
Tikhonov Regularization is Optimal Transport Robust under Martingale Constraints

Jiajin Li

Sirui Lin

José Blanchet

Stanford University

{jiajinli, siruilin, jose.blanchet}@stanford.edu

Viet Anh Nguyen

Chinese University of Hong Kong

nguyen@se.cuhk.edu.hk

Abstract

Distributionally robust optimization has been shown to offer a principled way to regularize learning models. In this paper, we find that Tikhonov regularization is distributionally robust in an optimal transport sense (i.e., if an adversary chooses distributions in a suitable optimal transport neighborhood of the empirical measure), provided that suitable martingale constraints are also imposed. Further, we introduce a relaxation of the martingale constraints which not only provides a unified viewpoint to a class of existing robust methods but also leads to new regularization tools. To realize these novel tools, tractable computational algorithms are proposed. As a byproduct, the strong duality theorem proved in this paper can be potentially applied to other problems of independent interest.

1 Introduction

Regularization is an important tool in machine learning which is used in, for instance, reducing overfitting [23]. Recently, ideas from distributionally robust optimization (DRO) have led to a fresh viewpoint on regularization precisely in connection to overfitting; see, e.g., [7, 2, 20, 10, 29, 24, 25, 5] and the references therein.

In these references it is shown that many standard regularization-based estimators arise as the solution of a min-max game in which one wishes to minimize a loss over a class of parameters against an adversary that maximizes the out-of-distribution impact of any given parameter choice, that is, the adversary perturbs the empirical distribution in a certain way. The choice of adversarial distributions or perturbations in DRO is often non-parametric thus providing reassurance that the decision is reasonably robust to a wide range of out-of-distribution perturbations. For example, one such non-parametric choice is given by employing optimal transport costs [31] to construct a so-called distributional uncertainty set (e.g. a Wasserstein ball around the empirical distribution) for the adversary to choose. [27] shows that optimal transport-based DRO (OT-DRO) is closely related to adversarial robustness in the sense of steepest gradient loss contamination. This can be further explained by OT-DRO's hidden connection with generalized Lipschitz regularization [7]. Thus, understanding if a well-known regularization technique is actually distributionally robust and in what sense, allows us to understand its out-of-distribution benefits and potentially introduce improvements.

In this paper, we introduce a novel set of regularization techniques which incorporate martingale constraints into the OT-DRO framework. Our starting point is the conventional OT-DRO formulation. The conventional OT-DRO formulation can generally be interpreted as perturbing each data point in such a way that the average size perturbation is less than a given budget. In addition to this

conventional formulation, we will impose a *martingale constraint* in the joint distribution of the empirical data and the resulting adversarially perturbed data.

Why do we believe that the martingale constraint makes sense as a regularization technique? It turns out that two random variables X and \bar{X} form a martingale in the sense that $\mathbb{E}[\bar{X}|X] = X$ if and only if the distribution of \bar{X} dominates X in convex order [30]. In this sense, the adversary \bar{X} will have higher dispersion in non-parametric sense than the observed data X but in a suitably constrained way so that the average locations are preserved. This novel OT-DRO constrained regularization, we believe, is helpful to potentially combat conservative solutions, see [16]. Moreover, by allowing a small amount of violation in the martingale property, we can control the regularization properties of this constraint, thus obtaining a natural interpolation towards the conventional OT-DRO formulation and potentially improved regularization performance. We point out that related optimal transport problems with martingale constraints have been studied in robust mathematical finance [1, 8].

Consider, for example, the linear regression setting with the exact martingale constraints, which means that for any given observed data point, the conditional expectation of the additive perturbation under the worst-case joint distribution equals zero. Surprisingly, we show that the resulting martingale DRO model is *exactly* equivalent to the ridge regression [18] with the Tikhonov regularization. To the best of our knowledge, this paper is the first work to interpret the Tikhonov regularization from a DRO perspective showing that it is distributionally robust in a precise non-parametric sense. In stark contrast, it is well-known that the conventional OT-based DRO model (*without* the martingale constraint) is identical to the regularized square-root regression problem [2]. Therefore, introducing an additional power in norm regularization (i.e., converting square-root regression to Tikhonov regularization) can be translated into adding martingale constraints in the adversarial perturbations thus reducing the adversarial power. A natural question that arises here is whether we can interpolate between the conventional DRO model and the Tikhonov regularization, and further improve them.

We will provide a comprehensive and positive answer to the above question in this paper. The key idea here is to relax the equality constraint on the conditional expectation of the adversarial violation and thus allow a small perturbation of the martingale property to gain more flexibility of the uncertainty set. This idea leads to another novel model, termed the perturbed martingale DRO in the sequel. Intuitively, if the relaxation is sufficiently loose, the perturbed martingale DRO model will reduce to the conventional DRO, which is formally equivalent to setting an infinite amount of possible violations for the martingale constraint. By contrast, if no violation is allowed, the perturbed martingale DRO will automatically reduce to the exact counterpart — Tikhonov regularization. As a result, we are able to introduce a new class of regularizers via the interpolation between the conventional DRO model and the Tikhonov regularization.

Furthermore, such insightful interpolation also works for a broad class of nonlinear learning models. Inspired by our extensive exploration of linear regression, the developed martingale DRO model can also provide a new principled adversarial training procedure for deep neural networks. Extensive experiments are conducted to demonstrate the effectiveness of the proposed perturbed martingale DRO model for both linear regression and deep neural network training under the adversarial setting.

We summarize our main contributions as below:

- We reveal a new hidden connection in this paper, that is, Tikhonov regularization is optimal transport robust when exact martingale constraints (i.e., convex order between the adversary and empirical data) are imposed.
- Upon this finding, we develop a new *perturbed* martingale DRO model, which not only provides a unified viewpoint of existing regularization techniques, but also leads to a new class of robust regularizers.
- We introduce an easy-to-implement computational approach to capitalize the theoretical benefits in practice, in both linear regression and neural network training under the adversarial setting.
- As a byproduct, the strong duality theorem, which is proved in this paper and is used as the main technical tool, can be applied to a wider spectrum of problems of independent interest.

2 Preliminaries

Let us introduce some basic definitions and concepts preparing for the subsequent analysis.

Definition 2.1 (Optimal transport costs and the Wasserstein distance [21, 31]). *Suppose that $c(\cdot, \cdot) : \mathbb{R}^d \times \mathbb{R}^d \rightarrow [0, \infty]$ is a lower semi-continuous cost function such that $c(X, X) = 0$ for every $X \in \mathbb{R}^d$. The optimal transport cost between two distributions \mathbb{Q} and \mathbb{P} supported on \mathbb{R}^d is defined as*

$$D(\mathbb{Q}, \mathbb{P}) \triangleq \min_{\pi \in \mathcal{P}(\mathcal{X} \times \mathcal{X})} \{ \mathbb{E}_\pi [c(\bar{X}, X)] : P_1\pi = \mathbb{Q}, P_2\pi = \mathbb{P} \}.$$

Here, $\mathcal{P}(\mathcal{X} \times \mathcal{X})$ is the set of joint probability distribution π of (\bar{X}, X) supported on $\mathcal{X} \times \mathcal{X}$ while $P_1\pi$ and $P_2\pi$ respectively refer to the marginals of \bar{X} and X under the joint distribution π .

If $c(X, \bar{X}) = \|\bar{X} - X\|$ is any given norm on \mathbb{R}^d , then D recovers the Wasserstein distance [31]. In this paper, we are interested in a flexible family of functions for the computational tractability, so called the Mahalanobis cost functions in the form of $c(\bar{X}, X) = (X - \bar{X})^\top M (X - \bar{X})$, where M is a d -by- d positive definite matrix.

Next, we consider the conventional OT-DRO problem:

$$\min_{\beta} \mathbb{L}_\beta(\hat{\mathbb{P}}, \rho), \quad \text{where} \quad \mathbb{L}_\beta(\hat{\mathbb{P}}, \rho) \triangleq \begin{cases} \sup_{\pi} \mathbb{E}_\pi[\ell(f_\beta(\bar{X}))] \\ \text{s. t.} & \pi \in \mathcal{P}(\mathcal{X} \times \mathcal{X}) \\ & \mathbb{E}_\pi [c(\bar{X}, X)] \leq \rho, P_2\pi = \hat{\mathbb{P}}, \end{cases} \quad (2.1)$$

where $\hat{\mathbb{P}} \triangleq \frac{1}{N} \sum_{i=1}^N \delta_{X_i}$ is the empirical distribution. Using Definition 2.1, we have $\mathbb{L}_\beta(\hat{\mathbb{P}}, \rho) = \max_{\mathbb{Q}: D(\mathbb{Q}, \hat{\mathbb{P}}) \leq \rho} \mathbb{E}_{\mathbb{Q}}[\ell(f_\beta(\bar{X}))]$, which is the worst-case expected loss under all possible distributions around the empirical measure $\hat{\mathbb{P}}$ at most ρ with respect to the OT distance. It is well-known that under appropriate assumptions, the DRO problem (2.1) is equivalent to the regularized square-root regression problem.

Proposition 2.2 ([2, Proposition 2.]). *Suppose that (i) the loss function $\ell(\cdot)$ is a convex quadratic function, i.e., $\nabla^2 \ell(\cdot) = \gamma > 0$, where γ is a constant, (ii) the feature mapping $f_\beta(\bar{X}) = \beta^\top \bar{X}$ is linear, and (iii) the ground cost c is the squared Euclidean norm on $\mathcal{X} = \mathbb{R}^d$. Then*

$$\mathbb{L}_\beta(\hat{\mathbb{P}}, \rho) = \left(\sqrt{\mathbb{E}_{\hat{\mathbb{P}}}[\ell(f_\beta(X))]} + \sqrt{\rho} \|\beta\|_2 \right)^2.$$

We also present the strong duality result for a general class of optimal transport based DRO models with martingale constraints. This result serves as our main technical tool for reformulating the DRO models, and it can also be applied to other semi-infinity structured DRO models, which could be of independent interests. We consider the primal problem

$$\begin{aligned} & \sup_{\pi} \int_{\mathcal{X} \times \mathcal{X}} f(\bar{X}) d\pi \\ \text{s. t.} & \quad \pi \in \mathcal{P}(\mathcal{X} \times \mathcal{X}) \\ & \quad \int_{\mathcal{X} \times \mathcal{X}} c(\bar{X}, X) d\pi \leq \rho, \quad P_2\pi = \hat{\mathbb{P}} \\ & \quad \mathbb{E}_\pi[\bar{X}|X] = X \quad \hat{\mathbb{P}}\text{-a.s.} \end{aligned} \quad (\text{Primal})$$

and its associated dual form

$$\inf_{\substack{\lambda \in \mathbb{R}_+ \\ \alpha_i \in \mathbb{R}^d \forall i}} \lambda \rho + \sum_{i=1}^N \alpha_i^\top X_i + \frac{1}{N} \sum_{i=1}^N \sup_{\bar{X}} [f(\bar{X}) - \alpha_i^\top \bar{X} - \lambda c(\bar{X}, X_i)]. \quad (\text{Dual})$$

Here, $f : \mathcal{X} \rightarrow \mathbb{R}$ is upper semi-continuous and L^1 -integrable. The next theorem states the strong duality result linking these two problems.

Theorem 2.3 (Strong duality). *Let $\hat{\mathbb{P}} \triangleq \frac{1}{N} \sum_{i \in [N]} \delta_{X_i}$ be the reference measure. Suppose that (i) every sample point is in the interior of the cone generated by \mathcal{X} , i.e., $X_i \in \text{int}(\text{cone}(\mathcal{X})) \forall i \in [N]$, and (ii) the ambiguity radius $\rho > 0$. Then the strong duality holds, i.e., $\text{Val}(\text{Primal}) = \text{Val}(\text{Dual})$.*

Remark 2.4. *At the heart of our analysis tools is the abstract semi-infinite duality theory for conic linear program [26, Proposition 3.4]. Notably, it will be tricky and subtle to reformulate our problem into the standard form and further carefully check the general Slater condition. Moreover, we indeed fill the technical gap in [19, Theorem 4.2].*

Notation. We use \mathbb{S}_{++}^d to denote the set of d -by- d positive definite matrices and $\|X\|_M \triangleq \sqrt{X^\top M X}$ for any $X \in \mathbb{R}^d, M \in \mathbb{S}_{++}^d$; if M is the identity matrix, then we omit M and write $\|X\| \triangleq \sqrt{X^\top X}$. We use δ_X to denote the Dirac measure at X and let $\widehat{\mathbb{P}} \triangleq \frac{1}{N} \sum_{i=1}^N \delta_{X_i}$ be the empirical measure constructed from sample $\{X_1, \dots, X_N\}$. We use \mathbb{E}_P to denote the integration over P : $\mathbb{E}_P[f(X)] = \int_{\mathcal{X}} f(X) dP$. Specifically, for (\bar{X}, X) following joint distribution π , $\mathbb{E}_\pi[c(\bar{X}, X)] = \int_{\mathcal{X} \times \mathcal{X}} c(\bar{X}, X) d\pi$, $\mathbb{E}_\pi[f(\bar{X})] = \int_{\mathcal{X} \times \mathcal{X}} f(\bar{X}) d\pi$. We use $\ell(\cdot)$ to denote the loss function applied to the parametrized feature mapping f_β . Let $\nabla \ell(\cdot), \nabla^2 \ell(\cdot)$ be the first and second order derivative of $\ell(\cdot)$ respectively. Notably, we use $\mathbb{L}_\beta(\widehat{\mathbb{P}}, \rho)$ to denote the objective function of the conventional DRO model (2.1); we use $L_\beta(\widehat{\mathbb{P}}, \rho)$ to refer to the exact martingale DRO model (3.1); we use $\mathcal{L}_\beta(\widehat{\mathbb{P}}, \rho, \epsilon)$ to refer to the perturbed martingale DRO model (3.2).

3 Tractable Reformulations

In this section, we introduce an optimal transport-based DRO model with the exact martingale constraint at first. That is, on top of the vanilla DRO model [3], we add an additional martingale equality constraint on its coupling. It is surprisingly interesting to find out that the resulting DRO approach is equivalent to empirical risk minimization with Tikhonov regularization. Naturally, we can relax the equality constraint and thus allow a small violation of the martingale property to enrich the uncertainty set. Formally, by sending violation size to infinity in the martingale constraint, our relaxation allows to interpolate between the conventional DRO formulation (i.e. with no martingale constraints) and Tikhonov regularization (which involves exact martingale constraints). Therefore, this relaxation further leads to a new class of regularizers in a principled way, which improves upon Tikhonov regularization as we show in our experiments.

