

---

# Online Stochastic Optimization with Multiple Objectives

---

**Mehrdad Mahdavi**  
Dept. of Computer Science  
Michigan State University  
mahdavam@cse.msu.edu

**Tianbao Yang**  
Machine Learning Lab  
GE Global Research  
tyang@ge.com

**Rong Jin**  
Dept. of Computer Science  
Michigan State University  
rongjin@cse.msu.edu

## Abstract

In this paper we propose a general framework to characterize and solve the stochastic optimization problems with multiple objectives underlying many real world learning applications. We first propose a projection based algorithm which attains an  $O(T^{-1/3})$  convergence rate. Then, by leveraging on the theory of Lagrangian in constrained optimization, we devise a novel primal-dual stochastic approximation algorithm which attains the optimal convergence rate of  $O(T^{-1/2})$  for general Lipschitz continuous objectives.

## 1 Introduction

Stochastic optimization algorithms such as stochastic gradient descent (SGD) [9, 1] and its online counterpart, online gradient descent (OGD) [13, 11], have been focus of intensive study in the last few years. In the traditional setup of stochastic optimization, the learner is evaluated by a single loss function at each iteration. However, in many real world applications, the learner needs to consider several performance measures simultaneously [3] and we are not aware of any work addressing multi-objective stochastic problems.

In this work, we generalize online convex optimization (OCO) to the case of multiple objectives. In particular, at each iteration, the learner is asked to present a solution  $\mathbf{x}_t$ , which will be evaluated by multiple loss functions  $f_t^0(\mathbf{x}), f_t^1(\mathbf{x}), \dots, f_t^m(\mathbf{x})$ . Since it is impossible to simultaneously minimize multiple loss functions and in order to avoid complications caused by handling more than one objective, we choose one function as the objective and try to bound other objectives by appropriate thresholds. Specifically, the goal of OCO with multiple objectives becomes to minimize  $\sum_{t=1}^T f_t^0(\mathbf{x}_t)$  and at the same time keep the other loss functions below a given threshold, i.e.

$$\frac{1}{T} \sum_{t=1}^T f_t^i(\mathbf{x}_t) \leq \gamma_i,$$

where  $\mathbf{x}_1, \dots, \mathbf{x}_T$  are the solutions generated by the online learner and  $\gamma_i$  specifies the level of loss that is acceptable to the  $i$ -th loss function. We refer to the above problem as online convex optimization with multiple objectives. The proposed problem is closely related to the classical study of multiple objective optimization [10]. The main difference is that all the objectives (i.e., the loss functions) are changing over the iterations, making it a substantially more difficult problem. The proposed problem is also closely related to online optimization with side constraints [6, 4], where the constraint introduced is essentially the second objective in multiple objective optimization. The proposed problem generalizes online optimization with side constraints by allowing more than one constraints.

Since the general setup (i.e., full adversarial setup) is challenging for online convex optimization even with two objectives [7], in this work, we consider a simple scenario where all the loss functions  $\{f_t^i(\cdot)\}_{i=1}^m$  are i.i.d samples from unknown distribution [12]. We also note that our goal is not to find a sample from the Pareto optimal set (i.e. the set of solutions that are not dominated in the Pareto sense in the decision space), instead we are trying to satisfy all the objectives below a pre-

specified threshold. We denote by  $\bar{f}_i(\cdot) = \mathbb{E}_t[f_t^i(\cdot)]$ ,  $i = 0, 1, \dots, m$  the expected loss function of sampled function  $f_t^i(\cdot)$ . To solve the problem, as is standard in stochastic optimization, we assume that we do not have direct access to the expected loss functions and only information available to the solver is through a noisy oracle which provides a stochastic realization of the expected loss function at each call. We assume that there exist a solution  $\mathbf{x}$  strictly satisfying all the constraints, i.e.  $\bar{f}_i(\mathbf{x}) < \gamma_i$ ,  $i \in [m]$ . We denote by  $\mathbf{x}_*$  the optimal solution to the multiple objectives, i.e.,

$$\mathbf{x}_* = \arg \min \{ \bar{f}_0(\mathbf{x}) : \bar{f}_i(\mathbf{x}) \leq \gamma_i, i = 1, \dots, m \}.$$

