

DATA-DRIVEN OPTIMAL TRANSPORT COST SELECTION FOR DISTRIBUTIONALLY ROBUST OPTIMIZATION

Jose Blanchet

Management Science and Engineering
Stanford University
475 Via Ortega, Suite 310
Stanford, CA, 94305, USA

Yang Kang

Department of Statistics
Columbia University
1255th Amsterdam Ave. RM1005
New York, NY, 10027, USA

Karthyek Murthy

Engineering Systems & Design
Singapore University of Technology & Design
8 Somapah Rd
487372, SINGAPORE

Fan Zhang

Management Science and Engineering
Stanford University
475 Via Ortega, Suite 310
Stanford, CA, 94305, USA

ABSTRACT

Some recent works showed that several machine learning algorithms, such as square-root Lasso, Support Vector Machines, and regularized logistic regression, among many others, can be represented exactly as distributionally robust optimization (DRO) problems. The distributional uncertainty set is defined as a neighborhood centered at the empirical distribution, and the neighborhood is measured by optimal transport distance. In this paper, we propose a methodology which learns such neighborhood in a natural data-driven way. We show rigorously that our framework encompasses adaptive regularization as a particular case. Moreover, we demonstrate empirically that our proposed methodology is able to improve upon a wide range of popular machine learning estimators.

1 INTRODUCTION

A Distributionally Robust Optimization (DRO) problem takes the general form

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

where β is a decision variable, (X, Y) is a random element, and $l(x, y; \beta)$ measures a suitable loss or cost incurred when $(X, Y) = (x, y)$ and the decision β is taken. The expectation $\mathbb{E}_P[\cdot]$ is taken under the probability model P . The set \mathcal{U}_{δ} is called the distributional uncertainty set and it is indexed by the parameter $\delta > 0$, which measures the size of the distributional uncertainty.

The DRO problem is said to be *data-driven* if the uncertainty set \mathcal{U}_{δ} is informed by empirical observations. One natural way to supply this information is by letting the “center” of the uncertainty region be placed at the empirical measure, P_n , induced by a data set $\{X_i, Y_i\}_{i=1}^n$, which represents an empirical sample of realizations of (X, Y) that follows a population distribution P_0 . We denote by β_* the optimal decision that minimizes the population risk $\mathbb{E}_{P_0}[l(X, Y; \beta)]$. In order to emphasize the data-driven nature of a DRO formulation such as (1), when the uncertainty region is informed by an empirical sample, we

write $\mathcal{U}_\delta = \mathcal{U}_\delta(P_n)$. To the best of our knowledge, the available data is utilized in the DRO literature only by defining the center of the uncertainty region $\mathcal{U}_\delta(P_n)$ as the empirical measure P_n .

Our goal in this paper is to discuss a data-driven framework to inform the *shape* of $\mathcal{U}_\delta(P_n)$. Throughout this paper, we assume that a sensible loss function $l(x, y; \beta)$ has been selected for the problem at hand. Our contribution concerns the construction of the uncertainty region in a fully data-driven way and the implications of this design in machine learning applications. Before providing our construction, let us discuss the significance of the data-driven DRO in the context of machine learning.

Recently, (Blanchet et al. 2016; Shafieezadeh-Abadeh et al. 2015) showed that many prevailing machine learning estimators can be represented exactly as a data-driven DRO formulation in (1). For example, suppose that $X \in \mathbb{R}^d$ and $Y \in \{-1, 1\}$. Further, let $l(x, y, \beta) = \log(1 + \exp(-y\beta^T x))$ be the log-exponential loss associated to a logistic regression model where $Y \sim \text{Ber}(1/(1 + \exp(-\beta_*^T x)))$, and β_* is the underlying parameter to learn. Then, given a set of empirical samples $\mathcal{D}_n = \{(X_i, Y_i)\}_{i=1}^n$, and a judicious choice of the distributional uncertainty set $\mathcal{U}_\delta(P_n)$, (Blanchet et al. 2016) shows that

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

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 $1/p + 1/q = 1$. If $p = 1$ we see that (2) recovers L_1 regularized logistic regression (see (Friedman et al. 2001)). Other estimators such as Support Vector Machines and sqrt-Lasso are shown in (Blanchet et al. 2016) to admit DRO representations analogous to (2) – provided that the loss function and the uncertainty region are judiciously chosen. Note that the parameter δ in $\mathcal{U}_\delta(P_n)$ is precisely the regularization parameter in the right hand side of (2). So, the data-driven DRO representation (2) provides a direct interpretation of the regularization parameter as the size of the probabilistic uncertainty around the empirical evidence.

An important element to all of the DRO representations obtained in (Blanchet et al. 2016) is that the design of the uncertainty region $\mathcal{U}_\delta(P_n)$ is based on optimal transport theory. In particular, we have that

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

and $D_c(P, P_n)$ is 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

$$c((x, y), (x', y')) = \|x - x'\|_q^2 I(y = y') + \infty \cdot I(y \neq y'). \tag{4}$$

In the end, as we discuss in Section 3, $D_c(P, P_n)$ can be easily computed as the solution of a linear programming (LP) problem which is known as Kantorovich’s problem (Villani 2008).

Other discrepancy notions between probability models have been considered, typically using the Kullback-Leibler divergence and other divergence based notions (Hu, Z., and L. J. Hong. ; Lam and Zhou 2017; Ghosh and Lam 2019). Using divergence (or likelihood ratio) based discrepancies to characterize the uncertainty region $\mathcal{U}_\delta(P_n)$ forces the models $P \in \mathcal{U}_\delta(P_n)$ to share the same support with P_n , which may restrict generalization properties of a DRO-based estimator, and such restriction may induce overfitting problem (Esfahani and Kuhn (2018) and Blanchet et al. (2016)).