Assumption 3.1. *The following assumptions hold throughout.*

- (i) *The ground cost $c(\cdot, \cdot)$ is the Mahalanobis cost with the weighting matrix $M \in \mathbb{S}_{++}^d$,*
- (ii) *The domain \mathcal{X} is unconstrained, i.e., $\mathcal{X} = \mathbb{R}^d$.*

3.1 Optimal Transport-based DRO with Martingale Constraints

To start with, we investigate the exact martingale DRO problem:

$$\min_{\beta} L_{\beta}(\widehat{\mathbb{P}}, \rho), \quad \text{where} \quad L_{\beta}(\widehat{\mathbb{P}}, \rho) \triangleq \begin{cases} \sup_{\pi} \mathbb{E}_{\pi}[\ell(f_{\beta}(\bar{X}))] \\ \text{s. t.} & \pi \in \mathcal{P}(\mathcal{X} \times \mathcal{X}) \\ & \mathbb{E}_{\pi}[c(\bar{X}, X)] \leq \rho, P_2 \pi = \widehat{\mathbb{P}} \\ & \mathbb{E}_{\pi}[\bar{X}|X] = X \quad \widehat{\mathbb{P}}\text{-a.s.}, \end{cases} \quad (3.1)$$

and $\rho \geq 0$ is the radius of uncertainty set centered at $\widehat{\mathbb{P}}$. Note that because $\widehat{\mathbb{P}}$ is the empirical measure, the martingale constraint implies that the *conditional* expected value of the perturbation obtained by modifying each observed data point equals the observed data point itself. The quantity $L_{\beta}(\widehat{\mathbb{P}}, \rho)$ is referred to as the worst-case expected loss of the model parameter β under the martingale DRO model. It is easy to see that $L_{\beta}(\widehat{\mathbb{P}}, \rho) \leq \mathbb{L}_{\beta}(\widehat{\mathbb{P}}, \rho)$, where \mathbb{L}_{β} is defined as in (2.1). This is because the adversary in (3.1) has a smaller feasible set, and thus is less powerful than the adversary in (2.1). Hence, the martingale DRO solution for problem (3.1) is considered to be *less* conservative than the conventional DRO solution for problem (2.1).

The next result asserts that the martingale DRO problem coincides with the Tikhonov regularization problem under similar conditions of Proposition 2.2.

Proposition 3.2 (Tikhonov equivalence). *Suppose that (i) the loss function $\ell(\cdot)$ is a convex quadratic function, i.e., $\nabla^2 \ell(\cdot) = \gamma > 0$, and (ii) the feature mapping $f_{\beta}(\bar{X}) = \beta^\top \bar{X}$ is linear. Then we have*

$$L_{\beta}(\widehat{\mathbb{P}}, \rho) = \mathbb{E}_{\widehat{\mathbb{P}}}[\ell(\beta^\top X)] + \frac{\gamma \rho}{2} \|\beta\|_{M^{-1}}^2.$$

If the Mahalanobis matrix M is the identity matrix, the martingale DRO model (3.1) recovers the Tikhonov regularization problem.

Proof of Proposition 3.2. By a change of the variable, let $\Delta = \bar{X} - X$ and we have

$$\begin{aligned} \sup_{\substack{\mathbb{E}_\pi[\|\Delta\|_M^2] \leq \rho \\ \mathbb{E}_\pi[\Delta|X]=0}} \mathbb{E}_\pi [\ell(\beta^\top (X + \Delta))] &= \sup_{\substack{\mathbb{E}_\pi[\|\Delta\|_M^2] \leq \rho \\ \mathbb{E}_\pi[\Delta|X]=0}} \mathbb{E}_\pi \left[\ell(\beta^\top X) + \nabla \ell(\beta^\top X) \beta^\top \Delta + \frac{\gamma}{2} \|\beta^\top \Delta\|^2 \right] \\ &= \mathbb{E}_{\hat{\mathbb{P}}} [\ell(\beta^\top X)] + \sup_{\substack{\mathbb{E}_\pi[\|\Delta\|_M^2] \leq \rho \\ \mathbb{E}_\pi[\Delta|X]=0}} \mathbb{E}_\pi \left[\frac{\gamma}{2} \|\beta^\top \Delta\|^2 \right] \\ &= \mathbb{E}_{\hat{\mathbb{P}}} [\ell(\beta^\top X)] + \frac{\gamma \rho}{2} \|\beta\|_{M^{-1}}^2. \end{aligned}$$

The last equality follows from the general Hölder's inequality. To achieve the equality, we can, for example, take a normally distributed random variable C with mean 0 and variance ρ and which is independent of X , and then let $\Delta = CM^{-1}\beta$. \square

Example 3.3 (Linear regression). Let $X^\top \triangleq (Y, Z^\top) \in \mathbb{R}^d$ and $\beta^\top \triangleq (1, -b^\top) \in \mathbb{R}^d$, we have $\beta^\top X = Y - b^\top Z$. For any $Q \in \mathbb{S}_{++}^{d-1}$, we take $M = \text{diag}(+\infty, Q)$, which implies that we do not allow transport of the response Y , then the problem (3.1) with $\gamma = 2$ becomes

$$\min_b \left\{ \mathbb{E}_{\hat{\mathbb{P}}} [(Y - b^\top Z)^2] + \rho \|b\|_{Q^{-1}}^2 \right\}.$$

In Appendix C.1, we give another instructive proof based on the strong duality result in Section 2. For general convex loss functions, we have the following certificate of robustness that provides upper and lower bound for the worst-case DRO loss in (3.1).

Corollary 3.4 (General convex loss functions). Suppose that (i) the loss function $\ell(\cdot)$ is μ -strongly convex and C -smooth, that is, $\ell(\theta) \geq \ell(\theta') + \nabla \ell(\theta')(\theta - \theta') + \frac{\mu}{2}(\theta - \theta')^2$ and $\ell(\theta) \leq \ell(\theta') + \nabla \ell(\theta')(\theta - \theta') + \frac{C}{2}(\theta - \theta')^2$ hold for all $\theta, \theta' \in \mathbb{R}$, and (ii) the feature mapping $f_\beta(\bar{X}) = \beta^\top \bar{X}$ is linear. Then we have for any $\rho \geq 0$,

$$\mathbb{E}_{\hat{\mathbb{P}}}[\ell(f_\beta(X))] + \frac{\mu\rho}{2} \|\beta\|_{M^{-1}}^2 \leq L_\beta(\hat{\mathbb{P}}, \rho) \leq \mathbb{E}_{\hat{\mathbb{P}}}[\ell(f_\beta(X))] + \frac{C\rho}{2} \|\beta\|_{M^{-1}}^2.$$

Example 3.5 (Logistic regression). Let $X = YZ \in \mathbb{R}^d$, where $Z \in \mathbb{R}^d, Y \in \{\pm 1\}$, and $\ell(t) = \log(1 + \exp(-t))$, where $\ell(\cdot)$ satisfies Assumption (i) in Corollary 3.4 with $C = \frac{1}{4}$. Then we have

$$L_\beta(\hat{\mathbb{P}}, \rho) \leq \mathbb{E}_{\hat{\mathbb{P}}}[\log(1 + \exp(-Y\beta^\top Z))] + \frac{\rho}{8} \|\beta\|_{M^{-1}}^2.$$

3.2 Optimal Transport-based DRO with Perturbed Martingale Constraints

Now we turn to the relaxation of the martingale constraint to improve upon Tikhonov regularization and gain more flexibility. We consider the perturbed martingale coupling based DRO model (perturbed martingale DRO):

$$\min_{\beta} \mathcal{L}_\beta(\hat{\mathbb{P}}, \rho, \epsilon), \quad \text{where } \mathcal{L}_\beta(\hat{\mathbb{P}}, \rho, \epsilon) \triangleq \begin{cases} \sup_{\pi} \mathbb{E}_\pi[\ell(f_\beta(\bar{X}))] \\ \text{s. t. } \pi \in \mathcal{P}(\mathcal{X} \times \mathcal{X}) \\ \mathbb{E}_\pi [c(\bar{X}, X)] \leq \rho, P_2\pi = \hat{\mathbb{P}} \\ \|\mathbb{E}_\pi[\bar{X}|X] - X\|_M \leq \epsilon \quad \hat{\mathbb{P}}\text{-a.s.} \end{cases} \quad (3.2)$$

The parameter ϵ controls the allowed violations of the martingale constraint for the adversary. It is trivial that if we set $\epsilon = 0$, then we obtain $\mathcal{L}_\beta(\hat{\mathbb{P}}, \rho, 0) = L_\beta(\hat{\mathbb{P}}, \rho)$ and we recover the exact martingale DRO model (3.1). If we set $\epsilon = +\infty$ then the martingale constraint becomes ineffective, thus we have $\mathcal{L}_\beta(\hat{\mathbb{P}}, \rho, +\infty) = \mathbb{L}_\beta(\hat{\mathbb{P}}, \rho)$ and model (3.2) collapses to the DRO formulation (2.1). We thus can think of ϵ as an interpolating parameter connecting two extremes: the conventional DRO model (2.1) (at $\epsilon = +\infty$) and the exact martingale DRO model (3.1) (at $\epsilon = 0$).

However, the resulting optimization problem (3.2) constitutes an infinite-dimensional optimization problem over probability distributions and thus appears to be computationally intractable. To overcome this issue, we leverage Theorem 2.3 and prove that the problem (3.2) is actually equivalent to a finite-dimensional problem. To begin with, we present one crucial proposition that can be applied to more general settings.

Theorem 3.6 (General loss functions and feature mappings). *Suppose that the loss function $\ell(\cdot)$ and the feature mapping $f_\beta(\cdot)$ are upper semi-continuous. Then, for any $\rho > 0$ and $\epsilon > 0$, the perturbed martingale DRO model (3.2) admits*

$$\mathcal{L}_\beta(\widehat{\mathbb{P}}, \rho, \epsilon) = \inf_{\lambda \geq 0, \alpha} \lambda \rho + \frac{\epsilon}{N} \sum_{i=1}^N \|\alpha_i\|_{M^{-1}} + \frac{1}{N} \sum_{i=1}^N \sup_{\Delta_i} \left[\ell(f_\beta(X_i + \Delta_i)) - \alpha_i^\top \Delta_i - \lambda \|\Delta_i\|_M^2 \right]. \quad (3.3)$$

Sketch of proof. The key step is to decouple (3.2) as a two-layer optimization problem:

$$\begin{aligned} \mathcal{L}_\beta(\widehat{\mathbb{P}}, \rho, \epsilon) = & \sup_{\|\eta_i\|_M \leq \epsilon \forall i} \sup_{\pi} \int_{\mathcal{X}} \ell(f_\beta(\bar{X})) d\pi \\ \text{s.t. } & \pi \in \mathcal{P}(\mathcal{X} \times \mathcal{X}) \\ & \int_{\mathcal{X} \times \mathcal{X}} c(\bar{X}, X) d\pi \leq \rho, P_2 \pi = \widehat{\mathbb{P}} \\ & \int_{\mathcal{X} \times \mathcal{X}} \mathbb{I}_{X_i}(X) \cdot \bar{X} d\pi = \frac{1}{N}(X_i + \eta_i) \quad \forall i \in [N]. \end{aligned} \quad (3.4)$$

Then we invoke Theorem 2.3 for the inner maximization problem over π and apply the Sion's minimax theorem [28] for the outer maximization over η . The desired result is obtained. \square

Due to the specific structure of the quadratic cost and the linear feature mapping, the dual variable λ admits a closed-form representation and $\alpha \in \mathbb{R}^{N \times d}$ in (3.3) can be reduced to N parallel one-dimensional optimization problems. Upon this observation, we thus conduct an instructive and intuitive reformulation in Theorem 3.7 for linear regression with proof detailed in Appendix C.2.

Theorem 3.7 (Linear regression). *Suppose that (i) the loss function $\ell(\cdot)$ is a convex quadratic function, i.e., $\nabla^2 \ell(\cdot) = \gamma > 0$ and (ii) the feature mapping $f_\beta(\bar{X}) = \beta^\top \bar{X}$ is linear. Then, the perturbed martingale DRO model (3.2) admits:*

$$\mathcal{L}_\beta(\widehat{\mathbb{P}}, \rho, \epsilon) = \mathbb{E}_{\widehat{\mathbb{P}}}[\ell(f_\beta(X))] + \frac{\rho \|\beta\|_{M^{-1}}^2}{2} + \boxed{R(\beta)}, \quad (3.5)$$

where the additional regularizer is defined as $R(\beta) \triangleq \|\beta\|_{M^{-1}} \min_{s \in \mathbb{R}^N} \left(\frac{\epsilon}{N} \|s\|_1 + \sqrt{\frac{\rho}{N}} \|G_\beta - s\|_2 \right)$ and $G_\beta = (\nabla \ell(\beta^\top X_1), \dots, \nabla \ell(\beta^\top X_N))^\top \in \mathbb{R}^N$.

Obviously, compared with the exact martingale constraint investigated in Proposition 3.2, the perturbed constraint we consider here involves an additional term $R(\beta)$. Fortunately, armed with Lemma D.3, we are able to shed light on its intuitive interpretation based on the quantitative relationship between modeling parameters ϵ and ρ . First, we discuss two extreme cases — there are interesting hidden connections between (3.2) and other existing regularization techniques.

