A STOCHASTIC MULTI-GRADIENT ALGORITHM FOR WASSERSTEIN DISTRIBUTIONALLY ROBUST MULTI-OBJECTIVE OPTIMIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Distributionally robust multi-objective optimization (DRMOO) has been proposed to improve the robustness of multi-task learning models under distribution shifts. However, existing DRMOO methods mainly construct ambiguity sets based on f-divergence, which may have limited capability in capturing the geometric structure of the sample space and changes in the support of distributions. To address this limitation, we propose a Wasserstein distributionally robust multi-objective optimization algorithm, which leverages the geometry of the sample space to characterize distributional uncertainty more flexibly. Specifically, by applying the dual reformulation of Wasserstein distributionally robust optimization, each robust objective is reformulated as an expectation involving an inner adversarial maximization problem. Based on this, we develop a double-loop optimization algorithm, where the inner loop employs K-step gradient ascent to approximately solve the inner maximization problem, and the outer loop adopts an MGDA-type method to construct a common descent direction among multiple objectives. Theoretically, under non-convex and generalized ℓ-smooth assumptions, we establish the convergence guarantees of the proposed algorithm, proving that it achieves an ϵ-Pareto stationary point with a total sample complexity of $\mathcal{O}(\epsilon^{-4})$.

1 INTRODUCTION

Multi-objective optimization (MOO) has become an important research direction in machine learning for problems involving multiple potential conflicting objectives, such as multi-task learning Dimitriadis et al. (2025); Ban & Ji (2024), recommender systems (Zhang et al., 2024a; Fan et al., 2025), and autonomous driving (He & Lv, 2023). In general, there does not exist a single solution that simultaneously minimizes all objectives in MOO problem. Therefore, the goal of MOO is to seek Pareto optimal solutions that achieve optimal trade-offs among objectives. Among gradient-based methods, the multiple-gradient descent algorithm (MGDA) is one of the most representative approaches, which constructs a common descent direction for all objectives at each iteration. It has subsequently been adapted to machine learning, particularly multi-task learning, and has demonstrated promising empirical performance across various applications (Sener & Koltun, 2018; Fernando et al., 2023; Zhang et al., 2026).

Despite the recent progress of MGDA-based methods, most existing approaches optimize multiple objectives under a fixed empirical training distribution, without explicitly considering potential distribution shifts. However, in real-world applications, models often encounter data that may deviate from the empirical training distribution due to domain mismatch, changes in class proportions, or adversarial perturbations (Yu et al., 2025; Yang et al., 2024; 2026; Watkins et al., 2024). Thus, a Pareto optimal solution obtained under the empirical data distribution may no longer provide an optimal trade-off among objectives when the underlying data distribution changes. Distributionally robust optimization (DRO) (Zhang et al., 2025b; Jin et al., 2021; Zhang et al., 2024b; Micheli et al.) provides a well-established framework for handling distribution shifts in single-objective optimization. Instead of optimizing solely with respect to the empirical training distribution, DRO constructs an ambiguity set containing plausible distributions and optimizes against the worst-case distribution within this set.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Table 1: Comparison of stochastic multi-objective optimization methods. Explanation on the upper footmarks: (LS) indicates that the objectives are standard L-smooth while (GS) the objectives are generalized ℓ-smooth as defined in Assumption 1. (Note: Reference Yang et al. (2026) assumes that the original objective functions are globally L-smooth.)

Methods	**Original-Objective Smoothness**	**Batch Size**	**Sample Complexity**
MoCo Fernando et al. (2023)	LS	$\mathcal{O}(1)$	$\mathcal{O}(\epsilon^{-4})$
SDMGrad Xiao et al. (2023)	LS	$\mathcal{O}(1)$	$\mathcal{O}(\epsilon^{-4})$
MoDo Chen et al. (2024)	LS	$\mathcal{O}(1)$	$\mathcal{O}(\epsilon^{-4})$
MGDA Zhang et al. (2025a)	GS	$\mathcal{O}(1)$	$\mathcal{O}(\epsilon^{-4})$
Double-Loop MGDA Yang et al. (2026)	LS	$\mathcal{O}(1)$	$\mathcal{O}(\epsilon^{-12})$
Double-Clip MGDA Yang et al. (2026)	LS	$\mathcal{O}(\epsilon^{-2})$	$\mathcal{O}(\epsilon^{-4})$
WDR-MOO	GS	$\mathcal{O}(1)$	$\mathcal{O}(\epsilon^{-4})$