In summary, data-driven DRO via optimal transport has been shown to encompass a wide range of prevailing machine learning estimators. However, so far the cost function $c(\cdot)$ is fixed, and not chosen in a data-driven way.

Our main contribution in this paper is to propose a comprehensive approach for designing the uncertainty region $\mathcal{U}_\delta(P_n)$ in a fully data-driven way, using the convenient role of $c(\cdot)$ in the definition of the optimal transport discrepancy $D_c(P, P_n)$. Our modeling approach further underscores, beyond the existence of representations such as (2), the convenience of working with an optimal transport discrepancy for the

design of data-driven DRO machine learning estimators. In other words, because one can select $c(\cdot)$ in a data driven way, it is sensible to use our data-driven DRO formulation even if one is not able to simplify the inner optimization in order to achieve a representation such as (2).

Our idea is to apply metric-learning procedures to estimate $c(\cdot)$ from the training data. Then, use such data-driven $c(\cdot)$ in the definition of $D_c(P, P_n)$ and the construction $\mathcal{U}_\delta(P_n)$ in (3). Finally, solve the DRO problem (1), using cross-validation to choose δ .

The intuition behind our proposal is the following. By using a metric learning procedure we are able to calibrate a cost function $c(\cdot)$ which attaches relatively high transportation costs to (u, w) if transporting mass between these locations substantially impacts performance (e.g. in the response variable, so increasing the expected risk). In turn, the adversary, with a given budget δ , will carefully choose the data which is to be transported. The mechanism will then induce enhanced out-of-sample performance focusing precisely on regions of relevance, while improving generalization error.

One of the challenges for the implementation of our idea is to efficiently solve (1). We address this challenge by proposing a stochastic gradient descent algorithm which takes advantage of a duality representation and fully exploits the nature of the linear programming structure embedded in the definition of $D_c(P, P_n)$, together with a smoothing technique.

Another challenge consists in selecting the type of cost $c(\cdot)$ to be used in practice, and the methodology to fit such cost. To cope with this challenge, we rely on metric-learning procedures. We do not contribute any novel metric learning methodology; rather, we discuss various parametric cost functions and methods developed in the metric-learning literature. In particular, we discuss their use in the context of fully data-driven DRO formulations for machine learning problems – which is what we propose in this paper. The choice of $c(\cdot)$ ultimately will be influenced by the nature of the data and the application at hand. For example, in the setting of image recognition, it might be natural to use a cost function related to similarity notions.

In addition to discuss intuitively the benefits of our approach in Section 2, we also show that our methodology provides a way to naturally estimate various parameters in the setting of adaptive regularization. For example, Theorem 1 below, shows that choosing $c(\cdot)$ using a suitable weighted norm, allows us to recover an adaptive regularized ridge regression estimator (Ishwaran and Rao 2014). In turn, using standard techniques from metric learning we can estimate $c(\cdot)$. Hence, our technique connects metric learning tools to estimate the parameters of adaptive regularized estimators.

More broadly, we compare the performance of our procedure with a number of alternatives in the setting of various data sets and show that our approach exhibits consistently superior performance.

2 DATA-DRIVEN DRO: INTUITION AND INTERPRETATIONS

One of the main benefits of the DRO formulations such as (1) and (2) is their interpretability. For example, we can readily see from the left hand side of (2) that the regularization parameter corresponds precisely to the size of the *data-driven* distributional uncertainty.

The data-driven aspect is important because we can employ statistical thinking to optimally characterize the size of the uncertainty, δ . This readily implies an optimal choice of the regularization parameter, as explained in (Blanchet et al. 2016), in settings such as (2). Elaborating, we can interpret $\mathcal{U}_\delta(P_n)$ as the set of plausible variations of the empirical data, P_n . Consequently, for instance, in the linear regression setting leading to (2), the estimate $\beta_P = \arg \min_\beta \mathbb{E}_P[l(X, Y, \beta)]$ is a plausible estimate of the regression parameter β_* as long as $P \in \mathcal{U}_\delta(P_n)$. Hence, the set

$$\Lambda_\delta(P_n) = \{\beta_P : P \in \mathcal{U}_\delta(P_n)\}$$

is a natural confidence region for β_* that shrinks when δ is decreasing. Thus, a statistically minded approach for choosing δ is to fix a confidence level, say $(1 - \alpha)$, and choose an optimal δ ($\delta_*(n)$) via

$$\delta_*(n) := \inf\{\delta : P(\beta_* \in \Lambda_\delta(P_n)) \geq 1 - \alpha\}. \tag{5}$$

Note that the random element in $P(\beta_* \in \Lambda_\delta(P_n))$ is given by P_n . In (Blanchet et al. 2016) this optimization problem is solved asymptotically as $n \rightarrow \infty$ under standard assumptions on the data generating process. If the underlying model is correct, one would typically obtain, as in (Blanchet et al. 2016), that $\delta_*(n) \rightarrow 0$ at a suitable rate. For instance, in the linear regression setting corresponding to (2), if the data is i.i.d. with finite variance and the linear regression model holds then $\delta_*(n) = \chi_{1-\alpha}(1 + o(1))/n$ as $n \rightarrow \infty$ (where χ_α is the α quantile of an explicitly characterized distribution).

In practice, one can also choose δ by cross-validation. The work of (Blanchet et al. 2016) compares the asymptotically optimal choice $\delta_*(n)$ against cross-validation, concluding that the performance is comparable in the experiments performed. In this paper, we use cross validation to choose δ , but the insights behind the limiting behavior of (5) are useful, as we shall see, to inform the design of our algorithms.