Remark 3.8. *Intuitively, if ϵ is relatively small, the optimal s^* is zero which means that all perturbed martingale constraints will be active. Precisely, if $\epsilon^2 \leq \rho$, the additional regularization term satisfies $R(\beta) = \epsilon \mathbb{E}_{\widehat{\mathbb{P}}}[\|\nabla_X \ell(f_\beta(X))\|_{M^{-1}}]$, which is so-called Jacobian or input gradient regularizer in the literature when M is an identity matrix. Recently, it has received much attention owing to its ability to improve adversarial robustness [32, 9]. Conversely, if ρ is relatively small, the optimal s^* is equal to G_β , implying that none of perturbed martingale constraints is active. In fact, if $\epsilon^2 \geq N\rho$, we have $R(\beta) = \sqrt{\rho \mathbb{E}_{\widehat{\mathbb{P}}}[\|\nabla_X \ell(f_\beta(X))\|_{M^{-1}}^2]}$. Then, we can easily figure out that $\mathcal{L}_\beta(\widehat{\mathbb{P}}, \rho, \epsilon)$ can be reduced to the conventional OT-based DRO model in Proposition 2.2.*

Then for the middle case, it is natural to infer that only part of constraints will be active. In Lemma D.3, we justify this conjecture rigorously. We refer the reader to Appendix D for the details. As such, the proposed martingale DRO model takes the first step bridging the input gradient regularization and regularized square-root regression problem in a unified framework. On the other hand, it also opens up an exciting brand new avenue of robustified regularizers. In the next section, we validate its effectiveness for the adversarial training task.

4 Optimization Algorithms

To take advantage of the proposed perturbed martingale DRO model, a natural question here is whether we can address problem (3.2) in a tractable manner. In this section, we answer the above question in the affirmative by developing two different computational paradigms for linear regression and deep neural network respectively.

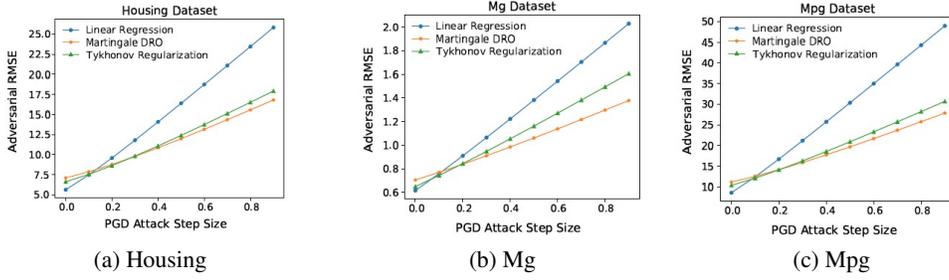


Figure 1: Compare the proposed martingale DRO with two standard benchmarks (linear regression and ridge regression) on three real world datasets — Housing, Mg and Mpg. The martingale DRO performs better than competing methods under large PGD step size.

4.1 Subgradient Method for Linear Regression

Unfortunately, the resulting formulation we conducted in Theorem 3.7 (i.e., (3.5)) for linear regression is potentially intractable for optimizing (β, s) jointly. One obvious computational challenge here is that the overall problem is not necessarily convex in (β, s) , even if the original problem is convex over β . This is because the minima of convex functions is not convex. However, under a mild assumption, it is possible to give a reliable computational routine for solving the resulting problem over β directly.

To start with, a key observation is $\mathcal{L}_\beta(\widehat{\mathbb{P}}, \rho, \epsilon)$ is convex over β . This is essentially from the fact that the pointwise supremum of a class of convex functions is still convex. To check the details, $\mathcal{L}_\beta(\widehat{\mathbb{P}}, \rho, \epsilon)$ is originally defined in (3.2) and $\mathbb{E}_\pi[\ell(f_\beta(\bar{X}))]$ is convex over β for all π . Thus, a natural yet simple algorithm is subgradient method. The main difficulty is to obtain the correct subgradient oracle. Since the first two terms are smooth and strongly convex, we know that $R(\beta)$ is weakly convex and thus subdifferentially regular. All the subdifferential concepts are coincide. As such, we may simply use the Clarke subdifferential [6] in the sequel. Moreover, the sum rule is hold for computing the subgradient of $\mathcal{L}_\beta(\widehat{\mathbb{P}}, \rho, \epsilon)$ due to the weakly convexity. We have

$$\partial_\beta \mathcal{L}_\beta(\widehat{\mathbb{P}}, \rho, \epsilon) = \nabla_\beta \mathbb{E}_{\widehat{\mathbb{P}}}[\ell(f_\beta(X))] + \nabla_\beta \frac{\rho \|\beta\|_{M-1}^2}{2} + \partial_\beta R(\beta).$$

The remaining question is how to compute the Clarke subdifferential of $R(\beta)$. When $\epsilon^2 \leq \rho$ and $\epsilon^2 \geq N\rho$, $R(\beta)$ will enjoy the convex composite structure. Thus, we can get the correct subgradient by invoking the chain rule developed in [22, Theorem 10.6] directly. The more subtle and tricky case is the middle one — $\rho < \epsilon^2 < N\rho$. Without of loss generality, we assume that $M = I$ for simplicity. Based on [6, Theorem 2.3.9], we have

$$\partial_\beta R(\beta) \subseteq \text{Conv} \left\{ \partial_\beta \left(\frac{\epsilon}{N} \sum_{i=1}^N \|s_i^* \beta\|_1 + \sqrt{\frac{\rho}{N} \sum_{i=1}^N \|(\nabla \ell(\beta^\top X_i) - s_i^*) \beta\|_2^2} \right), s^* \in \mathcal{S}(\beta) \right\},$$

where $\mathcal{S}(\beta)$ is the optimal solution set and $\text{Conv}\{\cdot\}$ denotes the convex hull. If we assume the inclusion here is tight, then the vanilla subgradient method will converge to the optimal solution with the rate $\mathcal{O}(1/\sqrt{K})$ [4]. Empirically, we find out that the resulting subgradient method works well and the violated case will never happen.

4.2 A New Principled Adversarial Training Procedure for Deep Learning

In this subsection, we develop inexact stochastic gradient-type methods for (3.3) and thus are able to realized the benefits of the proposed DRO model in adversarial learning tasks. As the nonconvexity of $f_\beta(\cdot)$, the inner maximization problem over Δ_i will be no longer tractable. Therefore, we leverage the methodology proposed in [27] to gain the computational efficiency, that is, regarding the dual variable as a modeling parameter. To proceed, we make the same smoothness assumption in [27, Assumption B] (i.e., see Assumption E.1 in appendix for details).

Lemma 4.1 (Convex-concave minimax theorem without compactness). *Suppose that $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is convex and level bounded and $g : \mathbb{R}^m \rightarrow \mathbb{R}$ is strongly convex. Then we have*

$$\min_{x \in \mathbb{R}^n} \max_{y \in \mathbb{R}^m} f(x) + x^\top Ay - g(y) = \max_{y \in \mathbb{R}^m} \min_{x \in \mathbb{R}^n} f(x) + x^\top Ay - g(y).$$

Leveraging Lemma 4.1 and the smoothness assumption, (3.3) leads to a simple and instructive form:

$$\min_{\beta} \frac{1}{N} \sum_{i=1}^N \max_{\|\Delta_i\|_M \leq \epsilon} [\ell(f_{\beta}(X_i + \Delta_i)) - \lambda \|\Delta_i\|_M^2]. \quad (4.1)$$

In contrast to the vanilla DRO model studied in [27], (4.1) further constrains the perturbation into a Euclidean ball with the correlation information M . Moreover, we can observe that the magnitude of ϵ decides how many martingale constraints will be active, which also perfectly matches our theoretical results and interpretations established for linear regression, see Theorem 3.7 and Remark 3.8.

From a computational viewpoint, if λ is large enough (see Lemma E.2 for details), the inner maximization problem is strongly concave and thus the outer minimization problem over β will be smooth. This motivates Algorithm 1, an inexact stochastic gradient method for Problem (4.1). The convergence guarantee is provided in [27, Theorem 2]. It is worthwhile mentioning that the resulting new principled adversarial training is extremely easy to implement by only adding three lines of *Pytorch* code based on [27]. We refer the interested readers to Appendix F for details.

Algorithm 1: Martingale Distributionally Robust Optimization with Adversarial Training

Input: Sampling distribution $\widehat{\mathbb{P}}$, stepsize sequence $\{t_k\}_{k=0}^{K-1}$;
for $k = 0, 1, 2, \dots, K - 1$ **do**

Sample $X^k \sim \widehat{\mathbb{P}}$ and find an η -approximate maximizer $\hat{\Delta}^k$ satisfying

$$\|\hat{\Delta}^k - \Delta_k^*\| \leq \eta, \quad \text{where} \quad \Delta_k^* = \arg \max_{\|\Delta\|_M \leq \epsilon} \{\ell(f_{\beta^k}(X^k + \Delta)) - \lambda \|\Delta\|_M^2\}.$$

Set $\beta^{k+1} \leftarrow \beta^k - t_k \nabla_{\beta} \ell(f_{\beta^k}(X^k + \hat{\Delta}^k))$.

5 Numerical Results

In this section, we validate the effectiveness of our methods (referred to as *martingale DRO*) on both linear regression and deep neural networks under the adversarial setting. All simulations are implemented using Python 3.8 on: (1) a computer running Windows 10 with a 2.80GHz, Intel(R) Core(TM) i7-1165G7 processor and 16 GB of RAM, and (2) Google Colab with NVIDIA Tesla P100 GPU and 16 GB of RAM. As for the adversarial setting, we consider three types of attack, the detailed definitions of which are collected in Appendix F.

5.1 Linear Regression

To start with, we demonstrate the effectiveness of the proposed martingale DRO model (3.5) with the quadratic loss function and linear feature mapping, i.e., $\ell(f_{\beta}(X)) = \frac{1}{2}(Y - b^{\top}Z)^2$ with $X^{\top} \triangleq (Y, Z^{\top})$ and $\beta^{\top} \triangleq (1, -b^{\top})$, where Z is the feature vector and Y is the target variable. In this experiment, we test our method on three LIBSVM regression real world datasets¹. More specifically, we randomly select 60% of the data to train the models and the rest as our test data. To showcase the effectiveness of martingale DRO model under adversarial setting, we apply one-step projected-gradient method (PGM) attack [17] on test data and report the performance in terms of the the root-mean-square error (RMSE) on adversarial test data, where $\text{RMSE} \triangleq \sqrt{\frac{1}{N} \sum_i (\hat{\beta}^{\top} x_{\text{adv}}^{(i)} - y_{\text{adv}}^{(i)})^2}$ and $\hat{\beta}$ is the estimator of β . All numerical results with different step sizes for PGM attack are collected in Figure 1. As we mentioned in Proposition 3.2, since the exact martingale DRO model (i.e., $\epsilon = 0$) is equivalent to Tikhonov regularization, we choose the same hyperparameter $\rho = 0.08$ for ridge regression and martingale DRO mode for fair comparison. We can observe that the martingale DRO can outperform other two benchmarks over three real world datasets consistently when the step size for the PGD attack is relatively large. This result also corroborates our theoretical intuition — the additional regularization $R(\beta)$ can further improve the adversarial robustness.

¹<https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/regression.html>

5.2 Deep Neural Network for Adversarial Training

We generate the synthetic training data $\{(Y_i, Z_i)\}_{i \in \mathcal{I}}$ with a wide margin as follows: generate i.i.d. $Z_j \sim N(0, I_2)$, where $Z \in \mathbb{R}^2$, I_2 is the identity matrix in \mathbb{R}^2 ; set $\mathcal{I} = \{j : \|Z_j\|_2 \notin (\sqrt{2}/\eta, \eta\sqrt{2})\}$, where $\eta = 1.6$; let $Y_i = \text{sign}(\|Z_i\|_2 - \sqrt{2})$, $\forall i \in \mathcal{I}$. We train a neural network with 3 hidden layers of size 4, 3 and 2 and ELU activations between layers. We compare our approach (cf. martingale DRO) with ERM and the conventional DRO approach developed in [27]. More details about the experiment setup are collected in Appendix F.

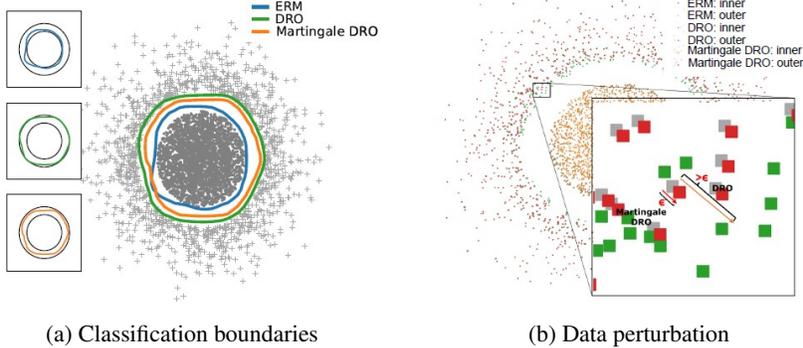


Figure 2: Illustration of the performance comparison between competing methods. ERM tends to overfit to the inner class while DRO becomes too conservative due to unconstrained perturbation. The martingale DRO leaves bigger margins from the sample points than the other methods.

Figure 2 show the experimental results on the synthetic dataset. Test data are shown in darkgray and gray with different shapes, which are generated by the above-mentioned procedure with a smaller margin ($\eta = 1.2$). Classification boundaries are shown in blue, green, and orange for ERM, DRO, and martingale DRO respectively, as well as with the true class boundaries of the test data. Intuitively, the boundary generated by ERM is too close to the true inner boundary since the majority of points are of darkgray class, while the DRO approach pushes the classification boundary outwards. However, as illustrated in Figure 2(a), the DRO approach suffers from over-conservativeness and becomes entangled with the boundary of the outer gray class. In contrast, our martingale DRO boundary lands in between the two extremes and it leaves bigger margin. Figure 2(b) explicitly shows the qualitative difference between these two methods in terms of the perturbation to the data: the Martingale perturbation is constrained below ϵ while the DRO perturbation is unconstrained. Moreover, previously shown, decreasing non-zero ϵ pushes the perturbed martingale constraints towards the exact martingale constraints and forces the classification boundary increasingly inward. More results are in Appendix F.

Then, we validate our method on the MNIST dataset [15]. For the classifier, we train a neural network equipped with 8×8 , 6×6 , 5×5 convolutional filter layers and ELU activations followed by a fully connected layer and softmax output. To show the robustness of our method, we test the performance of four methods (ERM, DRO, Jacobian regularization [13] and martingale DRO) under the PGD and FGSM attacks (Definition F.1) with test error defined to be: $1 - \text{classification accuracy}$.