Our goal is to compute a solution  $\widehat{\mathbf{x}}_T$  after  $T$  trials that (i) obeys all the constraints, i.e.  $\bar{f}_i(\widehat{\mathbf{x}}_T) \leq \gamma_i$ ,  $i \in [m]$  and (ii) minimizes the regret with respect to the optimal solution  $\mathbf{x}_*$ , i.e.  $\bar{f}_0(\widehat{\mathbf{x}}_T) - \bar{f}_0(\mathbf{x}_*)$ . For the convenience of discussion, we refer to  $f_t^0(\cdot)$  and  $\bar{f}_0(\cdot)$  as the objective function, and to  $f_t^i(\cdot)$  and  $\bar{f}_i(\cdot)$  as the constraint functions.

Before discussing the algorithms, we first describe a few assumptions made in our analysis. We assume that the final solution  $\mathbf{x}_*$  lives in a ball  $\mathcal{B}$  of radius  $R$ , i.e.,  $\mathcal{B} = \{ \mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\| \leq R \}$ . We also make the standard assumption that all the loss functions are Lipschitz continuous, i.e.,  $|f_t^i(\mathbf{x}) - f_t^i(\mathbf{x}')| \leq L\|\mathbf{x} - \mathbf{x}'\|$  for any  $\mathbf{x} \in \mathcal{B}$  and  $\mathbf{x}' \in \mathcal{B}$ .

In this extended abstract, we only sketch the results, and omit many important details, which appear in the full version of our paper. In section 2 we propose a projection based algorithm which reduces the problem into a standard optimization problem with changing solution space. Section 3 introduces our efficient primal-dual stochastic optimization algorithm achieving the optimal known bound.

## 2 Warmup: a Projection based Algorithm

The main challenge of the proposed problem is that the expected constraint functions  $\bar{f}_i(\cdot)$  are not given. Instead, only a sampled constraint function is provided at each trial  $t$ . Our naive approach is to turn the multiple objective optimization problem into a constrained optimization problem as

$$\min_{\mathbf{x} \in \mathcal{B} \cap \mathcal{K}} \bar{f}_0(\mathbf{x}) \quad (1)$$

where domain  $\mathcal{K}$  is defined as  $\mathcal{K} = \{ \mathbf{x} : \bar{f}_i(\mathbf{x}) \leq \gamma_i, i = 1, \dots, m \}$ .

This approach circumvents the problem of optimizing multiple objective into the original online convex optimization with complex projections. Since the domain  $\mathcal{K}$  is unknown in (1), a naive approach is to estimate the expected constraint functions based on the sampled constraints received so far, and project the updated solution into the domain constructed by the estimated constraint functions. More specifically, at trial  $t$ , given the current solution  $\mathbf{x}_t$  and received loss functions  $f_t^i(\mathbf{x})$ ,  $i = 0, 1, \dots, m$ , we first estimate the expected constraint functions as

$$\widehat{f}_t^i(\mathbf{x}) = \frac{1}{t} \sum_{k=1}^t f_k^i(\mathbf{x}), i \in [m]$$

and then update the solution by  $\mathbf{x}_{t+1} = \Pi_{\mathcal{K}_t}(\mathbf{x}_t - \eta \nabla f_t^i(\mathbf{x}_t))$  where  $\eta > 0$ ,  $\Pi_{\mathcal{K}}(\mathbf{x}) = \min_{\mathbf{z} \in \mathcal{K}} \|\mathbf{z} - \mathbf{x}\|$ , and  $\mathcal{K}_t$  is an approximate domain and is given by  $\mathcal{K}_t = \{ \mathbf{x} : \widehat{f}_t^i(\mathbf{x}) \leq \gamma_i, i = 1, \dots, m \}$ .

The problem with the above approach is that although it is feasible to satisfy all the constraints based on the true expected constraint functions, there is no guarantee that the approximate domain  $\mathcal{K}_t$  is not empty. One way to address this issue is to estimate the expected constraint functions by burning the first  $bT$  trials, where  $b \in (0, 1)$  is a constant that needs to be adjusted to obtain the optimal performance. Given the sampled constraint functions  $f_1^i, \dots, f_{bT}^i$  received in the first  $bT$  trials, we compute the approximate domain  $\mathcal{K}'$  as