Motivated by the success of DRO in single-objective optimization, recent studies have extended DRO to multi-objective settings. In particular, Yang et al. (2026) proposed a DR-MOO framework, which constructs an ambiguity set based on f-divergences and develops MGDA-type algorithms with provable convergence guarantees. However, it still leaves several limitations. First, f-divergences primarily characterize distributional discrepancies through changes in probability mass and typically rely on an absolute continuity relationship between the nominal and perturbed distributions (Ben-Tal et al., 2013; Kuhn et al., 2025), which may be less suitable for capturing distribution shifts involving geometric displacement or changes in distributional support (Gao & Kleywegt, 2023). In contrast, the Wasserstein distance measures distributional discrepancies through an optimal transport cost and explicitly incorporates the geometry of the underlying sample space (Blanchet & Murthy, 2019; Montesuma et al., 2025). These properties make Wasserstein-based ambiguity sets particularly suitable for modeling distributional perturbations involving geometric displacement or support variation. On the optimization side, (Yang et al., 2026) assumes that the original loss functions are globally L-smooth, although a generalized (L_0, L_1)-smooth property can be derived for its dual of DR-MOO formulation. Nevertheless, a recent study Zhang et al. (2020b) indicates that global constant smoothness assumptions may not necessarily be true for the training of neural networks. Generalized smoothness provides a more flexible framework by allowing the local smoothness to depend on the gradient magnitude. In particular, the generalized ℓ-smooth framework (Li et al., 2023) characterizes the local smoothness through a general non-decreasing function $\ell(\cdot)$, thereby covering a broader class of non-convex objectives. Both standard L-smoothness and (L_0, L_1)-smoothness can be recovered as special cases of this framework. Therefore, the following question naturally arises:

Can we develop a Wasserstein distributionally robust multi-objective optimization method with provable convergence guarantees under the more general ℓ-smooth condition?

To answer this question, we propose a Wasserstein distributionally robust multi-objective optimization (WDR-MOO) framework, in which each objective is optimized against distributional perturbations characterized by a Wasserstein ambiguity set. By exploiting the dual reformulation of Wasserstein DRO (Sinha et al., 2018; Li et al., 2024), each robust objective can be represented as an expectation involving an inner adversarial maximization problem in the sample space. Based on this formulation, we develop a stochastic multi-gradient algorithm with a double-loop structure. In the inner loop, a K-step gradient ascent procedure approximately solves the adversarial maximization problem for each objective and generates stochastic estimates of the corresponding robust gradients. In the outer loop, an MGDA-type update combines these estimated gradients to construct a common descent direction for all robust objectives. Table 1 presents a detailed comparison of stochastic multi-objective methods. Our main contributions are summarized as follows:

- We propose a WDR-MOO framework together with a double-loop stochastic multi-gradient algorithm. The inner loop uses K-step gradient ascent to approximate the worst-case perturbation, while the outer loop employs an MGDA-type update to construct a common descent direction for multiple robust objectives.

- We establish convergence guarantees for the proposed method under nonconvex settings and the generalized ℓ-smoothness. In particular, we prove that the proposed method achieves an ϵ-Pareto stationary solution with $\mathcal{O}(\epsilon^{-4})$ sample complexity.

2 PRELIMINARIES

In this section, we introduce the basic notions of MOO, multiple gradient descent algorithm (MGDA), and Wasserstein distributionally robust optimization.

2.1 MULTI-OBJECTIVE OPTIMIZATION

Unlike single-objective optimization, multi-objective optimization aims to simultaneously optimize multiple potentially conflicting objectives. Mathematically, multi-objective optimization problem can be formulated as

$$\min_{x\in\mathbb{R}^d} F(x) = [f_1(x), f_2(x), \ldots, f_m(x)], \tag{2.1}$$

where $f_i(x) = \mathbb{E}_{w\sim\mathbb{P}_i}[f_i(x,w)]$, x is a shared decision variable for all objectives, w is a random variable sampled from distribution $\mathbb{P}$, and m denotes the number of objectives.

Since the objectives may conflict with each other, there generally does not exist a single solution that minimizes all objectives simultaneously. Therefore, the goal of MOO is generally to find a Pareto optimal points.

Definition 2.1 (Pareto optimality)**.** *Given two solutions $x_1, x_2 \in \mathbb{R}^d$, x_1 is said to dominate x_2, if and only if $\forall i \in \{1,2,\ldots,m\}$, we have $f_i(x_1) \le f_i(x_2)$ and $F(x_1) \neq F(x_2)$. A solution x^* is called a Pareto optimal point if there does not exist another point that dominates it.*

Under non-convex settings, MOO methods commonly seek a Pareto stationary point, which is a necessary condition for Pareto optimality.

Definition 2.2 (Pareto stationary)**.** *A point x is called Pareto stationary, if*

$$range(\nabla F(x)) \cap (-\mathbb{R}^m_{++}) = \emptyset, \tag{2.2}$$

where $\nabla F(x) = [\nabla f_1(x), \nabla f_2(x), \ldots, \nabla f_m(x)]$, and $\mathbb{R}^m_{++}$ is the positive orthant cone.

