

Zeroth-Order Methods for Convex-Concave Minmax Problems: Applications to Decision-Dependent Risk Minimization

Chinmay Maheshwari^{1*}, Chih-Yuan Chiu^{1*}, Eric Mazumdar¹,
S. Shankar Sastry¹, Lillian J. Ratliff²

¹Electrical Engineering and Computer Sciences, University of California, Berkeley

²Electrical and Computer Engineering, University of Washington, Seattle

Abstract

Min-max optimization is emerging as a key framework for analyzing problems of robustness to strategically and adversarially generated data. We propose a random reshuffling-based gradient free Optimistic Gradient Descent-Ascent algorithm for solving convex-concave min-max problems with finite sum structure. We prove that the algorithm enjoys the same convergence rate as that of zeroth-order algorithms for convex minimization problems. We further specialize the algorithm to solve distributionally robust, decision-dependent learning problems, where gradient information is not readily available. Through illustrative simulations, we observe that our proposed approach learns models that are simultaneously robust against adversarial distribution shifts and strategic decisions from the data sources, and outperforms existing methods from the strategic classification literature.

1 Introduction

The deployment of learning algorithms in real-world scenarios necessitates versatile and robust algorithms that operate efficiently under mild information structures. Min-max optimization has been used as a tool ensure robustness in variety of domains e.g. robust optimization [BTEGN09], robust control [HÅBB13], to name a few. Recently, min-max optimization has emerged as a promising framework for framing problems of algorithmic robustness against adversaries [GPAM⁺14, SKL17, MMS⁺17], strategically generated data [DRS⁺18, BHK20], and distributional shifts in dynamic environments [YLMJ21].

Despite this, recent works in machine learning and robust optimization on designing and analyzing stochastic algorithms for min-max optimization problems have largely operated on a number of assumptions that preclude their application to a broad range of real-world problems e.g., access to first-order oracles that provide exact gradients [YKH20, NSH⁺19, JNJ20] or restrictive structural assumptions such as strong convexity [LLC⁺20, WBMR20, SBDG21]. Moreover, the developed theory is often not well-aligned with the practical implementation of these algorithms in real-world machine learning applications. For example, [BSG20] propose zeroth-order methods for convex-concave problems but the proposed algorithm may not be suitable for machine learning applications where the objective function is a sum of large numbers of component functions (depending on the size of dataset). Indeed, in order to compute the gradient estimate at any iteration Beznosikov et al requires perturbing **all** the functions which might not be suitable/possible for many applications. Furthermore, stochastic gradient methods are often used with random reshuffling (without replacement) in practice, yet their theoretical performance is usually characterized under the assumption of uniform sampling with replacement [Bot09, JNN19].

*Equal contribution

In this work, we do away with these assumptions and formulate a gradient-free (zeroth-order), random reshuffling-based algorithm with non-asymptotic convergence guarantees under mild structural assumptions on the underlying min-max objective. Our convergence guarantees are established by balancing the bias and variance of the zeroth-order gradient estimator [BLM18], using coupling-based arguments to analyze the correlations between iterates due to the random reshuffling procedure [JNN19], and exploiting the recent connections between the Optimistic Gradient Descent Ascent (OGDA) and Proximal Gradient algorithms [MOP20b].

One of the primary problem areas in which such an algorithm becomes necessary is in learning from strategically generated or decision-dependent data, a classical problem in operations research (see, e.g., [HBT18] and references therein). This problem has garnered a lot of attention of late in the machine learning community under the name “performative prediction” [PZMDH20, MPZ21, BHK20] due to the growing recognition that learning algorithms are increasingly dealing with data from strategic agents. In such problems, assuming access to the response map of strategic agents is often too restrictive, and the introduction of agent’s strategic responses into a convex loss function can often result in non-convex objectives.

As an example of such a decision-dependent problem, consider a scenario in which a ride-sharing platform seeks to devise an adaptive pricing strategy which is responsive to changes in supply and demand. The platform observes the current supply and demand in the environment and adjusts the price to increase the supply of drivers (and potentially decrease the demand) as needed. Drivers, however, have the ability to adjust their availability, and can strategically create dips in supply to trigger price increases. Such gaming has been observed in real ride-share markets (see, e.g., [Ham19, You19]) and results in negative externalities like higher prices for passengers. Importantly, in this situation, the platform does not observe precisely the decision making process of the drivers, only their strategically generated availability, and must learn to optimize through these agents’ responses. This lack of precise knowledge regarding the data generation process, and the reactive nature of the data, motivate the use of game theoretic abstractions for the decision problem, as well as algorithms for finding solutions in the absence of full information.

Previous work analyzing this problem studies this phenomenon through the lens of risk minimization in which the data distribution is decision-dependent, and seeks out settings in which the decision maker can optimize the decision-dependent risk [MPZ21]. These works, however, do not account for model misspecification in their analysis. In particular, if the data generation model is incorrect, the performance of the optimal solution returned by their training methods may potentially degrade rapidly, something we explore in our experiments.

We show that the decision-dependent learning (performative prediction) problem can be *robustified* by taking a distributional robustness perspective on the original problem. Moreover, we show that, under mild assumptions, the distributionally robust decision-dependent learning problem can be transformed to a min-max problem and hence our zeroth-order random reshuffling algorithm can be applied. The gradient-free nature of our algorithm is important for applications where data is generated by strategic users that one must query; in these scenarios, the decision-maker is unlikely to access the best response map (data generation mechanism) of the strategic users, and hence will lack access to precise gradients.

Contributions. In this paper, we analyze the class of convex-concave min-max problems given by

$$\min_{x \in \mathcal{X}} \max_{y \in \mathcal{Y}} L(x, y), \tag{1}$$

where $\mathcal{X} \subset \mathbb{R}^{d_x}$, $\mathcal{Y} \subset \mathbb{R}^{d_y}$, and $L : \mathbb{R}^{d_x} \times \mathbb{R}^{d_y} \rightarrow \mathbb{R}$ has the finite-sum structure given by $L := \frac{1}{n} \sum_{i=1}^n L_i$, where $L_1, \dots, L_n : \mathbb{R}^{d_x} \times \mathbb{R}^{d_y} \rightarrow \mathbb{R}$ denote n individual loss functions. This formulation is ubiquitous in machine learning applications, where the overall loss objective is often the average of the loss function evaluated over each data point in a dataset.

The contributions of this paper can be summarized as follows.

- I) We propose an efficient zeroth-order random reshuffling-based OGDA algorithm for a convex-concave min-max optimization problem, *without* assuming any other structure on the curvature of the min-max loss (e.g., strong convexity or strong concavity). We provide (to our knowledge) the *first* non-asymptotic analysis of OGDA algorithm *with* random reshuffling and zeroth-order gradient information.
- II) As an important application, we formulate the *Wasserstein distributionally robust learning with decision-dependent data* problem as a constrained finite-dimensional, smooth convex-concave min-max problem of the form (1). In particular, we consider the setting of learning from strategically generated data, where the goal is to fit a generalized linear model, and where an ambiguity set is used to capture model misspecification regarding the data generation process. This setting encapsulates a distributionally robust version of the recently introduced problem of strategic classification [HMPW16]. We show that this problem, under mild assumptions on data generation model and the ambiguity set, can be transformed into a convex-concave min-max problem to which our algorithm applies.
- III) We complement the theoretical contributions of this paper by presenting illustrative numerical examples.

2 Related Work

Our work draws upon the existing literature on zeroth-order methods for min-max optimization problems, decision-dependent learning (performative prediction), and distributionally robust optimization.

Zero-Order Methods for Min-Max Optimization. Zeroth-order methods provide a computationally efficient method for applications in which first-order or higher-order information is inaccessible or impractical to compute, e.g., when generating adversarial examples to test the robustness of black-box machine learning models [LLC⁺20, LCK⁺20, CZS⁺17, IEAL18, TTC⁺19, LLW⁺19, AH17]. Recently, Liu et al. and others [LLC⁺20, GJZ18, WBMR20] provided the first non-asymptotic convergence bounds for zeroth-order algorithms, based on analysis methods for gradient-free methods in convex optimization [NS17]. However, these works assume that the min-max objective is either strongly concave in the maximizing variable [LLC⁺20, WBMR20] or strongly convex [GJLJ17] in the minimizing variable, an assumption that fails to hold in many applications [DRS⁺18, YLMJ21]. In contrast, the zeroth-order algorithm presented in this work provides non-asymptotic guarantees under the less restrictive assumption that the objective function is convex-concave. In particular, we present the first (to our knowledge) zeroth-order variant of the Optimistic Gradient Ascent-Descent (OGDA) algorithm [MOP20b, MOP20a].

In single-variable optimization problems, first-order stochastic gradient descent algorithms are empirically observed to converge faster when random reshuffling (RR, or sampling without replacement) is deployed, compared to sampling with replacement [RR11, Bot09]. Although considerably more difficult to analyze theoretically, gradient-based RR methods have recently been shown to enjoy faster convergence when the underlying objective function is convex [Sha16, JNN19, MKR21, HS19, SS19, GOP21, RGP20]. Recently, these theoretical results have been extended to first-order methods for convex-concave min-max optimization problems [YLMJ21]; in this paper, we further extend these results to the case

of zero-order algorithms. Specifically, we present the first (to our knowledge) non-asymptotic convergence rates for zeroth-order random reshuffling min-max optimization algorithms.

Distributionally Robust Optimization. *Distributionally Robust Optimization (DRO)* seeks to find solutions to optimization problems (e.g., supervised learning tasks) robust against changes in the data distribution between training and test time [MMS⁺17, YLMJ21]. These distributional differences may arise due to imbalanced data, sample selection bias, or adversarial perturbations or deletions [CSSL09, MMS⁺17], and are often modeled as min-max optimization problems, in which the classifier and an adversarial noise component are respectively modeled as the minimizer and maximizer of a common min-max loss objective [Bag05, BBC10, GMT14, GYdH15, RM19]. In particular, the noise is assumed to generate the worst possible loss corresponding to a bounded training data distribution shift, with the bound given by either the f -divergence or Wasserstein distance. [YLMJ21, BTHW⁺13, ND16, HNSS18, SAEK15]. While this work considers adversarial noise in generated data, largely in a worst-case context, it has yet to capture strategically generated data wherein a data source generates data via a best response mapping.

Strategic Classification and Performative Prediction. *Strategic classification* [HMPW16, DRS⁺18, KTS⁺19, SBKK20] and *performative prediction* [PZMDH20, MDPZH20, MPZ21, DX20] concern supervised learning problems in which the training data distribution shifts in response to the deployed classifier or predictor more generally. This setting naturally arises in machine learning applications in which the selection of the deployed classifier either directly changes the training data (e.g., decisions based on credit scores, such as loan approvals, themselves change credit scores), or prompts the data source to artificially alter their attributes (e.g. withdrawals during bank runs spur worried clients to make more withdrawals) [PZMDH20, MDPZH20, MPZ21]. Here, the learner accesses only perturbed features representing the strategic agents' best responses to a deployed classifier, and not the true underlying features [DRS⁺18]. This is a recently introduced formulation to machine learning; the results in this body of literature (to our knowledge) have not introduced the concept of robustness to model misspecification or the data generation process, in the same manner as we capture in this work.

3 Preliminaries

Recall that in this paper, we consider the class of convex-concave min-max problems given by:

$$\min_{x \in \mathcal{X}} \max_{y \in \mathcal{Y}} L(x, y), \quad (2)$$

where $\mathcal{X} \subset \mathbb{R}^{d_x}$, $\mathcal{Y} \subset \mathbb{R}^{d_y}$, and $L := \frac{1}{n} \sum_{i=1}^n L_i$, where $L_1, \dots, L_n : \mathbb{R}^{d_x} \times \mathbb{R}^{d_y} \rightarrow \mathbb{R}$ denote n individual loss functions. For convenience, we denote $d := d_x + d_y$.

Assumption 3.1. *The following statements hold:*

- (i) *The sets $\mathcal{X} \subset \mathbb{R}^{d_x}$ and $\mathcal{Y} \subset \mathbb{R}^{d_y}$ are convex and compact.*
- (ii) *The functions $L_1, \dots, L_n : \mathbb{R}^d \rightarrow \mathbb{R}$ are convex in $x \in \mathbb{R}^{d_x}$ for each $y \in \mathbb{R}^{d_y}$, concave in $y \in \mathbb{R}^{d_y}$ for each $x \in \mathbb{R}^{d_x}$, and G -Lipschitz and ℓ -smooth in $(x, y) \in \mathbb{R}^d$ (which implies that $L : \mathbb{R}^d \rightarrow \mathbb{R}$, by definition, also possesses the same properties).*

For ease of exposition, we denote $u := (x, y)$, $M_L := \sup_{u \in \mathcal{X} \times \mathcal{Y}} |L(u)|$, $D := \sup_{u, u' \in \mathcal{X} \times \mathcal{Y}} \|u - u'\|_2$, and define the operators $F, F_i : \mathbb{R}^d \rightarrow \mathbb{R}^d$, for each $i \in [n]$, by:

$$F(u) := \begin{bmatrix} \nabla_x L(u) \\ -\nabla_y L(u) \end{bmatrix}, \quad F_i(u) := \begin{bmatrix} \nabla_x L_i(u) \\ -\nabla_y L_i(u) \end{bmatrix}, \quad \forall i \in [n]. \quad (3)$$

Observe that under Assumption 3.1, $M_L, D < \infty$, and F and each F_i are monotone¹. Finally, we define the *gap function* $\Delta : \mathbb{R}^d \rightarrow [0, \infty)$ associated with the loss L by

$$\Delta(x, y) := L(x, y^*) - L(x^*, y) \geq 0, \quad (4)$$

where $u^* := (x^*, y^*) \in \mathcal{X} \times \mathcal{Y}$ denotes any min-max saddle point of the overall loss $L(x, y)$, and $(x, y) \in \mathcal{X} \times \mathcal{Y}$ denotes any feasible point. This gap function allows us to measure the convergence rate of our proposed algorithm. To this end, we define the ϵ -optimal saddle-point of (2) as follows.

Definition 3.1 (ϵ -optimal saddle point solution). *A feasible point $(x, y) \in \mathcal{X} \times \mathcal{Y}$ is said to be an ϵ -optimal saddle-point solution of (2) if*

$$\Delta(x, y) = L(x, y^*) - L(x^*, y) \leq \epsilon,$$

4 Algorithms and Performance

In this section we introduce a gradient-free version of the well-studied Optimistic Gradient Descent Ascent (OGDA) algorithm, and give finite time rates showing that it can efficiently find the saddle point in constrained convex-concave problems.

4.1 Zero-Order Gradient Estimates

In our zero-order, random-reshuffling based variant of the OGDA algorithms, we use the one-shot randomized gradient estimator in [Spa97, FKM05, GJZ18, LLC⁺20]. In particular, given the current iterate $u \in \mathbb{R}^d$ and a *query radius* $\varepsilon > 0$, we sample a vector v uniformly from unit sphere \mathcal{S}^{d-1} (i.e. $v \sim \text{Unif}(\mathcal{S}^{d-1})$), and define the zeroth-order estimator $\hat{F}(u; \varepsilon, v) \in \mathbb{R}^d$ of the min-max loss $L(u)$ to be:

$$\hat{F}(u; \varepsilon, v) := \frac{d}{\varepsilon} L(u + \varepsilon v) v$$

Properties of this zeroth-order estimator, derived in [BLM18], are reproduced as Proposition A.4 in Appendix A.1.

4.2 Optimistic Gradient Descent Ascent with Random Reshuffling (OGDA-RR)

In this subsection, we formulate our main algorithm, Optimistic Gradient Descent Ascent with Random Reshuffling (OGDA-RR). In each epoch $t \in \{0, 1, \dots, T-1\}$, the algorithm generates a uniformly random permutation $\sigma^t := (\sigma_1^t, \dots, \sigma_n^t)$ of $[n] := \{1, \dots, n\}$ independently of any other randomness, and fixes a query radius $\varepsilon^t > 0$ and search direction $v_i^t \in \mathbb{R}^d$. (Note: query radii only depends on epoch indices t , and not on sample indices). For each index $i \in [n]$, we compute the OGDA-RR update as follows:

$$u_{i+1}^t = \text{Proj}_{\mathcal{X} \times \mathcal{Y}} \left(u_i^t - \eta^t \hat{F}_{\sigma_i^t}^t(u_i^t; \varepsilon^t, v_i^t) - \eta^t \hat{F}_{\sigma_{i-1}^t}^t(u_i^t; \varepsilon^t, v_i^t) + \eta^t \hat{F}_{\sigma_{i-1}^t}^t(u_{i-1}^t; \varepsilon^t, v_{i-1}^t) \right), \quad (5)$$

After repeating this process for T epochs, the algorithm returns the step-size-weighted average of the iterates, $\tilde{u}^T := \frac{1}{n \cdot \sum_{t=0}^{T-1} \eta^t} \sum_{t=0}^{T-1} \sum_{i=1}^n \eta^t u_i^t$. Roughly, the weighting described in Theorem 4.1 below optimally balances the bias and variance of the zero-order gradient estimator in Section 4.1.

¹A function $F : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is called *monotone* if $\langle F(x) - F(y), x - y \rangle \geq 0$ for all $x, y \in \mathbb{R}^d$.

Algorithm 1: OGDA-RR Algorithm

1 Input: stepsizes η^t, ε^t , data points $\{(x_i, y_i)\}_{i=1}^n \sim \mathcal{D}, u_0^{(0)}$, time horizon duration T ;
2 for $t = 0, 1, \dots, T - 1$ **do**
3 $\sigma^t = (\sigma_1^t, \dots, \sigma_n^t) \leftarrow$ a random permutation of set $[n]$;
4 **for** $i = 0, \dots, n - 1$ **do**
5 Sample $v_i^t \sim \text{Unif}(\mathcal{S}^{d-1})$
6 $u_{i+1}^t \leftarrow (5)$
7 **end**
8 $u_0^{(t+1)} \leftarrow u_n^t$
9 $u_{-1}^{(t+1)} \leftarrow u_{n-1}^t$
10 end
11 Output: $\tilde{u}^T := \frac{1}{n \cdot \sum_{t=0}^{T-1} \eta^t} \sum_{t=0}^{T-1} \sum_{i=1}^n \eta^t u_i^t$.

Theorem 4.1. *Let $L(u)$ denote the objective function in the constrained min-max optimization problem given by (1), and let $u^* = (x^*, y^*) \in \mathcal{X} \times \mathcal{Y}$ denote any saddle point of $L(u)$. Fix $\epsilon > 0$. Suppose Assumption 3.1 holds, and the number of epochs T , step sizes sequence $\{\eta^t\}_{t=0}^{T-1}$, and query radii sequence $\{\varepsilon^t\}_{t=0}^{T-1}$ satisfy:*

$$\begin{aligned} \eta^t &:= \eta^0 \cdot (t+1)^{-3/4+\chi}, & \forall t \in \{0, 1, \dots, T-1\}, \\ \varepsilon^t &:= \varepsilon^0 \cdot (t+1)^{-1/4}, & \forall t \in \{0, 1, \dots, T-1\}, \\ T &> \frac{1}{\varepsilon^4} \left(\frac{3}{16n} D + \frac{5}{4n} \cdot C \cdot \max \left\{ \varepsilon^0, \eta^0, \eta^0 \varepsilon^0, \frac{\eta^0}{\varepsilon^0}, \frac{\eta^0}{(\varepsilon^0)^2} \right\} \left(1 + \frac{1}{\chi} \right) \right)^{\frac{4}{1-4\chi}}, \end{aligned}$$

for some initial step size $\eta^0 \in \left(0, \frac{1}{2\ell}\right)$, initial query radius $\varepsilon^0 > 0$, parameter $\chi \in (0, 1/4)$, and constant:

$$C := \max \{ 3ndD\ell, 18ndDG\ell, 54ndG^2 + 18ndD\ell M_L, 90ndGM_L, 36ndM_L^2, 6dGn^2 + 14Gn^2 + 4nG, 6dM_Ln, 3dDG\ell, 3dD\ell M_L \} > 0.$$

Then the iterates $\{u_i^t\}$ generated by the OGDA-RR Algorithm (Alg. 1) satisfy:

$$\mathbb{E}[\Delta(\tilde{u}^T)] < \epsilon.$$

Remark. Note that our OGDA-RR algorithm is more computationally efficient than Alg. 2 in [YLMJ21], even if one replaces the gradient estimates with true gradient values. This is because Alg. 2 in [YLMJ21] requires $M \sim \mathcal{O}(\log(n))$ inner loop iterations to approximate a proximal point update. Here, we avoid this restriction by exploiting the recent perspective that the OGDA update is a perturbed proximal point update [MOP20a, MOP20b]. For more details, see Appendix A.2 for the proof of Theorem 4.1.

