OF STATISTICS AND DEPARTMENT OF INDUSTRIAL ENGINEERING & OPERATIONS RESEARCH. NEW YORK, NY 10027, UNITED STATES.

E-mail address: fh2293@columbia.edu

E-mail address: hu.zhangyi@gmail.com