Following the results in (Tanabe et al., 2019; Chen et al., 2024), a point x is a Pareto stationary point if and only if

$$\min_{\lambda\in\Delta^m} \|\nabla F(x)\lambda\| = 0 \tag{2.3}$$

where $\Delta^m = \{\lambda \in \mathbb{R}^m | \sum_{i=1}^m \lambda_i = 1, \lambda_i \ge 0\}$. Therefore, Eq. (2.3) provides a convergence criterion for MOO in non-convex settings.

Definition 2.3 (ϵ-Approximate Pareto stationary)**.** *Given a constant $\epsilon > 0$, if*

$$\min_{\lambda\in\Delta^m} \|\nabla F(x)\lambda\| \le \epsilon, \tag{2.4}$$

the point x is called an ϵ-Approximate Pareto stationary point.

2.2 MULTI-GRADIENT DESCENT ALGORITHM

The Multi-Gradient Descent Algorithm (MGDA) is a representative gradient-based method for solving multi-objective optimization problems. Its core idea is to construct a common descent direction that decreases all objectives simultaneously. At each iteration, MGDA first computes an optimal weight vector $\lambda^*(x_t)$ by solving

$$\lambda_t^* \in \arg\min_{\lambda\in\Delta^m} \frac{1}{2}\|\nabla F(x_t)\lambda\|^2 + \frac{\rho}{2}\|\lambda\|^2, \tag{2.5}$$

where ρ is a regularization parameter. The corresponding multi-gradient direction is defined as $\nabla F(x_t)\lambda^*(x_t)$. Then, the decision variable is updated according to

$$x_{t+1} = x_t - \eta\nabla F(x_t)\lambda_t^*, \tag{2.6}$$

where $\eta > 0$ is the learning rate.

2.3 WASSERSTEIN DISTRIBUTIONALLY ROBUST OPTIMIZATION

Distributionally robust optimization (DRO) considers optimization problems where the underlying data distribution is uncertain. Instead of optimizing the expected objective under a fixed distribution, DRO optimizes against the worst-case distribution within an ambiguity set.

For a given nominal distribution $\mathbb{P}$, the Wasserstein ambiguity set can be defined as

$$\mathcal{P} = \{\mathbb{Q} : W_c(\mathbb{Q}, \mathbb{P}) \leq \nu\}, \tag{2.7}$$

where ν is the ambiguity set radius and W_c denotes the Wasserstein distance, defined as

$$W_c(\mathbb{Q}, \mathbb{P}) := \inf_{\pi \in \Pi(\mathbb{Q}, \mathbb{P})} \left\{\mathbb{E}_{(w', w) \sim \pi} \left[c(w', w)\right]\right\}, \tag{2.8}$$

where $\Pi(\mathbb{Q}, \mathbb{P})$ represents the set of all couplings with marginals $\mathbb{P}$ and $\mathbb{Q}$, and $c(w', w)$ is the transportation cost function.

The Wasserstein DRO formulation is given by

$$\min_{x \in \mathbb{R}^d} \max_{\mathbb{Q} \in \mathcal{P}} \mathbb{E}_{w \sim \mathbb{Q}} \left[f(x, w)\right]. \tag{2.9}$$

This is a minmax optimization problem, where inner maximization searches for the worst-case distribution within the Wasserstein ambiguity set, the outer seeks a decision variable, which minimizes corresponding worst-case expected loss. Following Sinha et al. (2018), we consider the following penalized Wasserstein DRO problems:

$$\min_{x \in \mathbb{R}^d} \sup_{\mathbb{Q}} \mathbb{E}_w \left[f(x, w)\right] - \gamma W_c(\mathbb{Q}, \mathbb{P}), \tag{2.10}$$

where γ is a penalty parameter. This transformation uses Lagrangian relaxation to convert the hard constraint that the distribution $\mathbb{Q}$ lies within the Wasserstein ambiguity set into a penalty term. The penalized formulation avoids the additional optimization over the dual variable associated with the Wasserstein constraint, thereby simplifying the resulting optimization problem.