$$\widehat{f}_i(\mathbf{x}) = \frac{1}{bT} \sum_{t=1}^{bT} f_t^i(\mathbf{x}), i \in [m], \mathcal{K}' = \{ \mathbf{x} : \widehat{f}_i(\mathbf{x}) \leq \hat{\gamma}_i, i = 1, \dots, m \}$$

where  $\hat{\gamma}_i = \gamma_i + LR\sqrt{[2/(bT)] \ln(m/\delta)}$ . It is straightforward to show that with a probability  $1 - \delta$ , for any  $\mathbf{x} \in \mathcal{K}$ , we have  $\mathbf{x} \in \mathcal{K}'$ . We note that for projection onto the estimated domain, we only consider only a special solution and therefore the negative results of uniform convergence [12] does not apply. Using the approximate domain  $\mathcal{K}'$ , for trial  $t \in [bT + 1, T]$ , we update the solution by

$$\mathbf{x}_{t+1} = \Pi_{\mathcal{K}'}(\mathbf{x}_t - \eta \nabla f_t^i(\mathbf{x}_t)).$$

There are however several drawbacks with this approach. First, by simple counting, it is not difficult to see that the overall violation of constraints, given by  $\sum_{t=1}^T \widehat{f}_i(\mathbf{x}_t)$ , is  $O(bT + (1 - b)T/\sqrt{bT})$ :

---

**Algorithm 1** Online Convex Optimization with Multiple Objectives

---

- 1: **INPUT:** step size  $\eta$ ,  $\lambda_i^0 > 0, i \in [m]$  and total time  $T$
- 2:  $\mathbf{x}_1 = 0$ .
- 3: **for**  $t = 1, \dots, T$  **do**
- 4:   Submit the solution  $\mathbf{x}_t$
- 5:   Receive loss functions  $f_t^i, i = 0, 1, \dots, m$
- 6:   Compute the gradient  $\nabla f_t^i(\mathbf{x}_t), i = 0, 1, \dots, m$
- 7:   Update the solution  $\mathbf{x}$  and  $\lambda$  by

$$\begin{aligned}\mathbf{x}_{t+1} &= \Pi_{\mathcal{B}}(\mathbf{x}_t - \eta \nabla_{\mathbf{x}} \mathcal{L}_t(\mathbf{x}_t, \lambda_t)) = \Pi_{\mathcal{B}}\left(\mathbf{x}_t - \eta \left[ \nabla f_t^0(\mathbf{x}_t) + \sum_{i=1}^m \lambda_t^i \nabla f_t^i(\mathbf{x}_t) \right]\right) \\ \lambda_{t+1}^i &= \Pi_{[0, \lambda_i^0]}(\lambda_t^i + \eta \nabla_{\lambda_i} \mathcal{L}_t(\mathbf{x}_t, \lambda_t)) = \Pi_{[0, \lambda_i^0]}(\lambda_t^i + \eta [f_t^i(\mathbf{x}_t) - \gamma_i])\end{aligned}$$

- 8: **end for**
  - 9: **Return**  $\hat{\mathbf{x}}_T = \sum_{t=1}^T \mathbf{x}_t / T$
- 

$O(bT)$  comes from the first  $bT$  trials used to estimate the expected constraint functions, where the violation of each trial is about  $O(1)$  and  $(1-b)T/\sqrt{bT}$  comes from the rest  $(1-b)T$  trials where the violation is  $O(1/\sqrt{bT})$ . By minimizing the overall violation, we choose  $b = O(T^{-1/3})$ , leading to the overall violation of  $O(T^{2/3})$ . Using the same trick as in [5], we could obtain a solution with zero violation of constraints but with a regret bound of  $O(T^{2/3})$ , leading to unsatisfied result. Second, this approach requires memorizing the constraint functions of the first  $bT$  trials. This is in contrast to the typical assumption of online learning where only the solution is memorized. Third, even though the difference between  $\hat{f}_i(\cdot)$  and  $\bar{f}_i(\mathbf{x})$  is small, i.e.,  $\max_{\mathbf{x} \in \mathcal{B}} |\hat{f}_i(\mathbf{x}) - \bar{f}_i(\mathbf{x})| = O(1/\sqrt{bT})$ ,  $\hat{f}_i(\cdot)$  could be non-convex, leading to inefficient computation when performing the projection.