5 Applications to Decision-Dependent DRO

In this section we discuss a novel convex-concave min-max reformulation of a class of decision-dependent distributional robust risk minimization problems, which reflects the need for learning classifiers that are simultaneously robust to strategic data sources and adversarial model-specification. In particular, we present a distributionally robust formulation of strategic classification [DRS⁺18] with generalized linear loss, a semi-infinite optimization problem that can be reformulated to a finite-dimensional convex-concave min-max problem.

Strategic classification is an emerging paradigm in machine learning which attempts to “close the loop”—i.e., account for data (user) reaction at training time—while designing

classifiers to be deployed in strategic environments in the real world, where deploying *naïve* classifier (designed ignoring the distribution shift) can be catastrophic. Modeling the exact behavior of such strategic interactions is very complex, since the decision-maker (learner) does not have access to the strategic users' preferences and hence lacks access to their best response function. To overcome this difficulty, we use a natural model for these strategic behaviors that has been exploited in Dong et. al.(2018), and then impose robustness conditions (in the form of an ambiguity set on the decision-dependent data distribution) to capture model misspecification. To facilitate the discussion, we provide a primer on decision-dependent DRO in the next subsection.

5.1 Primer on decision-dependent distributionally robust optimization

Consider a *generalized linear problem*, where the goal is to estimate the parameter $\theta \in \Theta$, which is assumed to be a compact set, by solving the following convex optimization program:

$$\inf_{\theta \in \Theta} \mathbb{E}_{\mathcal{D}} [\phi(\langle \bar{x}, \theta \rangle) - \bar{y} \langle \bar{x}, \theta \rangle]$$

where $\phi : \mathbb{R} \rightarrow \mathbb{R}$ is a smooth convex function and the tuple $(\bar{x}, \bar{y}) \in \mathbb{R}^d \times \{-1, +1\}$ is sampled from an unknown distribution \mathcal{D} , often approximated by the empirical distribution of a set of observed data. The generalized linear model encompasses a wide range of machine learning formulations [MN19].

A distributionally robust generalized linear problem, on the other hand, minimizes the worst case expectation over an uncertainty set \mathcal{P} in the space of probability measures. This setup can be envisioned as a game between a learning algorithm and an adversary. Based on parameters chosen by the learning algorithm, the adversary then picks a probability measure from the uncertainty set which maximizes the risk for that choice of parameter:

$$\inf_{\theta} \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{E}_{\mathbb{P}} [\phi(\langle \bar{x}, \theta \rangle) - \bar{y} \langle \bar{x}, \theta \rangle],$$

where $(\bar{x}, \bar{y}) \sim \mathbb{P} \in \mathcal{P}$. Typically \mathcal{P} is chosen as a Wasserstein ball around the empirical distribution $\tilde{\mathcal{D}}_n$ of a set of n observed data points, $\{(\tilde{x}_i, \tilde{y}_i) \in \mathbb{R}^d \times \{-1, 1\}\}_{i=1}^n$, sampled independently from the data distribution \mathcal{D} . Then, for any $\delta > 0$ the uncertainty set \mathcal{P} is given by $\mathbb{B}_{\delta}(\tilde{\mathcal{D}}_n) = \{\mathbb{P} : \mathcal{W}(\mathbb{P}, \tilde{\mathcal{D}}_n) \leq \delta\}$.

A critique of the above problem formulation is that the underlying data distribution \mathcal{D} is considered fixed, while in many strategic settings underlying data distribution will depend on the classifier parameter θ . *Decision-dependent supervised learning* aims to tackle such distribution shifts. When specialized to the generalized linear model, the problem formulation becomes:

$$\inf_{\theta} \mathbb{E}_{\mathcal{D}(\theta)} [\phi(\langle \bar{x}, \theta \rangle) - \bar{y} \langle \bar{x}, \theta \rangle],$$

where $(\bar{x}, \bar{y}) \sim \mathcal{D}(\theta)$. In this work, we take a step forward and work with the *distributionally robust decision-dependent generalized linear model*, defined as:

$$\inf_{\theta} \sup_{\mathbb{P} \in \mathcal{P}(\theta)} \mathbb{E}_{\mathbb{P}} [\phi(\langle \bar{x}, \theta \rangle) - \bar{y} \langle \bar{x}, \theta \rangle], \quad (6)$$

where $(\bar{x}, \bar{y}) \sim \mathbb{P} \in \mathcal{P}(\theta)$ and $\mathcal{P}(\theta) = \mathbb{B}_{\delta}(\tilde{\mathcal{D}}_n(\theta))$. Here, the dependence of \mathbb{P} on the choice of classifier θ is captured by its inclusion in $\mathcal{P}(\theta) = \mathbb{B}_{\delta}(\tilde{\mathcal{D}}_n(\theta))$. To describe decision-dependent distribution shifts $\tilde{\mathcal{D}}_n(\theta)$, we restrict our focus to the setting of strategic classification. The following subsection formalizes our setting.

5.2 Model for strategic response

Below, we denote the data points sampled from *true distribution* by $(\tilde{x}_i, \tilde{y}_i) \sim \mathcal{D}$ where \mathcal{D} is a unknown, underlying distribution. For ease of presentation, we associate each data

point index i with an agent. For each agent $i \in [n]$, let $u_i(x; \theta, \tilde{x}_i, \tilde{y}_i) \in \mathbb{R}$ denote its utility function that a strategic agent seeks to maximize. In other words, when a classifier parametrized by $\theta \in \mathbb{R}^d$ is deployed, the agent $i \in [n]$ responds by reporting $b_i(\theta, \tilde{x}_i, \tilde{y}_i)$, defined as:

$$b_i(\theta, \tilde{x}_i, \tilde{y}_i) \in \arg \max_x u_i(x; \theta, \tilde{x}_i, \tilde{y}_i).$$

Note that we allow different agent to have different utility function.

We now impose the following assumptions on the utility functions; these are crucial for ensuring guaranteed convergence of our proposed algorithms.

Assumption 5.1. For each agent $i \in [n]$, define $u_i(x; \theta, \tilde{x}_i, \tilde{y}_i) := \frac{1-\tilde{y}_i}{2} \langle x, \theta \rangle - g_i(x - \tilde{x}_i)$, where $g_i : \mathbb{R}^d \rightarrow \mathbb{R}$ satisfies:

- (i) $g_i(x) > 0$ for all $x \neq 0$;
- (ii) g_i is convex on \mathbb{R}^d ;
- (iii) g_i is positive homogeneous² of degree $p > 1$;
- (iv) Its convex conjugate $g_i^*(\theta) := \sup_{x \in \mathbb{R}^d} \langle x, \theta \rangle - g_i(x)$ is G_i -Lipschitz and \bar{G}_i -smooth on Θ .

As is pointed out in Dong et. al. (2018), a large class of functions $g(\cdot)$ satisfy the requirements posited in Assumption 5.1. For example, for any arbitrary norm and any $p > 1$ the function $g(x) = \frac{1}{p} \|x\|^p$ is a candidate. Note that these assumptions are not very restrictive and capture a large variety of practical scenarios [DRS⁺18]. A natural consequence of the above modeling paradigm is that $b_i(\theta, \tilde{x}_i, +1) = \tilde{x}_i$. To wit, the agents act strategically only if their true label is -1 . This is a reasonable setting for many real world applications [DRS⁺18]. We now present a technical lemma which will be helpful in subsequent presentation.

Lemma 5.2 (Dong et. al. (2018)). Under Assumption 5.1, for each agent $i \in [n]$, the set of best responses $\arg \max_x u_i(x; \theta, \tilde{x}_i, \tilde{y}_i)$ is finite and bounded. The function $\theta \mapsto \langle b_i(\theta, \tilde{x}_i, \tilde{y}_i), \theta \rangle$ is convex. To wit, for any $i \in [n]$: $\langle b_i(\theta, \tilde{x}_i, \tilde{y}_i), \theta \rangle = \langle \tilde{x}_i, \theta \rangle + \frac{1-\tilde{y}_i}{2} q g_i^*(\theta)$ where $\frac{1}{p} + \frac{1}{q} = 1$

Against the preceding backdrop, we now present the convex-concave min-max reformulation of the Wasserstein Distributionally Robust Strategic Classification (WDRSC) problem.

5.3 Reformulation of the WDRSC Problem

The WDRSC problem formulation contains two main components—the *strategic component* that accounts for a distribution shift $\mathcal{D}(\theta)$ in response to the choice of classifier θ , and the *adversarial component* that accounts for the uncertainty set $\mathcal{P}(\theta)$. As per the modeling assumptions described in Section 5.2, we have $(\tilde{x}_i, \tilde{y}_i) \sim \mathcal{D}$ and $(b_i(\theta, \tilde{x}_i, \tilde{y}_i), \tilde{y}_i) \sim \mathcal{D}(\theta)$ for all $i \in [n]$. We now impose certain restrictions on the adversarial component that would enable us to reformulate the WDRSC problem as a convex-concave min-max optimization problem. Crudely speaking, we allow adversarial modifications on *features* for all data points, but adversarial modifications on *labels* only when the true label is $+1$.

For the distributionally robust strategic classification problem, we consider a specific form of uncertainty set $\mathcal{P}(\theta)$ that allows us to reformulate the infinite-dimensional optimization problem as a finite-dimensional convex-concave min-max problem. As described above, in our formulation, the features of a given data point i can be perturbed strategically if $\tilde{y}_i = -1$, but not if $\tilde{y}_i = +1$. On top of the strategic perturbations we also consider the

²A function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is *positive homogenous of degree r* if for any scalar $\alpha > 0$ and $x \in \mathbb{R}^d$ we have $f(\alpha x) = \alpha^r f(x)$

adversarial perturbations to the data points. Specifically, we also assume that the adversary can perturb both the features and label of a data point i if $\tilde{y}_i = 1$, but can only perturb the features and not the label if $\tilde{y}_i = -1$. A rigorous exposition of this restriction is deferred to Appendix B.1. Under these assumptions, we now present a convex-concave min-max reformulation of the WDRSC problem.

Theorem 5.3. *Let the strategic behavior of the agents be governed in accordance with Assumption 5.1. Suppose ϕ is convex and β -smooth. In addition, suppose $\mathbb{R} \ni x \mapsto \phi(x) + x \in \mathbb{R}$ is non-decreasing. Then the WDRSC problem (6) can be reformulated into the following convex-concave min-max problem:*

$$\begin{aligned} \min_{(\theta, \alpha)} \max_{\gamma \in \mathbb{R}^n} & \left\{ \alpha(\delta - \kappa) + \frac{1}{n} \sum_i \frac{1 + \tilde{y}_i}{2} (\phi(\langle b_i(\theta), \theta \rangle)) + \gamma_i (\langle b_i(\theta), \theta \rangle - \alpha\kappa) \right. \\ & \left. + \frac{1}{n} \sum_i \frac{1 - \tilde{y}_i}{2} (\phi(\langle b_i(\theta), \theta \rangle) + \langle b_i(\theta), \theta \rangle) \right\} \\ \text{s.t. } & \|\theta\| \leq \alpha/(\beta + 1), \quad \|\gamma\|_\infty \leq 1 \end{aligned} \quad (7)$$

where for any $i \in [n]$, we have concisely written $b_i(\theta, \tilde{x}_i, \tilde{y}_i)$ as $b_i(\theta)$.

The proof of Theorem 5.3 is presented in Appendix B.2.

Remark. The non-decreasing assumption on the map $\mathbb{R} \ni x \mapsto \phi(x) + x \in \mathbb{R}$ is not overly restrictive; in fact, it is satisfied by the logistic regression model in supervised learning (see Appendix C).

Remark. Note that we can convert the smooth convex-concave minmax problem (7) into a non-smooth convex minimization problem by explicitly taking maximization over γ . But we refrain from doing as it has been observed [YLMJ21] that solving the smooth minimax optimization problem is faster than solving the non-smooth problem. In fact, we have presented an experimental study in Appendix C which corroborates this observation.

Throughout the rest of this paper, we denote the min-max objective in (7) by $L(\alpha, \theta, \gamma)$.

6 Empirical Results

In this section we deploy zeroth-order OGDA algorithm with random reshuffling to solve the convex concave reformulation of WDRSC as presented in (7). We point out that in order to solve (7), the zeroth-order method should only be applied to estimate the gradient with respect to θ . This is because the gradient with respect to other variables, namely (α, γ) , can be exactly computed. Specifically, to compute derivative with respect to θ the designer must know the best response function which is often not available and it can only be queried.

We now present some illustrations of the empirical performance of our proposed algorithm, as well as empirical justification for solving the WDRSC problem over existing prior approaches to strategic classification.

6.1 Experimental Setup

Our first set of empirical results uses synthetic data to illustrate the effectiveness of our algorithms. The datasets used in this section are constructed as follows: the ground truth classifier θ^* and features \tilde{x}_i are sampled as $\theta^* \sim \mathcal{N}(0, I_d)$ and $\tilde{x}_i \sim$ i.i.d. $\mathcal{N}(0, I_d)$, for each $i \in [n]$, while the ground truth labels \tilde{y}_i are given by $\tilde{y}_i = \text{sign}(\langle \tilde{x}_i, \theta^* \rangle + z_i)$ for each $i \in [n]$, where $z_i \sim$ i.i.d. $\mathcal{N}(0, 0.1 \cdot I_d)$. We use $n \in \{500, 1000\}$ with $d = 10$. The first five of the

$d = 10$ features are chosen to be strategic. In all experiments, we take $\kappa = 0.5$ and $\delta = 0.4$. Each strategic agent $i \in [n]$ has a utility function given by:

$$u_i(x; \theta, \tilde{x}_i, \tilde{y}_i, \zeta_i) = \frac{1 - \tilde{y}_i}{2} \langle x, \theta \rangle - \frac{1}{2\zeta_i} \|x - \tilde{x}_i\|^2, \quad (8)$$

where ζ_i denote the perturbation ‘‘power’’ of agent i . For simplicity, we assume all agents are homogeneous, in the sense that $\zeta_i = \zeta > 0$ for all $i \in [n]$; in practice, one need not impose this assumption. Given this utility function, the best response of agents takes the form:

$$b_i(\theta, \tilde{x}_i, \tilde{y}_i; \zeta) = \begin{cases} \tilde{x}_i & \text{if } \tilde{y}_i = +1, \\ \tilde{x}_i + \zeta\theta & \text{if } \tilde{y}_i = -1 \end{cases} \quad (9)$$

where, in our simulations, we fix $\zeta = 0.05$. We reemphasize that our algorithm does not use the value of ζ in any of its computations. For purposes of illustration, we focus on the performance of the following algorithms:

- (A-I) Zeroth-order optimistic-GDA *with* random reshuffling (see Algorithm 1),
- (A-II) Zeroth-order optimistic-GDA *without* random reshuffling (see Appendix C),
- (A-III) Zeroth-order stochastic-GDA *with* random reshuffling (see Appendix C),
- (A-IV) Zeroth-order stochastic-GDA *without* random reshuffling (see Appendix C).

and we evaluate the proposed algorithms and model formulation on two criteria:

- (i) *Suboptimality*: To measure suboptimality, we use the *gap function* $\Delta(\alpha, \theta, \gamma) = L(\alpha, \theta, \gamma^*) - L(\alpha^*, \theta^*, \gamma)$ (Def. 4) where $(\alpha^*, \theta^*, \gamma^*)$ is a solution of the min-max reformulation (7) of the WDRSC problem. If the objective $L(\cdot)$ is convex-concave, $\Delta(\cdot)$ is non-negative, and equals zero at (and only at) saddle points.
- (ii) *Accuracy*: Given a data set $\{(\tilde{x}_i, \tilde{y}_i)\}_{i \in [n]}$, the accuracy of a classifier θ is measured as $\frac{1}{n} \sum_{i \in [n]} \tilde{y}_i \langle b_i(\theta, \tilde{x}_i, \tilde{y}_i; \zeta), \theta \rangle$. Under this criterion we compare the accuracy under different perturbations for different classifiers θ ;

To compute suboptimality, we first compute a true min-max saddle point $(\alpha^*, \theta^*, \gamma^*)$ via a first order gradient based algorithm (namely, GDA). All experiments were run using Python 3.7 on a standard MacBook Pro laptop (2.6 GHz Intel Core i7 and 16 GB of RAM).

6.2 Results

Simulation results presented in Figure (1a)-(1b) show that our proposed algorithm (i.e. (A-I)) outperforms algorithms without reshuffling (i.e. (A-II) and (A-IV)). However, its performance resembles that of zeroth-order stochastic-GDA with random reshuffling. More experimental studies need to be conducted to more conclusively determine whether (A-I) outperforms (A-III), or vice versa. In fact, there has been no theoretical investigations even for the *first order* stochastic-GDA algorithm with random reshuffling; this is an interesting future direction to explore.

In Figure 1, we also compare the robustness of the classifier obtained by using Algorithm (A-I) with that obtained from prior work on solving problems of strategic classification trained with $\zeta = 0.05$ (referred as *LogReg SC* in Figure 1). As expected, due to the formulation, the performance of the classifier obtained via (A-I) degrades gracefully even when subject to large perturbations, while the performance of existing approaches to strategic classification degrades rapidly. Further numerical results on synthetically generated and real world datasets are given in Appendix C.

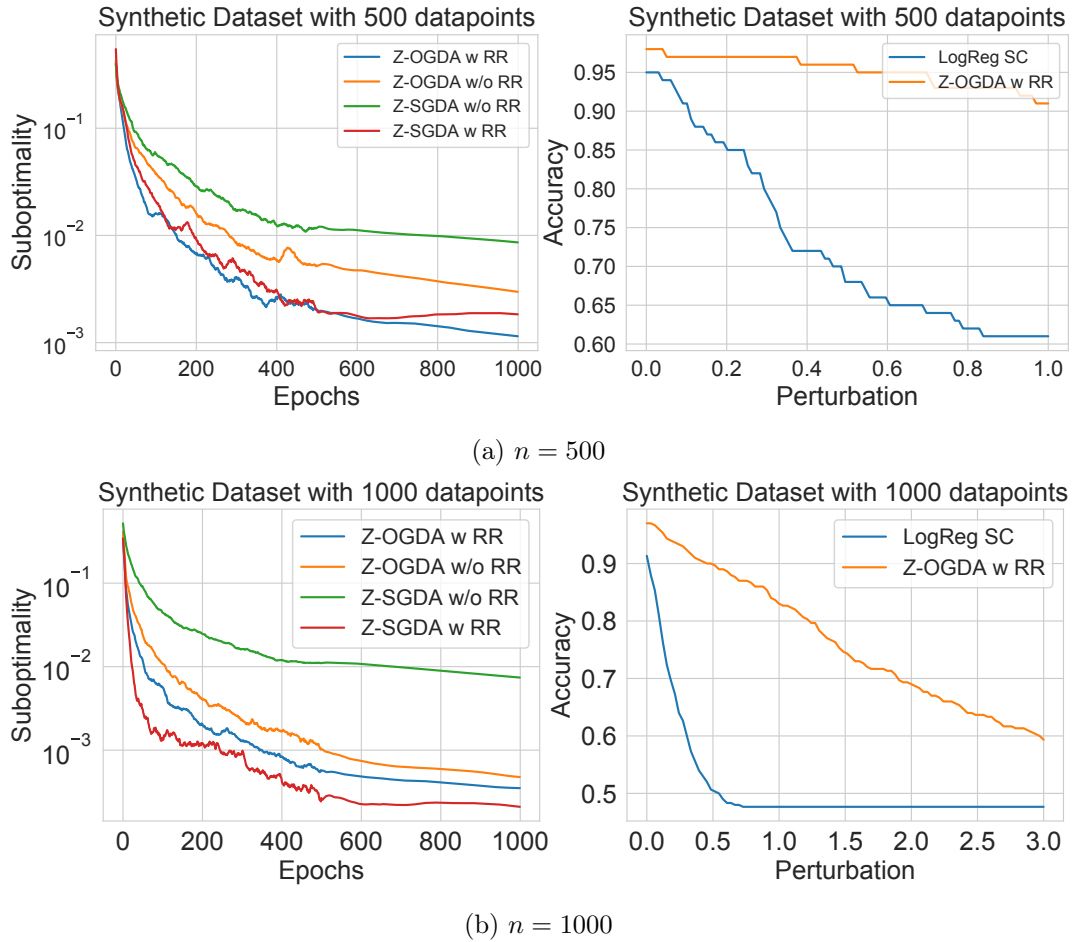


Figure 1: Experimental results for a synthetic dataset with $n = 500$ and $n = 1000$. (Left panes of (1a), (1b)) Suboptimality iterates generated by the four algorithms (A-I), (A-II), (A-III), (A-IV), respectively denoted as *Z-OGDA w RR*, *Z-OGDA w/o RR*, *Z-SGDA w RR*, *Z-SGDA w/o RR*. (Right panes of (1a), (1b)) Comparison between decay in accuracy of strategic classification with logistic regression (trained with $\zeta = 0.05$) and Alg. (A-I) with change in perturbation.