More generally, the DRO formulation (1) is appealing because the distributional uncertainty endows the estimation of β directly with a mechanism to enhance generalization properties. To wit, we can interpret (1) as a game in which we (the outer player) choose a decision β , while the adversary (the inner player) selects a model which is a perturbation, P , of the data (encoded by P_n). The amount of the perturbation is dictated by the size of δ which, as discussed earlier, is data driven. But the type of perturbation and how the perturbation is measured is dictated by $D_c(P, P_n)$. It makes sense to inform the design of $D_c(\cdot)$ using a data-driven mechanism, which is our goal in this paper. The intuition is to allow the types of perturbations which focus the effort and budget of the adversary mostly on out-of-sample exploration over regions of relevance.

The type of benefit that is obtained by informing $D_c(P, P_n)$ with data is illustrated in Figure 1 below, which illustrates a classification task. The data roughly lies on a lower dimensional non-linear manifold. Some data which is classified with a negative label is seen to be “close” to data which is classified with a positive label when seeing the whole space (i.e. \mathbb{R}^2) as the natural ambient domain of the data. However, if we use a distance similar to the geodesic distance intrinsic to the manifold we would see that the negative instances are actually far from the positive instances. So, the generalization properties of the algorithm would be enhanced relative to using a standard metric in the ambient space, because with a given budget δ the adversarial player would be allowed perturbations mostly along the intrinsic manifold where the data lies and instances which are surrounded (in the intrinsic metric) by similarly classified examples will naturally carry significant impact in testing performance. A quantitative example to illustrate this point will be discussed in the sequel.




Figure 1: Stylized examples illustrating the need for data-driven cost function.

3 BACKGROUND ON OPTIMAL TRANSPORT AND METRIC LEARNING PROCEDURES

In this section we quickly review basic notions on optimal transport and metric learning methods so that we can define $D_c(P, P_n)$ and explain how to calibrate the function $c(\cdot)$.

3.1 Defining Optimal Transport Distance 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

$$D_c(P, Q) = \inf \{ \mathbb{E}_\pi [c(U, V)] : \pi \in \mathcal{P}(\mathcal{S}_P \times \mathcal{S}_Q), \pi_U = P, \pi_V = Q \}, \quad (6)$$

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 $\rho \geq 1$ such that $c^{1/\rho}(u, w) \leq c^{1/\rho}(u, v) + c^{1/\rho}(v, w)$ (i.e. $c^{1/\rho}(\cdot)$ satisfies the triangle inequality) then it can be easily verified (see (Villani 2008)) that $D_c^{1/\rho}(P, Q)$ is a metric. For example, if $c(u, v) = \|u - v\|_q^p$ 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 (6) is a linear program in the variable π .

3.2 On Metric Learning Procedures

In order to keep the discussion focused, we use a few metric learning procedures, but we emphasize that our approach can be used in combination with virtually any method in the metric learning literature, see the survey paper (Bellet et al. 2013) that contains additional discussion on metric learning procedures. The procedures that we consider, as we shall see, can be seen to already improve significantly upon natural benchmarks. Moreover, as we shall see, these metric families can be related to adaptive regularization. This connection will be useful to further enhance the intuition of our procedure.

3.2.1 The Mahalanobis Distance

The Mahalanobis metric is defined as

$$d_\Lambda(x, x') = \left((x - x')^T \Lambda (x - x') \right)^{1/2},$$

where Λ is symmetric and positive semi-definite and we write $\Lambda \in PSD$. Note that $d_\Lambda(x, x')$ is the metric induced by the norm $\|x\|_\Lambda = \sqrt{x^T \Lambda x}$.

For a discussion, assume that 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

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

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 $(X_i, X_j) \in \mathcal{N}$.

The work of (Xing et al. 2002), one of the earlier reference on the subject, suggests considering

$$\min_{\Lambda \in PSD} \sum_{(X_i, X_j) \in \mathcal{M}} d_\Lambda^2(X_i, X_j) \quad s.t. \quad \sum_{(X_i, X_j) \in \mathcal{N}} d_\Lambda^2(X_i, X_j) \geq \bar{\lambda}. \quad (7)$$

In words, the previous optimization problem minimizes the total distance between pairs that should be connect, while keeping the pairs that should not connect well separated. The constant $\tilde{\lambda} > 0$ is somewhat arbitrary (given that Λ can be normalized by $\tilde{\lambda}$, we can choose $\tilde{\lambda} = 1$).

The optimization problem (7) is an LP problem on the convex cone PSD and it has been widely studied. Since $\Lambda \in PSD$, we can always write $\Lambda = LL^T$, and therefore $d_\Lambda(X_i, X_j) = \|X_i - X_j\|_\Lambda := \|LX_i - LX_j\|_2$. There are various techniques which can be used to exploit the PSD structure to efficiently solve (7); see, for example, (Xing et al. 2002) for a projection-based algorithm; or (Schultz and Joachims 2003), which uses a factorization-based procedure; or the survey paper (Bellet et al. 2013) for the discussion of a wide range of techniques.

We have chosen formulation (7) to estimate Λ because it is intuitive and easy to state, but the topic of learning Mahalanobis distances is an active area of research and there are different algorithms which can be implemented (Li et al. 2018).

3.2.2 Using Mahalanobis Distance In Data-Driven DRO

Let us assume that the underlying data takes the form $\mathcal{D}_n = \{(X_i, Y_i)\}_{i=1}^n$, where $X_i \in R^d$ and $Y_i \in R$ and the loss function, depending on a decision variable $\beta \in R^m$, is given by $l(x, y; \beta)$. Note that we are not imposing any linear structure on the underlying model or in the loss function. Then, motivated by the cost function (4), we may consider

$$c_\Lambda((x, y), (x', y')) = d_\Lambda^2(x, x')I(y = y') + \infty I(y \neq y'), \quad (8)$$

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

Even in this case, since the definitions of \mathcal{M} and \mathcal{N} depend on $W_i = (X_i, Y_i)$ (in particular, on the response variable), cost function $c_\Lambda(\cdot)$ (once Λ is calibrated using, for example, the method discussed in the previous subsection), will be informed by the Y_i s. Then, we estimate β via

$$\min_{\beta} \sup_{P: D_{c_\Lambda}(P, P_n) \leq \delta} \mathbb{E}_P[l(X, Y; \beta)]. \quad (9)$$

It is important to note that Λ has been applied only to the definition of the cost function.