As indicated by the above analysis, the main limitation of the naive approach is that it requires a projection step. To address this limitation, we present an algorithm that does not require projection when updating the solution. We show that with a high probability, the solution found by the proposed algorithm will exactly satisfy the expected constraints and achieves a regret bound of  $O(\sqrt{T})$ .

### 3 An Efficient Online Stochastic Primal Dual Algorithm

The main idea of the proposed algorithm is to design an appropriate objective that combines the loss function  $f_0$  with  $\{f_i\}_{i=1}^m$ . To this end, we define the following objective function

$$\bar{\mathcal{L}}(\mathbf{x}, \lambda) = \bar{f}_0(\mathbf{x}) + \sum_{i=1}^m \lambda_i (\bar{f}_i(\mathbf{x}) - \gamma_i)$$

Note that the objective function consists of both the primal variables  $\mathbf{x}$  and dual variables  $\lambda = (\lambda_1, \dots, \lambda_m)$ . In the proposed algorithm, we will simultaneously update solutions for both  $\mathbf{x}$  and  $\lambda$ . By exploring convex-concave optimization theory [8], we will show that with a high probability, the solution of regret  $O(\sqrt{T})$  that exactly obey the constraints.

As the first step, we consider a simple scenario where the obtained solution is allowed to violate the constraints. Algorithm 1 shows the detailed steps. It follows the same procedure as convex-concave optimization. Since at each iteration, we only observed a randomly sampled loss functions  $f_t^i(\cdot), i = 0, 1, \dots, m$ , the objective function given by

$$\mathcal{L}_t(\mathbf{x}, \lambda) = f_t^0(\mathbf{x}) + \sum_{i=1}^m \lambda_i (f_t^i(\mathbf{x}) - \gamma_i)$$

provides an unbiased estimate of  $\bar{\mathcal{L}}(\mathbf{x}, \lambda)$ . Given the approximate objective  $\mathcal{L}_t(\mathbf{x}, \lambda)$ , Algorithm 1 tries to minimize the objective  $\mathcal{L}_t(\cdot, \cdot)$  with respect to the primal variable  $\mathbf{x}$  and maximize the objective with respect to the dual variable  $\lambda$ . In the following theorem, we show that under appropriate conditions, the solution  $\hat{\mathbf{x}}_T$  output by Algorithm 1 will have a convergence rate of  $O(1/\sqrt{T})$  for both the regret and the violation of the constraints. To facilitate the analysis, we rewrite the constrained optimization problem in (1) as

$$\min_{\mathbf{x} \in \mathcal{B}} \max_{\lambda \in \mathbb{R}_+^m} \bar{f}_0(\mathbf{x}) + \sum_{i=1}^m \lambda_i (\bar{f}_i(\mathbf{x}) - \gamma_i) \quad (2)$$

We denote by  $\mathbf{x}_*$  and  $\lambda_* = (\lambda_1^*, \dots, \lambda_m^*)^\top$  as the optimal solution to the above convex-concave optimization problem, i.e.

$$\mathbf{x}_* = \arg \min_{\mathbf{x} \in \mathcal{B}} \bar{f}_0(\mathbf{x}) + \sum_{i=1}^m \lambda_i^* (\bar{f}_i(\mathbf{x}) - \gamma_i) \quad (3)$$

$$\lambda_* = \arg \max_{\lambda \in \mathbb{R}_+^m} \bar{f}_0(\mathbf{x}_*) + \sum_{i=1}^m \lambda_i (\bar{f}_i(\mathbf{x}_*) - \gamma_i) \quad (4)$$

We define two quantities that are useful for bounding the gradients  $\nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \lambda)$  and  $\nabla_{\lambda} \mathcal{L}(\mathbf{x}, \lambda)$ :

$$D^2 = \sum_{i=1}^m [\lambda_i^0]^2, \quad G^2 = L^2 \left( 1 + \sum_{i=1}^m \lambda_i^0 \right)^2 + \max_{\mathbf{x} \in \mathcal{B}} \sum_{i=1}^m \bar{f}_i^2(\mathbf{x}) \quad (5)$$