7 Conclusion

This paper presents the first (to our knowledge) non-asymptotic convergence rates for a gradient-free stochastic min-max optimization algorithm with random reshuffling. Our theoretical results, established for smooth convex-concave min-max objectives, do not require any additional, restrictive structural assumptions to hold. As a concrete application, we reformulate a distributionally robust strategic classification problem as a convex-concave min-max optimization problem that can be iteratively solved using our method. Empirical results on synthetic and real datasets demonstrate the efficiency and effectiveness of our algorithm, as well as its robustness against adversarial distributional shifts and strategic behavior of the data sources. Immediate directions for future work include establishing convergence results for the random-reshuffling based Stochastic Gradient Descent Ascent (SGDA-RR) algorithm, as well as performing more extensive experimental studies to better understand the empirical performance of our algorithm.

Acknowledgements

Research supported by NSF under grant DMS 2013985 “THEORINet: Transferable, Hierarchical, Expressive, Optimal, Robust and Interpretable Networks”. Chinmay Maheshwari, Chih-Yuan Chiu and S. Shankar Sastry were also supported in part by U.S. Office of Naval Research MURI grant N00014-16-1- 2710.

References

- [AH17] Charles Audet and Warren Hare. *Derivative-free and blackbox optimization*. Springer, 01 2017.
- [Bag05] J. A. Bagnell. Robust supervised learning. In *AAAI*, volume 2, pages 714–719, 2005.
- [BBC10] Dimitris Bertsimas, David Brown, and Constantine Caramanis. Theory and applications of robust optimization. *SIAM Review*, 53, 10 2010.
- [BHK20] Gavin Brown, Shlomi Hod, and Iden Kalemaj. Performative prediction in a stateful world. *arXiv preprint arXiv:2011.03885*, 2020.
- [BLM18] Mario Bravo, David Leslie, and Panayotis Mertikopoulos. Bandit learning in concave N-person games. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS’18, page 5666–5676, Red Hook, NY, USA, 2018.
- [Bot09] L. Bottou. Curiously fast convergence of some stochastic gradient descent algorithms. In *Proceedings of the Symposium on Learning and Data Science, Paris*, 2009.
- [BSG20] Aleksandr Beznosikov, Abdurakhmon Sadiev, and Alexander Gasnikov. Gradient-free methods with inexact oracle for convex-concave stochastic saddle-point problem. *ArXiv*, 2020.
- [BTEGN09] Aharon Ben-Tal, Laurent El Ghaoui, and Arkadi Nemirovski. *Robust optimization*. Princeton university press, 2009.
- [BTHW⁺13] A. Ben-Tal, D. D. Hertog, A. D. Waegenaere, B. Melenberg, and G. Rennen. Robust solutions of optimization problems affected by uncertain probabilities. *Manag. Sci.*, 59:341–357, 2013.

- [CSSL09] Joaquin Candela, Masashi Sugiyama, Anton Schwaighofer, and Neil D. Lawrence. *Dataset Shift in Machine Learning*. The MIT Press, 2009.
- [CZS⁺17] P. Chen, Huan Zhang, Yash Sharma, Jinfeng Yi, and Cho-Jui Hsieh. ZOO: Zeroth-Order Optimization-based Black-box Attacks to Deep Neural Networks Without Training Substitute Models. *Proceedings of the 10th ACM Workshop on Artificial Intelligence and Security*, 2017.
- [DRS⁺18] Jinshuo Dong, Aaron Roth, Zachary Schutzman, Bo Waggoner, and Zhiwei Steven Wu. Strategic classification from revealed preferences. In *Proceedings of the 2018 ACM Conference on Economics and Computation, EC '18*, page 55–70, New York, NY, USA, 2018. Association for Computing Machinery.
- [DX20] Dmitriy Drusvyatskiy and Lin Xiao. Stochastic optimization with decision-dependent distributions. *arXiv*, 2020.
- [FKM05] Abraham D. Flaxman, Adam Tauman Kalai, and H. Brendan McMahan. Online convex optimization in the bandit setting: Gradient descent without gradient. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '05*, page 385–394, USA, 2005. Society for Industrial and Applied Mathematics.
- [GJLJ17] Gauthier Gidel, Tony Jebara, and Simon Lacoste-Julien. Frank-wolfe algorithms for saddle point problems. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pages 362–371, Fort Lauderdale, FL, USA, 20–22 Apr 2017. PMLR.
- [GJZ18] Xiang Gao, B. Jiang, and S. Zhang. On the information-adaptive variants of the ADMM: An iteration complexity perspective. *Journal of Scientific Computing*, 76:327–363, 2018.
- [GMT14] Virginie Gabrel, Cécile Murat, and Aurélie Thiele. Recent advances in robust optimization: An overview. *European Journal of Operational Research*, 235(3):471–483, 2014.
- [GOP21] Mert Gürbüzbalaban, A. Ozdaglar, and P. Parrilo. Why random reshuffling beats stochastic gradient descent. *Math. Program.*, 186:49–84, 2021.
- [GPAM⁺14] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27, 2014.
- [GYdH15] Bram L. Gorissen, Ihsan Yanikoglu, and Dick den Hertog. A practical guide to robust optimization. *Omega*, 53(C):124–137, 2015.
- [HÅBB13] Martin Hast, Karl Johan Åström, Bo Bernhardsson, and Stephen Boyd. Pid design by convex-concave optimization. In *2013 European Control Conference (ECC)*, pages 4460–4465. IEEE, 2013.
- [Ham19] Isobel Hamilton. Uber drivers are reportedly colluding to trigger surge prices because they say the company is not paying them enough. *Business Insider*, June 14, 2019.

- [HBT18] Lars Hellemo, Paul I Barton, and Asgeir Tomasgard. Decision-dependent probabilities in stochastic programs with recourse. *Computational Management Science*, 15(3):369–395, 2018.
- [HMPW16] Moritz Hardt, Nimrod Megiddo, Christos Papadimitriou, and Mary Wootters. Strategic classification. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, ITCS ’16, page 111–122, New York, NY, USA, 2016. Association for Computing Machinery.
- [HNSS18] Weihua Hu, Gang Niu, Issei Sato, and Masashi Sugiyama. Does distributionally robust supervised learning give robust classifiers? In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 2029–2037. PMLR, 10–15 Jul 2018.
- [HS19] Jeff Z. HaoChen and S. Sra. Random shuffling beats SGD after finite epochs. In *ICML*, 2019.
- [IEAL18] Andrew Ilyas, Logan Engstrom, Anish Athalye, and Jessy Lin. Black-box adversarial attacks with limited queries and information. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 2137–2146. PMLR, 10–15 Jul 2018.
- [JNJ20] Chi Jin, Praneeth Netrapalli, and Michael I. Jordan. What is local optimality in nonconvex-nonconcave minimax optimization? In *ICML*, 2020.
- [JNN19] Prateek Jain, Dheeraj M. Nagaraj, and Praneeth Netrapalli. SGD without replacement: sharper rates for general smooth convex functions. In *ICML*, 2019.
- [KTS⁺19] M. Khajehnejad, Behzad Tabibian, B. Schölkopf, A. Singla, and M. Gomez-Rodriguez. Optimal Decision Making Under Strategic Behavior. *ArXiv*, abs/1905.09239, 2019.
- [LCK⁺20] Sijia Liu, Pin-Yu Chen, Bhavya Kailkhura, Gaoyuan Zhang, Alfred O. Hero III, and Pramod K. Varshney. A primer on zeroth-order optimization in signal processing and machine learning: Principals, recent advances, and applications. *IEEE Signal Processing Magazine*, 37(5):43–54, 2020.
- [Lee13] John M. Lee. *Introduction to Smooth Manifolds*. Springer Science+Business Media New York, 2013.
- [LLC⁺20] Sijia Liu, Songtao Lu, Xiangyi Chen, Yao Feng, Kaidi Xu, Abdullah Al-Dujaili, Mingyi Hong, and Una-May O’Reilly. Min-max optimization without gradients: Convergence and applications to adversarial ML. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 6282–6293. PMLR, 13–18 Jul 2020.
- [LLW⁺19] Yandong Li, Lijun Li, Liqiang Wang, Tong Zhang, and Boqing Gong. NAT-TACK: learning the distributions of adversarial examples for an improved black-box attack on deep neural networks. In *arXiv*, 2019.
- [MDPZH20] Celestine Mendler-Dünner, Juan Perdomo, Tijana Zrnic, and Moritz Hardt. Stochastic optimization for performative prediction. In *Advances in Neural Information Processing Systems*, volume 33, pages 4929–4939, 2020.

- [MKR21] Konstantin Mishchenko, Ahmed Khaled, and Peter Richtárik. Random reshuffling: Simple analysis with vast improvements, 2021.
- [MMS⁺17] Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *ArXiv*, 06 2017.
- [MN19] Peter McCullagh and John A Nelder. *Generalized Linear Models*. Routledge, 2019.
- [MOP20a] Aryan Mokhtari, A. Ozdaglar, and S. Pattathil. A Unified Analysis of Extra-gradient and Optimistic Gradient Methods for Saddle Point Problems: Proximal Point Approach. In *AISTATS*, 2020.
- [MOP20b] Aryan Mokhtari, A. Ozdaglar, and S. Pattathil. Convergence Rate of $O(1/k)$ for optimistic gradient and extragradient methods in smooth convex-concave saddle point problems. *SIAM J. Optim.*, 30:3230–3251, 2020.
- [MPZ21] John Miller, Juan Perdomo, and Tijana Zrnic. Outside the echo chamber: Optimizing the performative risk. *arXiv preprint arXiv:2102.08570*, 2021.
- [ND16] Hongseok Namkoong and John C Duchi. Stochastic gradient methods for distributionally robust optimization with f-divergences. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29, 2016.
- [Nes14] Yurii Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Springer Publishing Company, Incorporated, 1 edition, 2014.
- [NS17] Yurii Nesterov and Vladimir Spokoiny. Random gradient-free minimization of convex functions. *Found. Comput. Math.*, 17(2):527–566, April 2017.
- [NSH⁺19] Maher Nouiehed, Maziar Sanjabi, Tianjian Huang, J. Lee, and Meisam Razaviyayn. Solving a class of non-convex min-max games using iterative first order methods. In *NeurIPS*, 2019.
- [PZMDH20] Juan Perdomo, Tijana Zrnic, Celestine Mendler-Dünner, and Moritz Hardt. Performative prediction. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 7599–7609. PMLR, 13–18 Jul 2020.
- [RGP20] Shashank Rajput, Anant Gupta, and Dimitris Papailiopoulos. Closing the convergence gap of SGD without replacement. *arXiv e-prints*, page arXiv:2002.10400, February 2020.
- [RM19] Hamed Rahimian and Sanjay Mehrotra. Distributionally robust optimization: A review. In *SIAM*, 08 2019.
- [RR11] Benjamin Recht and Christopher Ré. Parallel stochastic gradient algorithms for large-scale matrix completion. *Mathematical Programming Computation*, 5, 04 2011.
- [Rud76] Walter Rudin. *Principles Of Mathematical Analysis*. McGraw-Hill, Inc., 1976.
- [RW21] Benjamin Recht and Stephen Wright. *Optimization for Data Analysis*. Cambridge University Press, 1 edition, 2021.
- [SAEK15] Soroosh Shafieezadeh-Abadeh, Peyman Mohajerin Esfahani, and D. Kuhn. Distributionally robust logistic regression. In *NeurIPS*, 2015.

- [SBDG21] Abdurakhmon Sadiev, Aleksandr Beznosikov, Pavel Dvurechensky, and Alexander Gasnikov. Zeroth-Order Algorithms for Smooth Saddle-Point Problems. *ArXiv*, 2021.
- [SBKK20] Pier Giuseppe Sessa, Ilija Bogunovic, M. Kamgarpour, and A. Krause. Learning to play sequential games versus unknown opponents. *ArXiv*, abs/2007.05271, 2020.
- [Sha16] O. Shamir. Without-Replacement Sampling for Stochastic Gradient Methods. In *NIPS*, 2016.
- [SKL17] Jacob Steinhardt, Pang Wei Koh, and Percy Liang. Certified defenses for data poisoning attacks. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS’17, page 3520–3532, Red Hook, NY, USA, 2017.
- [Spa97] J. Spall. A one-measurement form of simultaneous perturbation stochastic approximation. *Autom.*, 33:109–112, 1997.
- [SS19] Itay Safran and O. Shamir. How good is SGD with random shuffling? *ArXiv*, abs/1908.00045, 2019.
- [TTC⁺19] Chun-Chen Tu, Pai-Shun Ting, P. Chen, Sijia Liu, Huan Zhang, Jinfeng Yi, Cho-Jui Hsieh, and Shin-Ming Cheng. AutoZOOM: autoencoder-based zeroth order optimization method for attacking black-box neural networks. In *AAAI*, 2019.
- [WBMR20] Z. Wang, K. Balasubramanian, Shiqian Ma, and Meisam Razaviyayn. Zeroth-order algorithms for nonconvex minimax problems with improved complexities. *ArXiv*, abs/2001.07819, 2020.
- [YKH20] Junchi Yang, Negar Kiyavash, and Niao He. Global convergence and variance reduction for a class of nonconvex-nonconcave minimax problems. In *Advances in Neural Information Processing Systems*, volume 33, pages 1153–1165, 2020.
- [YLMJ21] Yaodong Yu, Tianyi Lin, Eric V. Mazumdar, and Michael I. Jordan. Fast distributionally robust learning with variance reduced min-max optimization. *ArXiv*, abs/2104.13326, 2021.
- [You19] Soo Youn. Uber, Lyft drivers coordinate to manipulate surge pricing at Virginia airport over pay concerns: Report. *ABC News*, May 18, 2019.

A Results for the Proof of Theorem 4.1

A.1 Lemmas for Theorem 4.1

First, we list some fundamental facts regarding projections onto convex, compact subsets of an Euclidean space. Below, for any fixed convex, compact subset $\Omega \subset \mathbb{R}^d$, we denote the projection operator onto Ω by $\text{Proj}_\Omega(x) := \operatorname{argmin}_{z \in \Omega} \|x - z\|_2$ for each $x \in \mathbb{R}^d$. Note that $\text{Proj}_\Omega(x)$ is well-defined (i.e., exists and is unique) for each $x \in \mathbb{R}^d$, if $\Omega \subset \mathbb{R}^d$ were convex and compact.

We begin by summarizing some fundamental properties of the projection operator $\text{Proj}_\Omega(\cdot)$.

Proposition A.1. *Let $\Omega \subset \mathbb{R}^d$ be compact and convex, and fix $x, y \in \mathbb{R}^d$ arbitrarily. Then:*

$$\begin{aligned} \|\text{Proj}_\Omega(x) - \text{Proj}_\Omega(y)\|_2^2 &\leq (\text{Proj}_\Omega(x) - \text{Proj}_\Omega(y))^\top (x - y), \\ \|\text{Proj}_\Omega(x) - \text{Proj}_\Omega(y)\|_2 &\leq \|x - y\|_2. \end{aligned}$$

Proof. From [Nes14], Lemma 3.1.4 (see also [RW21], Lemma 7.4), we have:

$$\begin{aligned} (\text{Proj}_\Omega(x) - \text{Proj}_\Omega(y))^\top (x - \text{Proj}_\Omega(x)) &\geq 0, \\ (\text{Proj}_\Omega(y) - \text{Proj}_\Omega(x))^\top (y - \text{Proj}_\Omega(y)) &\geq 0. \end{aligned}$$

Adding the two expressions and rearranging terms, we obtain:

$$\begin{aligned} (\text{Proj}_\Omega(x) - \text{Proj}_\Omega(y))^\top ((x - y) - (\text{Proj}_\Omega(x) - \text{Proj}_\Omega(y))) &\geq 0, \\ \Rightarrow \|\text{Proj}_\Omega(x) - \text{Proj}_\Omega(y)\|_2^2 &\leq (\text{Proj}_\Omega(x) - \text{Proj}_\Omega(y))^\top (x - y), \end{aligned}$$

as given in the first claim. The Cauchy Schwarz inequality then implies:

$$\begin{aligned} \|\text{Proj}_\Omega(x) - \text{Proj}_\Omega(y)\|_2^2 &\leq (\text{Proj}_\Omega(x) - \text{Proj}_\Omega(y))^\top (x - y) \\ &\leq \|\text{Proj}_\Omega(x) - \text{Proj}_\Omega(y)\|_2 \cdot \|x - y\|_2. \end{aligned}$$

If $\text{Proj}_\Omega(x) = \text{Proj}_\Omega(y)$, then the second claim becomes $0 \leq \|x - y\|_2$, which is clearly true. Otherwise, dividing both sides above by $\|\text{Proj}_\Omega(x) - \text{Proj}_\Omega(y)\|_2$ gives the second claim. \square

Lemma A.2. *Let $\Omega \subset \mathbb{R}^d$ be a compact, convex subset of \mathbb{R}^d , and consider the update $z_{k+1} = \text{Proj}_\Omega(z_k - \eta F(z_{k+1}) + \gamma_k)$, where $z_k, z_{k+1}, \gamma_k \in \mathbb{R}^d$. Then, for each $z \in \Omega$:*

$$\begin{aligned} &\langle F(z_{k+1}), z_{k+1} - z \rangle \\ &\leq \frac{1}{2\eta} \|z_k - z\|^2 - \frac{1}{2\eta} \|z_{k+1} - z\|^2 - \frac{1}{2\eta} \|z_{k+1} - z_k\|^2 + \frac{1}{\eta} \langle \gamma_k, z_{k+1} - z \rangle. \end{aligned}$$

Proof. Note that:

$$\begin{aligned} \|z_{k+1} - z\|^2 &= \|z_{k+1} - z_k + z_k - z\|^2 \\ &= \|z_{k+1} - z_k\|^2 + \|z_k - z\|^2 + 2 \langle z_{k+1} - z_k, z_k - z \rangle \\ &= \|z_{k+1} - z_k\|^2 + \|z_k - z\|^2 + 2 \langle z_{k+1} - z_k, z_k - z_{k+1} + z_{k+1} - z \rangle \\ &= \|z_k - z\|^2 - \|z_{k+1} - z_k\|^2 + 2 \langle z_{k+1} - z_k, z_{k+1} - z \rangle \end{aligned}$$

By definition of z_{k+1} , and optimality conditions for the projection operator:

$$\begin{aligned} \langle z_{k+1} - z, z_{k+1} - z_k + \eta F(z_{k+1}) - \gamma_k \rangle &\leq 0, \\ \Rightarrow \langle z_{k+1} - z_k, z_{k+1} - z \rangle &\leq \langle \gamma_k, z_{k+1} - z \rangle - \eta \cdot \langle F(z_{k+1}), z_{k+1} - z \rangle. \end{aligned}$$

Substituting back, we obtain:

$$\begin{aligned}\|z_{k+1} - z\|^2 &= \|z_k - z\|^2 - \|z_{k+1} - z_k\|^2 + 2\langle z_{k+1} - z_k, z_{k+1} - z \rangle \\ &\leq \|z_k - z\|^2 - \|z_{k+1} - z_k\|^2 + 2\langle \gamma_k, z_{k+1} - z \rangle - 2\eta \cdot \langle F(z_{k+1}), z_{k+1} - z \rangle.\end{aligned}$$

Rearranging and dividing by η gives the claim in the lemma. \square

Next, we state the properties of the mean and variance of the zeroth-order gradient estimator defined in Section ?? ([BLM18], Lemma C.1). Below, we define the R -smoothed loss function $L^R : \mathbb{R}^d \rightarrow \mathbb{R}$ by $L^R(u) := \mathbb{E}_{\bar{v} \sim \text{Unif}(B^d)}[L(u + R\bar{v})]$, where \mathcal{S}^{d-1} denotes the $(d-1)$ -dimensional unit sphere in \mathbb{R}^d , B^d denotes the d -dimensional unit open ball in \mathbb{R}^d , and $\text{Unif}(\cdot)$ denotes the continuous uniform distribution over a set. Similarly, we define $L_i^R : \mathbb{R}^d \rightarrow \mathbb{R}$ by $L_i^R(u) := \mathbb{E}_{\bar{v} \sim \text{Unif}(B^d)}[L_i(u + R\bar{v})]$, for each $i \in [n] := \{1, \dots, n\}$. We further define $R \cdot \mathcal{S}^{d-1} := \{Rv : v \in \mathcal{S}^{d-1}\}$ and $R \cdot B^d := \{R\bar{v} : \bar{v} \in B^d\}$. Finally, we use $\text{vol}_d(\cdot)$ to denote the volume of a set in d dimensions.