In Figures 3a and 3b, our martingale DRO model outperforms the other methods and still provides robustness under the ∞ -norm FGSM attacks. In Figure 3c, we show the performance of our model with different ϵ . As expected, when ϵ is relatively small, the model is not flexible enough and shows large test error. Alternatively, as ϵ becomes large, our model will behave similarly to the original DRO model since the ϵ -constraint is almost inactive in this case.

Figure 4 visualizes the different levels of robustness for the four methods. For each test data point, we perturb the image using the DRO attacks (Definition F.2) with decaying level of perturbation and respectively record the first perturbed images that each model correctly classifies. In Figure 4, the original label is 6 and all methods output the correct prediction, whereas in the adversarial example that the DRO model predicts 6, the correct classification seems unreasonable to human eyes (see Appendix F for more examples). This observation shows an insight that the original DRO model is too conservative in predicting and our model puts more constraints on the perturbation when training thus providing a model that is more consistent to human eyes.

PGD Attack	ERM	DRO	Jacobian Regularization	Martingale DRO
$\epsilon = 0$	84.16%	84.02%	81.73%	85.48%
$\epsilon = 0.04$	77.50%	82.87%	78.78%	83.25%
$\epsilon = 0.08$	70.20%	80.68%	73.85%	80.86%

Table 1: Top-1 accuracy results with different levels of perturbation on CIFAR-10.

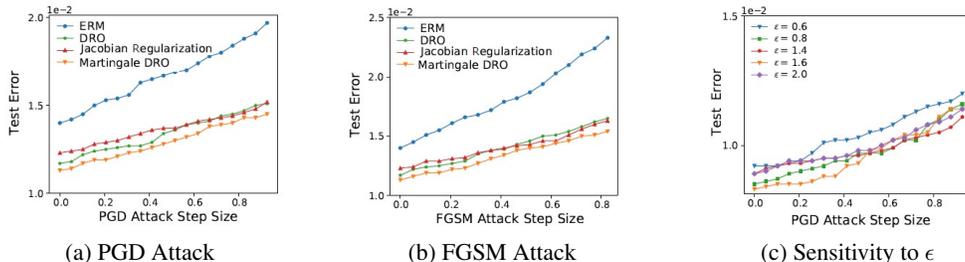


Figure 3: Compare the proposed martingale DRO with ERM and DRO on the MNIST datasets under PGD and FGSM attack; compare the proposed martingale DRO with different values of ϵ .

Experimental setup for CIFAR-10 [14]: For the classifier, we train a ResNet with the architecture in [12]. We optimize using Adam with a batch size of 128 for all methods. The learning rate starts from 0.01 and shrinks by $0.1 \frac{\text{epoch}}{\text{total epochs}}$, and each model is trained for 100 epochs. The simulations are implemented using Python 3.8 on Google Colab with TPU v2 and 16GB RAM. Similarly, we test the performance of four methods (ERM, DRO, Jacobian regularization and martingale DRO) under the PGD attack with different levels of perturbation, the results shown in Table 1 are consistent with those from the MNIST dataset.

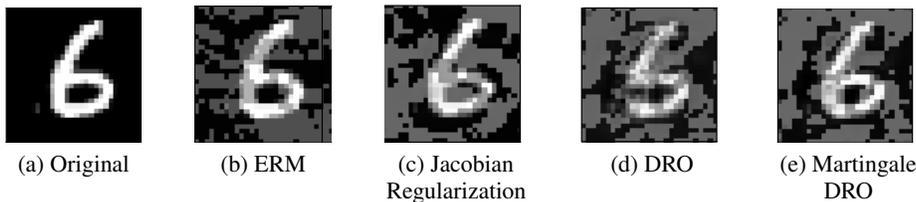


Figure 4: The largest DRO perturbation such that each model makes correct prediction.

6 Closing Remarks

In this paper, we find that the OT-based DRO model is equivalent to Tikhonov regularization when exact martingale constraints are imposed. Upon this interesting hidden connection, we introduce a new model called the perturbed martingale DRO model, which not only provides a unified viewpoint to several common robust methods but also leads to new regularization tools. Empirically, we validate the effectiveness of our model in addressing the conservativeness issue for the conventional DRO model. From the statistical perspective, how to optimally select the size of uncertainty regions and the perturbation size of the martingale constraint, is a natural problem to be further explored.

Acknowledgements Material in this paper is based upon work supported by the Air Force Office of Scientific Research under award number FA9550-20-1-0397. Additional support is gratefully acknowledged from NSF grants 1915967 and 2118199. Viet Anh Nguyen acknowledges the support from the CUHK’s Improvement on Competitiveness in Hiring New Faculties Funding Scheme.

References

- [1] Mathias Beiglböck, Pierre Henry-Labordere, and Friedrich Penkner. Model-independent bounds for option prices—a mass transport approach. *Finance and Stochastics*, 17(3):477–501, 2013.
- [2] Jose Blanchet, Yang Kang, and Karthyek Murthy. Robust Wasserstein profile inference and applications to machine learning. *Journal of Applied Probability*, 56(3):830–857, 2019.
- [3] Jose Blanchet, Karthyek Murthy, and Fan Zhang. Optimal transport-based distributionally robust optimization: Structural properties and iterative schemes. *Mathematics of Operations Research*, 2021.
- [4] Stephen Boyd, Lin Xiao, and Almir Mutapcic. Subgradient methods. *lecture notes of EE392o, Stanford University, Autumn Quarter, 2004:2004–2005*, 2003.
- [5] Ruidi Chen and Ioannis C Paschalidis. A robust learning approach for regression models based on distributionally robust optimization. *Journal of Machine Learning Research*, 19(13), 2018.
- [6] Frank H Clarke. *Optimization and nonsmooth analysis*. SIAM, 1990.
- [7] Zac Cranko, Zhan Shi, Xinhua Zhang, Richard Nock, and Simon Kornblith. Generalised Lipschitz regularisation equals distributional robustness. In *International Conference on Machine Learning*, pages 2178–2188. PMLR, 2021.
- [8] Yan Dolinsky and H Mete Soner. Martingale optimal transport and robust hedging in continuous time. *Probability Theory and Related Fields*, 160(1):391–427, 2014.
- [9] Chris Finlay and Adam M Oberman. Scaleable input gradient regularization for adversarial robustness. *arXiv preprint arXiv:1905.11468*, 2019.
- [10] Rui Gao, Xi Chen, and Anton J Kleywegt. Wasserstein distributionally robust optimization and variation regularization. *arXiv preprint arXiv:1712.06050*, 2017.
- [11] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [13] Judy Hoffman, Daniel A Roberts, and Sho Yaida. Robust learning with jacobian regularization. *arXiv preprint arXiv:1908.02729*, 2019.
- [14] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- [15] Yann LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.
- [16] Jiashuo Liu, Zheyang Shen, Peng Cui, Linjun Zhou, Kun Kuang, and Bo Li. Distributionally robust learning with stable adversarial training. *arXiv preprint arXiv:2106.15791*, 2021.
- [17] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
- [18] Gary C McDonald. Ridge regression. *Wiley Interdisciplinary Reviews: Computational Statistics*, 1(1):93–100, 2009.
- [19] Peyman Mohajerin Esfahani and Daniel Kuhn. Data-driven distributionally robust optimization using the wasserstein metric: Performance guarantees and tractable reformulations. *Mathematical Programming*, 171(1):115–166, 2018.
- [20] Hongseok Namkoong and John C Duchi. Variance-based regularization with convex objectives. *Advances in Neural Information Processing Systems*, 30, 2017.

- [21] Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- [22] R Tyrrell Rockafellar and Roger J-B Wets. *Variational Analysis*, volume 317. Springer Science & Business Media, 2009.
- [23] Stephan R Sain. *The Nature of Statistical Learning Theory*. Taylor & Francis, 1996.
- [24] Soroosh Shafieezadeh-Abadeh, Daniel Kuhn, and Peyman Mohajerin Esfahani. Regularization via mass transportation. *Journal of Machine Learning Research*, 20(103):1–68, 2019.
- [25] Soroosh Shafieezadeh Abadeh, Peyman M Mohajerin Esfahani, and Daniel Kuhn. Distributionally robust logistic regression. *Advances in Neural Information Processing Systems*, 28, 2015.
- [26] Alexander Shapiro. On duality theory of conic linear problems. In *Semi-infinite programming*, pages 135–165. Springer, 2001.
- [27] Aman Sinha, Hongseok Namkoong, and John Duchi. Certifying some distributional robustness with principled adversarial training. In *International Conference on Learning Representations*, 2018.
- [28] Maurice Sion. On general minimax theorems. *Pacific Journal of Mathematics*, 8(1):171–176, 1958.
- [29] Matthew Staib and Stefanie Jegelka. Distributionally robust optimization and generalization in kernel methods. *Advances in Neural Information Processing Systems*, 32, 2019.
- [30] Volker Strassen. The existence of probability measures with given marginals. *The Annals of Mathematical Statistics*, 36(2):423–439, 1965.
- [31] Cédric Villani. *Optimal Transport: Old and New*, volume 338. Springer, 2009.
- [32] Wei Emma Zhang, Quan Z Sheng, Ahoud Alhazmi, and Chenliang Li. Adversarial attacks on deep-learning models in natural language processing: A survey. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 11(3):1–41, 2020.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes]
 - (c) Did you discuss any potential negative societal impacts of your work? [N/A]
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes] See section 3
 - (b) Did you include complete proofs of all theoretical results? [Yes] See Appendix
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [Yes]
 - (b) Did you mention the license of the assets? [Yes]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [No]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

Broader impact

This work does not present any foreseeable societal consequence. While our contribution has a theoretical orientation, we believe that the structure of our method holds significant promise in the adversarial learning and robust optimization as we mentioned in the body context.

A Organization of the Appendix

We organize the appendix as follows:

- The proof details of Theorem 2.3 (cf. **Strong Duality Result**) is given in Section B.
- We collect all proof details of tractable reformulation results in Section C, including Proposition 3.2, Theorem 3.6 and Theorem 3.7.
- All useful technical lemmas are summarized in D.
- Convergence analysis of optimization algorithms are provided in Section E.
- Supplementary materials of numerical results are included in Section F.

B Strong Duality Result

To obtain the tractable reformulation result, we start to prove the strong duality theorem for a general class of martingale constraints-based Wasserstein DRO optimization problems (i.e., in the main context, we just provide the simplified version for simplicity):

$$\begin{aligned}
 & \sup_{\mathbb{Q}, \pi} \int_{\mathcal{X}} f(\bar{X}) d\mathbb{Q} \\
 & \text{s. t. } \mathbb{Q} \in \mathcal{P}(\mathcal{X}), \pi \in \mathcal{P}(\mathcal{X} \times \mathbb{X}) \\
 & P_1 \pi = \mathbb{Q}, P_2 \pi = \hat{\mathbb{P}} \\
 & \int_{\mathcal{X} \times \mathbb{X}} c(\bar{X}, X) d\pi \leq \rho \\
 & \mathbb{E}_\pi[\tilde{X}|X] = \tilde{X} \quad \hat{\mathbb{P}}\text{-a.s.}
 \end{aligned} \tag{Primal}$$

Here,

- $f : \mathcal{X} \rightarrow \mathbb{R}$ is assumed to be upper semi-continuous and $\hat{\mathbb{P}}$ -integrable i.e., $f \in L^1(\hat{\mathbb{P}})$.
- $\mathcal{P}(\mathcal{X})$ denotes the set of all Borel probability measures supported on \mathcal{X} .
- The cost function $c : \mathcal{X} \times \mathcal{X} \rightarrow [0, \infty]$ is a lower semicontinuous function satisfying $c(X, X) = 0$ for every $X \in \mathcal{X}$.
- $P_1 \pi$ and $P_2 \pi$ refer to the first and second marginal probability measure of π , that is, $(P_1 \pi)(S) = \pi(S \times \mathcal{X})$ and $(P_2 \pi)(S) = \pi(\mathcal{X} \times S)$ for any Borel subset S of \mathcal{X} .
- For simplicity, let the reference measure be the empirical distribution $\hat{\mathbb{P}} = \frac{1}{N} \sum_{i=1}^N \delta_{X_i}$ and $\mathbb{X} \triangleq \{X_1, X_2, \dots, X_N\} \subset \mathcal{X}$.
- \tilde{X} is the perturbed discrete distribution based on the empirical distribution $\hat{\mathbb{P}}$ supported on $\{X_1 + \eta_1, \dots, X_N + \eta_N\}$, i.e., $\tilde{\mathbb{P}} = \frac{1}{N} \sum_{i=1}^N \delta_{X_i + \eta_i}$.