The intuition behind the formulation can be gained in the context of a logistic regression setting, see the example in Figure 2: Suppose that $d = 2$, and that Y depends only on $X(1)$ (i.e. the first coordinate of X). Then, the metric learning procedure in (7) will induce a relatively low transportation cost across the $X(2)$ direction and a relatively high transportation cost in the $X(1)$ direction, whose contribution, being highly informative, is reasonably captured by the metric learning mechanism. Since the $X(1)$ direction is most impactful, from the standpoint of expected loss estimation, the adversarial player will reach a compromise, between transporting (which is relatively expensive) and increasing the expected loss (which is the adversary's objective). Out of this compromise the DRO procedure localizes the out-of-sample enhancement, and yet improves generalization.

3.2.3 Mahalanobis Metric on a Non-Linear Feature Space

In this section, in order to capture the non-linear structure of the data, we apply a non-linear feature map $\Phi: R^d \rightarrow R^l$ to the data and then learn the Mahalanobis metric on the feature space. Assume that the data takes the form $\mathcal{D}_n = \{(X_i, Y_i)\}_{i=1}^n$, where $X_i \in R^d$ and $Y_i \in R$ and the loss function, depending on decision variable $\beta \in R^m$, is given by $l(x, y; \beta)$. Once again, motivated by the cost function (4), we may define

$$c_\Lambda^\Phi((x, y), (x', y')) = d_\Lambda^2(\Phi(x), \Phi(x'))I(y = y') + \infty I(y \neq y'), \quad (10)$$

for $\Lambda \in PSD$. To preserve the properties of a cost function (i.e. non-negativity, lower semicontinuity and $c_\Lambda^\Phi(u, w) = 0$ implies $u = w$), we assume that $\Phi(\cdot)$ is continuous and that $\Phi(w) = \Phi(u)$ implies that $w = u$.




Figure 2: Motivation for Mahalanobis cost function.

Then we can apply a metric learning procedure, such as the one described in (7), to calibrate Λ . The intuition is the same as the one provided in the linear case in Section 3.2.2.

4 DATA-DRIVEN COST SELECTION AND ADAPTIVE REGULARIZATION

In this section we establish a direct connection between our fully data-driven DRO procedure and adaptive regularization. Moreover, our main result here, together with our discussion from the previous section, provides a direct connection between the metric learning literature and adaptive regularized estimators. As a consequence, the methods from the metric learning literature can be used to estimate the parameter of adaptively regularized estimators.

Throughout this section we consider again a data set of the form $\mathcal{D}_n = \{(X_i, Y_i)\}_{i=1}^n$ with $X_i \in \mathbb{R}^d$ and $Y_i \in \mathbb{R}$. Motivated by the cost function (4) we define the cost function $c_\Lambda(\cdot)$ as in (8). Using (8) we obtain the following result.

Theorem 1 (DRO Representation for Generalized Adaptive Regularization) Assume that $\Lambda \in \mathbb{R}^{d \times d}$ in (8) is positive definite. Given the data set \mathcal{D}_n , we obtain the following representation

$$\min_{\beta} \max_{P: D_{c_\Lambda}(P, P_n) \leq \delta} \mathbb{E}_P^{1/2} \left[(Y - X^T \beta)^2 \right] = \min_{\beta} \sqrt{\frac{1}{n} \sum_{i=1}^n (Y_i - X_i^T \beta)^2} + \sqrt{\delta} \|\beta\|_{\Lambda^{-1}}. \tag{11}$$

Moreover, if $Y \in \{-1, +1\}$ in the context of adaptive regularized logistic regression, we obtain the following representation

$$\min_{\beta} \max_{P: D_{c_\Lambda}(P, P_n) \leq \delta} \mathbb{E}_P \left[\log \left(1 + e^{-Y(X^T \beta)} \right) \right] = \min_{\beta} \frac{1}{n} \sum_{i=1}^n \log \left(1 + e^{-Y_i(X_i^T \beta)} \right) + \delta \|\beta\|_{\Lambda^{-1}}. \tag{12}$$

In order to recover a more familiar setting in adaptive regularization, assume that Λ is a diagonal positive definite matrix. In which case we obtain, in the setting of (11),

$$\min_{\beta} \max_{P: D_{c_\Lambda}(P, P_n) \leq \delta} \mathbb{E}_P^{1/2} \left[(Y - X^T \beta)^2 \right] = \min_{\beta} \sqrt{\frac{1}{n} \sum_{i=1}^n (Y_i - X_i^T \beta)^2} + \sqrt{\delta} \sqrt{\sum_{i=1}^d \beta_i^2 / \Lambda_{ii}}. \tag{13}$$

The proof of Theorem 1 follows similar dual problem derivation as for Theorem 1 in (Blanchet and Kang 2017a). We refer reader to Appendix A.2 in Blanchet and Kang (2017a) for technique details.

The adaptive regularization method was first derived as a generalization for ridge regression in (Hoerl and Kennard 1970b; Hoerl and Kennard 1970a). Recent works show that adaptive regularization can improve the predictive power of its non-adaptive counterpart, specially in high-dimensional settings (Zou 2006; Ishwaran and Rao 2014).