3 WASSERSTEIN DISTRIBUTIONALLY ROBUST MULTI-OBJECTIVE OPTIMIZATION ALGORITHM

We extend Wasserstein DRO to the multi-objective setting. Specifically, we formulate Wasserstein distributionally robust multi-objective optimization problem as

$$\min_{x \in \mathbb{R}^d} \sup_{\mathbb{Q}_1, \mathbb{Q}_2, \ldots, \mathbb{Q}_m} \left[H_1(x, \mathbb{Q}_1), H_2(x, \mathbb{Q}_2), \ldots, H_m(x, \mathbb{Q}_m)\right], \tag{3.1}$$

where $H_i(x, \mathbb{Q}_i) = \mathbb{E}_{w \sim \mathbb{Q}_i} \left[f(x, w)\right] - \gamma W_c(\mathbb{Q}_i, \mathbb{P}_i)$. Directly solving the Wasserstein problem involves optimization over probability distributions and is generally computationally challenging. Following the lemma 4.16 in (Kuhn et al., 2025), the penalized Wasserstein DRO objective can be reformulated as

$$\sup_{\mathbb{Q}_i} H_i(x, \mathbb{Q}_i) = \mathbb{E}_{w \sim \mathbb{P}_i} \left[\sup_{w'} \left\{f_i(x, w') - \frac{\gamma}{2} \left\|w' - w\right\|^2\right\}\right], \tag{3.2}$$

where we adopt the transport cost function as $c(w', w) = \frac{1}{2} \left\|w' - w\right\|^2$. This reformulation converts the optimization over probability distributions into a sample-wise inner maximization problem, which can be approximately solved by gradient-based methods.

To simplify the notation, we define

$$h_i(x, w', w) = f_i(x, w') - \frac{\gamma}{2} \left\|w' - w\right\|^2, \tag{3.3}$$

and its corresponding inner maximization function

$$h_i^*(x, w) = \max_{w'} h_i(x, w', w). \tag{3.4}$$

Then, the i-the objective function can be written as

$$H_i(x) = \mathbb{E}_{w \sim \mathbb{P}_i} \left[h_i^*(x, w)\right]. \tag{3.5}$$

Algorithm 1 WDR-MOO

Init: x_0, $\lambda_0 \in \Delta^m$, learning rate $\alpha, \beta, \eta > 0$, regularization parameter $\rho > 0$, and inner iteration number K.
for $t = 0, \ldots, T-1$ **do**
  for $i = 1, \ldots, m$ **do**
    Draw $w_{t,i,j} \sim \mathbb{P}_i$ independently.
    Initialize $w'^{,0}_{t,i,j} = w_{t,i,j}$.
    Perform K-step, according to $w'^{,k+1}_{t,i,j} = w'^{,k}_{t,i,j} + \alpha \nabla_{w'} h_i(x_t, w'^{,k}_{t,i,j}, w_{t,j})$.
    Evaluate $y_{t,i,j} = \nabla_x f_i(x_t, w'^{,K}_{t,i,j})$.
    Construct $Y_{t,j} = [y_{t,1,j}, \ldots, y_{t,m,j}]$.
  end for
  Compute weight $\lambda_t = \Pi_{\Delta^m}\left[\lambda_{t-1} - \beta\left(Y_{t,2}^\top Y_{t,3}\lambda_{t-1} + \rho\lambda_{t-1}\right)\right]$.
  Update decision variable $x_{t+1} = x_t - \eta Y_{t,1}\lambda_t$.
end for

The original Wasserstein distributionally robust multi-objective optimization problem is transformed as

$$\min_{x\in\mathbb{R}^d} H(x) = [H_1(x), H_2(x), \ldots, H_m(x)]\,. \tag{3.6}$$

Although, Eq. (3.6) has a standard MOO form, each objective H_i involves an inner maximization problem, which makes its gradient difficult to compute directly. Therefore, we adopt a double-loop optimization framework. In the inner loop, we employ a K-step gradient ascent to approximately solve the inner maximization problem. Specifically, for the i-th objective and the j-th independent sample $w_{t,j}, j \in \{1,2,3\}$, the variable w' is updated as

$$w'^{,k+1}_{t,i,j} = w'^{,k}_{t,i,j} + \alpha \nabla_{w'} h_i(x_t, w'^{,k}_{t,i,j}, w_{t,j}), \tag{3.7}$$

where $\alpha > 0$ is the learning rate. After K-step, $w'^{,K}_{t,i,j}$ is used as approximation of

$$w'^{,*}_{t,i,j} \in \arg\max_{w'} h_i(x_t, w', w_{t,j}). \tag{3.8}$$

Therefore, the gradient of the i-th objective is approximated by

$$y_{t,i,j} = \nabla_x h_i(x_t, w'^{K}_{t,i,j}, w_{t,j}) = \nabla_x f_i(x_t, w'^{K}_{t,i,j}). \tag{3.9}$$

To address the conflicts among objectives, we employ an MGDA-type method in the outer loop to construct a common descent direction. However, the exact gradient matrix $\nabla H(x_t)$ is unavailable because each $H_i(x)$ involves both an expectation and an inner maximization. Directly using stochastic gradient estimators in Eq. (2.5) to obtain optimal weight λ generally introduces an additional bias Fernando et al. (2023). To reduce this bias, we adopt a double-sampling stochastic gradient scheme to update the weight vector:

$$\lambda_{t+1} = \Pi_{\Delta^m}\left(\lambda_t - \beta\left(Y_{t,2}^\top Y_{t,3}\lambda_t + \rho\lambda_t\right)\right), \tag{3.10}$$

where $Y_{t,j} = [y_{t,1,j}, y_{t,2,j}, \ldots, y_{t,m,j}]$, β is the learning rate for the weight update, ρ is a regularization parameter to ensure the optimal λ is unique, and $\Pi_{\Delta^m}(\cdot)$ denotes the projection onto the probability simplex.