Proposition A.3. *Let $\hat{F}(u; R, v) = \frac{d}{R} \cdot L(u + Rv)v$ and $F(u) = \nabla L(u)$. Then the following holds:*

$$\mathbb{E}_{v \sim \text{Unif}(\mathcal{S}^{d-1})}[\hat{F}(u; R, v)] = \nabla L^R(u), \quad (10)$$

$$\|\nabla L^R(u) - F(u)\|_2 \leq \ell R, \quad (11)$$

$$\|\hat{F}(u; R, v)\|_2 \leq dG + \frac{dM_L}{R}, \quad (12)$$

$$\|\hat{F}(u; R, v) - F(u)\| \leq \min \left\{ (d+1)G + \frac{dM_L}{R}, \ell R + 2dG + \frac{2dM_L}{R} \right\}. \quad (13)$$

Proof. First, to establish (10), observe that since $L^R(u) = \mathbb{E}_{\bar{v} \sim \text{Unif}(B^d)}[L(u + R\bar{v})]$ and $\hat{F}(u; R, v) = \frac{d}{R} \cdot L(u + Rv)v$ for each $u \in \mathbb{R}^d$, $R > 0$, and $v \in \mathcal{S}^{d-1}$:

$$\begin{aligned}\nabla L^R(u) &= \nabla \mathbb{E}_{\bar{v} \sim \text{Unif}(B^d)}[L(u + R\bar{v})] \\ &= \nabla \mathbb{E}_{\bar{v} \sim \text{Unif}(R \cdot B^d)}[L(u + \bar{v})] \\ &= \frac{1}{\text{vol}_d(R \cdot B^d)} \cdot \nabla \left(\int_{R \cdot B^d} L(u + \bar{v}) d\bar{v} \right) \\ &= \frac{1}{\text{vol}_d(R \cdot B^d)} \cdot \int_{R \cdot \mathcal{S}^{d-1}} L(u + v) \cdot \frac{v}{\|v\|_2} dv, \quad (14) \\ \mathbb{E}_{v \sim \text{Unif}(\mathcal{S}^{d-1})}[\hat{F}(u; R, v)] &= \frac{d}{R} \cdot \mathbb{E}_{v \sim \text{Unif}(\mathcal{S}^{d-1})}[L(u + Rv)v] \\ &= \frac{d}{R} \cdot \mathbb{E}_{v \sim \text{Unif}(R \cdot \mathcal{S}^{d-1})} \left[L(u + v) \cdot \frac{v}{\|v\|_2} \right] \\ &= \frac{d}{R} \cdot \frac{1}{\text{vol}_{d-1}(R \cdot \mathcal{S}^{d-1})} \cdot \int_{R \cdot \mathcal{S}^{d-1}} L(u + v) \cdot \frac{v}{\|v\|_2} dv,\end{aligned}$$

where (14) follows because Stokes' Theorem (see, e.g., Lee, Theorem 16.11 [Lee13]) implies that:

$$\nabla \int_{R \cdot B^d} L(u + \bar{v}) d\bar{v} = \int_{R \cdot \mathcal{S}^{d-1}} L(u + v) \cdot \frac{v}{\|v\|_2} dv.$$

The equality (10) now follows by observing that the surface-area-to-volume ratio of $R \cdot B^d$ is d/R .

Next, to establish (11), we note that:

$$\|\nabla L^R(u) - F(u)\|_2 = \|\nabla \mathbb{E}_{\bar{v} \sim \text{Unif}(B^d)}[L^R(u) - L(u)]\|_2$$

$$\begin{aligned}
&= \frac{1}{\text{vol}_d(B^d)} \cdot \left\| \nabla \left(\int_{B^d} [L(u + R\bar{v}) - L(u)] d\bar{v} \right) \right\|_2 \\
&\leq \frac{1}{\text{vol}_d(B^d)} \cdot \left\| \int_{B^d} [F(u + R\bar{v}) - F(u)] d\bar{v} \right\|_2 \quad (15) \\
&\leq \frac{1}{\text{vol}_d(B^d)} \cdot \int_{B^d} \|F(u + R\bar{v}) - F(u)\|_2 d\bar{v} \\
&\leq \frac{1}{\text{vol}_d(B^d)} \cdot \int_{B^d} \ell R \cdot \|\bar{v}\|_2 d\bar{v} \\
&\leq \ell R,
\end{aligned}$$

where (15) follows by differentiating under the integral sign (see, e.g., Rudin, Theorem 9.42 [Rud76]), and the remaining inequalities follow from the fact that F is ℓ -Lipschitz.

Next, we establish (12) by using the triangle inequality and the M_L -boundedness of $L(\cdot)$ on $\mathcal{X} \times \mathcal{Y}$, and the G -Lipschitzness of $L(\cdot)$:

$$\begin{aligned}
|\hat{F}(u; R, v)| &= \frac{d}{R} |L(u + Rv)| \cdot \|v\|_2 \\
&\leq \frac{d}{R} \cdot (|L(u)| + |L(u + Rv) - L(u)|) \cdot 1 \\
&\leq \frac{d}{R} \cdot (M_L + RG).
\end{aligned}$$

We can then use (12) to establish (13) by observing that:

$$|\hat{F}(u; R, v) - F(u)| \leq |\hat{F}(u; R, v)| + |F(u)| \leq (d+1)G + \frac{dM_L}{R}.$$

and, from (12):

$$\begin{aligned}
&|\hat{F}(u; R, v) - F(u)| \\
&\leq |\hat{F}(u; R, v) - \mathbb{E}_v[\hat{F}(u; R, v)|u]| + |\mathbb{E}_v[\hat{F}(u; R, v)|u] - F(u)| \\
&\leq |\hat{F}(u; R, v) - \mathbb{E}_v[\hat{F}(u; R, v)|u]| + |\nabla L^R(u) - F(u)| \\
&\leq 2 \left(dG + \frac{dM_L}{R} \right) + \ell R
\end{aligned}$$

This concludes the proof. \square

Below, we present technical lemmas that allow us to analyze the convergence rate of the correlated iterates $\{u_i^t\}$ in our random reshuffling-based OGD Algorithm (Alg. 1).

Let $\sigma^0, \dots, \sigma^{t-1}$ denote the permutations drawn from epoch 0 to epoch $t-1$, and let $\{u_i^t(\sigma^t)\}_{1 \leq i \leq n}$ and $\{u_i^t(\tilde{\sigma}^t)\}_{1 \leq i \leq n}$ denote the iterates obtained at epoch t , when the permutations σ^t and $\tilde{\sigma}^t$ are used for the epoch t , respectively. Moreover, let $\mathcal{D}_{i,t}$ denote the distribution of $\{u_i^t(\sigma^t)\}_{1 \leq i \leq n}$ under σ^t , and for $1 \leq r \leq n$ let $\mathcal{D}_{i,t}^{(r)}$ denote the distribution of $\{u_i^t(\sigma^t)\}_{1 \leq i \leq n}$ with σ^t conditioned on the event $\{\sigma_{i-1}^t = r\}$.

We use the p -Wasserstein distance between probability distributions on \mathbb{R}^d , defined below, to characterize the distance between $\mathcal{D}_{i,t}$ and $\mathcal{D}_{i,t}^{(r)}$. This is used in the coupling-based techniques employed to establish non-asymptotic convergence results for our random reshuffling algorithm. Note the difference between the p -Wasserstein distance for probability distributions on \mathbb{R}^d , and the Wasserstein distance on $\mathcal{Z} := \mathbb{R}^d \times \{+1, -1\}$ associated with a metric $c : \mathcal{Z} \times \mathcal{Z} \rightarrow [0, \infty)$, defined in Appendix B.2 (Definition B.1).

Definition A.1 (p -Wasserstein distance between distributions on \mathbb{R}^d). *Let μ, ν be probability distributions over \mathbb{R}^d with finite p -th moments, for some $p \geq 1$, and let $\Pi(\mu, \nu)$*

denote the set of all couplings (joint distributions) between μ and ν . The p -Wasserstein distance between μ and ν , denoted $\mathcal{W}_p(\mu, \nu)$, is defined by:

$$\mathcal{W}_p(\mu, \nu) = \inf_{(X, X') \sim \pi \in \Pi(\mu, \nu)} \left(\mathbb{E}_\pi [\|X - X'\|^p] \right)^{1/p}.$$

The following proposition characterizes the 1-Wasserstein distance as a measure of the gap between Lipschitz functions of random variables.

Proposition A.4 (Kantorovich Duality). *If μ, ν are probability distributions over \mathbb{R}^d with finite second moments, then:*

$$\mathcal{W}_1(\mu, \nu) = \sup_{g \in \text{Lip}(1)} \mathbb{E}_{X \sim \mu}[g(X)] - \mathbb{E}_{Y \sim \nu}[g(Y)],$$

where $\text{Lip}(1) := \{g : \mathbb{R}^d \rightarrow \mathbb{R} : g \text{ is 1-Lipschitz}\}$.

Using [YLMJ21, Lemma C.2], we now bound the difference between the unbiased gap $\mathbb{E}[\Delta(u_i^t)]$ and the biased gap $\mathbb{E}[L_{\sigma_i^t}(x_{i+1}^t, y^*) - L_{\sigma_i^t}(x^*, y_{i+1}^t)]$ using the Wasserstein metric.

Lemma A.5. *Let $u^* := (x^*, y^*) \in \mathbb{R}^{d_x} \times \mathbb{R}^{d_y} = \mathbb{R}^d$ denote a saddle point of the min-max optimization problem (2). Then, for each $t \in [T]$ and $i \in [n]$, the iterates $\{u_i^t\} = \{(x_i^t, y_i^t)\}$ of the OGD-RR algorithm satisfy:*

$$\left| \mathbb{E}[\Delta(u_{i+1}^t)] - \mathbb{E}[L_{\sigma_i^t}(x_{i+1}^t, y^*) - L_{\sigma_i^t}(x^*, y_{i+1}^t)] \right| \leq \frac{G}{n} \sum_{r=1}^n \mathcal{W}_2(\mathcal{D}_{i+1,t}, \mathcal{D}_{i+1,t}^r)$$

Proof. Since σ^t and $\tilde{\sigma}^t$ are independently generated permutations of $[n]$, the iterates $\{u_i^t\}_{1 \leq i \leq n} = \{u_i^t(\sigma^t)\}_{1 \leq i \leq n}$ and $\{u_i^t(\tilde{\sigma}^t)\}_{1 \leq i \leq n}$ are i.i.d. Thus, we have:

$$\mathbb{E}[\Delta(u_{i+1}^t)] = \mathbb{E}[L_{\sigma_i^t}(x_{i+1}^t(\tilde{\sigma}^t), y^*) - L_{\sigma_i^t}(x^*, y_{i+1}^t(\tilde{\sigma}^t))],$$

and thus:

$$\begin{aligned} & \left| \mathbb{E}[\Delta(u_{i+1}^t)] - \mathbb{E}[L_{\sigma_i^t}(x_{i+1}^t, y^*) - L_{\sigma_i^t}(x^*, y_{i+1}^t)] \right| \\ &= \left| \mathbb{E}[L_{\sigma_i^t}(x_{i+1}^t(\tilde{\sigma}^t), y^*) - L_{\sigma_i^t}(x^*, y_{i+1}^t(\tilde{\sigma}^t))] - \mathbb{E}[L_{\sigma_i^t}(x_{i+1}^t, y^*) - L_{\sigma_i^t}(x^*, y_{i+1}^t)] \right| \\ &= \left| \frac{1}{n} \sum_{r=1}^n \mathbb{E}[L_r(x_{i+1}^t(\tilde{\sigma}^t), y^*) - L_r(x^*, y_{i+1}^t(\tilde{\sigma}^t))] \right. \end{aligned} \quad (16)$$

$$\left. - \frac{1}{n} \sum_{r=1}^n \mathbb{E}[L_r(x_{i+1}^t, y^*) - L_r(x^*, y_{i+1}^t) | \sigma_i^t = r] \right|$$

$$\leq \frac{1}{n} \sum_{r=1}^n \left| \mathbb{E}[L_r(x_{i+1}^t(\tilde{\sigma}^t), y^*) - L_r(x^*, y_{i+1}^t(\tilde{\sigma}^t))] - \mathbb{E}[L_r(x_{i+1}^t, y^*) - L_r(x^*, y_{i+1}^t) | \sigma_i^t = r] \right|$$

$$\leq \frac{1}{n} \sum_{r=1}^n \sup_{g \in \text{Lip}(G)} \left(\mathbb{E}[g(x_{i+1}^t(\tilde{\sigma}^t), y_{i+1}^t(\tilde{\sigma}^t))] - \mathbb{E}[g(x_{i+1}^t, y_{i+1}^t) | \sigma_i^t = r] \right) \quad (17)$$

$$\leq \frac{1}{n} \sum_{r=1}^n G \cdot \mathcal{W}_1(\mathcal{D}_{i+1,t}, \mathcal{D}_{i+1,t}^{(r)}) \quad (18)$$

$$\leq \frac{1}{n} \sum_{r=1}^n G \cdot \mathcal{W}_2(\mathcal{D}_{i+1,t}, \mathcal{D}_{i+1,t}^{(r)}), \quad (19)$$

where (16) follows by properties of the conditional expectation on $\{\sigma_i^t = r\}$ and the fact that σ^t and $\tilde{\sigma}^t$ are independent, (17) follows from the fact that L is Lipschitz, (18) follows from Proposition A.4, and (19) follows from the fact that $\mathcal{W}_1(\mu, \nu) \leq \mathcal{W}_2(\mu, \nu)$ for any two probability distributions μ, ν . \square

The next lemma bounds the difference in the iterates $\{u_i^t(\sigma^t)\}$ and $\{u_i^t(\tilde{\sigma}^t)\}$ (assuming, as before, that $\sigma^0, \dots, \sigma^{t-1}$ were fixed and identical for both sequences.)

Lemma A.6. *Denote, with a slight abuse of notation, $u_i^t := u_i^t(\sigma^t)$ and $\tilde{u}_i^t := u_i^t(\tilde{\sigma}^t)$. Then:*

$$\|u_{i+1}^t - \tilde{u}_{i+1}^t\|_2 \leq \left(6nd + 14n + 2 \cdot \sum_{i=1}^n \mathbf{1}\{\sigma_i^t \neq \tilde{\sigma}_i^t\} \right) G \cdot \eta^t + 6ndM_L \cdot \frac{\eta^t}{R^t}.$$

Proof. Our proof strategy is to bound the differences between zeroth-order and first-order OGD updates, and between the OGD and proximal point updates. To this end, we define:

$$\begin{aligned} u_{i+1}^t &= \text{Proj}_{\mathcal{X} \times \mathcal{Y}} \left(u_i^t - \eta^t \hat{F}_{\sigma_i^t}(u_i^t; R^t, v_i^t) - \eta^t \hat{F}_{\sigma_{i-1}^t}(u_i^t; R^t, v_i^t) + \eta^t \hat{F}_{\sigma_{i-1}^t}(u_{i-1}^t; R^t, v_{i-1}^t) \right), \\ \tilde{u}_{i+1}^t &= \text{Proj}_{\mathcal{X} \times \mathcal{Y}} \left(\tilde{u}_i^t - \eta^t \hat{F}_{\tilde{\sigma}_i^t}(\tilde{u}_i^t; R^t, v_i^t) - \eta^t \hat{F}_{\tilde{\sigma}_{i-1}^t}(\tilde{u}_i^t; R^t, v_i^t) + \eta^t \hat{F}_{\tilde{\sigma}_{i-1}^t}(\tilde{u}_{i-1}^t; R^t, v_{i-1}^t) \right), \\ v_{i+1}^t &= \text{Proj}_{\mathcal{X} \times \mathcal{Y}} \left(u_i^t - \eta^t F_{\sigma_i^t}(u_i^t) - \eta^t F_{\sigma_{i-1}^t}(u_i^t) + \eta^t F_{\sigma_{i-1}^t}(u_{i-1}^t) \right), \\ \tilde{v}_{i+1}^t &= \text{Proj}_{\mathcal{X} \times \mathcal{Y}} \left(\tilde{u}_i^t - \eta^t F_{\tilde{\sigma}_i^t}(\tilde{u}_i^t) - \eta^t F_{\tilde{\sigma}_{i-1}^t}(\tilde{u}_i^t) + \eta^t F_{\tilde{\sigma}_{i-1}^t}(\tilde{u}_{i-1}^t) \right), \\ w_{i+1}^t &= \text{Proj}_{\mathcal{X} \times \mathcal{Y}} \left(u_i^t - \eta^t F_{\sigma_i^t}(w_{i+1}^t) \right), \\ \tilde{w}_{i+1}^t &= \text{Proj}_{\mathcal{X} \times \mathcal{Y}} \left(\tilde{u}_i^t - \eta^t F_{\tilde{\sigma}_i^t}(\tilde{w}_{i+1}^t) \right). \end{aligned}$$

By the triangle inequality:

$$\begin{aligned} \|u_{i+1}^t - \tilde{u}_{i+1}^t\|_2 &\leq \|u_{i+1}^t - v_{i+1}^t\|_2 + \|v_{i+1}^t - w_{i+1}^t\|_2 + \|w_{i+1}^t - \tilde{w}_{i+1}^t\|_2 \\ &\quad + \|\tilde{w}_{i+1}^t - \tilde{v}_{i+1}^t\|_2 + \|\tilde{v}_{i+1}^t - \tilde{u}_{i+1}^t\|_2. \end{aligned} \quad (20)$$

Observe that bounding the fourth term is equivalent to bounding the second term, and bounding the fifth term is equivalent to bounding the first term.