In view of (13), our discussion in Section 3.2.1 uncovers tools which can be used to estimate the coefficients $\{1/\Lambda_{ii} : 1 < i \leq d\}$ using the connection to metric learning procedures. To complement the intuition given in Figure 2, note that in the adaptive regularization literature one often choose $\Lambda_{ii} \approx 0$ to induce $\beta_i \approx 0$ (i.e., there is a high penalty to variables with low explanatory power). This, in our setting, would correspond to transport costs which are low in such low explanatory directions.

5 SOLVING DATA-DRIVEN DRO BASED ON OPTIMAL TRANSPORT DISCREPANCIES

In order to fully take advantage of the combination synergies between metric learning methodology and our DRO formulation, it is crucial to have a methodology which allows us to efficiently estimate β in DRO problems such as (1). In the presence of a simplified representation such as (2) or (13), we can apply standard stochastic optimization results (Lei and Jordan 2016).

Our objective in this section is to study algorithms which can be applied for more general loss and cost functions, for which a simplified representation might not be accessible.

Throughout this section, once again we assume that the data is given in the form $\mathcal{D}_n = \{(X_i, Y_i)\}_{i=1}^n \subset \mathbb{R}^{d+1}$. The loss function is written as $\{l(x, y; \beta) : (x, y) \in \mathbb{R}^{d+1}, \beta \in \mathbb{R}^m\}$. We assume that for each (x, y) , the function $l(x, y, \cdot)$ is convex and continuously differentiable. Further, we shall consider cost functions of the form

$$\bar{c}((x, y), (x', y')) = c(x, x')I(y = y') + \infty I(y \neq y'),$$

as this will simplify the form of the dual representation in the inner optimization of our DRO formulation. To ensure boundedness of our DRO formulation, we impose the following assumption.

Assumption 1. There exists $\Gamma(\beta, y) \in (0, \infty)$ such that $l(u, y; \beta) \leq \Gamma(\beta, y) \cdot (1 + c(u, x))$, for all $(x, y) \in \mathcal{D}_n$. Under Assumption 1, we can guarantee that

$$\max_{P: D_c(P, P_n) \leq \delta} \mathbb{E}_P[l(X, Y; \beta)] \leq (1 + \delta) \max_{i=1, \dots, n} \Gamma(\beta, Y_i) < \infty.$$

Using the strong duality theorem for semi-infinite linear programming problem in Appendix B of (Blanchet et al. 2016),

$$\max_{P: D_c(P, P_n) \leq \delta} \mathbb{E}_P[l(X, Y; \beta)] = \min_{\lambda \geq 0} \frac{1}{n} \sum_{i=1}^n \phi(X_i, Y_i, \beta, \lambda),$$

where $\psi(u, X, Y, \beta, \lambda) := l(u, Y; \beta) - \lambda(c(u, X) - \delta)$, $\phi(X, Y, \beta, \lambda) := \max_{u \in \mathbb{R}^d} \psi(u, X, Y, \beta, \lambda)$. Therefore,

$$\min_{\beta} \max_{P: D_c(P, P_n) \leq \delta} \mathbb{E}_P[l(X, Y; \beta)] = \min_{\lambda \geq 0, \beta} \{\mathbb{E}_{P_n}[\phi(X, Y, \beta, \lambda)]\}. \quad (14)$$

The optimization in (14) is minimize over β and λ , which we can consider stochastic approximation algorithm if the gradient of $\phi(\cdot)$ with respect to β and λ exist. However, $\phi(\cdot)$ is given in the form of the value function of a maximization problem, of which the gradient is not easily accessible. We will discuss the detailed algorithm and the validity of the smoothing approximation below.

We consider a smoothing approximation technique to remove the maximization problem $\phi(\cdot)$ using soft-max counterpart, $\phi_{\varepsilon, f}(\cdot)$. The smoothing soft-max approximation has been explored and applied to approximately solve the DRO problem for the discrete case, where we restrict the distributionally uncertainty set only contains probability measures support on finite set (i.e., labeled training data and unlabeled training data with pseudo labels), we refer Blanchet and Kang (2017b) for further details.

However, due to the continuous-infinite support constraint, the soft-max approximation is a non-trivial generalization of the finite-discrete analogue. The smoothing approximation for $\phi(\cdot)$ is defined as,

$$\phi_{\varepsilon, f}(X, Y, \beta, \lambda) = \varepsilon \log \left(\int_{\mathbb{R}^d} \exp([\psi(u, X, Y, \beta, \lambda)]/\varepsilon) f(u) du \right),$$

where $f(\cdot)$ is a probability density in \mathbb{R}^d ; for example, we can consider a multivariate normal distribution and ε is a small positive number regarded as smoothing parameter.

Theorem 2 below allows to quantify the error due to smoothing approximation.

Theorem 2 Under mild technical assumptions (see Assumption 1-4 in Appendix A), there exists $\varepsilon_0 > 0$ such that for every $\varepsilon < \varepsilon_0$, we have

$$\phi(X, Y, \beta, \lambda) \geq \phi_{\varepsilon, f}(X, Y, \beta, \lambda) \geq \phi(X, Y, \beta, \lambda) - d\varepsilon \log(1/\varepsilon)$$

The proof of Theorem 2 is given in Appendix A.