The Lagrangian dual problem is derived as

$$\min_{\substack{\lambda \in \mathbb{R}_+ \\ \alpha_i \in \mathbb{R}^d \forall i}} \lambda \rho + \sum_{i=1}^N \alpha_i^\top (X_i + \eta_i) + \frac{1}{N} \sum_{i=1}^N \max_{\bar{X}} [f(\bar{X}) - \alpha_i^\top \bar{X} - \lambda c(\bar{X}, X_i)]. \tag{Dual}$$

Theorem B.1 (Restate Theorem 2.3 in a more general fashion). *Suppose that (i) the reference measure $\hat{\mathbb{P}}$ is the empirical distribution, i.e., $\hat{\mathbb{P}} = \frac{1}{N} \sum_{i \in [N]} \delta_{X_i}$, (ii) \tilde{X} follows from the perturbed empirical distribution, i.e., $\tilde{\mathbb{P}} = \frac{1}{N} \sum_{i \in [N]} \delta_{X_i + \eta_i}$ satisfying $X_i + \eta_i \in \text{int}(\text{cone}(\mathcal{X})) \forall i \in [N]$, and (iii) the ambiguity radius satisfies $\rho > 0$. Then strong duality holds, i.e., $\text{Val}(\text{Primal}) = \text{Val}(\text{Dual})$.*

Proof of Theorem B.1. Since $\mathbb{Q} = P_1\pi$ a change of variables allows us to rewrite the objective function as

$$\int_{\mathcal{X} \times \mathbb{X}} f(\bar{X}) d\pi.$$

Then, as the reference measure $\widehat{\mathbb{P}} = \frac{1}{N} \sum_{i=1}^N \delta_{X_i}$, we can recast the marginal constraint $P_2\pi = \widehat{\mathbb{P}}$ as

$$\int_{\mathcal{X} \times \mathbb{X}} \mathbb{I}_{\mathcal{X} \times \{X_i\}}(\bar{X}, X) d\pi = \frac{1}{N} \quad \forall i \in [N],$$

where $\mathbb{I}_{\mathcal{S}}$ is the indicator function of the set \mathcal{S} . Similarly, we can also reformulate the martingale constraint via further exploiting the discrete structure of the reference measure $\widehat{\mathbb{P}}$:

$$\begin{aligned} \int_{\mathcal{X} \times \mathbb{X}} \bar{X} \mathbb{I}_{X_i}(X) d\pi &= \frac{1}{N} \sum_{i=1}^N \int_{\mathcal{X}} \bar{X} \mathbb{I}_{X_i}(X_i) \mathbb{Q}^i(d\bar{X}) \\ &= \frac{1}{N} \int_{\mathcal{X}} \mathbb{Q}^i(d\bar{X}) = \frac{1}{N} (X_i + \eta_i). \end{aligned}$$

Thus, we have

$$\int_{\mathcal{X} \times \mathbb{X}} \mathbb{I}_{X_i}(X) \cdot \bar{X} d\pi = \frac{1}{N} (X_i + \eta_i) \quad \forall i \in [N],$$

where $X_i + \eta_i \in \mathcal{X}$. If we make the normalization of π explicit, we obtain the following equivalent reformulation of Problem (Primal):

$$\begin{aligned} \sup_{\pi \in M_+(\mathcal{X} \times \mathbb{X})} \quad & \int_{\mathcal{X} \times \mathbb{X}} f(\bar{X}) d\pi. \\ \text{s. t.} \quad & \int_{\mathcal{X} \times \mathbb{X}} \mathbb{I}_{\mathcal{X} \times \{X_i\}}(\bar{X}, X) d\pi = \frac{1}{N} \quad \forall i \in [N] \\ & \int_{\mathcal{X} \times \mathbb{X}} \mathbb{I}_{X_i}(X) \cdot \bar{X} d\pi = \frac{1}{N} (X_i + \eta_i) \quad \forall i \in [N] \\ & \int_{\mathcal{X} \times \mathbb{X}} c(\bar{X}, X) d\pi \leq \rho. \end{aligned} \tag{B.1}$$

Here, $M_+(\mathcal{X} \times \mathbb{X})$ is the set of all non-negative Borel measures supported on $\mathcal{X} \times \mathbb{X}$ and the first integral constraint ensures that π is a probability measure. As $M_+(\mathcal{X} \times \mathbb{X})$ is a convex cone and all of constraints regarding π are linear, problem (B.1) can be fitted into the standard primal problem in [26, (3.2)]. That is,

$$\begin{aligned} \min_{\pi \in M_+(\mathcal{X} \times \mathbb{X})} \quad & \langle f, \pi \rangle \\ \text{s. t.} \quad & \mathcal{A}(\pi) - b \in K, \end{aligned} \tag{B.2}$$

where

$$K = \{0\}^{N+Nd} \times \mathbb{R}_{\leq 0}, \quad b = \left(\frac{1}{N} \mathbf{e}_N, X_1 + \eta_1, \dots, X_N + \eta_N, \rho \right),$$

and \mathcal{A} is the linear mapping defined through the left hand side of the constraints in (B.1):

$$\mathcal{A} : \pi \mapsto \begin{bmatrix} \int_{\mathcal{X} \times \mathbb{X}} \mathbb{I}_{\mathcal{X} \times \{X_1\}}(\bar{X}, X) d\pi \\ \vdots \\ \int_{\mathcal{X} \times \mathbb{X}} \mathbb{I}_{\mathcal{X} \times \{X_N\}}(\bar{X}, X) d\pi \\ \int_{\mathcal{X} \times \mathbb{X}} \mathbb{I}_{X_1}(X) \cdot \bar{X} d\pi \\ \vdots \\ \int_{\mathcal{X} \times \mathbb{X}} \mathbb{I}_{X_N}(X) \cdot \bar{X} d\pi \\ \int_{\mathcal{X} \times \mathbb{X}} c(\bar{X}, X) d\pi \end{bmatrix}$$

Next, we aim at invoking Proposition 3.4 in [26] to prove the strong duality. A sufficient condition is the generalized Slater condition, see (3.12) in [26]. That is, we have to check

$$b \in \text{int}[\mathcal{A}(M_+(\mathcal{X} \times \mathbb{X})) - K], \tag{B.3}$$

where $\text{int}(\cdot)$ is the interior of a set. As such,

$$\mathcal{A}(M_+(\mathcal{X} \times \mathbb{X})) = [0, +\infty]^N \times \text{Range}(F)^N \times [0, \infty],$$

where $F : M_+(\mathcal{X} \times \mathbb{X}) \rightarrow \mathbb{R}^d$ satisfying $F(\pi) = \int_{\mathcal{X} \times \mathbb{X}} \mathbb{I}_{X_i}(X) \cdot \bar{X} d\pi$. Then,

$$\mathcal{A}(M_+(\mathcal{X} \times \mathbb{X})) - K = [0, +\infty]^N \times \text{Range}(F)^N \times [0, \infty].$$

To check the Slater condition, we validate each constraint separately.

- $\frac{1}{N} \in \text{int}([0, +\infty])$, for all $i \in [N]$;
- As $X_i + \eta_i \in \text{int}(\text{cone}(\mathcal{X}))$, $\forall i \in [N]$ and Lemma D.1 holds (i.e., $\text{cone}(\mathcal{X}) \subseteq \text{Range}(F)$), then $\frac{1}{N}(X_i + \eta_i) \in \text{int}(\text{Range}(F))$, $\forall i \in [N]$.
- Due to $\rho > 0$, we have $\rho \in \text{int}([0, +\infty])$.

Then, we obtained the desirable result. At last, we derive the dual problem by the standard Lagrangian method following [26],

$$\begin{aligned} \mathcal{L}(\pi; \lambda, s, \alpha) = & \lambda\rho + \frac{1}{N} \sum_{i=1}^N s_i + \frac{1}{N} \sum_{i=1}^N \alpha_i^\top (X_i + \eta_i) \\ & + \int_{\mathcal{X} \times \mathbb{X}} \left[f(\bar{X}) - \sum_{i=1}^N \mathbb{I}_{X_i}(X) \cdot \alpha_i^\top \bar{X} - \lambda c(\bar{X}, X) - \sum_{i=1}^N s_i \mathbb{I}_{\mathcal{X} \times \{X_i\}}(\bar{X}, X) \right] d\pi. \end{aligned}$$

Due to the strong duality result, we have

$$\sup_{\pi \in M_+(\mathcal{X} \times \mathbb{X})} \min_{\lambda \geq 0, s, \alpha} \mathcal{L}(\pi; \lambda, s, \alpha) = \min_{\lambda \geq 0, s, \alpha} \sup_{\pi \in M_+(\mathcal{X} \times \mathbb{X})} \mathcal{L}(\pi; \lambda, s, \alpha).$$

Moreover, since $\mathbb{X} = \{X_1, X_2, \dots, X_N\}$, the nonnegative measure $\pi \in M_+(\mathcal{X} \times \mathbb{X})$ can be decomposed as $\pi(\bar{X}, X) = \sum_{i=1}^N w_i \mathbb{I}_{X_i}(X) \mathbb{Q}^i(\bar{X})$ where $w_i \geq 0$, $\forall i \in [N]$. Then, we have

$$\begin{aligned} & \min_{\lambda \geq 0, s, \alpha} \sup_{\pi \in M_+(\mathcal{X} \times \mathbb{X})} \mathcal{L}(\pi; \lambda, s, \alpha) \\ = & \min_{\lambda \geq 0, s, \alpha} \sup_{w_i \geq 0} \lambda\rho + \frac{1}{N} \sum_{i=1}^N s_i + \frac{1}{N} \sum_{i=1}^N \alpha_i^\top (X_i + \eta_i) \\ & + \sum_{i=1}^N w_i \max_{\bar{X}} [f(\bar{X}) - \alpha_i^\top \bar{X} - \lambda c(\bar{X}, X_i) - s_i] \\ = & \min_{\lambda \geq 0, s, \alpha} \sup_{w_i \geq 0} \lambda\rho + \sum_{i=1}^N \left(\frac{1}{N} - w_i \right) s_i + \frac{1}{N} \sum_{i=1}^N \alpha_i^\top (X_i + \eta_i) \\ & + \sum_{i=1}^N w_i \max_{\bar{X}} [f(\bar{X}) - \alpha_i^\top \bar{X} - \lambda c(\bar{X}, X_i)] \\ = & \min_{\lambda \geq 0, \alpha} \lambda\rho + \frac{1}{N} \sum_{i=1}^N \alpha_i^\top (X_i + \eta_i) + \frac{1}{N} \sum_{i=1}^N \max_{\bar{X}} [f(\bar{X}) - \alpha_i^\top \bar{X} - \lambda c(\bar{X}, X_i)]. \end{aligned} \tag{B.4}$$

We complete the proof. □

C Proof Details of Tractable Reformulation Results

C.1 Proof of Proposition 3.2

Proof of Proposition 3.2. Problem (3.1) can be recast into

$$\begin{aligned} & \sup \int_{\mathcal{X}} \ell(\beta^\top \bar{X}) d\mathbb{Q} \\ \text{s. t. } & \mathbb{Q} \in \mathcal{P}(\mathcal{X}), \pi \in \mathcal{P}(\mathcal{X} \times \mathcal{X}) \\ & P_1 \pi = \mathbb{Q}, P_2 \pi = \hat{\mathbb{P}} \\ & \int_{\mathcal{X} \times \mathbb{X}} c(X, \bar{X}) d\pi \leq \rho \\ & \int_{\mathcal{X} \times \mathbb{X}} \mathbb{I}_{X_i}(X) \cdot \bar{X} d\pi = \frac{1}{N} X_i \quad \forall i \in [N]. \end{aligned} \tag{C.1}$$

Because $\hat{\mathbb{P}}$ is the empirical measure and because any feasible measure π satisfy the constraint $P_2 \pi = \hat{\mathbb{P}}$, the integral in the last two constraints of (C.1) is restricted to $\mathcal{X} \times \mathbb{X}$ (instead of $\mathcal{X} \times \mathcal{X}$) without any loss of optimality.

The trivial case $\rho = 0$ is easy to verify. To begin with, we focus on the case where $\rho > 0$. Here, we want to invoke the strong result, i.e., Theorem 2.3. Before getting into details, we check the conditions at first. $\ell(\cdot)$ is quadratic and thus upper semi-continuous; $\mathcal{X} = \mathbb{R}^d$ can help us to get rid of the mild regularity condition, that is, $X_i \in \text{int}(\mathbb{R}^d)$ automatically holds. Then, we get

$$L_\beta(\widehat{\mathbb{P}}, \rho) = \min_{\lambda \geq 0, \alpha} \lambda \rho + \frac{1}{N} \sum_{i=1}^N \alpha_i^\top X_i + \frac{1}{N} \sum_{i=1}^N \max_{\bar{X}} [\ell(\beta^\top \bar{X}) - \alpha_i^\top \bar{X} - \lambda c(\bar{X}, X_i)].$$

By a change of the variables, i.e., $\Delta_i = \bar{X} - X_i \forall i \in [N]$, we have

$$\begin{aligned} & L_\beta(\widehat{\mathbb{P}}, \rho) \\ &= \min_{\lambda \geq 0, \alpha} \lambda \rho + \frac{1}{N} \sum_{i=1}^N \max_{\Delta_i} [\ell(\beta^\top (X_i + \Delta_i)) - \alpha_i^\top \Delta_i - \lambda \|\Delta_i\|_M^2] \\ &= \min_{\lambda \geq 0, \alpha} \lambda \rho + \frac{1}{N} \sum_{i=1}^N \max_{\Delta_i} \left[\ell(\beta^\top X_i) + \nabla \ell(\beta^\top X_i) \beta^\top \Delta_i + \frac{\gamma}{2} \|\beta^\top \Delta_i\|^2 - \alpha_i^\top \Delta_i - \lambda \|\Delta_i\|_M^2 \right] \\ &= \frac{1}{N} \sum_{i=1}^N \ell(\beta^\top X_i) + \min_{\lambda \geq 0, \alpha} \lambda \rho + \frac{1}{N} \sum_{i=1}^N \max_{\Delta_i} \left[(\nabla \ell(\beta^\top X_i) \beta - \alpha_i)^\top \Delta_i + \frac{\gamma}{2} \|\beta^\top \Delta_i\|^2 - \lambda \|\Delta_i\|_M^2 \right]. \end{aligned}$$

Thus, the crux is the inner maximization problem. To proceed, we exhaust all possible cases. When $\lambda < \|\beta\|_{M^{-1}}^2 \gamma / 2$, it is easy to check that the inner maximization problem will go to $+\infty$ due to the general Cauchy-Schwarz inequality for the normed space $\|\beta^\top \Delta_i\|^2 \leq \|\beta\|_{M^{-1}}^2 \|\Delta_i\|_M^2$. When $\lambda = \|\beta\|_{M^{-1}}^2 \gamma / 2$ and $\alpha_i \neq \nabla \ell(\beta^\top X_i) \beta$, the inner maximization problem will also go to $+\infty$. As such, we have

$$\begin{aligned} L_\beta(\widehat{\mathbb{P}}, \rho) &\leq \frac{1}{N} \sum_{i=1}^N \ell(\beta^\top X_i) + \min_{\lambda \geq 0, \alpha} \lambda \rho + \frac{1}{N} \sum_{i=1}^N \max_{\Delta_i} \left[\frac{\gamma}{2} \|\beta^\top \Delta_i\|_M^2 - \lambda \|\Delta_i\|_M^2 \right] \\ &= \frac{1}{N} \sum_{i=1}^N \ell(\beta^\top X_i) + \min_{\lambda \geq 0, \alpha} \lambda \rho + \frac{1}{N} \sum_{i=1}^N \max_{\Delta_i} \left[\frac{\gamma}{2} \|\beta\|_{M^{-1}}^2 \|\Delta_i\|_M^2 - \lambda \|\Delta_i\|_M^2 \right] \\ &= \frac{1}{N} \sum_{i=1}^N \ell(\beta^\top X_i) + \frac{\gamma \rho}{2} \|\beta\|_{M^{-1}}^2, \end{aligned}$$

if $\lambda = \|\beta\|_{M^{-1}}^2 \gamma / 2$ and $\alpha_i = \nabla \ell(\beta^\top X_i) \beta$. At last, we focus on the left case and further prove the above inequality is the equality. If $\lambda > \|\beta\|_{M^{-1}}^2 \gamma / 2$, we have

$$\lambda \rho + \frac{1}{N} \sum_{i=1}^N \max_{\Delta_i} \left[(\nabla \ell(\beta^\top X_i) \beta - \alpha_i)^\top \Delta_i + \frac{\gamma}{2} \|\beta^\top \Delta_i\|^2 - \lambda \|\Delta_i\|_M^2 \right] > \|\beta\|_{M^{-1}}^2 \gamma / 2 + 0.$$