To bound the first term on the right hand side, we use Proposition A.3 to conclude that:

$$\begin{aligned} \|u_{i+1}^t - v_{i+1}^t\|_2 &\leq \eta^t \cdot \|\hat{F}_{\sigma_i^t}(u_i^t; R^t, v_i^t) - F_{\sigma_i^t}(u_i^t)\| + \eta^t \cdot \|\hat{F}_{\sigma_{i-1}^t}(u_i^t; R^t, v_i^t) - F_{\sigma_{i-1}^t}(u_i^t)\| \\ &\quad + \eta^t \cdot \|\hat{F}_{\sigma_{i-1}^t}(u_{i-1}^t; R^t, v_{i-1}^t) - F_{\sigma_{i-1}^t}(u_{i-1}^t)\| \\ &\leq 3(d+1)G\eta^t + 3dM_L \cdot \frac{\eta^t}{R^t} \end{aligned} \quad (21)$$

For the second term, we use the G -Lipschitzness of L_r , for each $r \in [n]$ to conclude that:

$$\begin{aligned} \|v_{i+1}^t - w_{i+1}^t\|_2 &\leq \eta^t \cdot |F_{\sigma_i^t}(u_i^t)| + \eta^t \cdot |F_{\sigma_{i-1}^t}(u_i^t)| + \eta^t \cdot |F_{\sigma_{i-1}^t}(u_{i-1}^t)| + \eta^t \cdot |F_{\sigma_i^t}(w_{i+1}^t)| \\ &\leq 4G \cdot \eta^t. \end{aligned} \quad (22)$$

For the third term, we observe that if $\sigma_i^t \neq \tilde{\sigma}_i^t$, then:

$$\begin{aligned} \|w_{i+1}^t - \tilde{w}_{i+1}^t\|_2 &\leq \|u_i^t - \tilde{u}_i^t\|_2 + \eta^t \cdot \|F_{\sigma_i^t}(w_{i+1}^t) - F_{\tilde{\sigma}_i^t}(\tilde{w}_{i+1}^t)\|_2 \\ &\leq \|u_i^t - \tilde{u}_i^t\|_2 + 2G \cdot \eta^t. \end{aligned} \quad (23)$$

On the other hand, if $\sigma_i^t = \tilde{\sigma}_i^t$, then:

$$\begin{aligned} w_{i+1}^t &= \text{Proj}_{\mathcal{X} \times \mathcal{Y}} \left(u_i^t - \eta^t F_{\sigma_i^t}(w_{i+1}^t) \right), \\ \tilde{w}_{i+1}^t &= \text{Proj}_{\mathcal{X} \times \mathcal{Y}} \left(\tilde{u}_i^t - \eta^t F_{\sigma_i^t}(\tilde{w}_{i+1}^t) \right), \end{aligned}$$

so we have:

$$\|w_{i+1}^t - \tilde{w}_{i+1}^t\|_2^2$$

$$\leq (w_{i+1}^t - \tilde{w}_{i+1}^t)^\top ((u_i^t - \eta \cdot F_{\sigma_i^t}(w_{i+1}^t)) - (\tilde{u}_i^t - \eta \cdot F_{\sigma_i^t}(\tilde{w}_{i+1}^t))) \quad (24)$$

$$= (w_{i+1}^t - \tilde{w}_{i+1}^t)^\top (u_i^t - \tilde{u}_i^t) - \eta (w_{i+1}^t - \tilde{w}_{i+1}^t)^\top (F_{\sigma_i^t}(w_{i+1}^t) - F_{\sigma_i^t}(\tilde{w}_{i+1}^t))$$

$$\leq (w_{i+1}^t - \tilde{w}_{i+1}^t)^\top (u_i^t - \tilde{u}_i^t) \quad (25)$$

$$\leq \|w_{i+1}^t - \tilde{w}_{i+1}^t\|_2 \cdot \|u_i^t - \tilde{u}_i^t\|_2, \quad (26)$$

so $\|w_{i+1}^t - \tilde{w}_{i+1}^t\|_2 \leq \|u_i^t - \tilde{u}_i^t\|_2$. Here, (24) follows from the definitions of w_{i+1}^t and \tilde{w}_{i+1}^t , as well as Proposition A.1, while (25) holds because the monotonicity of F_i , for each $i \in [n]$, implies that $(w_{i+1}^t - \tilde{w}_{i+1}^t)^\top (F_{\sigma_i^t}(w_{i+1}^t) - F_{\sigma_i^t}(\tilde{w}_{i+1}^t)) \geq 0$. Putting together (21), (22), (23), (26), we have:

$$\begin{aligned} \|u_{i+1}^t - \tilde{u}_{i+1}^t\|_2 &\leq \|u_i^t - \tilde{u}_i^t\|_2 + (6d + 14)G \cdot \eta^t + 6dM_L \cdot \frac{\eta^t}{R^t} \\ &\quad + 2G \cdot \mathbf{1}\{\sigma_i^t \neq \tilde{\sigma}_i^t\} \cdot \eta^t, \end{aligned}$$

where the indicator $\mathbf{1}(A)$ returns 1 if the given event A occurs, and 0 otherwise.

Since $u_0^t = \tilde{u}_0^t$, we can iteratively apply the above inequality to obtain that, for any and epoch t and $i \in [n]$:

$$\|u_{i+1}^t - \tilde{u}_{i+1}^t\|_2 \leq (6d + 14)nG \cdot \eta^t + 6ndM_L \cdot \frac{\eta^t}{R^t} + 2\eta_t G \cdot \sum_{i=1}^n \mathbf{1}\{\sigma_i^t \neq \tilde{\sigma}_i^t\},$$

□

Remark. In the theorems and lemmas below, we will be concerned with the case where σ^t and $\tilde{\sigma}^t$ have the following specific relationship. Let \mathcal{R}_n denote the set of all random permutations over the set $[n]$. For each $l, m \in [n]$, let $S_{l,m} : \mathcal{R}_n \rightarrow \mathcal{R}_n$ denote the map that swaps, for each input permutation σ , the l -th and m -th entries. For each $r, i \in [n]$, define the map $\omega_{r,i} : \mathcal{R}_n \rightarrow \mathcal{R}_n$ as follows:

$$\omega_{r,i}(\sigma) = \begin{cases} \sigma, & \text{if } \sigma_{i-1} = r, \\ S_{i-1,j}(\sigma), & \text{if } \sigma_j = r \text{ and } j \neq i-1. \end{cases}$$

Intuitively, $\omega_{r,i}$ performs a single swap such that the $(i-1)$ -th position of the permutation is r . Clearly, if σ^t is a random permutation (i.e., selected from a uniform distribution over \mathcal{R}_n), then $\omega_{r,i}(\sigma^t)$ has the same distribution as $\sigma^t | (\sigma_{i-1}^t = r)$. Based on this construction, we have $u_i(\sigma^t) \sim \mathcal{D}_{i,t}$ and $u_i(\omega_{r,i}(\sigma^t)) \sim \mathcal{D}_{i,t}^{(r)}$. This gives a coupling between $\mathcal{D}_{s,t}$ and $\mathcal{D}_{s,t}^{(r)}$. Since σ^t and $\tilde{\sigma}^t$ differ by at most two entries, by iteratively applying Lemma A.6, we have:

$$\begin{aligned} \|u_{i+1}^t - \tilde{u}_{i+1}^t\|_2 &\leq n \left((6d + 14)G \cdot \eta^t + 6dM_L \cdot \frac{\eta^t}{R^t} \right) + 4G \cdot \eta^t \\ &= (6nd + 14n + 4)G \cdot \eta^t + 6ndM_L \cdot \frac{\eta^t}{R^t}, \end{aligned}$$

as claimed.

Lemma A.7. *If $\eta^t \leq 1/(2\ell)$ for each $t \in \{0, 1, \dots, T-1\}$, the iterates $\{u_i^t\} = \{(x_i^t, y_i^t)\}$ of the OGD-RR algorithm satisfy, for each $u \in \mathcal{X} \times \mathcal{Y}$:*

$$\begin{aligned} &2\eta^t \cdot \mathbb{E} \left[\left\langle F_{\sigma_i^t}(u_{i+1}^t), u_{i+1}^t - u \right\rangle \right] \\ &\leq \mathbb{E} [\|u_i^t - u\|_2^2] - \mathbb{E} [\|u_{i+1}^t - u\|_2^2] - \frac{1}{2} \mathbb{E} [\|u_{i+1}^t - u_i^t\|_2^2] + \frac{1}{2} \mathbb{E} [\|u_i^t - u_{i-1}^t\|_2^2] \\ &\quad + 2\eta^t \cdot \mathbb{E} \left[\left\langle F_{\sigma_i^t}(u_{i+1}^t) - F_{\sigma_i^t}(u_i^t), u_{i+1}^t - u \right\rangle \right] \end{aligned}$$

$$\begin{aligned}
& -2\eta^t \cdot \mathbb{E} \left[\left\langle F_{\sigma_{i-1}^t}(u_i^t) - F_{\sigma_{i-1}^t}(u_{i-1}^t), u_i^t - u \right\rangle \right] \\
& + 6C_1 \cdot \left(\eta^t R^t + (\eta^t)^2 R^t + (\eta^t)^2 + \frac{(\eta^t)^2}{R^t} + \frac{(\eta^t)^2}{(R^t)^2} \right),
\end{aligned}$$

where $C_1 := d^2 \max \{6G\ell D, 18G^2 + 6M_L\ell D, 30M_L G, 12M_L^2\}$ is a constant independent of the sequences $\{\eta^t\}$ and $\{R^t\}$.

Proof. The iterates of the OGD-RR algorithm are given by:

$$\begin{aligned}
u_{i+1}^t &= \text{Proj}_{\mathcal{X} \times \mathcal{Y}} \left(u_i^t - \eta^t \hat{F}_{\sigma_i^t}(u_i^t; R^t, v_i^t) - \eta^t \hat{F}_{\sigma_{i-1}^t}(u_i^t; R^t, v_i^t) \right. \\
& \quad \left. - \eta^t \hat{F}_{\sigma_{i-1}^t}(u_{i-1}^t; R^t, v_{i-1}^t) \right) \\
&= \text{Proj}_{\mathcal{X} \times \mathcal{Y}} \left(u_i^t - \eta^t F_{\sigma_i^t}(u_{i+1}^t) + \eta^t (\gamma_i^t + E_{i,1}^t + E_{i,2}^t + E_{i,3}^t) \right), \tag{27}
\end{aligned}$$

where we have defined:

$$\begin{aligned}
\gamma_i^t &:= F_{\sigma_i^t}(u_{i+1}^t) - F_{\sigma_i^t}(u_i^t) - F_{\sigma_{i-1}^t}(u_i^t) + F_{\sigma_{i-1}^t}(u_{i-1}^t), \\
E_{i,1}^t &:= F_{\sigma_i^t}(u_i^t) - \hat{F}_{\sigma_i^t}(u_i^t; R^t, v_i^t), \\
E_{i,2}^t &:= F_{\sigma_{i-1}^t}(u_i^t) - \hat{F}_{\sigma_{i-1}^t}(u_i^t; R^t, v_i^t), \\
E_{i,3}^t &:= F_{\sigma_{i-1}^t}(u_{i-1}^t) - \hat{F}_{\sigma_{i-1}^t}(u_{i-1}^t; R^t, v_{i-1}^t).
\end{aligned}$$

First, by applying Lemma A.2 we have:

$$\begin{aligned}
& 2\eta^t \cdot \mathbb{E} \left[\left\langle F_{\sigma_i^t}(u_{i+1}^t), u_{i+1}^t - u \right\rangle \right] \tag{28} \\
& \leq \mathbb{E} [\|u_i^t - u\|_2^2] - \mathbb{E} [\|u_{i+1}^t - u\|_2^2] - \mathbb{E} [\|u_{i+1}^t - u_i^t\|_2^2] \\
& \quad + 2\eta^t \cdot \mathbb{E} \left[\left\langle \gamma_i^t, u_{i+1}^t - u \right\rangle \right] + \sum_{k=1}^3 2\eta^t \cdot \mathbb{E} \left[\left\langle E_{i,k}^t, u_{i+1}^t - u \right\rangle \right].
\end{aligned}$$

Below, we proceed to bound the inner product terms on the right-hand-side of (28). First, we bound $\langle \gamma_i^t, u_{i+1}^t - u \rangle$:

$$\begin{aligned}
\left\langle \gamma_i^t, u_{i+1}^t - u \right\rangle &= \left\langle F_{\sigma_i^t}(u_{i+1}^t) - F_{\sigma_i^t}(u_i^t), u_{i+1}^t - u \right\rangle \\
& \quad - \left\langle F_{\sigma_{i-1}^t}(u_i^t) - F_{\sigma_{i-1}^t}(u_{i-1}^t), u_{i+1}^t - u \right\rangle \\
&= \left\langle F_{\sigma_i^t}(u_{i+1}^t) - F_{\sigma_i^t}(u_i^t), u_{i+1}^t - u \right\rangle \\
& \quad - \left\langle F_{\sigma_{i-1}^t}(u_i^t) - F_{\sigma_{i-1}^t}(u_{i-1}^t), u_i^t - u \right\rangle \\
& \quad - \left\langle F_{\sigma_{i-1}^t}(u_i^t) - F_{\sigma_{i-1}^t}(u_{i-1}^t), u_{i+1}^t - u_i^t \right\rangle \\
&\leq \left\langle F_{\sigma_i^t}(u_{i+1}^t) - F_{\sigma_i^t}(u_i^t), u_{i+1}^t - u \right\rangle \tag{29} \\
& \quad - \left\langle F_{\sigma_{i-1}^t}(u_i^t) - F_{\sigma_{i-1}^t}(u_{i-1}^t), u_i^t - u \right\rangle \\
& \quad + \frac{1}{2} \ell \cdot \|u_i^t - u_{i-1}^t\|_2^2 + \frac{1}{2} \ell \cdot \|u_{i+1}^t - u_i^t\|_2^2.
\end{aligned}$$

Note that the final inequality follows by applying Young's inequality, and noting that F is ℓ -Lipschitz. Next, we bound $\langle E_{i,1}^t, u_{i+1}^t - u \rangle$:

$$\mathbb{E} \left[\left\langle E_{i,1}^t, u_{i+1}^t - u \right\rangle \right]$$

$$\begin{aligned}
&= \mathbb{E} \left[\left\langle F_{\sigma_i^t}(u_i^t) - \hat{F}_{\sigma_i^t}(u_i^t, R^t, v_i^t), u_{i+1}^t - u \right\rangle \right] \\
&= \mathbb{E} \left[\left\langle F_{\sigma_i^t}(u_i^t) - \nabla L_{\sigma_i^t}^{R^t}(u_i^t), u_{i+1}^t - u \right\rangle \right] \\
&\quad + \mathbb{E} \left[\left\langle \mathbb{E}[\hat{F}_{\sigma_i^t}(u_i^t; R^t, v_i^t | u_i^t)] - \hat{F}_{\sigma_i^t}(u_i^t, R^t, v_i^t), u_{i+1}^t - u \right\rangle \right] \\
&= \mathbb{E} \left[\left\langle F_{\sigma_i^t}(u_i^t) - \nabla L_{\sigma_i^t}^{R^t}(u_i^t), u_{i+1}^t - u \right\rangle \right] \\
&\quad + \mathbb{E} \left[\left\langle \mathbb{E}_v[\hat{F}_{\sigma_i^t}(u_i^t; R^t, v | u_i^t)] - \hat{F}_{\sigma_i^t}(u_i^t, R^t, v_i^t), u_i^t - u \right\rangle \right] \\
&\quad + \mathbb{E} \left[\left\langle \mathbb{E}_v[\hat{F}_{\sigma_i^t}(u_i^t; R^t, v | u_i^t)] - \hat{F}_{\sigma_i^t}(u_i^t, R^t, v_i^t), u_{i+1}^t - u_i^t \right\rangle \right],
\end{aligned} \tag{30}$$

where the first equality above follows by applying Proposition A.3, (10), and we have used the shorthand $\mathbb{E}_v := \mathbb{E}_{v \sim \text{Unif}(\mathcal{S}^{d-1})}$. (Recall that $L^R(u) := \mathbb{E}_{v \sim \text{Unif}(\mathcal{S}^{d-1})}[L(u + Rv)]$) Next, we upper bound each of the three quantities in (30). First, by Proposition A.3, (11), we have:

$$\begin{aligned}
&\mathbb{E} \left[\left\langle F_{\sigma_i^t}(u_i^t) - \nabla L_{\sigma_i^t}^{R^t}(u_i^t), u_{i+1}^t - u \right\rangle \right] \\
&\leq \mathbb{E} \left[\|F_{\sigma_i^t}(u_i^t) - \nabla L_{\sigma_i^t}^{R^t}(u_i^t)\|_2 \cdot \|u_{i+1}^t - u\|_2 \right] \\
&\leq \ell D \cdot R^t,
\end{aligned} \tag{31}$$

with $C_1 > 0$ as given in Lemma A.7. Meanwhile, the law of iterated expectations can be used to bound the second quantity:

$$\begin{aligned}
&\mathbb{E} \left[\left\langle \mathbb{E}_v[\hat{F}_{\sigma_i^t}(u_i^t; R^t, v | u_i^t)] - \hat{F}_{\sigma_i^t}(u_i^t, R^t, v_i^t), u_i^t - u \right\rangle \right] \\
&= \mathbb{E} \left[\mathbb{E}_v \left[\left\langle \hat{F}_{\sigma_i^t}(u_i^t, R^t, v_i^t), u_i^t - u \right\rangle | u_i^t \right] \right] - \mathbb{E} \left[\left\langle \hat{F}_{\sigma_i^t}(u_i^t, R^t, v_i^t), u_i^t - u \right\rangle \right] \\
&= 0,
\end{aligned} \tag{32}$$

and we can upper-bound the third quantity as shown below. By using the compactness of $\mathcal{X} \times \mathcal{Y}$ and the continuity of L , we have:

$$\begin{aligned}
&\mathbb{E} \left[\left\langle \mathbb{E}_v[\hat{F}_{\sigma_i^t}(u_i^t; R^t, v | u_i^t)] - \hat{F}_{\sigma_i^t}(u_i^t, R^t, v_i^t), u_{i+1}^t - u_i^t \right\rangle \right] \\
&\leq \left(\|\mathbb{E}_v[\hat{F}_{\sigma_i^t}(u_i^t; R^t, v | u_i^t)]\|_2 + \|\hat{F}_{\sigma_i^t}(u_i^t, R^t, v_i^t)\| \right) \cdot \|u_{i+1}^t - u_i^t\|_2 \\
&\leq 2 \cdot \frac{d}{R^t} \cdot \sup_{\substack{u \in \mathcal{X} \times \mathcal{Y} \\ v \sim \text{Unif}(\mathcal{S}^{d-1})}} |L(u_i^t + R^t v)| \cdot \|u_{i+1}^t - u_i^t\|_2, \\
&\leq 2 \cdot \frac{d}{R^t} \cdot (M_L + R^t G) \cdot \|u_{i+1}^t - u_i^t\|_2,
\end{aligned} \tag{33}$$

and using (31) and the bound for each $\|\hat{F}_{\sigma_i^t}\|_2$ given in (33), we have:

$$\begin{aligned}
&\|u_{i+1}^t - u_i^t\|_2 \\
&\leq \eta^t \cdot \|\hat{F}_{\sigma_i^t}(u_i^t; R^t, v_i^t) + \hat{F}_{\sigma_{i-1}^t}(u_i^t; R^t, v_i^t) - \hat{F}_{\sigma_{i-1}^t}(u_{i-1}^t; R^t, v_{i-1}^t)\| \\
&\leq \eta^t \cdot \|F_{\sigma_i^t}(u_i^t) + F_{\sigma_{i-1}^t}(u_i^t) - F_{\sigma_{i-1}^t}(u_{i-1}^t)\|_2 \\
&\quad + \eta^t d \cdot \|\hat{F}_{\sigma_i^t}(u_i^t; R^t, v_i^t) - F_{\sigma_i^t}(u_i^t)\|_2 \\
&\quad + \eta^t d \cdot \|\hat{F}_{\sigma_{i-1}^t}(u_i^t; R^t, v_i^t) - F_{\sigma_{i-1}^t}(u_i^t)\|_2 \\
&\quad + \eta^t d \cdot \|\hat{F}_{\sigma_{i-1}^t}(u_{i-1}^t; R^t, v_{i-1}^t) - F_{\sigma_{i-1}^t}(u_{i-1}^t)\|_2 \\
&\leq 3G\eta^t + 3\eta^t d \cdot \left(2(M_L + GR^t) \cdot \frac{1}{R^t} + \ell D \cdot R^t \right).
\end{aligned} \tag{34}$$