After applying smooth approximation, the optimization problem turns into a standard stochastic optimization problem and we can apply mini-batch based stochastic approximation (SA) algorithm to solve it. As we can notice, as a function and β and λ , the gradient of $\phi_{\varepsilon, f}(\cdot)$ satisfies

$$\begin{aligned} \nabla_{\beta} \phi_{\varepsilon, f}(X, Y, \beta, \lambda) &= \frac{\mathbb{E}_{U \sim f} [\exp(\psi(U, X, Y, \beta, \lambda) / \varepsilon) \nabla_{\beta} l(f_{\beta}(U), Y)]}{\mathbb{E}_{U \sim f} [\exp(\psi(U, X, Y, \beta, \lambda) / \varepsilon)]}, \\ \nabla_{\lambda} \phi_{\varepsilon, f}(X, Y, \beta, \lambda) &= \frac{\mathbb{E}_{U \sim f} [\exp(\psi(u, X, Y, \beta, \lambda) / \varepsilon) (\delta - c_{\mathcal{D}_n}(u, X))]}{\mathbb{E}_{U \sim f} [\exp(\psi(U, X, Y, \beta, \lambda) / \varepsilon)]}. \end{aligned}$$

However, since the gradients are still given in the form of expectation, we can apply a simple Monte Carlo sampling algorithm to approximate the gradient, i.e., we sample U_i 's from $f(\cdot)$ and evaluate the numerators and denominators of the gradient using Monte Carlo separately. For more details of the SA algorithm, please see in Algorithm 1.

Algorithm 1 Stochastic Gradient Descent with Continuous State

- 1: **Initialize** $\lambda = 0$, and β to be empirical risk minimizer, $\varepsilon = 0.5$, tracking error $Error = 100$.
- 2: **while** $Error > 10^{-3}$ **do**
- 3: **Sample** a mini-batch uniformly from observations $\{X_{(j)}, Y_{(j)}\}_{j=1}^M$, with $M \leq n$.
- 4: For each $j = 1, \dots, M$, sample i.i.d. $\{U_k^{(j)}\}_{k=1}^L$ from $\mathcal{N}(0, \sigma^2 I_{d \times d})$.
- 5: We denote f_L^j as empirical distribution for $U_k^{(j)}$'s, and estimate the batched as

$$\nabla_{\beta} \phi_{\varepsilon, f} = \frac{1}{M} \sum_{j=1}^M \nabla_{\beta} \phi_{\varepsilon, f_L^j}(X_{(j)}, Y_{(j)}, \beta, \lambda), \nabla_{\lambda} \phi_{\varepsilon, f} = \frac{1}{M} \sum_{j=1}^M \nabla_{\lambda} \phi_{\varepsilon, f_L^j}(X_{(j)}, Y_{(j)}, \beta, \lambda).$$

- 6: Update β and λ using $\beta = \beta + \alpha_{\beta} \nabla_{\beta} \phi_{\varepsilon, f}$ and $\lambda = \lambda + \alpha_{\lambda} \nabla_{\lambda} \phi_{\varepsilon, f}$.
 - 7: Update tracking error $Error$ as the norm of difference between latest parameter and average of last 50 iterations.
 - 8: **Output** β .
-

6 NUMERICAL EXPERIMENTS

We validate our data-driven cost function based DRO using 6 real data examples from the UCI machine learning database (Lichman, M. 2013), with BC for breast-cancer, BN for, QSAR for QSAR biodegration, Magic for MAGIC Gamma Telescope, MB for MiniBooNE particle identification, and SB for Spambase. Those are standard binary classification task with multivariate predictors. And in the table, we show the average log-exponential loss function and its 1-standard deviation. We focus on a DRO formulation based on the log-exponential loss for a linear model. We use the linear metric learning framework explained in equation (7), which we feed into the cost function, c_{Λ} , as in (8), denoting by DRO-L. In addition, we also fit a cost function c_{Λ}^{Φ} , as explained in (10) using linear and quadratic transformations, in order to

capture the nonlinear structure of the data; the outcome is denote as (DRO-NL). We compare our DRO-L and DRO-NL with logistic regression (LR), and regularized logistic regression (LRL1). For each iteration and each data set, the data is split randomly into training and test sets. In order to demonstrate the power of DRO formulation, the empirical distribution P_n should be deviate from the underlying distribution P_0 to some extent. Thus, we keep the size of training set relatively small to demonstate the prediction power of DRO formulation. We fit the models on the training set and evaluate the performance on test set. The regularization parameter is chosen via 5–fold cross-validation for LRL1, DRO-L and DRO-NL. We report the mean and standard deviation for training and testing log-exponential error and testing accuracy for 200 independent experiments for each data set. The details of the numerical results and basic information of the data is summarized in Table 1.

		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	$.929 \pm .023$	$.930 \pm .042$	$.646 \pm .036$	$.665 \pm .045$	$.717 \pm .041$	$.811 \pm .034$
DRO-L	Train	$.022 \pm .019$	$.197 \pm .112$	$.402 \pm .039$	$.469 \pm .064$	$.294 \pm .046$	$.166 \pm .031$
	Test	$.126 \pm .034$	$.275 \pm .093$	$.557 \pm .023$	$.571 \pm .043$	$.613 \pm .053$	$.333 \pm .018$
	Accur	$.954 \pm .015$	$.919 \pm .050$	$.733 \pm .026$	$.727 \pm .039$	$.714 \pm .032$	$.887 \pm .011$
DRO-NL	Train	$.032 \pm .015$	$.113 \pm .035$	$.339 \pm .044$	$.381 \pm .084$	$.287 \pm .049$	$.195 \pm .034$
	Test	$.119 \pm .044$	$.194 \pm .067$	$.554 \pm .032$	$.576 \pm .049$	$.607 \pm .060$	$.332 \pm .015$
	Accur	$.955 \pm .016$	$.931 \pm .036$	$.736 \pm .027$	$.730 \pm .043$	$.716 \pm .054$	$.889 \pm .009$
Num Predictors		30	4	30	10	20	56
Train Size		40	20	80	30	30	150
Test Size		329	752	475	9,990	125,034	2,951

Table 1: Numerical results for real data sets.

We can observe from Table 1 that DRO-L performs better than the benchmark model LRL1 in terms of the test error, and the DRO-NL model further improves the average test error in all except Magic dataset.