Then, the decision variable is updated following as

$$x_{t+1} = x_t - \eta Y_{t,1}\lambda_t, \tag{3.11}$$

where η is a learning rate. The complete procedure of our algorithm is summarized in Algorithm 1.

4 Convergence Analysis

In this section, we establish the convergence guarantee of the proposed algorithm under generalized ℓ-smoothness. We first introduce the assumptions required in the analysis and then present our main convergence result. The detailed proofs are deferred to Appendix B.3.

Assumption 1 (Generalized Smoothness). *For each $i \in \{1, 2, \ldots, m\}$, the function $f_i(x, w)$ satisfies generalized ℓ_x-smoothness with respect to x and global L_w-smoothness with respect to w. Where ℓ_x is a continuous non-decreasing and sub-quadratic function, defined as $\ell_x(u) = L_0 + L_1 u^v$, where constant $L_0, L_1 > 0, \quad 0 \leq v < 2$. For any $z_0 = (x_0, w_0)$ and $z_1 = (x_1, w_1), z_2 = (x_2, w_2) \in \mathcal{B}(z_0, r_x(z_0))$, we have*

$$\|\nabla_x f_i(z_1) - \nabla_x f_i(z_2)\| \leq \ell_x \left(\|\nabla_x f_i(z_0)\| + G\right) \|z_1 - z_2\|, \tag{4.1}$$

where $G > 0$ and $r_x(z_0) = \frac{G}{\ell_x(\|\nabla_x f_i(z_0)\| + G)}$. Moreover, for any z_1, z_2, we have

$$\|\nabla_w f_i(z_1) - \nabla_w f_i(z_2)\| \leq L_w \|z_1 - z_2\|. \tag{4.2}$$

This assumption generalizes the classical smoothness condition by allowing the local smoothness with respect to x to depend on the gradient norm. In particular, $v = 0$ recovers the standard global Lipschitz-smooth setting, while $v = 1$ corresponds to the (L_0, L_1)-smoothness.

Assumption 2 (Strong concavity). *For each $i \in \{1, 2, \ldots, m\}$, function $h_i(x, w', w)$ is μ_h-strongly concave with respect to w', where $\mu_h = \gamma - L_w > 0$.*

This assumption ensures that the inner adversarial maximization admits a unique maximizer and enables a geometric contraction of the K-step gradient ascent procedure.

Assumption 3 (Bounded Fourth Moment). *For each $i \in \{1, 2, \ldots, m\}$ and any $x \in \mathbb{R}^d$, the gradient satisfies*

$$\mathbb{E}_w \left[\|\nabla_x h_i^*(x, w) - \nabla H_i(x)\|^4\right] \leq \sigma^4, \tag{4.3}$$

where $\sigma > 0$ is a constant. By Jensen's inequality, we have

$$\mathbb{E}_{w\sim\mathbb{P}} \left[\|\nabla_x h_i^*(x, w) - \nabla H_i(x)\|^2\right] \leq \sqrt{\mathbb{E}_{w\sim\mathbb{P}} \left[\|\nabla_x h_i^*(x, w) - \nabla H_i(x)\|^4\right]} \leq \sigma^2. \tag{4.4}$$

Assumption 4 (Bounded Inner Initialization Gap). *There exists a constant $C_h > 0$ such that, for any iteration t, objective i, and sample index j, we have*

$$\frac{2}{\mu_h} \left[h_i(x_t, w'^{,*}_{t,i,j}, w_{t,j}) - h_i(x_t, w'^{,0}_{t,i,j}, w_{t,j})\right] \leq C_h, \tag{4.5}$$

where $w'^{,*}_{t,i,j}$ denotes the exact maximizer of the corresponding inner problem.

This assumption uniformly bounds the initial optimality gap of the inner maximization. Together with the strong concavity of the inner problem, it provides a uniform bound on the initial distance to the inner maximizer and is used to control the approximation error after a finite number of gradient-ascent steps.