**Theorem 1.** Set  $\lambda_0^i \geq \lambda_i^* + \theta, i \in [m]$ , where  $\theta > 0$  is constant. Let  $\hat{\mathbf{x}}_T$  be the solution obtained by Algorithm 1 obtained after  $T$  iterations. Then, with a probability  $1 - (2m + 1)\delta$ , we have

$$\bar{f}_0(\hat{\mathbf{x}}_T) - \bar{f}_0(\mathbf{x}_0) \leq \frac{\mu(\delta)}{\sqrt{T}} \quad \text{and} \quad \bar{f}_i(\hat{\mathbf{x}}_T) - \gamma_i \leq \frac{\mu(\delta)}{\theta\sqrt{T}}, i \in [m]$$

where

$$\mu(\delta) = G\sqrt{R^2 + D^2} + 2G(R + D)\sqrt{2\ln \frac{1}{\delta}}. \quad (6)$$

We now develop an algorithm that allows the solution to exactly satisfy all the constraints. To this end, we define  $\hat{\gamma}_i = \gamma_i - \frac{\mu(\delta)}{\theta\sqrt{T}}$ . We will run Algorithm 1 but with  $\gamma_i$  replaced by  $\hat{\gamma}_i$ . The following theorem shows the property of the obtained solution.

**Theorem 2.** Let  $\hat{\mathbf{x}}_T$  be the solution obtained by Algorithm 1 with  $\gamma_i$  replaced by  $\hat{\gamma}_i$  and  $\lambda_0^i = \lambda_i^* + \theta, i \in [m]$ . Then, with a probability  $1 - (2m + 1)\delta$ , we have

$$\bar{f}_0(\hat{\mathbf{x}}_T) - \bar{f}_0(\mathbf{x}_0) \leq \frac{(1 + \sum_{i=1}^m \lambda_i^0)\mu(\delta)}{\sqrt{T}}, \quad \bar{f}_i(\hat{\mathbf{x}}_T) \leq \gamma_i, i \in [m]$$

where  $\mu(\delta)$  is defined (6).

In order to run Algorithm 1, we need to estimate the parameter  $\lambda_i^0$ , which requires estimating the upper bound for  $\lambda_i^*$ . To this end, we consider an alternative problem to the convex-concave optimization problem in (2), i.e.

$$\min_{\mathbf{x} \in \mathcal{B}} \max_{\lambda \geq 0} \bar{f}_0(\mathbf{x}) + \lambda \max_{1 \leq i \leq m} (\bar{f}_i(\mathbf{x}) - \gamma_i) \quad (7)$$

Evidently  $\mathbf{x}_*$  is the optimal primal solution to (7). Let  $\lambda_a$  be the optimal dual solution to the problem in (7). We have the following proposition that links  $\lambda_i^*, i \in [m]$ , the optimal dual solution to (2), with  $\lambda_a$ , the optimal dual solution to (7).

**Proposition 1.** Let  $\lambda_a$  be the optimal dual solution to (7) and  $\lambda_i^*, i \in [m]$  be the optimal solution to (2). We have  $\lambda_a = \sum_{i=1}^m \lambda_i^*$ .

Given the result from Proposition 1, it is sufficiently to bound  $\lambda_a$ . In order to bound  $\lambda_a$ , we need to make certain assumption about  $\bar{f}_i, i \in [m]$ .

**Assumption 1.** We assume  $\min_{\alpha \in \Delta_m} |\sum_{i=1}^m \alpha_i \nabla \bar{f}_i(\mathbf{x})| \geq \tau$ , where  $\tau > 0$  is a constant and domain  $\Delta_m$  is defined as  $\Delta_m = \{\alpha \in \mathbb{R}_+^m : \sum_{i=1}^m \alpha_i = 1\}$ .

The following lemma bounds  $\lambda_a$  by  $\tau$ .

**Lemma 1.** Under Assumption 1, we have  $\lambda_a \leq \frac{L}{\tau}$ .

Combining Proposition 1 with Lemma 1, we have, under Assumption 1,  $\lambda_i^* \leq \frac{L}{\tau}, i = 1, \dots, m$ .

## 4 Conclusion

In this paper we have addressed the problem of online stochastic optimization with multiple objectives and presented an efficient primal-dual algorithm which attains the optimal convergence rate  $O(1/\sqrt{T})$  for all the objectives.