7 CONCLUSION AND DISCUSSION

Our fully data-driven DRO procedure combines a semiparametric approach (i.e., the metric learning procedure) with a parametric procedure (expected loss minimization) to enhance the generalization performance of the underlying parametric model. We emphasize that our approach is applicable to any DRO formulation and is not restricted to classification tasks. An interesting research avenue that might be considered is the development of a semisupervised framework as in (Blanchet and Kang 2017b), in which unlabeled data is used to inform the support of the elements in $\mathcal{U}_\delta(P_n)$.

ACKNOWLEDGMENTS

We gratefully acknowledge support from the following NSF grants 1915967, 1820942, 1838676 as well as DARPA award N660011824028.

A PROOF OF THEOREM 2

Let us begin by listing the assumptions required to prove Theorem 2. First, we begin by recalling Assumption 1 from Section 5.

Assumption 1. There exists $\Gamma(\beta, y) \in (0, \infty)$ such that $l(u, y; \beta) \leq \Gamma(\beta, y) \cdot (1 + c(u, x))$, for all $(x, y) \in \mathcal{D}_n$,

We now introduce Assumptions 2-4 below.

Assumption 2. $\psi(\cdot, X, Y, \beta, \lambda)$ is twice continuously differentiable and the Hessian of $\psi(\cdot, X, Y, \beta, \lambda)$ evaluated at u^* , $D_u^2 \psi(u^*, X, Y, \beta, \lambda)$, is positive definite. In particular, we can find $\theta > 0$ and $\eta > 0$, such that

$$\psi(u, X, Y, \beta, \lambda) \geq \psi(u^*, X, Y, \beta, \lambda) - \frac{\theta}{2} \|u - u^*\|_2^2, \quad \forall u \text{ s.t. } \|u - u^*\|_2 \leq \eta.$$

Assumption 3. For a constant $\lambda_0 > 0$ such that $\phi(X, Y, \beta, \lambda_0) < \infty$, let $K = K(X, Y, \beta, \lambda_0)$ be any upper bound for $\phi(X, Y, \beta, \lambda_0)$.

Assumption 4. In addition to the lower semicontinuity of $c(\cdot) \geq 0$, we assume that $c(\cdot, X)$ is coercive in the sense that $c(u, X) \rightarrow \infty$ as $\|u\|_2 \rightarrow \infty$.

For any set S , the r -neighborhood of S is defined as the set of all points in \mathbb{R}^d that are at distance less than r from S , i.e. $S_r = \cup_{u \in S} \{\bar{u} : \|\bar{u} - u\|_2 \leq r\}$.

Proof of Theorem 2. The first part of the inequality is easy to derive. For the second part, we proceed as follows: Under Assumptions 3 and 4, we can define the compact set

$$\mathcal{C} = \mathcal{C}(X, Y, \beta, \lambda) = \{u : c(u, X) \leq l(X, Y; \beta) - K + \lambda_0 / (\lambda - \lambda_0)\}.$$

It is easy to check that $\arg \max \{\psi(u, X, Y, \lambda)\} \subset \mathcal{C}$. Owing to optimality of u^* and from Assumption 2 that $K \geq \phi(X, Y, \beta, \lambda_0)$, we can see that

$$l(X, Y) \leq l(u^*, Y) - \lambda c(u, X) = l(u^*, Y) - \lambda_0 c(u^*, X) - (\lambda - \lambda_0) c(u^*, X) \leq K - \lambda_0 - (\lambda - \lambda_0) c(u^*, X).$$

Thus by definition of $\mathcal{C} = \mathcal{C}(X, Y, \beta, \lambda)$, it follows easily that $u^* \in \mathcal{C}$, which further implies $\{u : \|u - u^*\|_2 \leq \eta\} \subset \mathcal{C}_\eta$. Then we combine the strongly convexity assumption in Assumption 2 and the definition of $\phi_{\varepsilon, f}(u, X, Y, \beta, \lambda)$, which yields

$$\begin{aligned} \phi_{\varepsilon, f}(X, Y, \beta, \lambda) &\geq \varepsilon \log \left(\int_{\|u - u^*\|_2 \leq \eta} \exp \left(\left[\phi(X, Y, \beta, \lambda) - \frac{\theta}{2} \|u - u^*\|_2^2 \right] / \varepsilon \right) f(u) du \right) \\ &= \varepsilon \log \left(\exp(\phi(X, Y, \beta, \lambda) / \varepsilon) \int_{\|u - u^*\|_2 \leq \eta} \exp \left(-\frac{\theta}{2} \|u - u^*\|_2^2 / \varepsilon \right) f(u) du \right) \\ &= \phi(X, Y, \beta, \lambda) + \varepsilon \log \int_{\|u - u^*\|_2 \leq \eta} \exp \left(-\frac{\theta \|u - u^*\|_2^2}{2\varepsilon} \right) f(u) du. \end{aligned}$$

As $\{u : \|u - u^*\|_2 \leq \eta\} \subset \mathcal{C}_\eta$, we can use the lower bound of $f(\cdot)$ to deduce that

$$\begin{aligned} \int_{\|u - u^*\|_2 \leq \eta} \exp \left(-\frac{\theta \|u - u^*\|_2^2}{2\varepsilon} \right) f(u) du &\geq \inf_{u \in \mathcal{C}_\eta} f(u) \times \int_{\|u - u^*\|_2 \leq \eta} \exp \left(-\frac{\theta \|u - u^*\|_2^2}{2\varepsilon} \right) du \\ &= \inf_{u \in \mathcal{C}_\eta} f(u) \times (2\pi\varepsilon/\theta)^{d/2} P(Z_d \leq \eta^2\theta/\varepsilon), \end{aligned}$$

where Z_d is a chi-squared random variable of d degrees of freedom. To conclude, recall that $\varepsilon \in (0, \eta^2\theta\chi_\alpha)$, the lower bound of $\phi_{\varepsilon, f}(\cdot)$ can be written as

$$\phi_{\varepsilon, f}(X, Y, \beta, \lambda) \geq \phi(X, Y, \beta, \lambda) - \frac{d}{2} \varepsilon \log(1/\varepsilon) + \frac{d}{2} \varepsilon \log \left((2\pi\alpha/\theta) \inf_{u \in \mathcal{C}_\eta} f(u) \right).$$