Theorem 4.1. *Suppose that Assumptions 1–4 hold. For any $\epsilon > 0$ and $0 < \delta < 1/2$, let the parameters $\eta = \mathcal{O}(\epsilon^2)$, $\beta = \mathcal{O}(\epsilon^2)$, $\rho = \mathcal{O}(\epsilon^2)$, $K = \Theta(\log(1/\epsilon))$, and $T = \Theta(\epsilon^{-4})$. Then, with probability at least $1 - \delta$, we have*

$$\frac{1}{T} \sum_{t=0}^{T-1} \|\nabla H(x_t)\lambda_t^*\|^2 \leq 2\epsilon^2, \tag{4.6}$$

where $\lambda_t^* \in \arg\min_{\lambda\in\Delta^m} \|\nabla H(x_t)\lambda\|^2$.

Theorem 4.1 provides a high-probability convergence guarantee for the proposed WDR-MOO algorithm under non-convex settings and generalized ℓ-smoothness. Specifically, to achieve ϵ-Pareto stationarity, the algorithm requires $T = \Theta(\epsilon^{-4})$ outer iterations and only $K = \Theta(\log(1/\epsilon))$ inner gradient-ascent steps. Since each outer iteration uses a constant number of stochastic samples, the resulting sample complexity is $\mathcal{O}(\epsilon^{-4})$. Compared with existing DR-MOO methods Yang et al. (2026), which either incur $\mathcal{O}(\epsilon^{-12})$ total sample complexity or require a batch size of $\mathcal{O}(\epsilon^{-2})$. Our method achieves the standard$\mathcal{O}(\epsilon^{-4})$ sample complexity while retaining a constant batch size. Moreover, our convergence analysis is established directly under generalized ℓ-smoothness of the original objectives, rather than requiring global L-smoothness of the original loss functions.

Proof sketch. The proof consists of three main steps. First, by the strong concavity of the inner maximization problem, the K-step gradient-ascent yields a geometric contraction of the inner approximation error. This allows us to decompose the stochastic gradient error as

$$E_{t,j} = R_{t,j} + S_{t,j}, \tag{4.7}$$

where $R_{t,j}$ denotes the error caused by the inexact inner solution and decays geometrically with K, while $S_{t,j}$ denotes the stochastic sampling error and has bounded moments.

Second, we introduce the stopping time $\tau = \min\{\tau_1, \tau_2, \tau_3, \tau_4\}$ to simultaneously control the objective gap, stochastic gradient errors, error-product terms, and the local region required by generalized ℓ-smoothness. On the event $t < \tau$, the gradients of the robust objectives remain bounded, so the generalized smoothness condition reduces to a uniform local smoothness bound along the optimization trajectory. Combining this descent property with the update of weight and the independent double-sampling scheme, and summing over $t = 0, \ldots, \tau - 1$, yields the stopped descent lemma:

$$\mathbb{E}[H(x_\tau)\lambda] - \inf_x H(x)\lambda \leq \frac{\delta F}{16} - \eta \mathbb{E}\left[\sum_{t=0}^{\tau-1} \|Q_t \lambda_{t-1}\|^2\right]. \tag{4.8}$$

Finally, we complete the high-probability analysis in two steps. First, we bound the probability that the stopping time occurs before T. By separately controlling the events associated with τ_1, τ_2, τ_3, and τ_4, and applying the union bound, we show that

$$\mathbb{P}(\tau < T) \leq \frac{\delta}{2}. \tag{4.9}$$

Therefore,

$$\mathbb{P}(\tau = T) \geq 1 - \frac{\delta}{2}. \tag{4.10}$$

Second, on the event $\{\tau = T\}$, all the required bounds remain valid throughout the first T iterations, and the stopped descent inequality reduces to the full-trajectory bound

$$\frac{1}{T}\sum_{t=0}^{T-1} \|Q_t \lambda_{t-1}\|^2. \tag{4.11}$$

Using the descent lemma together with the parameter choices in Theorem 4.1, we further show that

$$\mathbb{P}\left(\frac{1}{T}\sum_{t=0}^{T-1} \|Q_t \lambda_{t-1}\|^2 > \epsilon^2,\ \tau = T\right) \leq \frac{\delta}{2}. \tag{4.12}$$

Combining the two probability bounds gives

$$\mathbb{P}\left(\frac{1}{T}\sum_{t=0}^{T-1} \|Q_t \lambda_{t-1}\|^2 > \epsilon^2\right) \leq \mathbb{P}(\tau < T) + \mathbb{P}\left(\frac{1}{T}\sum_{t=0}^{T-1} \|Q_t \lambda_{t-1}\|^2 > \epsilon^2,\ \tau = T\right) \leq \delta. \tag{4.13}$$

Hence, with probability at least $1 - \delta$,

$$\frac{1}{T}\sum_{t=0}^{T-1} \|Q_t \lambda_{t-1}\|^2 \leq \epsilon^2. \tag{4.14}$$

5 Conclude

In this paper, we propose a distributionally robust multi-objective optimization algorithm with a double-loop structure. Under nonconvex and generalized ℓ-smoothness assumptions, we establish a high-probability convergence guarantee and show that the proposed method reaches an ϵ-Pareto-stationary point with a total sample complexity of $\mathcal{O}(\epsilon^{-4})$. We further validate the effectiveness of the proposed algorithm through adversarial-attack experiments in multi-task learning. Future work may investigate more computationally efficient inner solvers and single-step approximations for the Wasserstein adversarial maximization, as well as extensions to broader multi-objective learning applications.