## References

- [1] F. Bach and E. Moulines. Non-asymptotic analysis of stochastic approximation algorithms for machine learning. In *NIPS*, pages 451–459, 2011.
- [2] S. Boucheron, G. Lugosi, and O. Bousquet. Concentration inequalities. In *Advanced Lectures on Machine Learning*, pages 208–240, 2003.
- [3] Y. Jin and B. Sendhoff. Pareto-based multiobjective machine learning: An overview and case studies. *IEEE Transactions on Systems, Man, and Cybernetics, Part C*, 38(3):397–415, 2008.
- [4] B. Kveton, J. Y. Yu, G. Theodorou, and S. Mannor. Online learning with expert advice and finite-horizon constraints. In *AAAI*, pages 331–336, 2008.
- [5] M. Mahdavi, R. Jin, and T. Yang. Trading regret for efficiency: online convex optimization with long term constraints. *JMLR*, 13:2465–2490, 2012.
- [6] S. Mannor and J. N. Tsitsiklis. Online learning with constraints. In *COLT*, pages 529–543, 2006.
- [7] S. Mannor, J. N. Tsitsiklis, and J. Y. Yu. Online learning with sample path constraints. *Journal of Machine Learning Research*, 10:569–590, 2009.
- [8] A. Nemirovski. Efficient methods in convex programming. Lecture Notes, Available at <http://www2.isye.gatech.edu/nemirovs>, 1994.
- [9] A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM J. on Optimization*, 19:1574–1609, 2009.
- [10] Y. Sawaragi, H. Nakayama, and T. Tanino. *Theory of Multiobjective Optimization*. Academic Press Inc., 1985.
- [11] S. Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends in Machine Learning*, 4(2):107–194, 2012.
- [12] S. Shalev-Shwartz, O. Shamir, N. Srebro, and K. Sridharan. Stochastic convex optimization. In *COLT*, 2009.
- [13] M. Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *ICML*, pages 928–936, 2003.

## Appendix A. Proof of Theorem 1

Using the standard analysis of convex-concave optimization, for any  $\mathbf{x} \in \mathcal{B}$  and  $\lambda_i \in [0, \lambda_i^0], i \in [m]$ , we have

$$\begin{aligned}
\bar{\mathcal{L}}(\mathbf{x}_t, \lambda) - \bar{\mathcal{L}}(\mathbf{x}, \lambda_t) &\leq (\mathbf{x}_t - \mathbf{x})^\top \nabla_{\mathbf{x}} \bar{\mathcal{L}}(\mathbf{x}_t, \lambda_t) - (\lambda_t - \lambda) \nabla_{\lambda} \bar{\mathcal{L}}(\mathbf{x}_t, \lambda_t) \\
&= (\mathbf{x}_t - \mathbf{x})^\top \nabla_{\mathbf{x}} \mathcal{L}_t(\mathbf{x}_t, \lambda_t) - (\lambda_t - \lambda) \nabla_{\lambda} \mathcal{L}_t(\mathbf{x}_t, \lambda_t) \\
&\quad + (\mathbf{x}_t - \mathbf{x})^\top (\nabla_{\mathbf{x}} \bar{\mathcal{L}}(\mathbf{x}_t, \lambda_t) - \nabla_{\mathbf{x}} \mathcal{L}_t(\mathbf{x}_t, \lambda_t)) - (\lambda_t - \lambda)^\top (\nabla_{\lambda} \bar{\mathcal{L}}(\mathbf{x}_t, \lambda_t) - \nabla_{\lambda} \mathcal{L}_t(\mathbf{x}_t, \lambda_t)) \\
&\leq \frac{\|\mathbf{x}_t - \mathbf{x}\|^2 - \|\mathbf{x}_{t+1} - \mathbf{x}\|^2}{2\eta} + \frac{\|\lambda_t - \lambda\|^2 - \|\lambda_{t+1} - \lambda\|^2}{2\eta} + \frac{\eta}{2} (\|\nabla_{\mathbf{x}} \mathcal{L}_t(\mathbf{x}_t, \lambda_t)\|^2 + \|\nabla_{\lambda} \mathcal{L}_t(\mathbf{x}_t, \lambda_t)\|^2) \\
&\quad + \sum_{t=1}^T (\mathbf{x}_t - \mathbf{x})^\top (\nabla_{\mathbf{x}} \bar{\mathcal{L}}(\mathbf{x}_t, \lambda_t) - \nabla_{\mathbf{x}} \mathcal{L}_t(\mathbf{x}_t, \lambda_t)) - (\lambda_t - \lambda)^\top (\nabla_{\lambda} \bar{\mathcal{L}}(\mathbf{x}_t, \lambda_t) - \nabla_{\lambda} \mathcal{L}_t(\mathbf{x}_t, \lambda_t))
\end{aligned}$$