This completes the proof of Theorem 2. □

REFERENCES

- Bellet, A., A. Habrard, and M. Sebban. 2013. "A Survey on Metric Learning for Feature Vectors and Structured Data". *arXiv preprint arXiv:1306.6709*.
- Blanchet, J., and Y. Kang. 2017a. "Distributionally Robust Groupwise Regularization Estimator". *arXiv preprint arXiv:1705.04241*.
- Blanchet, J., and Y. Kang. 2017b. "Distributionally Robust Semi-supervised Learning". *arXiv preprint arXiv:1702.08848*.
- Blanchet, J., Y. Kang, and K. Murthy. 2016. "Robust Wasserstein Profile Inference and Applications to Machine Learning". *arXiv preprint arXiv:1610.05627*.
- Esfahani, P. M., and D. Kuhn. 2018. "Data-Driven Distributionally Robust Optimization Using the Wasserstein Metric: Performance Guarantees and Tractable Reformulations". *Mathematical Programming* 171(1-2):115–166.
- Friedman, J., T. Hastie, and R. Tibshirani. 2001. *The Elements of Statistical Learning*, Volume 1. Berlin: Springer.
- Ghosh, S., and H. Lam. 2019. "Robust Analysis in Stochastic Simulation: Computation and Performance Guarantees". *Operations Research* 67(1):232–249.
- Hoerl, A. E., and R. W. Kennard. 1970a. "Ridge Regression: Applications to Nonorthogonal Problems". *Technometrics* 12(1):69–82.
- Hoerl, A. E., and R. W. Kennard. 1970b. "Ridge Regression: Biased Estimation for Nonorthogonal Problems". *Technometrics* 12(1):55–67.
- Hu, Z., and L. J. Hong. "Kullback-Leibler Divergence Constrained Distributionally Robust Optimization". http://www.optimization-online.org/DB_FILE/2012/11/3677.pdf, accessed 3rd August.
- Ishwaran, H., and J. S. Rao. 2014. "Geometry and Properties of Generalized Ridge Regression in High Dimensions". *Contemporary Mathematics* 622:81–93.
- Lam, H., and E. Zhou. 2017. "The Empirical Likelihood Approach to Quantifying Uncertainty in Sample Average Approximation". *Operations Research Letters* 45(4):301–307.
- Lei, L., and M. I. Jordan. 2016. "Less than a Single Pass: Stochastically Controlled Stochastic Gradient Method". *arXiv preprint arXiv:1609.03261*.
- Li, L., C. Sun, L. Lin, J. Li, and S. Jiang. 2018. "A Mahalanobis Metric Learning-Based Polynomial Kernel for Classification of Hyperspectral Images". *Neural Computing and Applications* 29(4):1103–1113.
- Lichman, M. 2013. "UCI Machine Learning Repository". <http://archive.ics.uci.edu/ml>, accessed 3rd August.
- Schultz, M., and T. Joachims. 2003. "Learning a Distance Metric from Relative Comparisons". In *Proceedings of the 16th International Conference on Neural Information Processing Systems*, 41–48. Cambridge: MIT Press.
- Shafieezadeh-Abadeh, S., P. M. Esfahani, and D. Kuhn. 2015. "Distributionally Robust Logistic Regression". In *Proceedings of the 28th International Conference on Neural Information Processing Systems*, 1576–1584. Cambridge: MIT Press.
- Villani, C. 2008. *Optimal Transport: Old and New*, Volume 338. Berlin: Springer.
- Xing, E. P., A. Y. Ng, M. I. Jordan, and S. Russell. 2002. "Distance Metric Learning with Application to Clustering with Side Information". In *Proceedings of the 15th International Conference on Neural Information Processing Systems*, 521–528. Cambridge: MIT Press.
- Zou, H. 2006. "The Adaptive Lasso and its Oracle Properties". *Journal of the American Statistical Association* 101(476):1418–1429.

AUTHOR BIOGRAPHIES

JOSE BLANCHET is a faculty member in the Department of Management Science and Engineering at Stanford University. He holds a Ph.D. in Operations Research from Stanford University. Jose is a recipient of the 2009 Best Publication Award given by the INFORMS Applied Probability Society and of the 2010 Erlang Prize. He also received a PECASE award given by NSF in 2010. He has research interests in applied probability and Monte Carlo methods. His email is jose.blanchet@stanford.edu

YANG KANG got his Ph.D. in Statistics from Columbia University. He has research interests in applied probability and robust statistic inference. Now he is working in the financial industry. His email is yang.kang@columbia.edu. His website is <https://sites.google.com/view/yangkang1026/>.

KARTHYEK MURTHY is an Assistant Professor in the Engineering Systems and Design pillar of Singapore University of Technology and Design. His research interests lie in decision-making under uncertainty and broadly in applied probability. Specific topics of his research include include distributionally robust optimization and estimation of tail risks / rare event probabilities. He holds a PhD in Computer and System Sciences from Tata Institute of Fundamental Research, Mumbai. His email is karthyek_murthy@sutd.edu.sg.

FAN ZHANG is a Ph.D. candidate in the Department of Management Science and Engineering at Stanford University. His research interests include applied probability, stochastic simulation and robust optimization. His email is fzh@stanford.edu.