AI USE STATEMENT

In this work, we used generative AI tools to polish and improve the presentation of the manuscript, assist with literature search, and verify the correctness of mathematical derivations. All AI-assisted content, references, and mathematical arguments were reviewed and checked by the authors before being included in the final manuscript. We take responsibility for the final content of this work, including text, claims, and mathematical results produced with the aid of generative AI.

ETHICS STATEMENT

(This section is **recommended** and does not count toward the page limit.) If authors feel that their paper submission raises questions regarding the Code of Ethics, they are encouraged to include a paragraph of Ethics Statement (at the end of the main text before references) to address potential concerns where appropriate. Topics include, but are not limited to, studies that involve human subjects, practices to data set releases, potentially harmful insights, methodologies and applications, potential conflicts of interest and sponsorship, discrimination/bias/fairness concerns, privacy and security issues, legal compliance, and research integrity issues (e.g., IRB, documentation, research ethics). This statement should not be more than 1 page.

REPRODUCIBILITY STATEMENT

(This section is **recommended** and does not count toward the page limit.)

AUTHOR CONTRIBUTIONS

If you'd like to, you may include a section for author contributions as is done in many journals. This is optional and at the discretion of the authors.

ACKNOWLEDGMENTS

Use unnumbered third level headings for the acknowledgments. All acknowledgments, including those to funding agencies, go at the end of the paper.

REFERENCES

Hao Ban and Kaiyi Ji. Fair resource allocation in multi-task learning. In *Proceedings of the 41st International Conference on Machine Learning, ICML*, volume 235, pp. 2715–2731, 21–27 Jul 2024.

Aharon Ben-Tal, Dick den Hertog, Anja De Waegenaere, Bertrand Melenberg, and Gijs Rennen. Robust solutions of optimization problems affected by uncertain probabilities. *Management Science*, 59(2):341–357, 2013.

Jose Blanchet and Karthyek Murthy. Quantifying distributional model risk via optimal transport. *Mathematics of Operations Research*, 44(2):565–600, 2019.

Lisha Chen, Heshan Devaka Fernando, Yiming Ying, and Tianyi Chen. Three-way trade-off in multi-objective learning: Optimization, generalization and conflict-avoidance. *Journal of Machine Learning Research*, 25:193:1–193:53, 2024.

Nikos Dimitriadis, Pascal Frossard, and François Fleuret. Pareto low-rank adapters: Efficient multi-task learning with preferences. In *Proceedings of the 13th the International Conference on Learning Representations, ICLR*, volume 2025, pp. 22323–22346, 2025.

Chunxiao Fan, Dan Guo, Ziqi Wang, and Meng Wang. Multi-objective convex quantization for efficient model compression. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47 (4):2313–2329, 2025.

Heshan Devaka Fernando, Han Shen, Miao Liu, Subhajit Chaudhury, Keerthiram Murugesan, and Tianyi Chen. Mitigating gradient bias in multi-objective learning: A provably convergent approach. In *Proceedings of the 11th International Conference on Learning Representations, ICLR*, 2023.

Rui Gao and Anton J. Kleywegt. Distributionally robust stochastic optimization with wasserstein distance. *Mathematics of Operations Research*, 48(2):603–655, 2023.

Xiangkun He and Chen Lv. Toward personalized decision making for autonomous vehicles: A constrained multi-objective reinforcement learning technique. *Transportation Research Part C: Emerging Technologies*, 156:104352, 2023.

Jikai Jin, Bohang Zhang, Haiyang Wang, and Liwei Wang. Non-convex distributionally robust optimization: Non-asymptotic analysis. In *Proceedings of the 35th Advances in Neural Information Processing Systems, NeurIPS*, pp. 2771–2782, 2021.

Daniel Kuhn, Soroosh Shafiee, and Wolfram Wiesemann. Distributionally robust optimization. *Acta Numerica*, 34:579–804, 2025.

Haochuan Li, Jian Qian, Yi Tian, Alexander Rakhlin, and Ali Jadbabaie. Convex and non-convex optimization under generalized smoothness. In *Proceedings of the 37th Advances in Neural Information Processing Systems, NeurIPS*, volume 36, pp. 40238–40271, 2023.