Substituting (34) back into (33), we have:

$$\begin{aligned}
& \mathbb{E} \left[\left\langle \mathbb{E}_v [\hat{F}_{\sigma_i^t}(u_i^t; R^t, v) | u_i^t] - \hat{F}_{\sigma_i^t}(u_i^t, R^t, v_i^t), u_{i+1}^t - u_i^t \right\rangle \right] \\
& \leq d^2 \ell D 6 \eta^t G \cdot R^t + 6d^2 \eta^t (3G^2 + M_L \ell D) + 30d^2 \eta^t M_L G \cdot \frac{1}{R^t} + 12d^2 \eta^t M_L^2 \cdot \left(\frac{1}{R^t} \right)^2 \\
& \leq C_1 \cdot \left(\eta^t R^t + \eta^t + \frac{\eta^t}{R^t} + \frac{\eta^t}{(R^t)^2} \right), \tag{35}
\end{aligned}$$

where $C_1 := d^2 \cdot \max \{6G\ell D, 18G^2 + 6M_L \ell D, 30M_L G, 12M_L^2\}$ is a constant independent of the sequences $\{\eta^t\}$ and $\{R^t\}$. The quantities $\mathbb{E}[\langle E_{i,2}^t, u_{i+1}^t - u \rangle]$ and $\mathbb{E}[\langle E_{i,3}^t, u_{i+1}^t - u \rangle]$ can be similarly bounded. Substituting (31), (32), (35) back into (30), and substituting (30) and (29) into (28), we find that:

$$\begin{aligned}
& 2\eta^t \cdot \mathbb{E} \left[\left\langle F_{\sigma_i^t}(u_{i+1}^t), u_{i+1}^t - u \right\rangle \right] \\
& = \mathbb{E}[\|u_i^t - u\|_2^2] - \mathbb{E}[\|u_{i+1}^t - u\|_2^2] - \mathbb{E}[\|u_{i+1}^t - u_i^t\|_2^2] \\
& \quad + 2\eta^t \cdot \mathbb{E} \left[\left\langle \gamma_i^t, u_{i+1}^t - u \right\rangle \right] + 2\eta^t \cdot \sum_{k=1}^3 \mathbb{E} \left[\left\langle E_{i,k}^t, u_{i+1}^t - u \right\rangle \right] \\
& \leq \mathbb{E}[\|u_i^t - u\|_2^2] - \mathbb{E}[\|u_{i+1}^t - u\|_2^2] - \mathbb{E}[\|u_{i+1}^t - u_i^t\|_2^2] \\
& \quad + 2\eta^t \cdot \mathbb{E} \left[\left\langle F_{\sigma_i^t}(u_{i+1}^t) - F_{\sigma_i^t}(u_i^t), u_{i+1}^t - u \right\rangle \right] \\
& \quad - 2\eta^t \cdot \mathbb{E} \left[\left\langle F_{\sigma_{i-1}^t}(u_i^t) - F_{\sigma_{i-1}^t}(u_{i-1}^t), u_i^t - u \right\rangle \right] \\
& \quad + \eta^t \ell \cdot \mathbb{E}[\|u_i^t - u_{i-1}^t\|_2^2] + \eta^t \ell \cdot \mathbb{E}[\|u_{i+1}^t - u_i^t\|_2^2] \\
& \quad + 6C_1 \cdot \left(\eta^t R^t + (\eta^t)^2 R^t + (\eta^t)^2 + \frac{(\eta^t)^2}{R^t} + \frac{(\eta^t)^2}{(R^t)^2} \right),
\end{aligned}$$

In particular, since by assumption $\eta^t \leq 1/(2\ell)$ for each $t \in \{0, 1, \dots, T-1\}$, then:

$$\begin{aligned}
& 2\eta^t \cdot \mathbb{E} \left[\left\langle F_{\sigma_i^t}(u_{i+1}^t), u_{i+1}^t - u \right\rangle \right] \\
& \leq \mathbb{E}[\|u_i^t - u\|_2^2] - \mathbb{E}[\|u_{i+1}^t - u\|_2^2] - \frac{1}{2} \mathbb{E}[\|u_{i+1}^t - u_i^t\|_2^2] + \frac{1}{2} \mathbb{E}[\|u_i^t - u_{i-1}^t\|_2^2] \\
& \quad + 2\eta^t \cdot \mathbb{E} \left[\left\langle F_{\sigma_i^t}(u_{i+1}^t) - F_{\sigma_i^t}(u_i^t), u_{i+1}^t - u \right\rangle \right] \\
& \quad - 2\eta^t \cdot \mathbb{E} \left[\left\langle F_{\sigma_{i-1}^t}(u_i^t) - F_{\sigma_{i-1}^t}(u_{i-1}^t), u_i^t - u \right\rangle \right] \\
& \quad + 6C_1 \cdot \left(\eta^t R^t + (\eta^t)^2 R^t + (\eta^t)^2 + \frac{(\eta^t)^2}{R^t} + \frac{(\eta^t)^2}{(R^t)^2} \right),
\end{aligned}$$

□

Finally, to bound the step size terms above, we require the following lemma, which follows from standard calculus arguments.

Lemma A.8.

$$\begin{aligned}
\sum_{t=1}^T t^{-\beta} & \geq \frac{1}{1-\beta} T^{1-\beta}, & \forall \beta < 1, \\
\sum_{t=1}^T t^{-(1+\beta)} & \leq \frac{1}{\beta} + 1, & \forall \beta > 0.
\end{aligned}$$

A.2 Proof of Theorem 4.1

Proof. (Proof of Theorem 4.1) By applying Lemma A.7 (note that $\eta^t \leq \eta^0 \leq \frac{1}{2\ell}$, for each $t \in \{0, 1, \dots, T-1\}$) and using convex-concave nature of L_r (refer Proposition 1 in [MOP20b]), for each $r \in \{1, \dots, n\}$, we have:

$$\begin{aligned}
& 2\eta^t \cdot \mathbb{E}[L_{\sigma_i^t}(x_{i+1}^t, y^*) - L_{\sigma_i^t}(x^*, y_{i+1}^t)] \\
& \leq 2\eta^t \cdot \mathbb{E}[\langle F_{\sigma_i^t}(u_{i+1}^t), u_{i+1}^t - u^* \rangle] \\
& \leq \mathbb{E}[\|u_i^t - u^*\|_2^2] - \mathbb{E}[\|u_{i+1}^t - u^*\|_2^2] - \frac{1}{2}\mathbb{E}[\|u_{i+1}^t - u_i^t\|_2^2] + \frac{1}{2}\mathbb{E}[\|u_i^t - u_{i-1}^t\|_2^2] \\
& \quad + 2\eta^t \cdot \mathbb{E}\left[\left\langle F_{\sigma_i^t}(u_{i+1}^t) - F_{\sigma_i^t}(u_i^t), u_{i+1}^t - u^* \right\rangle\right] \\
& \quad - 2\eta^t \cdot \mathbb{E}\left[\left\langle F_{\sigma_{i-1}^t}(u_i^t) - F_{\sigma_{i-1}^t}(u_{i-1}^t), u_i^t - u^* \right\rangle\right] \\
& \quad + 6C_1 \cdot \left(\eta^t R^t + (\eta^t)^2 R^t + (\eta^t)^2 + \frac{(\eta^t)^2}{R^t} + \frac{(\eta^t)^2}{(R^t)^2} \right). \tag{36}
\end{aligned}$$

Meanwhile, Lemma A.5, Proposition A.4 (Kantorovich Duality), and Lemma A.6 imply that:

$$\begin{aligned}
& \left| \mathbb{E}[\Delta(u_{i+1}^t)] - \mathbb{E}[L_{\sigma_i^t}(x_{i+1}^t, y^*) - L_{\sigma_i^t}(x^*, y_{i+1}^t)] \right| \leq \frac{G}{n} \sum_{r=1}^n \mathcal{W}_2(\mathcal{D}_{i+1,t}, \mathcal{D}_{i+1,t}^r) \\
& \leq \frac{G}{n} \sum_{r=1}^n \sqrt{\mathbb{E}[\|u_{i+1}^t(\sigma^t) - u_{i+1}^t(\bar{\sigma}^t)\|_2^2]} \\
& \leq G \cdot \left((6nd + 14n + 4)G \cdot \eta^t + 6ndM_L \cdot \frac{\eta^t}{R^t} \right).
\end{aligned}$$

Substituting back into (36), we have:

$$\begin{aligned}
& 2\eta^t \cdot \mathbb{E}[\Delta(u_i^t)] \\
& \leq 2\eta^t \cdot \mathbb{E}[L_{\sigma_i^t}(x_{i+1}^t, y^*) - L_{\sigma_i^t}(x^*, y_{i+1}^t)] \\
& \quad + G \cdot \left((12nd + 28n + 8)G \cdot (\eta^t)^2 + 12ndM_L \cdot \frac{(\eta^t)^2}{R^t} \right) \\
& \leq \mathbb{E}[\|u_i^t - u^*\|_2^2] - \mathbb{E}[\|u_{i+1}^t - u^*\|_2^2] - \frac{1}{2}\mathbb{E}[\|u_{i+1}^t - u_i^t\|_2^2] + \frac{1}{2}\mathbb{E}[\|u_i^t - u_{i-1}^t\|_2^2] \\
& \quad + 2\eta^t \cdot \mathbb{E}\left[\left\langle F_{\sigma_i^t}(u_{i+1}^t) - F_{\sigma_i^t}(u_i^t), u_{i+1}^t - u \right\rangle\right] \\
& \quad - 2\eta^t \cdot \mathbb{E}\left[\left\langle F_{\sigma_{i-1}^t}(u_i^t) - F_{\sigma_{i-1}^t}(u_{i-1}^t), u_i^t - u \right\rangle\right] \\
& \quad + 6C_1 \cdot \left(\eta^t R^t + (\eta^t)^2 R^t + (\eta^t)^2 + \frac{(\eta^t)^2}{R^t} + \frac{(\eta^t)^2}{(R^t)^2} \right) \\
& \quad + G \cdot \left((12nd + 28n + 8)G \cdot (\eta^t)^2 + 12ndM_L \cdot \frac{(\eta^t)^2}{R^t} \right). \tag{37}
\end{aligned}$$

We can now sum the above telescoping terms across the t -th epoch, as shown below:

$$\begin{aligned}
& 2 \cdot \sum_{i=1}^n \eta^t \cdot \mathbb{E}[\Delta(u_i^t)] \\
& \leq \mathbb{E}[\|u_1^t - u^*\|_2^2] - \mathbb{E}[\|u_1^{t+1} - u^*\|_2^2] + \frac{1}{2}\mathbb{E}[\|u_1^t - u_0^t\|_2^2] - \frac{1}{2}\mathbb{E}[\|u_1^{t+1} - u_0^{t+1}\|_2^2] \\
& \quad + 2\eta^t \cdot \mathbb{E}\left[\left\langle F_{\sigma_0^t}(u_1^t) - F_{\sigma_0^t}(u_0^t), u_1^t - u^* \right\rangle\right]
\end{aligned}$$

$$\begin{aligned}
& -2\eta^t \cdot \mathbb{E} \left[\left\langle F_{\sigma_0^{t+1}}(u_1^{t+1}) - F_{\sigma_0^{t+1}}(u_0^{t+1}), u_1^{t+1} - u^* \right\rangle \right] \\
& + 6nC_1 \cdot \left(\eta^t R^t + (\eta^t)^2 R^t + (\eta^t)^2 + \frac{(\eta^t)^2}{R^t} + \frac{(\eta^t)^2}{(R^t)^2} \right) \\
& + nG \cdot \left((12nd + 28n + 8)G \cdot (\eta^t)^2 + 12ndM_L \cdot \frac{(\eta^t)^2}{R^t} \right).
\end{aligned}$$

Meanwhile, we have for each $t = 0, 1, \dots, T-1$, $i \in [n]$:

$$\begin{aligned}
& \mathbb{E} \left[\left\langle F_{\sigma_i^t}(u_{i+1}^t) - F_{\sigma_i^t}(u_i^t), u_{i+1}^t - u^* \right\rangle \right] \\
& \leq \mathbb{E} \left[\|F_{\sigma_i^t}(u_{i+1}^t) - F_{\sigma_i^t}(u_i^t)\| \cdot \|u_{i+1}^t - u^*\| \right] \\
& = \ell \cdot \mathbb{E} [\|u_{i+1}^t - u_i^t\|] \cdot D \\
& \leq \ell D \cdot \mathbb{E} \left[\left\| -\eta^t \hat{F}_{\sigma_i^t}(u_i^t; R^t, v_i^t) - \eta^t \hat{F}_{\sigma_{i-1}^t}(u_i^t; R^t, v_i^t) + \eta^t \hat{F}_{\sigma_{i-1}^t}(u_{i-1}^t; R^t, v_{i-1}^t) \right\| \right] \\
& \leq 3\ell D \cdot \eta^t \cdot \left(dG + \frac{dM_L}{R^t} \right) \\
& = 3\ell D dG \cdot \eta^t + 3\ell D dM_L \cdot \frac{\eta^t}{R^t},
\end{aligned}$$

where the final inequality follows from Proposition A.3, (12). We can upper bound $\mathbb{E} \left[\left\langle F_{\sigma_{i-1}^t}(u_i^t) - F_{\sigma_{i-1}^t}(u_{i-1}^t), u_i^t - u^* \right\rangle \right]$ in a similar fashion. Substituting back into (37), we have:

$$\begin{aligned}
& 2 \cdot \sum_{i=1}^n \eta^t \cdot \mathbb{E} [\Delta(u_i^t)] \\
& \leq \mathbb{E} [\|u_1^t - u^*\|_2^2] - \mathbb{E} [\|u_1^{t+1} - u^*\|_2^2] + \frac{1}{2} \mathbb{E} [\|u_1^t - u_0^t\|_2^2] - \frac{1}{2} \mathbb{E} [\|u_1^{t+1} - u_0^{t+1}\|_2^2] \\
& \quad + 6nC_1 \cdot \left(\eta^t R^t + (\eta^t)^2 R^t + (\eta^t)^2 + \frac{(\eta^t)^2}{R^t} + \frac{(\eta^t)^2}{(R^t)^2} \right) \\
& \quad + nG \cdot \left((12nd + 28n + 8)G \cdot (\eta^t)^2 + 12ndM_L \cdot \frac{(\eta^t)^2}{R^t} \right) \\
& \quad + 6\ell D dG \cdot (\eta^t)^2 + 6\ell D dM_L \cdot \frac{(\eta^t)^2}{R^t} \\
& \leq \mathbb{E} [\|u_1^t - u^*\|_2^2] - \mathbb{E} [\|u_1^{t+1} - u^*\|_2^2] + \frac{1}{2} \mathbb{E} [\|u_1^t - u_0^t\|_2^2] - \frac{1}{2} \mathbb{E} [\|u_1^{t+1} - u_0^{t+1}\|_2^2] \quad (38) \\
& \quad + 2C \cdot \left(\eta^t R^t + (\eta^t)^2 R^t + (\eta^t)^2 + \frac{(\eta^t)^2}{R^t} + \frac{(\eta^t)^2}{(R^t)^2} \right),
\end{aligned}$$

where $C := \max\{3nC_1, (6nd + 14n + 4)nG, 6ndM_L, 3\ell D dG, 3\ell D dM_L\}$.

Finally, summing the above telescoping terms over $i \in [n]$ and $t \in \{0, 1, \dots, T-1\}$, and removing non-positive terms, we obtain:

$$\begin{aligned}
& \frac{\sum_{t=0}^{T-1} \sum_{i=1}^n \eta^t \cdot \mathbb{E} [\Delta(u_i^t)]}{\sum_{t=0}^{T-1} \sum_{i=1}^n \eta^t} \\
& \leq \frac{1}{2 \cdot \sum_{t=0}^{T-1} \sum_{i=1}^n \eta^t} \left(\|u_0^0 - u^*\|_2 - \mathbb{E} [\|u_n^{T-1} - u^*\|_2] + \frac{1}{2} \|u_1^0 - u_0^0\|_2 - \frac{1}{2} \mathbb{E} [\|u_n^{T-1} - u_{n-1}^{T-1}\|_2] \right) \\
& \quad + C \cdot \frac{1}{\sum_{t=0}^{T-1} \sum_{i=1}^n \eta^t} \cdot \sum_{t=0}^{T-1} \left(\eta^t R^t + (\eta^t)^2 R^t + (\eta^t)^2 + \frac{(\eta^t)^2}{R^t} + \frac{(\eta^t)^2}{(R^t)^2} \right)
\end{aligned}$$

$$\leq \frac{1}{\sum_{t=0}^{T-1} \eta^t} \cdot \frac{3D}{4n} + C \cdot \frac{1}{n \sum_{t=0}^{T-1} \eta^t} \cdot \sum_{t=0}^{T-1} \left(\eta^t R^t + (\eta^t)^2 R^t + (\eta^t)^2 + \frac{(\eta^t)^2}{R^t} + \frac{(\eta^t)^2}{(R^t)^2} \right), \quad (39)$$

By definition, $\eta^t = \eta^0 \cdot (t+1)^{-3/4-\chi}$ and $R^t = R^0 \cdot (t+1)^{-1/4}$, so by Lemma A.8, we have:

$$\begin{aligned} \sum_{t=0}^{T-1} \eta^t &= \eta^0 \cdot \sum_{t=1}^T t^{-3/4-\chi} \geq 4\eta^0 \cdot T^{1/4-\chi}, \\ \sum_{t=0}^{T-1} \eta^t R^t &= \eta^0 R^0 \cdot \sum_{t=1}^T t^{-(1+\chi)} \leq \eta^0 R^0 \cdot \left(1 + \frac{1}{\chi}\right), \\ \sum_{t=0}^{T-1} (\eta^t)^2 &= (\eta^0)^2 \cdot \sum_{t=1}^T t^{-3/2-2\chi} \leq (\eta^0)^2 \cdot \left(1 + \frac{1}{\frac{1}{2} + 2\chi}\right) \leq 3 \cdot (\eta^0)^2, \\ \sum_{t=0}^{T-1} (\eta^t)^2 R^t &= (\eta^0)^2 R^0 \cdot \sum_{t=1}^T t^{-7/4-2\chi} \leq (\eta^0)^2 R^0 \cdot \left(1 + \frac{1}{\frac{3}{4} + 2\chi}\right) \leq \frac{7}{4} \cdot (\eta^0)^2 \epsilon^0, \\ \sum_{t=0}^{T-1} \frac{(\eta^t)^2}{R^t} &= \frac{(\eta^0)^2}{R^0} \cdot \sum_{t=1}^T t^{-5/4-2\chi} \leq \frac{(\eta^0)^2}{R^0} \cdot \left(1 + \frac{1}{\frac{1}{4} + 2\chi}\right) \leq 5 \cdot \frac{(\eta^0)^2}{\epsilon^0}, \\ \sum_{t=0}^{T-1} \frac{(\eta^t)^2}{(R^t)^2} &= \frac{(\eta^0)^2}{(R^0)^2} \cdot \sum_{t=1}^T t^{-1-2\chi} \leq \frac{(\eta^0)^2}{(R^0)^2} \cdot \left(1 + \frac{1}{2\chi}\right). \end{aligned}$$

Substituting back into (39) and using the convexity of the gap function $\Delta(\cdot)$, we have:

$$\begin{aligned} &\mathbb{E}[\Delta(u^T)] \\ &\leq \frac{\sum_{t=0}^{T-1} \sum_{i=1}^n \eta^t \cdot \mathbb{E}[\Delta(u_i^t)]}{\sum_{t=0}^{T-1} \sum_{i=1}^n \eta^t} \\ &\leq \frac{1}{\sum_{t=0}^{T-1} \eta^t} \cdot \frac{3}{4n} D + C \cdot \frac{1}{\sum_{t=0}^{T-1} \eta^t} \cdot \sum_{t=0}^{T-1} \left(\eta^t R^t + (\eta^t)^2 R^t + (\eta^t)^2 + \frac{(\eta^t)^2}{R^t} + \frac{(\eta^t)^2}{(R^t)^2} \right) \\ &\leq \left(\frac{3}{16n} D + \frac{47}{4n} \cdot C \max \left\{ R^0, \eta^0, \eta^0 R^0, \frac{\eta^0}{R^0}, \frac{\eta^0}{(R^0)^2} \right\} \left(1 + \frac{1}{\chi}\right) \right) T^{-1/4+\chi} \\ &\leq R. \end{aligned}$$

where the final inequality follows by definition of T . □

B Wasserstein Distributionally Robust Strategic Classification

B.1 Model of Adversary

In this subsection, we formally define our model for the adversary, and the uncertainty set of distributions for the resulting strategically and adversarially perturbed data. For better exposition, in this section we summarize the various distributions used in the main article in Table 1 below.