The desirable result is obtained, that is,

$$L_\beta(\widehat{\mathbb{P}}, \rho) = \frac{1}{N} \sum_{i=1}^N \ell(\beta^\top X_i) + \frac{\gamma \rho}{2} \|\beta\|_{M^{-1}}^2 = \mathbb{E}_{\widehat{\mathbb{P}}}[\ell(\beta^\top X)] + \frac{\gamma \rho}{2} \|\beta\|_{M^{-1}}^2.$$

This completes the proof. \square

C.2 Proof of Theorem 3.6 and Theorem 3.7

Proof of Theorem 3.6. To start with, we recast problem (3.2) into a two-layer optimization problem:

$$\begin{aligned} \mathcal{L}_\beta(\widehat{\mathbb{P}}, \rho, \epsilon) &= \sup_{\|\eta_i\|_M \leq \epsilon \forall i} \sup_{\pi} \int_{\mathcal{X}} \ell(f_\beta(\bar{X})) d\mathbb{Q} \\ &\quad \text{s.t. } \mathbb{Q} \in \mathcal{P}(\mathcal{X}), \pi \in \mathcal{P}(\mathcal{X} \times \mathcal{X}), \\ &\quad P_1 \pi = \mathbb{Q}, P_2 \pi = \widehat{\mathbb{P}} \\ &\quad \int_{\mathcal{X} \times \mathbb{X}} c(\bar{X}, X) d\pi \leq \rho \\ &\quad \int_{\mathcal{X} \times \mathbb{X}} \mathbb{I}_{X_i}(X) \cdot \bar{X} d\pi = \frac{1}{N} (X_i + \eta_i) \quad \forall i \in [N], \end{aligned} \tag{C.2}$$

where $\mathbb{X} = \{X_1, X_2, \dots, X_N\}$. Then, we apply Theorem 2.3 (i.e., strong duality) to the inner maximization problem, i.e.,

$$\begin{aligned} & \mathcal{L}_\beta(\widehat{\mathbb{P}}, \rho, \epsilon) \\ &= \sup_{\|\eta_i\|_M \leq \epsilon} \inf_{\lambda \geq 0, \alpha} \lambda \rho + \frac{1}{N} \sum_{i=1}^N \alpha_i^\top (X_i + \eta_i) + \frac{1}{N} \sum_{i=1}^N \sup_{\bar{X}_i} [\ell(f_\beta(\bar{X}_i)) - \alpha_i^\top \bar{X}_i - \lambda c(\bar{X}_i, X_i)] \\ &= \sup_{\|\eta_i\|_M \leq \epsilon} \inf_{\lambda \geq 0, \alpha} \lambda \rho + \frac{1}{N} \sum_{i=1}^N \alpha_i^\top \eta_i + \frac{1}{N} \sum_{i=1}^N \sup_{\Delta_i} [\ell(f_\beta(X_i + \Delta_i)) - \alpha_i^\top \Delta_i - \lambda \|\Delta_i\|_M^2]. \end{aligned}$$

The second equality follows by setting $\bar{X}_i = X_i + \Delta_i$. As $0 < \epsilon < +\infty$ and M is a positive definite matrix, the set $\{(\eta_1, \dots, \eta_N) : \|\eta_i\|_M \leq \epsilon \forall i \in [N]\}$ is a compact set. Consider the following mapping

$$(\eta, \lambda, \alpha) \mapsto \lambda \rho + \frac{1}{N} \sum_{i=1}^N \alpha_i^\top \eta_i + \frac{1}{N} \sum_{i=1}^N \sup_{\Delta_i} [\ell(f_\beta(X_i + \Delta_i)) - \alpha_i^\top \Delta_i - \lambda \|\Delta_i\|_M^2].$$

It is easy to see that this mapping is linear, and thus concave, in η . Moreover, it is convex in (λ, α) as the pointwise supremum of a class of convex functions (i.e., the inner function over (λ, α) is linear) is always convex. From Sion's minimax theorem [28], we can interchange the outer supremum and infimum operators to obtain

$$\begin{aligned} & \mathcal{L}_\beta(\widehat{\mathbb{P}}, \rho, \epsilon) \\ &= \inf_{\lambda \geq 0, \alpha} \sup_{\|\eta_i\|_M \leq \epsilon \forall i} \lambda \rho + \frac{1}{N} \sum_{i=1}^N \alpha_i^\top \eta_i + \frac{1}{N} \sum_{i=1}^N \sup_{\Delta_i} [\ell(f_\beta(X_i + \Delta_i)) - \alpha_i^\top \Delta_i - \lambda \|\Delta_i\|_M^2]. \end{aligned}$$

For any feasible value of (λ, α) , the optimal solution in η_i is either

$$\eta_i^* = M^{-1} \alpha_i \quad \text{or} \quad \eta_i^* = -M^{-1} \alpha_i.$$

We thus have

$$\mathcal{L}_\beta(\widehat{\mathbb{P}}, \rho, \epsilon) = \inf_{\lambda \geq 0, \alpha} \lambda \rho + \frac{\epsilon}{N} \sum_{i=1}^N \|\alpha_i\|_{M^{-1}} + \frac{1}{N} \sum_{i=1}^N \sup_{\Delta_i} [\ell(f_\beta(X_i + \Delta_i)) - \alpha_i^\top \Delta_i - \lambda \|\Delta_i\|_M^2].$$

We complete the proof. \square

Proof of Theorem 3.7. Taking $\ell(f_\beta(X)) = \ell(\beta^\top X)$ with the second derivative of $\nabla^2 \ell(\cdot) = \gamma$ in (3.3), we have

$$\begin{aligned} \mathcal{L}_\beta(\widehat{\mathbb{P}}, \rho, \epsilon) &= \inf_{\lambda \geq 0, \alpha} \lambda \rho + \frac{\epsilon}{N} \sum_{i=1}^N \|\alpha_i\|_{M^{-1}} + \frac{1}{N} \sum_{i=1}^N \sup_{\Delta_i} [\ell(\beta^\top (X_i + \Delta_i)) - \alpha_i^\top \Delta_i - \lambda \|\Delta_i\|_M^2] \\ &= \inf_{\lambda \geq 0, \alpha} \lambda \rho + \frac{\epsilon}{N} \sum_{i=1}^N \|\alpha_i\|_{M^{-1}} + \\ &\quad \frac{1}{N} \sum_{i=1}^N \sup_{\Delta_i} \left[\ell(\beta^\top X_i) + \nabla \ell(\beta^\top X_i) \beta^\top \Delta_i + \frac{\gamma}{2} \|\beta^\top \Delta_i\|^2 - \alpha_i^\top \Delta_i - \lambda \|\Delta_i\|_M^2 \right] \\ &= \frac{1}{N} \sum_{i=1}^N \ell(\beta^\top X_i) + \inf_{\lambda \geq 0, \alpha} \lambda \rho + \frac{\epsilon}{N} \sum_{i=1}^N \|\alpha_i\|_{M^{-1}} + \\ &\quad \frac{1}{N} \sum_{i=1}^N \max_{\Delta_i} \left[(\nabla \ell(\beta^\top X_i) \beta - \alpha_i)^\top \Delta_i + \frac{\gamma}{2} \|\beta^\top \Delta_i\|^2 - \lambda \|\Delta_i\|_M^2 \right] \\ &= \frac{1}{N} \sum_{i=1}^N \ell(\beta^\top X_i) + \inf_{\lambda \geq 0, \alpha} \lambda \rho + \frac{\epsilon}{N} \sum_{i=1}^N \|\alpha_i\|_{M^{-1}} + \end{aligned}$$

$$\frac{1}{N} \sum_{i=1}^N \sup_{\Delta_i} \left[(\nabla \ell(\beta^\top X_i) \beta - \alpha_i)^\top \Delta_i + \frac{\gamma}{2} \|\beta^\top \Delta_i\|^2 - \lambda \|\Delta_i\|_M^2 \right].$$

Similar with the argument to proof Proposition 3.2 in the appendix, see section ?? for details, we can conclude that $0 \leq \lambda \leq \frac{\gamma}{2} \|\beta\|_{M^{-1}}^2$. As such, we analyze two cases separately.

Case 1: suppose that the optimal value of $\lambda^* = \frac{\gamma}{2} \|\beta\|_{M^{-1}}^2$. As we discussed the exact martingale DRO mode in the last subsection, we have $\alpha_i^* = \nabla \ell(\beta^\top X_i) \beta$ and

$$\mathcal{L}_1^*(\hat{\mathbb{P}}, \rho, \epsilon) = \frac{1}{N} \sum_{i=1}^N \ell(\beta^\top X_i) + \frac{\rho\gamma}{2} \|\beta\|_{M^{-1}}^2 + \frac{\epsilon}{N} \sum_{i=1}^N \|\nabla \ell(\beta^\top X_i) \beta\|_{M^{-1}}. \quad (\text{C.3})$$

Case 2: suppose that the optimal value of $\lambda^* > \frac{\gamma}{2} \|\beta\|_{M^{-1}}^2$. For any fixed $i = 1, \dots, N$. Define

$$F(\lambda, \alpha) = \max_{\rho} \left[(\nabla \ell(\beta^\top X_i) \beta - \alpha_i)^\top \Delta_i + \frac{\gamma}{2} \|\beta^\top \Delta_i\|^2 - \lambda \|\Delta_i\|_M^2 \right].$$

As $\lambda^* > \frac{\gamma}{2} \|\beta\|_{M^{-1}}^2$, the inner maximization with respect to Δ_i is strongly convex. Consequently, it is necessary and sufficient to study its first-order optimality condition:

$$(\nabla \ell(\beta^\top X_i) \beta - \alpha_i) + (\gamma \beta \beta^\top - 2\lambda M) \Delta_i = 0. \quad (\text{C.4})$$

Then, we obtain the optimal solution and the optimal value,

$$\Delta_i^* = (2\lambda M - \gamma \beta \beta^\top)^{-1} (\nabla \ell(\beta^\top X_i) \beta - \alpha_i), \quad (\text{C.5})$$

where the matrix inversion is valid as $\lambda^* > \frac{\gamma}{2} \|\beta\|_{M^{-1}}^2$ and

$$\begin{aligned} F(\lambda, \alpha) &= \lambda \|\Delta_i^*\|_M^2 - \frac{\gamma}{2} \|\beta^\top \Delta_i^*\|^2 \\ &= (\Delta_i^*)^\top \left(\lambda M - \frac{\gamma}{2} \beta \beta^\top \right)^{-1} \Delta_i^* \\ &= \frac{1}{4} (\nabla \ell(\beta^\top X_i) \beta - \alpha_i)^\top \left(\lambda M - \frac{\gamma}{2} \beta \beta^\top \right)^{-1} (\nabla \ell(\beta^\top X_i) \beta - \alpha_i). \end{aligned}$$

For simplicity, let us ignore the empirical loss at first, which is the constant w.r.t. the dual variables λ and α .

$$\begin{aligned} &\min_{\lambda, \alpha} \lambda \rho + \frac{\epsilon}{N} \sum_{i=1}^N \|\alpha_i\|_{M^{-1}} + \frac{1}{N} \sum_{i=1}^N \max_{\Delta_i} \left[(\nabla \ell(\beta^\top X_i) \beta - \alpha_i)^\top \Delta_i + \frac{\gamma}{2} \|\beta^\top \Delta_i\|^2 - \lambda \|\Delta_i\|_M^2 \right] \\ &= \min_{\lambda > \frac{\gamma}{2} \|\beta\|_{M^{-1}}^2, \alpha} \lambda \rho + \frac{\epsilon}{N} \sum_{i=1}^N \|\alpha_i\|_{M^{-1}} + \\ &\quad \frac{1}{4N} \sum_{i=1}^N (\nabla \ell(\beta^\top X_i) \beta - \alpha_i)^\top \left(\lambda M - \frac{\gamma}{2} \beta \beta^\top \right)^{-1} (\nabla \ell(\beta^\top X_i) \beta - \alpha_i). \end{aligned}$$

The resulting structure of (λ, α) is still quite complicated. To further characterize the structure of the optimal solution, we utilize the parallel structure of α and focus on the corresponding subproblem as follow:

$$\min_{\alpha_i} \epsilon \|\alpha_i\| + \frac{1}{4} (\nabla \ell(\beta^\top X_i) \beta - \alpha_i)^\top \left(\lambda I - \frac{\gamma}{2} \beta \beta^\top \right)^{-1} (\nabla \ell(\beta^\top X_i) \beta - \alpha_i). \quad (\text{C.6})$$

By the Sherman–Morrison Formula (i.e., see Fact D.2), we have

$$\begin{aligned} \left(\lambda M - \frac{\gamma}{2} \beta \beta^\top \right)^{-1} &= (\lambda M)^{-1} + \frac{(\lambda M)^{-1} (\frac{\gamma}{2} \beta \beta^\top) (\lambda M)^{-1}}{1 - \frac{\gamma}{2} \beta^\top (\lambda M)^{-1} \beta} \\ &= \frac{1}{\lambda} M^{-1} + \frac{\gamma M^{-1} \beta \beta^\top M^{-1}}{\lambda (2\lambda - \gamma \|\beta\|_{M^{-1}}^2)}. \end{aligned}$$

Together with (C.6),

$$\min_{\alpha_i} \epsilon \|\alpha_i\|_{M^{-1}} + \frac{1}{4} (\nabla \ell(\beta^\top X_i) \beta - \alpha_i)^\top \left(\frac{1}{\lambda} M^{-1} + \frac{\gamma M^{-1} \beta \beta^\top M^{-1}}{\lambda(2\lambda - \gamma \|\beta\|_{M^{-1}}^2)} \right) (\nabla \ell(\beta^\top X_i) \beta - \alpha_i).$$