By adding all the inequalities together, we have

$$\begin{aligned}
&\sum_{t=1}^T \bar{\mathcal{L}}(\mathbf{x}_t, \lambda) - \bar{\mathcal{L}}(\mathbf{x}, \lambda_t) \\
&\leq \frac{\|\mathbf{x} - \mathbf{x}_1\|^2 + \|\lambda - \lambda_1\|^2}{2\eta} + \frac{\eta}{2} \sum_{t=1}^T \|\nabla_{\mathbf{x}} \mathcal{L}_t(\mathbf{x}_t, \lambda_t)\|^2 + \|\nabla_{\lambda} \mathcal{L}_t(\mathbf{x}_t, \lambda_t)\|^2 \\
&\quad + \sum_{t=1}^T (\mathbf{x}_t - \mathbf{x})^\top (\nabla_{\mathbf{x}} \bar{\mathcal{L}}(\mathbf{x}_t, \lambda_t) - \nabla_{\mathbf{x}} \mathcal{L}_t(\mathbf{x}_t, \lambda_t)) - (\lambda_t - \lambda)^\top (\nabla_{\lambda} \bar{\mathcal{L}}(\mathbf{x}_t, \lambda_t) - \nabla_{\lambda} \mathcal{L}_t(\mathbf{x}_t, \lambda_t)) \\
&\leq \frac{R^2 + D^2}{2\eta} + \frac{\eta G^2 T}{2} \\
&\quad + \sum_{t=1}^T (\mathbf{x}_t - \mathbf{x})^\top (\nabla_{\mathbf{x}} \bar{\mathcal{L}}(\mathbf{x}_t, \lambda_t) - \nabla_{\mathbf{x}} \mathcal{L}_t(\mathbf{x}_t, \lambda_t)) - (\lambda_t - \lambda)^\top (\nabla_{\lambda} \bar{\mathcal{L}}(\mathbf{x}_t, \lambda_t) - \nabla_{\lambda} \mathcal{L}_t(\mathbf{x}_t, \lambda_t)) \\
&\leq \frac{R^2 + D^2}{2\eta} + \frac{\eta G^2 T}{2} + 2G(R + D) \sqrt{2T \ln \frac{1}{\delta}} \quad (\text{w.p. } 1 - \delta),
\end{aligned}$$

where the last step uses the Hoeffding inequality for martingales [2]. For any fixed  $\lambda_i \in [0, \lambda_i^0], i \in [m]$  and  $\mathbf{x} \in \mathcal{B}$ , with a probability  $1 - \delta$ , we have

$$\begin{aligned}
\bar{f}_0(\hat{\mathbf{x}}_T) + \sum_{i=1}^m \lambda_i (\bar{f}_i(\hat{\mathbf{x}}_T) - \gamma_i) - \bar{f}_0(\mathbf{x}) - \sum_{i=1}^m \hat{\lambda}_T^i (\bar{f}_i(\mathbf{x}) - \gamma_i) \\
\leq G \sqrt{\frac{R^2 + D^2}{T}} + 2G(R + D) \sqrt{\frac{2}{T} \ln \frac{1}{\delta}}
\end{aligned} \tag{8}$$

By fixing  $\mathbf{x} = \mathbf{x}_*$  and  $\lambda = \mathbf{0}$  in (9), we have  $\bar{f}_i(\mathbf{x}_*) \leq \gamma_i, i \in [m]$ , and therefore, with a probability  $1 - \delta$ , have

$$\bar{f}_0(\hat{\mathbf{x}}_T) \leq \bar{f}_0(\mathbf{x}_*) + G \sqrt{\frac{R^2 + D^2}{T}} + 2G(R + D) \sqrt{\frac{2}{T} \ln \frac{1}{\delta}}$$