Table 1: Table of notations

Notation	Explanation
\mathcal{D}	Unknown underlying distribution
$\mathcal{D}(\theta)$	Unknown underlying distribution strategically perturbed by θ
$\tilde{\mathcal{D}}_n(\theta)$	Empirical distribution of strategically perturbed data
\mathbb{P}	An element of uncertainty set $\mathcal{P}(\theta)$
\mathbb{P}_θ^i	Conditional distribution of adversarially generated data given i^{th} data point

The WDRSL problem formulation contains two main components—the *strategic component* that accounts for the distribution shift $\mathcal{D}(\theta)$ in response to the choice of classifier θ , and the *adversarial component*, which accounts for the uncertainty set $\mathcal{P}(\theta)$. As per the modeling assumptions put forth in Section 5.2, we have $(\tilde{x}_i, \tilde{y}_i) \sim \mathcal{D}$ and $(b_i(\theta, \tilde{x}_i, \tilde{y}_i), \tilde{y}_i) \sim \mathcal{D}(\theta)$ for all $i \in [n]$. For the sake of brevity, we shall use $b_i(\theta)$ in place of $b_i(\theta, \tilde{x}_i, \tilde{y}_i)$ for all $i \in [n]$.

As per the standard formulation of distributionally robust optimization, we restrict $\mathcal{P}(\theta)$ to be a Wasserstein neighborhood of $\tilde{\mathcal{D}}_n(\theta)$ (the empirical distribution of strategic responses $\{(b_i(\theta), \tilde{y}_i)\}_{i=1}^n$), i.e., we set $\mathcal{P}(\theta) \subset \mathbb{B}_\delta(\tilde{\mathcal{D}}_n(\theta))$ for some $\delta > 0$. However, to ensure that the min-max problem reformulated from the WDRSC problem is convex-concave, we further require the adversary to modify the *label* of a data point i in the empirical distribution only when the true label \tilde{y}_i is +1, although they are still always allowed to modify the *feature* $b_i(\theta)$. As a consequence, this imposes some restrictions on the conditional distribution \mathbb{P}_θ^i of (dx, y) , as generated by the adversary, given a data point i in the empirical distribution. In particular:

$$\mathbb{P}_\theta^i(dx, +1|b_i(\theta), -1) = 0, \quad \forall i \in [n].$$

By definition of conditional distributions, we obtain that any distribution \mathbb{P} can be expressed as the average of the conditional distribution \mathbb{P}_θ^i . That is,

$$\mathbb{P}(dx, y) = \frac{1}{n} \sum_{i=1}^n \mathbb{P}_\theta^i(dx, y|b_i(\theta), \tilde{y}_i).$$

Below, we formally state the restriction described above.

Assumption B.1. *We assume that $\mathbb{P} \in \mathbb{B}_\delta(\tilde{\mathcal{D}}_n(\theta))$ and $\mathbb{P}_\theta^i(dx, +1|b_i(\theta), -1) = 0$ for all $i \in [n]$. As a direct result, the uncertainty set $\mathcal{P}(\theta)$ is characterized as:*

$$\mathcal{P}(\theta) = \mathbb{B}_\delta(\tilde{\mathcal{D}}_n(\theta)) \cap \left\{ \frac{1}{n} \sum_{i=1}^n \mathbb{P}_\theta^i(dx, y|b_i(\theta), \tilde{y}_i) \mid \mathbb{P}_\theta^i(dx, +1|b_i(\theta), -1) = 0, \forall i \in [n] \right\}. \quad (40)$$

In the following subsection, we reformulate the WDRSC problem with a generalized linear model and with the uncertainty set defined in (40).

B.2 Proof of Theorem 5.3

The proof takes inspirations from [SAEK15, Theorem 1]. First, we define the *Wasserstein distance between distributions on \mathcal{Z} with cost function c* ; note that this is different from the p -Wasserstein distance between probability distributions on \mathbb{R}^d defined in Appendix A.1.

Definition B.1. (*Wasserstein distance between distributions on \mathcal{Z} with cost Function c*) Let μ, ν be probability distributions over $\mathcal{Z} := \mathbb{R}^d \times \{+1, -1\}$ with finite second moments, and let $\Pi(\mu, \nu)$ denote the set of all couplings (joint distributions) between μ and ν . Given a metric $c : \mathcal{Z} \times \mathcal{Z} \rightarrow [0, \infty)$ on \mathcal{Z} , we define:

$$\mathcal{W}_c(\mu, \nu) = \inf_{(Z, Z') \sim \pi \in \Pi(\mu, \nu)} \mathbb{E}_\pi [c(Z, Z')].$$

In Theorem 6 and in our proof below, we use the cost function $c(z, z') := \|x - x'\|_2^2 + \kappa \cdot |y - y'|$, with a fixed constant $\kappa > 0$, for each $z := (x, y) \in \mathcal{Z}$ and $z' := (x', y') \in \mathcal{Z}$.

Proof. (Proof of Theorem 5.3) Fix a $\theta \in \Theta$. Note that $b_i(\theta, \tilde{x}_i, +1) = \tilde{x}_i$. For any $(x, y) \in \mathbb{R}^d \times \{-1, 1\}$, let $\ell((x, y), \theta) := \phi(\langle x, \theta \rangle) - y \langle x, \theta \rangle$. We first analyze the inner supremum term, i.e.

$$\begin{aligned} & \sup_{\mathbb{P} \in \mathcal{P}(\theta)} \mathbb{E}_\mathbb{P} [\phi(\langle x, \theta \rangle) - y \langle x, \theta \rangle] \\ &= \sup_{\mathbb{P} \in \mathcal{P}(\theta)} \int_{\mathcal{Z}} \ell(z, \theta) \mathbb{P}(z) dz \\ &= \begin{cases} \sup_{\pi_\theta \in \Pi(\mathbb{P}, \tilde{\mathcal{D}}_n(\theta))} & \int_{\mathcal{Z}} \ell(z, \theta) \pi_\theta(dz, \mathcal{Z}) \\ \text{s.t.} & \int_{\mathcal{Z} \times \mathcal{Z}} \|z - \tilde{z}\| \pi_\theta(dz, d\tilde{z}) \leq \delta \end{cases} \end{aligned}$$

Here, $\Pi(\mathbb{P}, \tilde{\mathcal{D}}_n(\theta))$ denotes the set of all joint distributions that couple $\mathbb{P} \in \mathcal{P}(\theta)$ and $\tilde{\mathcal{D}}_n(\theta)$. Since the marginal distribution $\tilde{\mathcal{D}}_n(\theta)$ of \tilde{z} is discrete, such couplings π_θ are completely determined by the conditional distribution \mathbb{P}_θ^i of z given $\tilde{z}_i = (\tilde{x}_i(\theta), \tilde{y}_i)$ for each $i \in \{1, \dots, n\}$. That is:

$$\pi_\theta(dz, d\tilde{z}) = \frac{1}{n} \sum_{i \in [n]} \vartheta_{(b_i(\theta), \tilde{y}_i)}(d\tilde{z}) \mathbb{P}_\theta^i(dz)$$

where for any $(x, y) \in \mathcal{Z}$, $\vartheta_{(x, y)}$ is a *Dirac delta* distribution with its support at point (x, y) .

We introduce some notations. Let $\mathcal{I}_{+1} = \{i \in [n] : \tilde{y}_i = +1\}$ and $\mathcal{I}_{-1} = \{i \in [n] : \tilde{y}_i = -1\}$. Let's introduce two distributions μ_θ^i and ν_θ^i such that

$$\mathbb{P}_\theta^i = \begin{cases} \mu_\theta^i & \text{if } i \in \mathcal{I}_{+1} \\ \nu_\theta^i & \text{if } i \in \mathcal{I}_{-1} \end{cases}$$

Due to the constraint (40), we have $\nu_\theta^i(dx, +1) = 0$ at every x . This implies:

$$\pi_\theta(dz, d\tilde{z}) = \frac{1}{n} \left(\sum_{i \in \mathcal{I}_{+1}} \vartheta_{(b_i(\theta), 1)}(d\tilde{z}) \mu_\theta^i(dz) + \sum_{i \in \mathcal{I}_{-1}} \vartheta_{(b_i(\theta), -1)}(d\tilde{z}) \nu_\theta^i(dz) \right)$$

With a slight abuse of notation, we denote $\mu_{\theta, +1}^i(dx) = \mu_\theta^i(dx, +1)$, $\mu_{\theta, -1}^i(dx) = \mu_\theta^i(dx, -1)$ and $\nu_\theta^i(dx) = \nu_\theta^i(dx, -1)$. The optimization problem of concern then simplifies to:

$$\sup_{\mu_{\theta, \pm 1}^i, \nu_\theta^i} \frac{1}{n} \sum_{i \in \mathcal{I}_{+1}} \int_{\mathbb{R}^d} \ell((x, +1), \theta) \mu_{\theta, +1}^i(dx) + \frac{1}{n} \sum_{i \in \mathcal{I}_{-1}} \int_{\mathbb{R}^d} \ell((x, -1), \theta) \mu_{\theta, -1}^i(dx)$$

$$\begin{aligned}
& + \frac{1}{n} \sum_{i \in \mathcal{I}_{-1}} \int_{\mathbb{R}^d} \ell((x, -1), \theta) \nu_{\theta}^i(dx) \\
\text{s.t. } & \frac{1}{n} \sum_{i: \tilde{y}_i = +1} \int_{\mathbb{R}^d} \|(x, +1) - (b_i(\theta), \tilde{y}_i)\| \mu_{\theta, +1}^i(dx) \\
& + \frac{1}{n} \sum_{i: \tilde{y}_i = -1} \int_{\mathbb{R}^d} \|(x, -1) - (b_i(\theta), \tilde{y}_i)\| \mu_{\theta, -1}^i(dx) \\
& \int_{\mathbb{R}^d} \mu_{\theta, +1}^i(dx) + \int_{\mathbb{R}^d} \mu_{\theta, -1}^i(dx) = 1, \quad \forall i \in \mathcal{I}_{+1} \\
& \int_{\mathbb{R}^d} \nu_{\theta}^i(dx) = 1, \quad \forall i \in \mathcal{I}_{-1}
\end{aligned}$$

First, we rewrite the inequality constraint above as follows. Recall that:

$$\begin{aligned}
& \frac{2\kappa}{n} \int_{\mathbb{R}^d} \sum_{i \in \mathcal{I}_{+1}} \mu_{\theta, -1}^i(dx) + \frac{1}{n} \int_{\mathbb{R}^d} \sum_{i \in \mathcal{I}_{+1}} \|x - b_i(\theta)\| \mu_{\theta, +1}^i(dx) \\
& + \frac{1}{n} \int_{\mathbb{R}^d} \sum_{i \in \mathcal{I}_{+1}} \|x - b_i(\theta)\| \mu_{\theta, -1}^i(dx) + \frac{1}{n} \int_{\mathbb{R}^d} \sum_{i \in \mathcal{I}_{-1}} \|x - b_i(\theta)\| \nu_{\theta}^i(dx) \leq \delta.
\end{aligned}$$

Hence,

$$\begin{aligned}
& \sup_{\mu_{\theta, \pm 1}^i, \nu_{\theta}^i} \frac{1}{n} \sum_{i \in \mathcal{I}_{+1}} \int_{\mathbb{R}^d} \ell((x, +1), \theta) \mu_{\theta, +1}^i(dx) + \frac{1}{n} \sum_{i \in \mathcal{I}_{+1}} \int_{\mathbb{R}^d} \ell((x, -1), \theta) \mu_{\theta, -1}^i(dx) \\
& + \frac{1}{n} \sum_{\tilde{y}_i = -1} \int_{\mathbb{R}^d} \ell((x, -1), \theta) \nu_{\theta}^i(dx) \\
\text{s.t. } & \frac{2\kappa}{n} \int_{\mathbb{R}^d} \sum_{i \in \mathcal{I}_{+1}} \mu_{\theta, -1}^i(dx) + \frac{1}{n} \int_{\mathbb{R}^d} \sum_{i \in \mathcal{I}_{+1}} \|x - b_i(\theta)\| \mu_{\theta, +1}^i(dx) \\
& + \frac{1}{n} \int_{\mathbb{R}^d} \sum_{i \in \mathcal{I}_{+1}} \|x - b_i(\theta)\| \mu_{\theta, -1}^i(dx) + \frac{1}{n} \int_{\mathbb{R}^d} \sum_{i \in \mathcal{I}_{-1}} \|x - b_i(\theta)\| \nu_{\theta}^i(dx) \leq \delta \\
& \int_{\mathbb{R}^d} \mu_{\theta, +1}^i(dx) + \int_{\mathbb{R}^d} \mu_{\theta, -1}^i(dx) = 1, \quad \forall i \in \mathcal{I}_{+1} \\
& \int_{\mathbb{R}^d} \nu_{\theta}^i(dx) = 1, \quad \forall i \in \mathcal{I}_{-1}
\end{aligned}$$

Now, we can use duality to reformulate the infinite-dimensional optimization problem into a finite-dimensional problem:

$$\begin{aligned}
& \sup_{\mathbb{P} \in \mathcal{P}(\theta)} \mathbb{E}_{\mathbb{P}}[\phi(\langle x, \theta \rangle) - y \langle x, \theta \rangle] \\
= & \begin{cases} \inf_{\alpha, s_i} & \alpha \delta + \frac{1}{n} \sum_{i \in \mathcal{I}_{+1}} s_i + \frac{1}{n} \sum_{i \in \mathcal{I}_{-1}} t_i \\ \text{s.t.} & \sup_x \ell((x, +1), \theta) - \alpha \cdot \frac{1 + \tilde{y}_i}{2} \|x - b_i(\theta)\| \leq s_i \quad \forall i \in \mathcal{I}_{+1} \\ & \sup_x \ell((x, -1), \theta) - \alpha \cdot \frac{1 + \tilde{y}_i}{2} \|x - b_i(\theta)\| - \alpha \kappa (1 + \tilde{y}_i) \leq s_i \quad \forall i \in \mathcal{I}_{+1} \\ & \sup_x \ell((x, -1), \theta) - \alpha \cdot \frac{1 - \tilde{y}_i}{2} \|x - b_i(\theta)\| \leq t_i \quad \forall i \in \mathcal{I}_{-1} \\ & \alpha \geq 0 \end{cases}
\end{aligned}$$

which is equivalent to:

$$\sup_{\mathbb{P} \in \mathcal{P}(\theta)} \mathbb{E}_{\mathbb{P}}[\phi(\langle x, \theta \rangle) - y \langle x, \theta \rangle]$$

$$= \begin{cases} \inf_{\alpha, s_i} & \alpha\delta + \frac{1}{n} \sum_{i \in \mathcal{I}_{+1}} s_i + \frac{1}{n} \sum_{i \in \mathcal{I}_{-1}} t_i \\ \text{s.t.} & \sup_x \ell((x, +1), \theta) - \alpha \|x - b_i(\theta)\| \leq s_i \quad \forall i \in \mathcal{I}_{+1} \\ & \sup_x \ell((x, -1), \theta) - \alpha \|x - b_i(\theta)\| - 2\alpha\kappa \leq s_i \quad \forall i \in \mathcal{I}_{+1} \\ & \sup_x \ell((x, -1), \theta) - \alpha \|x - b_i(\theta)\| \leq t_i \quad \forall i \in \mathcal{I}_{-1} \\ & \alpha \geq 0 \end{cases}$$

We now invoke [YLMJ21, Lemma A.1], which claims that for any $\tilde{y} \in \{+1, -1\}$ and $\tilde{x} \in \mathbb{R}^d$:

$$\sup_x \ell((x, \tilde{y}), \theta) - \alpha \|x - \tilde{x}\| = \begin{cases} \ell((\tilde{x}, \tilde{y}), \theta) & \text{if } \|\theta\| \leq \alpha/(L+1) \\ -\infty & \text{otherwise.} \end{cases}$$

We now have:

$$\begin{aligned} & \sup_{\mathbb{P} \in \mathcal{P}(\theta)} \mathbb{E}_{\mathbb{P}}[\phi(\langle x, \theta \rangle) - y\langle x, \theta \rangle] \\ &= \begin{cases} \inf_{\alpha, s_i} & \alpha\delta + \frac{1}{n} \sum_{i \in \mathcal{I}_{+1}} s_i + \frac{1}{n} \sum_{i \in \mathcal{I}_{-1}} t_i \\ \text{s.t.} & \ell((b_i(\theta), +1), \theta) \leq s_i \quad \forall i \in \mathcal{I}_{+1} \\ & \ell((b_i(\theta), -1), \theta) - 2\alpha\kappa \leq s_i \quad \forall i \in \mathcal{I}_{+1} \\ & \ell((b_i(\theta), -1), \theta) \leq t_i \quad \forall i \in \mathcal{I}_{-1} \\ & \alpha \geq 0 \\ & \|\theta\| \leq \alpha/(L+1) \end{cases} \end{aligned}$$

In the above presented optimization problem we can conclude that:

$$\begin{aligned} t_i &= \phi(\langle b_i(\theta), \theta \rangle) + \langle b_i(\theta), \theta \rangle & \forall i \in \mathcal{I}_{-1} \\ s_i &= \max\{\ell((b_i(\theta), +1), \theta), \ell((b_i(\theta), -1), \theta) - 2\alpha\kappa)\} & \forall i \in \mathcal{I}_{+1}. \end{aligned}$$

To further simplify the s_i expression, note that:

$$\begin{aligned} s_i &= \max\{\phi(\langle b_i(\theta), \theta \rangle) - \langle b_i(\theta), \theta \rangle, \phi(\langle b_i(\theta), \theta \rangle) + \langle b_i(\theta), \theta \rangle - 2\alpha\kappa\} \\ &= \phi(\langle b_i(\theta), \theta \rangle) - \langle b_i(\theta), \theta \rangle + \max\{0, 2\langle b_i(\theta), \theta \rangle - 2\alpha\kappa\} \\ &= \phi(\langle b_i(\theta), \theta \rangle) - \alpha\kappa + \max_{\gamma_i: |\gamma_i| \leq 1} \gamma_i (\langle b_i(\theta), \theta \rangle - \alpha\kappa), \end{aligned}$$

so the overall objective can be written as:

$$\begin{aligned} & \sup_{\mathbb{P} \in \mathcal{P}(\theta)} \mathbb{E}_{\mathbb{P}}[\phi(\langle x, \theta \rangle) - y\langle x, \theta \rangle] \\ &= \begin{cases} \inf_{\alpha} \max_{\gamma: \|\gamma\|_{\infty} \leq 1} & \alpha(\delta - \kappa) + \frac{1}{n} \sum_i \frac{1+\tilde{y}_i}{2} (\phi(\langle b_i(\theta), \theta \rangle) + \gamma_i (\langle b_i(\theta), \theta \rangle - \alpha\kappa)) \\ & + \frac{1}{n} \sum_i \frac{1-\tilde{y}_i}{2} (\phi(\langle b_i(\theta), \theta \rangle) + \langle b_i(\theta), \theta \rangle) \\ \text{s.t.} & \|\theta\| \leq \alpha/(L+1) \end{cases} \end{aligned}$$

We claim that the minimax objective above is convex in θ . There are mainly two cases to analyze:

1. **Case I** ($i \in \mathcal{I}_{+1}$): We have $b_i(\theta) = \tilde{x}_i$ as per the strategic classification model. Therefore $\langle b_i(\theta), \theta \rangle$ is a linear function. For every γ, α , we claim that the mapping $\theta \mapsto \phi(\langle b_i(\theta), \theta \rangle) + \gamma_i (\langle b_i(\theta), \theta \rangle - \alpha\kappa)$ is convex. Indeed, the assumption that ϕ is convex and the observation that $\langle b_i(\theta), \theta \rangle$ is affine in θ ensures the convexity.
2. **Case II** ($i \in \mathcal{I}_{-1}$): We know from Lemma 5.2 that $\langle b_i(\theta), \theta \rangle$ is convex in θ . Moreover, the convexity of ϕ and the assumption that $z \mapsto \phi(z) + z$ is non-decreasing ensures that $\phi(\langle b_i(\theta), \theta \rangle) + \langle b_i(\theta), \theta \rangle$ is convex for every i .