Similarly, as the minimization problem w.r.t. β is strongly convex, it is sufficient to study its first-order optimality condition. WLOG, we can assume the optimal solution $\alpha_i \neq 0$ to get rid of the non-smooth point. Then, we have

$$\begin{aligned} 0 &= \frac{\epsilon M^{-1} \alpha_i}{\|\alpha_i\|_{M^{-1}}} + \frac{1}{2} \left(\frac{1}{\lambda} M^{-1} + \frac{\gamma M^{-1} \beta \beta^\top M^{-1}}{\lambda(2\lambda - \gamma \|\beta\|_{M^{-1}}^2)} \right) (\nabla \ell(\beta^\top X_i) \beta - \alpha_i) \\ &= \left(\frac{\epsilon}{\|\alpha_i\|_{M^{-1}}} - \frac{1}{2\lambda} \right) M^{-1} \alpha_i + \left(\frac{\nabla \ell(\beta^\top X_i)}{2\lambda} + \frac{\gamma (\nabla \ell(\beta^\top X_i) \|\beta\|_{M^{-1}}^2 - \beta^\top M^{-1} \alpha_i)}{\lambda(2\lambda - \gamma \|\beta\|_{M^{-1}}^2)} \right) M^{-1} \beta. \end{aligned}$$

It is easy to observe that the optimal solution α_i is parallel to β . The conclusion is also valid for the corner case $\alpha_i = 0$. Consequently, problem (C.6) can be reduced to a one-dimensional problem, i.e., $\alpha_i = s_i \beta$,

$$\min_{s_i} \epsilon \|\beta\|_{M^{-1}} |s_i| + \frac{1}{4} \left(\frac{\|\beta\|_{M^{-1}}^2}{\lambda} + \frac{\gamma \|\beta\|_{M^{-1}}^4}{\lambda(2\lambda - \gamma \|\beta\|_{M^{-1}}^2)} \right) (\nabla \ell(\beta^\top X_i) - s_i)^2. \quad (\text{C.7})$$

Putting all pieces together, we get

$$\begin{aligned} & \min_{\lambda > \frac{\gamma}{2}, s} \lambda \rho + \frac{\epsilon \|\beta\|_{M^{-1}}}{N} \|s\|_1 + \frac{1}{4N} \left(\frac{\|\beta\|_{M^{-1}}^2}{\lambda} + \frac{\gamma \|\beta\|_{M^{-1}}^4}{\lambda(2\lambda - \gamma \|\beta\|_{M^{-1}}^2)} \right) \sum_{i=1}^N (\nabla \ell(\beta^\top X_i) - s_i)^2 \\ &= \min_{\lambda > \frac{\gamma}{2}, s} \lambda \rho + \frac{\epsilon \|\beta\|_{M^{-1}}}{N} \|s\|_1 + \frac{1}{4N} \left(\frac{\|\beta\|_{M^{-1}}^2}{\lambda} + \frac{\gamma \|\beta\|_{M^{-1}}^4}{\lambda(2\lambda - \gamma \|\beta\|_{M^{-1}}^2)} \right) \|G_\beta - s\|_2^2 \\ &= \min_{\lambda > \gamma/2, s} \lambda \rho \|\beta\|_{M^{-1}}^2 + \frac{\epsilon \|\beta\|_{M^{-1}}}{N} \|s\|_1 + \frac{1}{4N} \left(\frac{1}{\lambda} + \frac{\gamma}{\lambda(2\lambda - \gamma)} \right) \|G_\beta - s\|_2^2, \\ &= \min_{\lambda > \gamma/2, s} \lambda \rho \|\beta\|^2 + \frac{\epsilon \|\beta\|}{N} \|s\|_1 + \frac{2}{4N(2\lambda - \gamma)} \|G_\beta - s\|_2^2, \end{aligned}$$

where $G_\beta = (\nabla \ell(\beta^\top X_1), \dots, \nabla \ell(\beta^\top X_N))$. By changing the variables $\theta = \frac{2}{2\lambda - \gamma}$ and $\lambda = \frac{1}{\theta} + \frac{\gamma}{2}$ where $\theta > 0$,

$$\min_s \frac{\gamma \rho}{2} \|\beta\|_{M^{-1}}^2 + \frac{\epsilon}{N} \|\beta\|_{M^{-1}} \|s\|_1 + \min_{\theta > 0} \frac{\rho}{\theta} \|\beta\|_{M^{-1}}^2 + \frac{\theta}{4N} \|G_\beta - s\|_2^2.$$

Here, the optimal solution θ^* is

$$\theta^* = \frac{2\sqrt{N\rho} \|\beta\|_{M^{-1}}}{\|G_\beta - s\|}.$$

Consequently, we have

$$\mathcal{L}_2^*(\widehat{\mathbb{P}}, \rho, \epsilon) = \mathbb{E}_{\widehat{\mathbb{P}}}[\ell(\beta^\top X)] + \frac{\gamma \rho}{2} \|\beta\|_{M^{-1}}^2 + \|\beta\|_{M^{-1}} \min_s \left(\frac{\epsilon}{N} \|s\|_1 + \sqrt{\frac{\rho}{N}} \|G_\beta - s\|_2 \right). \quad (\text{C.8})$$

By applying Lemma D.3, we have

$$\min_s \left(\frac{\epsilon}{N} \|s\|_1 + \sqrt{\frac{\rho}{N}} \|G_\beta - s\|_2 \right) = \frac{\epsilon}{N} \|G_\beta\|_1$$

when $\epsilon \leq \sqrt{\rho}$. Then, combining these two cases, we can obtain,

$$\mathcal{L}_\beta(\widehat{\mathbb{P}}, \rho, \epsilon) = \mathbb{E}_{\widehat{\mathbb{P}}}[\ell(\beta^\top X)] + \frac{\gamma \rho}{2} \|\beta\|_{M^{-1}}^2 + \|\beta\|_{M^{-1}} \min_s \left(\frac{\epsilon}{N} \|s\|_1 + \sqrt{\frac{\rho}{N}} \|G_\beta - s\|_2 \right).$$

We complete the proof. \square

D Useful Technical Lemmas

Lemma D.1. *Suppose that $F : \mathcal{M}_+(A) \rightarrow \mathbb{R}^d$ defined by $F(\mu) = \int_A X\mu(dx)$, then we have $\text{cone}(A) \subseteq \text{Range}(F)$.*

Proof of Lemma D.1. Recall that

$$\text{cone}(A) = \left\{ \sum_{i=1}^k w_i x_i : x_i \in A, w_i \in \mathbb{R}_{\geq 0}, k \in \mathbb{N} \right\}. \quad (\text{D.1})$$

For any $x \in \text{cone}(A)$, then there exists $x_1, x_2, \dots, x_k \in A$ and $\{w_i\}_{i=1}^k \geq 0$ such that $x = \sum_{i=1}^k w_i x_i$. Pick $\mu = \sum_{i=1}^k w_i \delta_{x_i}$, where δ_{x_i} are Dirac's delta measure at x_i . Then $\mu \in \mathcal{M}_+(A)$, and

$$x = \sum_{i=1}^k w_i x_i = \sum_{i=1}^k w_i \int_A x \delta_{x_i}(dx) = \int_A x \mu(dx),$$

which leads to the postulated claim. \square

Fact D.2 (Sherman–Morrison Formula). *Suppose $A \in \mathbb{R}^{n \times n}$ is an invertible square matrix and $u, v \in \mathbb{R}^n$ are column vectors. Then $A + uv^\top$ is invertible if and only if $1 + v^\top A^{-1}u \neq 0$. In this case,*

$$(A + uv^\top)^{-1} = A^{-1} - \frac{A^{-1}uv^\top A^{-1}}{1 + v^\top A^{-1}u}.$$

Here, uv^\top is the outer product of two vectors u and v .

Lemma D.3. *Suppose that $y \in \mathbb{R}^d$ satisfying $|y_1| \leq |y_2| \leq \dots \leq |y_d|$ and $\vartheta > 0$. Then, there exist $1 < j < d$ and $\alpha > 0$ such that the problem*

$$\min_{x \in \mathbb{R}^d} \|x\|_1 + \vartheta \|y - x\|_2$$

admits the following optimal solution

$$x^*(\vartheta) = \begin{cases} \mathbf{0} & \text{if } \vartheta \leq \frac{\|y\|_2}{\|y\|_\infty}, \\ \begin{bmatrix} \mathbf{0}_{1:j}, y_{j+1:d} - \alpha \text{sign}(y_{j+1:d}) \\ y \end{bmatrix} & \text{if } \frac{\|y\|_2}{\|y\|_\infty} < \vartheta < \sqrt{d}, \\ y & \text{if } \vartheta \geq \sqrt{d}. \end{cases} \quad (\text{D.2})$$

Proof of Lemma D.3. The basic strategy here is to check the first-order optimality condition.

- If $\vartheta \leq \frac{\|y\|_2}{\|y\|_\infty}$ we have

$$0 \in \partial \|x\|_1|_{x=0} - \vartheta \frac{y}{\|y\|}$$

holds. Thus, 0 is the optimal solution.

- Moreover, if $\vartheta \geq \sqrt{d}$, we have

$$0 \in \text{sign}(y) + \vartheta \partial \|x - y\|_2|_{x=y}$$

holds as $v \in \partial \|x - y\|_2|_{x=y}$ satisfies $\|v\|_2 \leq 1$.

- The most complicated case is the middle one, i.e., $\frac{\|y\|_2}{\|y\|_\infty} < \vartheta < \sqrt{d}$. Here, we are trying to characterize the structure of the optimal solution. Without of loss generality, we assume that $y = \text{sort}(y, 'abs')$, i.e., sorted by its absolute value. Still, we focus on its first-order optimality condition:

$$0 \in \partial \|x\|_1 + \vartheta \frac{x - y}{\|x - y\|_2},$$

as x cannot equal to y derived from the condition $\frac{\|y\|_2}{\|y\|_\infty} < \vartheta < \sqrt{d}$. Furthermore, the optimal solution x shares the same sign of y and $|x_i| \leq |y_i|, \forall i \in [d]$, otherwise you can always decrease the objective value by changing the sign. Next, we will argue that the

optimal solution admits $x_i^* = 0$ for some index i . We prove it by contradiction. If we assume $x^* \neq 0$, there exists a constant $\alpha > 0$ such that

$$x^* = y - \alpha \text{sign}(y).$$

Then, the first-order optimality condition will not hold, i.e.,

$$\text{sign}(y) + \vartheta \frac{\alpha \text{sign}(y)}{\|\alpha \text{sign}(y)\|_2} \neq 0,$$

as $\vartheta < \sqrt{d}$. As such, there exist $1 < j < d$ and a constant $\alpha > 0$ such that $x^* = [\mathbf{0}_{1:j}, y_{j+1:d} - \alpha \text{sign}(y_{j+1:d})]$.

□

E Convergence Analysis of Optimization Algorithms

Denote $\ell(f_\beta(X)) = h(\beta, X)$ and we make the following blanket assumption:

Assumption E.1. *The loss function $h : \Omega \times \mathcal{X} \rightarrow \mathbb{R}$ satisfies the Lipschitzian smoothness conditions*

$$\begin{aligned} \|\nabla_\beta h(\beta_1, X) - \nabla_\beta h(\beta_2, X)\| &\leq C_{\beta\beta} \|\beta_1 - \beta_2\|, \\ \|\nabla_X h(\beta, X_1) - \nabla_X h(\beta, X_2)\| &\leq C_{XX} \|X_1 - X_2\|, \\ \|\nabla_\beta h(\beta, X_1) - \nabla_\beta h(\beta, X_2)\| &\leq C_{\beta X} \|X_1 - X_2\|, \\ \|\nabla_X h(\beta_1, X) - \nabla_X h(\beta_2, X)\| &\leq C_{X\beta} \|\beta_1 - \beta_2\|, \end{aligned}$$

where $\Omega \subset \mathbb{R}^d$ is a closed convex set.

Derivation of (4.1)

$$\begin{aligned} &\min_\beta \frac{1}{N} \sum_{i=1}^N \min_{\alpha_i} \max_{\Delta_i} [\ell(f_\beta(X_i + \Delta_i)) - \alpha_i^\top \Delta_i - \lambda \|\Delta_i\|_M^2 + \epsilon \|\alpha_i\|_{M-1}] \\ &\stackrel{(a)}{=} \min_\beta \frac{1}{N} \sum_{i=1}^N \max_{\Delta_i} \min_{\alpha_i} [\epsilon \|\alpha_i\|_{M-1} - \alpha_i^\top \Delta_i + \ell(f_\beta(X_i + \Delta_i)) - \lambda \|\Delta_i\|_M^2] \quad (\text{E.1}) \\ &= \min_\beta \frac{1}{N} \sum_{i=1}^N \max_{\|\Delta_i\|_M \leq \epsilon} [\ell(f_\beta(X_i + \Delta_i)) - \lambda \|\Delta_i\|_{M-1}^2], \end{aligned}$$

where equality (a) follows from the following minimax theorem as the inner maximization over Δ_i is strongly concave and $\|\alpha_i\|_{M-1}$ is level bounded.

Proof of Lemma 4.1. By invoking the general best-case primal-dual relations given in [22, Corollary 11.40 (d)], the key ingredient is to check the boundedness of

$$\left\{ x \in \mathbb{R}^n : x = \arg \min_x \max_y \{f(x) + x^\top Ay - g(y)\} \right\}$$

and

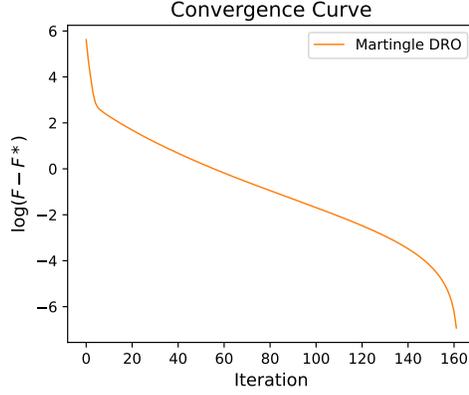
$$\left\{ y \in \mathbb{R}^m : y = \arg \max_y \min_x \{f(x) + x^\top Ay - g(y)\} \right\}.$$

For the purpose of this proof, we use f^* to denote the convex conjugate of f , formally defined as

$$f^*(z) = \max_{x \in \mathbb{R}^n} z^\top x - f(x).$$

Similarly, g^* is the conjugate of g . We have:

- As $f(x)$ is level-bounded, $\max_y f(x) + x^\top Ay - g(y) = f(x) + g^*(A^\top x)$ is also level bounded.
- As $g(y)$ is strongly convex, $\min_x f(x) + x^\top Ay - g(y) = -f^*(-Ay) - g(y)$ is strongly concave and thus its optimal solution set is compact.