Zhaoen Li, Maoqi Liu, and Zhi-Hai Zhang. Achieving robust data-driven contextual decision making in a data augmentation way, 2024.

Francesco Micheli, Efe C. Balta, Anastasios Tsiamis, and John Lygeros. Wasserstein distributionally robust bayesian optimization with continuous context. In *Proceedings of the 28 th International Conference on Artificial Intelligence and Statistics, AISTATS*, volume 258, pp. 4978–4986.

Eduardo Fernandes Montesuma, Fred Maurice Ngolè Mboula, and Antoine Souloumiac. Recent advances in optimal transport for machine learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(2):1161–1180, 2025.

Ozan Sener and Vladlen Koltun. Multi-task learning as multi-objective optimization. In *Proceedings of the 32nd Advances in Neural Information Processing Systems, NeurIPS*, volume 31, 2018.

Aman Sinha, Hongseok Namkoong, and John C. Duchi. Certifying some distributional robustness with principled adversarial training. In *Proceedings of the 6th International Conference on Learning Representations, ICLR*, 2018.

Hiroki Tanabe, Ellen H. Fukuda, and Nobuo Yamashita. Proximal gradient methods for multiobjective optimization and their applications. *Computational Optimization and Applications*, 72(2):339–361, 2019.

Austin Watkins, Thanh Nguyen-Tang, Enayat Ullah, and Raman Arora. In *Proceedings of the 38th Advances in Neural Information Processing Systems, NeurIPS*, volume 37, pp. 139213–139245, 2024.

Wenhan Xian, Ziyi Chen, and Heng Huang. Delving into the convergence of generalized smooth minimax optimization. In *Proceedings of the 51st International Conference on Machine Learning, ICML*, volume 235, pp. 54191–54211, 2024.

Peiyao Xiao, Hao Ban, and Kaiyi Ji. Direction-oriented multi-objective learning: Simple and provable stochastic algorithms. In *Proceedings of 37th Advances in Neural Information Processing Systems, NeurIPS*, 2023.

Enneng Yang, Zhenyi Wang, Li Shen, Shiwei Liu, Guibing Guo, Xingwei Wang, and Dacheng Tao. Adamerging: Adaptive model merging for multi-task learning. In *Proceedings of the 12th International Conference on Learning Representations, ICLR*, 2024.

Yufeng Yang, Fangning Zhuo, Ziyi Chen, Heng Huang, and Yi Zhou. Distributionally robust multi-objective optimization. *CoRR*, abs/2605.05660, 2026.

Jun Yu, Xiaokang Liu, Chongliang Luo, Jin Huang, Rong Zhou, Yixin Liu, Jie Hu, Jianmin Chen, Ke Zhang, Dazheng Zhang, et al. Multitask learning 1997–2024: Part i fundamentals. *Harvard Data Science Review*, 7(3), 2025.

Jingzhao Zhang, Tianxing He, Suvrit Sra, and Ali Jadbabaie. Why gradient clipping accelerates training: A theoretical justification for adaptivity. In *Prpceedings of the 8th International Conference on Learning Representations, ICLR*, 2020a.

Jingzhao Zhang, Tianxing He, Suvrit Sra, and Ali Jadbabaie. Why gradient clipping accelerates training: A theoretical justification for adaptivity. In *Proceedings of the 8th International Conference on Learning Representations, ICLR*, 2020b.

Lei Zhang, Sibo Liu, Haipeng Yang, Zihao Chen, and Hongke Zhao. MOREM: An evolutionary multitasking optimization algorithm for multi-objective recommendations. *Information Sciences*, 678:120921, 2024a.

Qi Zhang, Yi Zhou, Ashley Prater-Bennette, Lixin Shen, and Shaofeng Zou. Large-scale non-convex stochastic constrained distributionally robust optimization. In *Proceedings of the 39th AAAI Conference on Artificial Intelligence, AAAI*, pp. 8217–8225, 2024b.

Qi Zhang, Peiyao Xiao, Shaofeng Zou, and Kaiyi Ji. MGDA converges under generalized smoothness, provably. In *Proceedings of the 13th International Conference on Learning Representations, ICLR*, 2025a.

Qi Zhang, Yi Zhou, Simon Khan, Ashley Prater-Bennette, Lixin Shen, and Shaofeng Zou. Revisiting large-scale non-convex distributionally robust optimization. In *Proceedings of the 13th International Conference on Learning Representations, ICLR*, 2025b.

Xiaoyuan Zhang, Xi Lin, and Qingfu Zhang. PMGDA: A preference-based multiple gradient descent algorithm. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 10(2): 1098–1110, 2026.