This concludes the proof. \square

C Details on experimental study

Code used to reproduce the results in the main paper is available at <https://drive.google.com/drive/folders/1spuB3R6vEU2AqaXxAxeeXo9z5QMVdtdl?usp=sharing>

C.1 Algorithms

In our experiments, we compare the OGDA-RR algorithm (Alg. 1) with three other zeroth-order algorithms—Optimistic Gradient Descent Ascent with Sampling with Replacement (OGDA-WR), Stochastic Gradient Descent Ascent with Random Reshuffling (SGDA-RR), and Stochastic Gradient Descent Ascent with Sampling with Replacement (SGDA-WR)—characterized by the update equations (42), (43), (44), respectively. For convenience, we have reproduced (5), the update equation for the OGDA-RR algorithm (Algorithm 1), as (41) below:

$$u_{i+1}^t = \text{Proj}_{\mathcal{X} \times \mathcal{Y}} \left(u_i^t - \eta^t \hat{F}_{\sigma_i^t}(u_i^t; R^t, v_i^t) - \eta^t \hat{F}_{\sigma_{i-1}^t}(u_i^t; R^t, v_i^t) + \eta^t \hat{F}_{\sigma_{i-1}^t}(u_{i-1}^t; R^t, v_{i-1}^t) \right), \quad (41)$$

$$u_{i+1}^t = \text{Proj}_{\mathcal{X} \times \mathcal{Y}} \left(u_i^t - \eta^t \hat{F}_{j_i^t}(u_i^t; R^t, v_i^t) - \eta^t \hat{F}_{j_{i-1}^t}(u_i^t; R^t, v_i^t) + \eta^t \hat{F}_{j_{i-1}^t}(u_{i-1}^t; R^t, v_{i-1}^t) \right), \quad (42)$$

$$u_{i+1}^t = \text{Proj}_{\mathcal{X} \times \mathcal{Y}} \left(u_i^t - \eta^t \hat{F}_{\sigma_i^t}(u_i^t; R^t, v_i^t) \right), \quad (43)$$

$$u_{i+1}^t = \text{Proj}_{\mathcal{X} \times \mathcal{Y}} \left(u_i^t - \eta^t \hat{F}_{j_i^t}(u_i^t; R^t, v_i^t) \right), \quad (44)$$

where the indices σ_i^t and j_i^t are as defined in Algorithms 2, 3, and 4.

Algorithm 2: OGDA-WR Algorithm

```

1 Input: stepsizes  $\eta^t, R^t$ , data points  $\{(x_i, y_i)\}_{i=1}^n \sim \mathcal{D}, u_0^{(0)}$ , time horizon duration  $T$ ;
2 for  $t = 0, 1, \dots, T - 1$  do
3   for  $i = 0, \dots, n - 1$  do
4     Sample  $j_i^t \sim \text{Unif}(\{1, \dots, n\})$ 
5     Sample  $v_i^t \sim \text{Unif}(\mathcal{S}^{d-1})$ 
6      $u_{i+1}^t \leftarrow (42)$ 
7   end
8    $u_0^{(t+1)} \leftarrow u_n^t$ 
9    $u_{-1}^{(t+1)} \leftarrow u_{n-1}^t$ 
10 end
11 Output:  $\tilde{u}^T := \frac{1}{n \cdot \sum_{t=0}^{T-1} \eta^t} \sum_{t=0}^{T-1} \sum_{i=1}^n \eta^t u_i^t$ 

```

C.2 Additional Experimental Results

In this section, we present more experimental findings, on both synthetic and real-world datasets, that reinforces the utility of the proposed algorithm. In all experimental results throughout this subsection, we take $\delta = 0.4$, $\kappa = 0.5$ and $\zeta = 0.05$.

C.2.1 Experimental Study On Synthetic Datasets

Figure 2 compares the performance of (A-I)-(A-IV) on a synthetic dataset (whose generating process is the same as that described in Section 6), with 4000 training points and 800 test points. Our proposed algorithm performs better empirically compared to most of

Algorithm 3: SGDA-RR Algorithm

1 **Input:** stepsizes η^t, R^t , data points $\{(x_i, y_i)\}_{i=1}^n \sim \mathcal{D}, u_0^{(0)}$, time horizon duration T ;
2 **for** $t = 0, 1, \dots, T - 1$ **do**
3 **for** $i = 0, \dots, n - 1$ **do**
4 Sample $j_i^t \sim \text{Unif}(\{1, \dots, n\})$
5 Sample $v_i^t \sim \text{Unif}(\mathcal{S}^{d-1})$
6 $u_{i+1}^t \leftarrow (43)$
7 **end**
8 $u_0^{(t+1)} \leftarrow u_n^t$
9 $u_{-1}^{(t+1)} \leftarrow u_{n-1}^t$
10 **end**
11 **Output:** $\tilde{u}^T := \frac{1}{n \cdot \sum_{t=0}^{T-1} \eta^t} \sum_{t=0}^{T-1} \sum_{i=1}^n \eta^t u_i^t$.

Algorithm 4: SGDA-WR Algorithm

1 **Input:** stepsizes η^t, R^t , data points $\{(x_i, y_i)\}_{i=1}^n \sim \mathcal{D}, u_0^{(0)}$, time horizon duration T ;
2 **for** $t = 0, 1, \dots, T - 1$ **do**
3 $\sigma^t = (\sigma_1^t, \dots, \sigma_n^t) \leftarrow$ a random permutation of set $[n]$;
4 **for** $i = 0, \dots, n - 1$ **do**
5 Sample $v_i^t \sim \text{Unif}(\mathcal{S}^{d-1})$
6 $u_{i+1}^t \leftarrow (44)$
7 **end**
8 $u_0^{(t+1)} \leftarrow u_n^t$
9 $u_{-1}^{(t+1)} \leftarrow u_{n-1}^t$
10 **end**
11 **Output:** $\tilde{u}^T := \frac{1}{n \cdot \sum_{t=0}^{T-1} \eta^t} \sum_{t=0}^{T-1} \sum_{i=1}^n \eta^t u_i^t$.

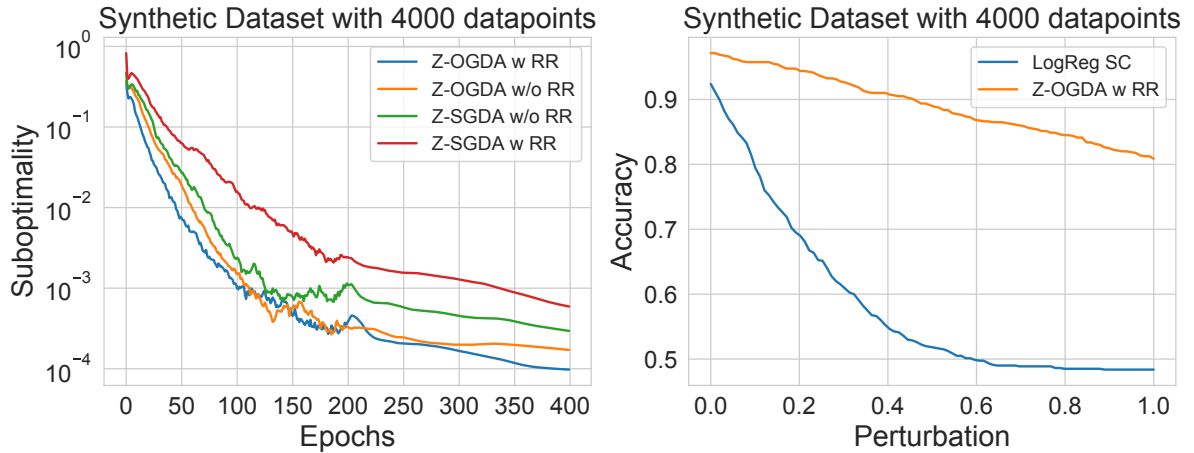


Figure 2: Experimental results for a synthetic dataset with $n = 4000$. (Left pane) Suboptimality iterates generated by the four algorithms (A-I), (A-II), (A-III), (A-IV), respectively denoted as *Z-OGDA w RR*, *Z-OGDA w/o RR*, *Z-SGDA w RR*, *Z-SGDA w/o RR*. (Right pane) Comparison between decay in accuracy of strategic classification with logistic regression (trained with $\zeta = 0.05$) and Alg. (A-I) with changes in perturbation.

its counterparts. Moreover, the proposed classifier, (A-I), is significantly more robust than a classifier obtained without considering adversarial perturbations. Note, however, that we cannot make any conclusive claims yet, because of the inherent randomness in these algorithms. Indeed, even if we fix the initialization, then there are two sources of randomness—the construction of the zeroth-order gradient estimator, and the sampling process that generates the data points.

To illustrate the variability in these algorithms’ performance, we run each algorithm repeatedly on a data set with 500 synthetically generated data points, using the same initialization, and present confidence interval plots with ± 2 standard deviations for the resulting performance (Figure 3). On average, our proposed algorithm (A-I) outperforms the other algorithms (A-II)-(A-IV). It is also interesting to point out that the performance of algorithms with random reshuffling is generally higher, and fluctuate less, compared to the performance of algorithms without random reshuffling.

We now illustrate the performance of our algorithm on two real-world data sets—the “GiveMeSomeCredit” dataset³, and the “Porto Bank” data set⁴.

C.2.2 Experimental Study on Credit Dataset

In modern times, banks use machine learning to determine whether or not to finance a customer. This process can be encoded into a classification framework, by using features such as age, debt ratio, monthly income to classify a customer as either likely or unlikely to default. However, those algorithms generally do not account for strategic or adversarial behavior on the part of the agents.

To illustrate the effect of our algorithm on datasets of practical significance, we deploy our algorithms on the “GiveMeSomeCredit”(GMSC) dataset, while assuming that the underlying features are subject to strategic or adversarial perturbations. We use a subset of the dataset of size 2000 with balanced labels. In Figure 4, we compare the empirical performance of our algorithm (A-I) with that of (A-II)-(A-IV). The left pane shows that (A-I) performs well, and the right pane illustrates that our classifier is significantly more robust to adversarial perturbations in data, compared to the strategic classification-based

³This dataset can be found at <https://www.kaggle.com/c/GiveMeSomeCredit>

⁴This dataset can be found at <https://archive.ics.uci.edu/ml/datasets/bank+marketing>




Figure 3: Experimental results for a synthetic dataset with $n = 500$. Suboptimality iterates generated by the four algorithms (A-I), (A-II), (A-III), and (A-IV) are respectively denoted as *Z-OGDA w RR*, *Z-OGDA w/o RR*, *Z-SGDA w RR*, and *Z-SGDA w/o RR*.

logistic regression algorithm developed recently in the literature [DRS⁺18].

C.2.3 Experimental Study on Porto-Bank Dataset

Next, we present empirical results obtained by applying our algorithm to the “Porto-Bank” dataset, which describes marketing campaigns of term deposits at Portuguese financial institutions. The classification task in this scenario aims to predict whether a customer with given features (eg. age, job, marital status etc.) would enroll for term deposits.

In Figure 5, we present the performance of our proposed algorithm (A-I) on the Porto-Bank dataset. For ease of illustration, we consider a subset of the dataset with 2000 training data points, 800 test data points, and balanced labels. In Figure 5, we compare the empirical performance of our algorithm (A-I) with that of (A-II)-(A-IV). The left pane shows that (A-I) performs well, while the right pane illustrates that our classifier is significantly more robust to adversarial perturbations in data, compared to the strategic classification-based logistic regression developed recently in the literature [DRS⁺18].

C.2.4 Effect of n, d on sample complexity

In this part, we demonstrate the empirical results that corroborates the theoretical dependence of sample complexity on n, d . For this purpose, we use synthetic dataset which is generated as per the method described in Section 6.1. Here we work in the setting where $n \in \{500, 1000, 1500, 2000\}$ and $d \in \{10, 15, 20, 25\}$. We fix the suboptimality to $\epsilon = 0.1$ and compute the number of samples required in each of the settings of n and d so that the iterates reach the ϵ -suboptimality. We present the results in Figure 6.

C.3 Logistic regression as a Generalized linear model

The goal in logistic regression is to maximize the log-likelihood of the conditional probability of y (the *label*) given x (the *feature*). In this model, it is assumed that:

$$P(Y = 1|x, \theta) = \frac{1}{1 + \exp(-\langle x, \theta \rangle)}$$

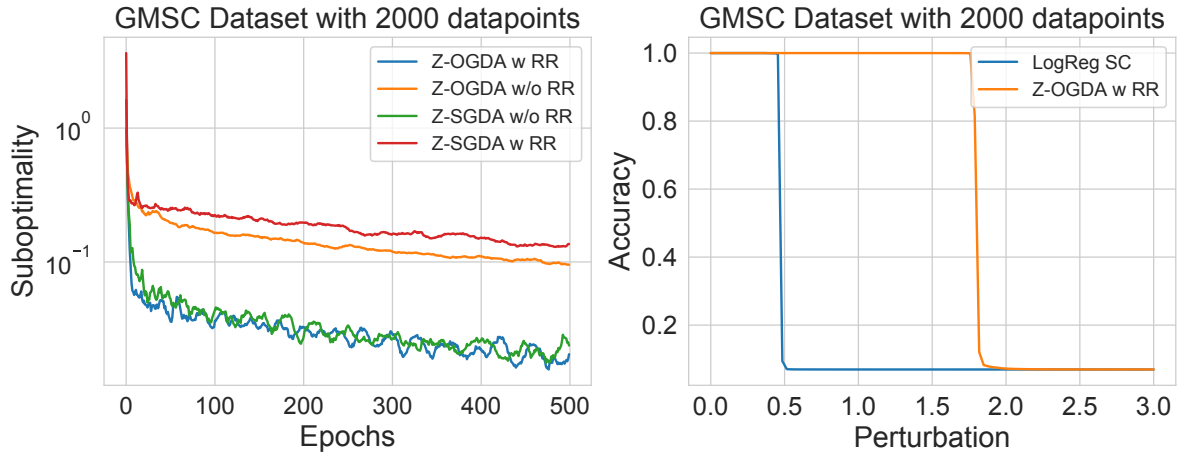


Figure 4: Experimental results for a balanced GiveMeSomeCredit dataset with $n = 2000$. (Left pane) Suboptimality iterates generated by the four algorithms (A-I), (A-II), (A-III), (A-IV), respectively denoted as $Z\text{-OGDA } w \text{ RR}$, $Z\text{-OGDA } w/o \text{ RR}$, $Z\text{-SGDA } w \text{ RR}$, $Z\text{-SGDA } w/o \text{ RR}$. (Right pane) Comparison between decay in accuracy of strategic classification with logistic regression (originally trained with $\zeta = 0.05$) and Alg. (A-I) with changes in perturbation.

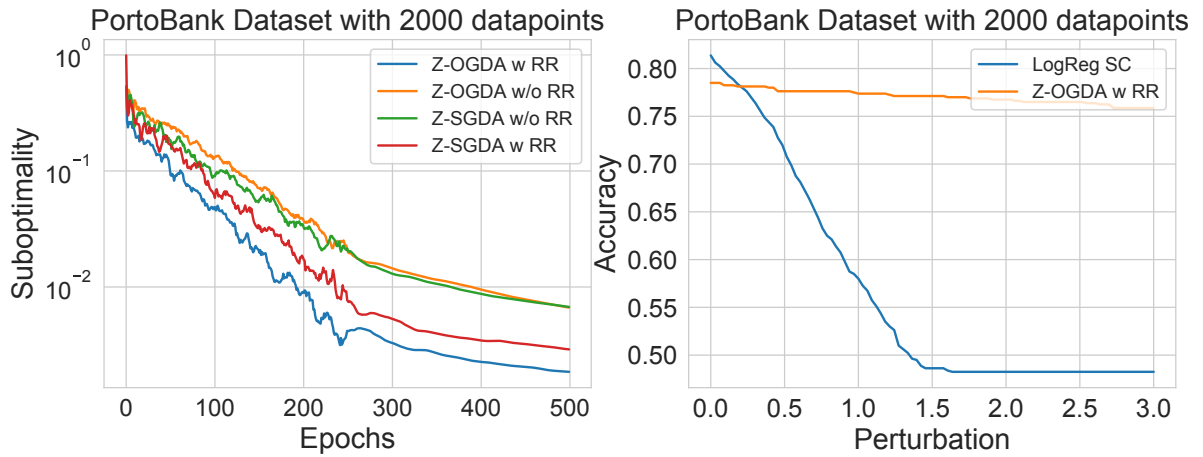


Figure 5: Experimental results for a balanced PortoBank dataset with $n = 2000$. (Left pane) Suboptimality iterates generated by the four algorithms (A-I), (A-II), (A-III), (A-IV), respectively denoted as $Z\text{-OGDA } w \text{ RR}$, $Z\text{-OGDA } w/o \text{ RR}$, $Z\text{-SGDA } w \text{ RR}$, $Z\text{-SGDA } w/o \text{ RR}$. (Right pane) Comparison between decay in accuracy of strategic classification with logistic regression (originally trained with $\zeta = 0.05$) and Alg. (A-I) with change in perturbation.




Figure 6: Experimental results presenting the number of samples required to reach ϵ -suboptimality, with $\epsilon = 0.1$, for our algorithm (A-I) on synthetic dataset with varying values of $n \in \{500, 1000, 1500, 2000\}$ and $d \in \{10, 15, 20, 25\}$.

This implies that:

$$P(Y = -1|x, \theta) = \frac{\exp(-\langle x, \theta \rangle)}{1 + \exp(-\langle x, \theta \rangle)}$$

Given a data point (x, y) the logistic loss is log-likelihood of observing y given x . For any θ and $y \in \{-1, 1\}$:

$$P(Y = y|x; \theta) = (P(Y = 1|x, \theta))^{\frac{1+y}{2}} (P(Y = -1|x, \theta))^{\frac{1-y}{2}}$$

Now, the log-likelihood is given by:

$$\begin{aligned} L(x, y; \theta) &= \log(P(Y = y|x; \theta)) \\ &= \frac{1+y}{2} \log\left(\frac{1}{1 + \exp(-\langle x, \theta \rangle)}\right) + \frac{1-y}{2} \log\left(\frac{\exp(-\langle x, \theta \rangle)}{1 + \exp(-\langle x, \theta \rangle)}\right) \\ &= -\frac{1-y}{2} \langle x, \theta \rangle + \left(\frac{1+y}{2} + \frac{1-y}{2}\right) \log\left(\frac{1}{1 + \exp(-\langle x, \theta \rangle)}\right) \\ &= -\frac{1-y}{2} \langle x, \theta \rangle - \log(1 + \exp(-\langle x, \theta \rangle)) \\ &= \frac{y}{2} \langle x, \theta \rangle - \frac{1}{2} \langle x, \theta \rangle + \langle x, \theta \rangle - \log(1 + \exp(\langle x, \theta \rangle)) \\ &= \frac{y}{2} \langle x, \theta \rangle + \frac{1}{2} \langle x, \theta \rangle - \log(1 + \exp(\langle x, \theta \rangle)) \end{aligned}$$

The goal is to maximize the log likelihood, which is equivalent to minimizing the negative log likelihood. Thus the logistic regression minimizes the following loss:

$$\tilde{L}(x, y; \theta) = -L(x, y; \theta) = -\frac{y}{2} \langle x, \theta \rangle + \phi(\langle x, \theta \rangle)$$

where $\phi(\beta) = \log(1 + \exp(\beta)) - \frac{\beta}{2}$. If $y = 1$, then the above loss becomes:

$$\log(1 + \exp(\langle x, \theta \rangle)) - \langle x, \theta \rangle = \log(1 + \exp(-\langle x, \theta \rangle))$$

Otherwise, if $y = -1$, then the above loss becomes $\log(1 + \exp(\langle x, \theta \rangle))$. Thus, the above loss is equivalent to $\log(1 + \exp(-y \langle x, \theta \rangle))$.