To bound the violation of constraints, for each  $i \in [m]$ , we set  $\mathbf{x} = \mathbf{x}_*, \lambda_i = \lambda_i^0$ , and  $\lambda_j = \lambda_j^*, j \neq i$  in (9). We have

$$\begin{aligned}
&\bar{f}_0(\hat{\mathbf{x}}_T) + \lambda_i^0 (\bar{f}_i(\hat{\mathbf{x}}_T) - \gamma_i) + \sum_{j \neq i} \lambda_j^* (\bar{f}_j(\hat{\mathbf{x}}_T) - \gamma_j) - \bar{f}_0(\mathbf{x}_*) - \sum_{i=1}^m \hat{\lambda}_T^i (\bar{f}_i(\mathbf{x}_*) - \gamma_i) \\
&\geq \bar{f}_0(\hat{\mathbf{x}}_T) + \lambda_i^0 (\bar{f}_i(\hat{\mathbf{x}}_T) - \gamma_i) + \sum_{j \neq i} \lambda_j^* (\bar{f}_j(\hat{\mathbf{x}}_T) - \gamma_j) - \bar{f}_0(\mathbf{x}_*) - \sum_{i=1}^m \lambda_i^* (\bar{f}_i(\mathbf{x}_*) - \gamma_i) \\
&\geq \theta (\bar{f}_i(\hat{\mathbf{x}}_T) - \gamma_i)
\end{aligned}$$

where the first inequality utilizes (4) and the second inequality utilizes (3). We thus have, with a probability  $1 - \delta$ ,

$$\bar{f}_i(\hat{\mathbf{x}}_T) - \gamma_i \leq \frac{G}{\theta} \sqrt{\frac{R^2 + D^2}{T}} + \frac{2G(R + D)}{\theta} \sqrt{\frac{2}{T} \ln \frac{1}{\delta}}$$

We complete the proof by taking the union bound over all the random events.

## Appendix B. Proof of Theorem 2

Following the proof of Theorem 1, with a probability  $1 - \delta$ , we have

$$\begin{aligned} \bar{f}_0(\hat{\mathbf{x}}_T) + \sum_{i=1}^m \lambda_i(\bar{f}_i(\hat{\mathbf{x}}_T) - \hat{\gamma}_i) - \bar{f}_0(\mathbf{x}) - \sum_{i=1}^m \hat{\lambda}_T^i(\bar{f}_i(\mathbf{x}) - \hat{\gamma}_i) \\ \leq G\sqrt{\frac{R^2 + D^2}{T}} + 2G(R + D)\sqrt{\frac{2}{T} \ln \frac{1}{\delta}} \end{aligned} \quad (9)$$

Using the definition of  $\hat{\gamma}_i$ , we have

$$\begin{aligned} \bar{f}_0(\hat{\mathbf{x}}_T) + \sum_{i=1}^m \lambda_i(\bar{f}_i(\hat{\mathbf{x}}_T) - \gamma_i) - \bar{f}_0(\mathbf{x}) - \sum_{i=1}^m \hat{\lambda}_T^i(\bar{f}_i(\mathbf{x}) - \gamma_i) \\ \leq G\sqrt{\frac{R^2 + D^2}{T}} + 2G(R + D)\sqrt{\frac{2}{T} \ln \frac{1}{\delta}} + \sum_{i=1}^m \hat{\lambda}_T^i \frac{\mu(\delta)}{T} \end{aligned} \quad (10)$$

We complete the proof by following the same steps as Theorem 1 and the fact  $\hat{\lambda}_T^i \leq \lambda_0^i$ .

## Appendix C. Proof of Proposition 1

We can rewrite (7) as

$$\min_{\mathbf{x} \in \mathcal{B}} \max_{\lambda \geq 0, \mathbf{p} \in \Delta} \bar{f}_0(\mathbf{x}) + \sum_{i=1}^m p_i \lambda (\bar{f}_i(\mathbf{x}) - \gamma_i)$$

where  $\Delta$  is a simplex. By redefining  $\lambda_i = p_i \lambda$ , we have the problem in (7) equivalent to (2) with  $\lambda = \sum_{i=1}^m \lambda_i$ .

## Appendix D. Proof of Lemma 1

Using the first order optimality condition, we have  $\lambda_a = \frac{|\nabla \bar{f}(\mathbf{x}_*)|}{|\partial g(\mathbf{x})|}$  where  $g(\mathbf{x}) = \max_{1 \leq i \leq m} \bar{f}_i(\mathbf{x})$ . Since  $\partial g(\mathbf{x}) \in \{\sum_{i=1}^m \alpha_i \nabla \bar{f}_i(\mathbf{x}) : \alpha \in \Delta_m\}$ , we complete the proof using Assumption 1.