The proof is complete. \square

Lemma E.2. *Let $h : \Omega \times \mathcal{X} \rightarrow \mathbb{R}$ be differentiable and $\phi_\lambda(\beta, X) = \sup_{\|\Delta\|_M \leq \epsilon} \{h(\beta, X + \Delta) - \lambda\|\Delta\|_M^2\}$. Suppose that Assumption E.1 holds and $\lambda > \sigma_{\min}(M)C_{XX}$, where $\sigma_{\min}(M)$ is the minimum eigenvalue of M . Then, $\phi_\lambda(\cdot, X)$ is differentiable.*

Proof. As the set $\|\Delta\|_M \leq \epsilon$ is a compact set whenever $0 < \epsilon < \infty$, we know the function $\phi_\lambda(\beta, X)$ is subsmooth function, see Definition 10.29 in [22]. Furthermore, since $h(\beta, \cdot)$ is L -smooth and $\lambda > \sigma_{\min}(M)C_{XX}$, we know that $h(\beta, X + \Delta) - \lambda\|\Delta\|_M^2$ is $(\lambda - \sigma_{\min}(M)C_{XX})$ -strongly concave with respect to Δ . Thus, the inner maximizer is unique and we can invoke [22, Theorem 10.31] (i.e., an extension of envelope theorem) to obtain the differentiability.

Compared with Lemma 1 in [27], our proof here is simpler as we utilize the compactness condition. \square

Recall that

$$\min_{\beta} F(\beta) := \frac{1}{N} \sum_{i=1}^N \underbrace{\max_{\|\Delta_i\|_M \leq \epsilon} [\ell(f_{\beta}(X_i + \Delta_i))] - \lambda\|\Delta_i\|_M^2}_{\phi_\lambda(\beta, X_i)}. \quad (\text{E.2})$$

Theorem E.3 (Convergence of nonconvex SGD; Adopted from Theorem 2 in [27]). *Suppose that $\Delta_F \geq F(\beta^0) - \inf_{\beta} F(\beta)$ and $\mathbb{E} [\|\nabla F(\beta) - \nabla_{\beta} \phi_{\gamma}(\beta, X)\|_2^2] \leq \sigma^2$ and we take constant stepsizes $\alpha = \sqrt{\frac{\Delta_F}{L_{\phi} K \sigma^2}}$ where $L_{\phi} := C_{\beta\beta} + \frac{C_{\beta X} C_{X\beta}}{\lambda - \sigma_{\min}(M)C_{XX}}$. For $K \geq \frac{L_{\phi} \Delta_F}{\sigma^2}$, Algorithm 1 satisfies*

$$\frac{1}{K} \sum_{k=0}^{K-1} \mathbb{E} [\|\nabla F(\beta^k)\|_2^2] - \frac{4C_{\beta X}^2}{\lambda - \sigma_{\min}(M)C_{XX}} \epsilon \leq 4\sigma \sqrt{\frac{L_{\phi} \Delta_F}{K}}.$$

F Supplementary Experiments

First we introduce the attack methods we use in the experiments of adversarial training.

Definition F.1 (PGD/FGSM attack). *For any model parameter β , let*

$$\Delta z_i(\beta) \triangleq \arg \max_{\|\eta\|_p \leq \xi} \{\nabla_z \ell(f_{\beta}(z_i, y_i))^{\top} \eta\} \quad \text{and} \quad \tilde{z}_i \triangleq \Pi_{\mathcal{B}_{\xi, p}(z_i)} \{z_i + \alpha \Delta z_i(\beta)\},$$

where ξ is the attack step size, α is a pre-specified hyperparameter and Π denotes the projection onto $\mathcal{B}_{\xi, p}(z_i) \triangleq \{z : \|z - z_i\|_p \leq \xi\}$. When $p = \infty$, the attack is reduced to the Fast Gradient Sign Method (FGSM) [11]. As in [27], we also consider the Euclidean case $p = 2$, which is a general version of Projected Gradient Descent (PGD) [17] with one step.

Definition F.2 (DRO attack). *For any parameter β , let*

$$\bar{z}_i \triangleq \arg \max_{z \in \mathbb{R}^d} \{ \ell(f_\beta(z_i, y_i)) - \gamma \|z - z_i\|_2^2 \},$$

where γ is a pre-specified hyperparameter.

As for the setup of the experiments on the synthetic data (i.e., Figure 2), we use $\lambda = 2$ for both the DRO approach and our approach. Further in our approach, we use a sequence of $\epsilon \in \{0.2, 0.22, \dots, 1.5\}$ to demonstrate the effectiveness of our approach for being less conservative than the conventional DRO approach. Full results are shown in Figure 5.

As for the setup of the experiments on the MNIST dataset (Figures 3, 4), $\mathbb{E}_{\mathbb{P}} \|X\|_2 = 9.21$ and we choose $\lambda = 0.04 \mathbb{E}_{\mathbb{P}} \|X\|_2$ for training the original DRO and Martingale DRO model. Additionally, we choose $\epsilon = 1.2$ in our model, which is smaller than the average L^2 norm of the perturbations suggested by the original DRO model when training on the MNIST dataset. As for the DRO attack, we choose an increasing sequence of γ (corresponding to a decaying sequence of perturbation) and collect the images after the largest perturbation so that these methods can classify correctly. Full results are shown in Figures 6, 7.

All the hyperparameters conducted in this section have been fine-tuned via grid search for optimal performance.

In the following, we want to highlight that our method can be applied within three lines of *PyTorch* code modification based on the original DRO approach, which is due to the formulation (4.1). The interested reader is referred to our code to see the details.

```
with torch.no_grad():
    delta_norm = delta.norm(p = 2, dim = (2,3))
    delta_index = delta_norm > eps
    delta[delta_index] /= (delta_norm[delta_index][:, None, None]/eps)
```

Simple implementation based on the original DRO approach results in little extra computation complexity, which is also shown in the following track of time during experiments on MNIST dataset.

Training time per epoch (s)	DRO	Martingale DRO
Average	1.66	1.73
Variance	1.90×10^{-3}	2.10×10^{-3}

Table 2: Per-iteration wall-clock time comparison between the vanilla DRO model and Martingale DRO model.

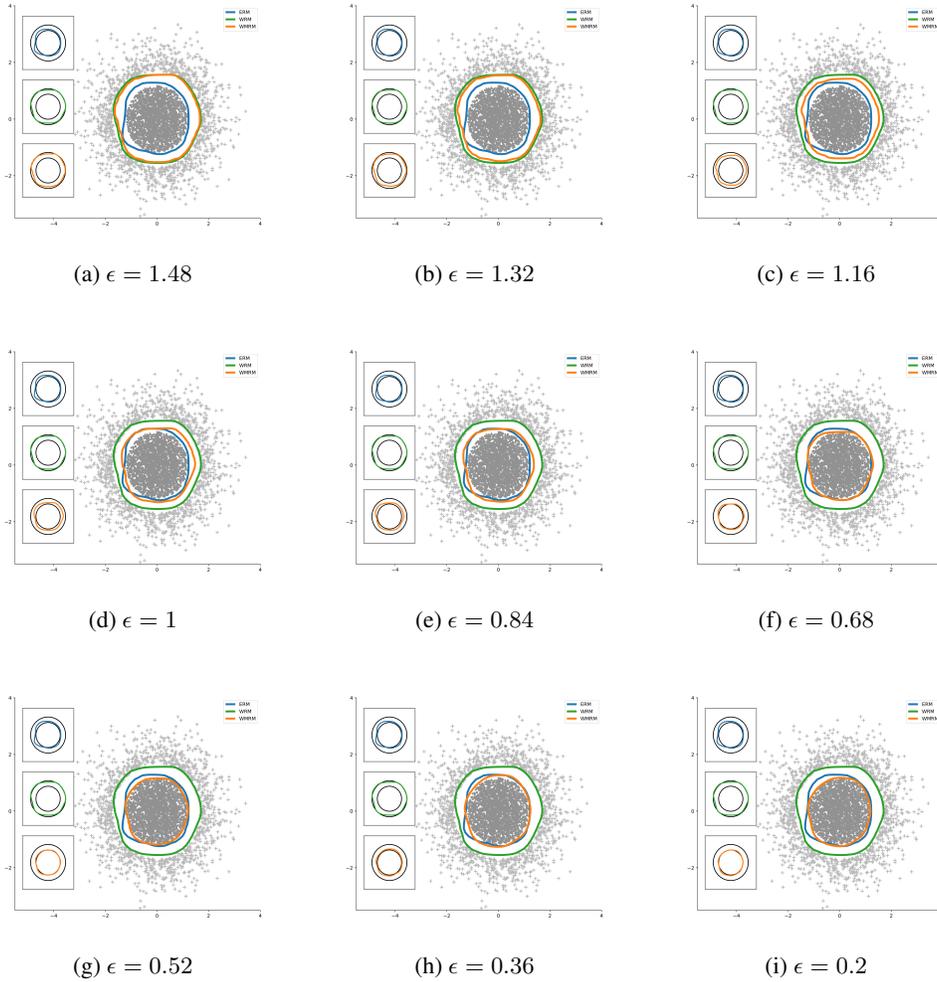


Figure 5: Synthetic data. Decreasing non-zero ϵ 's push the perturbed martingale constraints towards the exact martingale constraints and force the classification boundary increasingly inward.

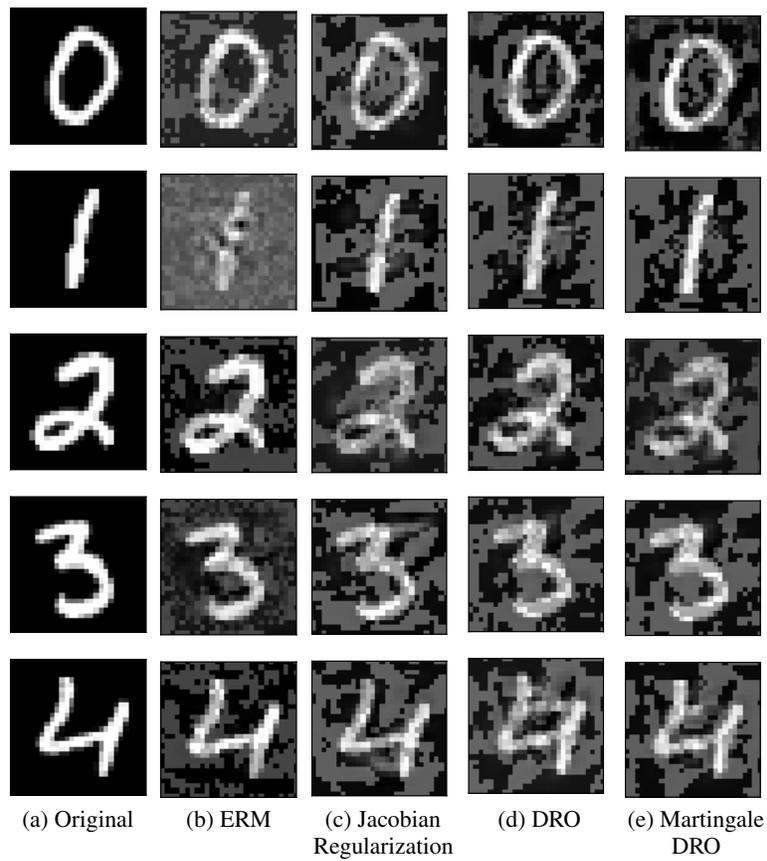


Figure 6: The largest DRO perturbations such that each model makes correct prediction.

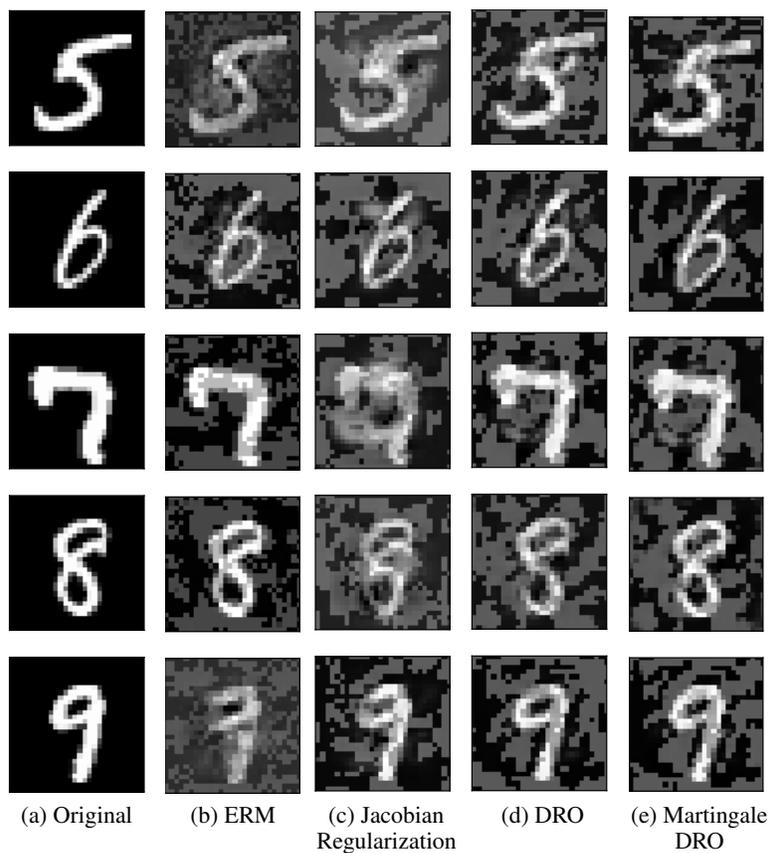


Figure 7: The largest DRO perturbations such that each model makes correct prediction.