2023 62nd IEEE Conference on Decision and Control (CDC)
December 13-15, 2023. Marina Bay Sands, Singapore

Data-Driven Stochastic Optimization Using Upper Confidence Bounds: Performance Guarantees and Distributional Robustness

Youngchae Cho Insoon Yang

***Abstract*— A data-driven approach for stochastic optimization should preferably offer a finite-sample performance guarantee as well as asymptotic optimality. However, existing approaches that enjoy both, mostly based on distributionally robust optimization (DRO), are often difficult to use as they require an ambiguity set containing the true distribution of uncertainty with confidence, which is difficult to obtain. The main contributions of this paper are two-fold. First, we propose a data-driven stochastic optimization approach that offers a finite-sample guarantee and asymptotic optimality without resorting to an ambiguity set. The core idea of our approach is to minimize an approximation to the true expected cost function derived using only a half of the sample data, and compute an upper confidence bound for the true expected cost using the other half. Second, we identify a case in which the proposed approach is tractable and accordingly design an algorithm for it. Simulation results demonstrate that our approach can be less conservative than existing DRO methods in terms of not only the finite-sample guarantee but also the true expected cost for a wide range of frequently used confidence levels.**

I. Introduction

The ultimate goal of many practical decision-making problems under uncertainty is to obtain a high-quality solution to

$$\inf_{x\in\mathcal{X}} \mu(x) := \mathbb{E}_{\xi\sim\mathbb{P}}[f(x,\xi)] \tag{1}$$

where x, $\mathcal{X}$, ξ, $\mathbb{P}$, $f:\mathcal{X}\times\Xi\to\mathbb{R}$, and $\mu:\mathcal{X}\to\mathbb{R}$ represent a decision vector, the feasible set of x, a random vector, the true distribution of ξ, a cost function, and the true expected cost function, respectively. We denote by Ξ and $\mathbb{R}$ the support of ξ and the real number set, respectively. Solving (1) is not easy regardless of the conditions for $\mathcal{X}$ and f, as $\mathbb{P}$ is mostly unknown in practice. To address (1) under the obscurity of $\mathbb{P}$, various data-driven approaches have been developed.

The sample average approximation (SAA) is one of the easiest data-driven approaches to understand and implement, which replaces $\mathbb{P}$ in (1) with an empirical distribution [1], [2]. Since it is consistent with solving (1) asymptotically as the sample size approaches infinity under mild assumptions, SAA has been extensively studied in numerous application domains. However, SAA generally fails to provide a probabilistic performance guarantee for finitely many samples.

Preferably, data-driven approaches for (1) should be consistent with solving (1) and offer a finite-sample performance guarantee. If a method is inconsistent, then it is implied that the method does not exploit data effectively to minimize the true expected cost and may yield solutions too far from any optimum no matter how many data are considered by the method. Moreover, if no performance guarantee is made for a finite number of samples, we may incur unexpectedly high costs arbitrarily often when there are only a few samples.[1]

Data-driven approaches for (1) offering a finite-sample guarantee while asymptotically consistent with solving (1) include distributionally robust optimization (DRO), or more specifically, Wasserstein DRO [3]–[6] and robust SAA [7]. The basic concept of DRO is to minimize the worst-case expected cost over a family of distributions of uncertainty, called an *ambiguity set*. Performances of DRO greatly depend on how an ambiguity set has been constructed. Although there is an extensive body of research on DRO with different ambiguity sets (e.g., [8]), we introduce only [3]–[7], focusing on the two properties of our interest: a finite-sample performance guarantee and asymptotic consistency.

In [3], an ambiguity set is defined as a ball, referred to as a Wasserstein ball, in the space of probability distributions centered at an empirical distribution of uncertainty. The distance from a distribution to its center is calculated using the Wasserstein metric. The authors also proved the existence of a Wasserstein ball's radius that acquires both finite-sample and asymptotic performance guarantees. However, computing such a radius is a challenging task in practice.

To obtain a theoretically valid radius of a Wasserstein ball in a less demanding way, explicit formulae presented in [4] and [5, Proposition 24] may be used. However, the former is often too conservative [9], while the latter requires an assumption on the dimension of uncertainty and the order of the Wasserstein metric. [6] designs a better structured ambiguity set with the Wasserstein metric. However, it is effective only for multiple independent uncertainties.

Meanwhile, robust SAA [7] defines an ambiguity set as a family of all distributions that pass a goodness-of-fit (GoF) test, which examines the hypothesis that a set of sample data is drawn from a distribution. If a statistic associated with a GoF test, which depends on both the sample data and the hypothetical distribution, is greater than a threshold called a critical value, then the hypothesis is rejected. The critical value is defined as a quantile of the statistic for the true distribution. As such, an ambiguity set for robust SAA contains the true distribution with probability encoded by the critical value. The authors showed that this approach leads to

This work was supported in part by the National Research Foundation of Korea (MSIT2020R1C1C1009766), the Information and Communications Technology Planning and Evaluation (IITP) grants (MSIT2022-0-00124, MSIT2022-0-00480), and Samsung Electronics.

The authors are with the Department of Electrical and Computer Engineering, Seoul National University, Seoul, South Korea. {youngchaecho,insoonyang}@snu.ac.kr

[1]Throughout the paper, we adopt the frequentist interpretation of probability in discussing and analyzing any property of a data-driven stochastic optimization approach.

the two performance guarantees for a wide class of problems. However, obtaining a critical value is burdensome unless it is distribution-free. Furthermore, a critical value should be computed via simulations in general, whose precision can affect the quality of a robust SAA problem.

In essence, although Wasserstein DRO [3]–[6] and robust SAA [7] can provide both performance guarantees, they may be hard to implement in practice. Thus far, the rules presented for constructing a theoretically meaningful ambiguity set are either limitedly applicable depending on the true distribution of uncertainty or overly conservative.

To resolve this issue, we present a novel data-driven approach for (1) without resorting to an ambiguity set. The proposed approach consists of two steps. As the first step, we construct an approximation to μ using a half of the sample data and minimize it. The data-driven approximation builds on an existing confidence bound on the tails of the difference between the true and sample means of a random variable. By design, the minimization problem we solve is asymptotically consistent with (1) under mild assumptions on $\mathcal{X}$, Ξ, and f. In the second step, we compute an upper confidence bound (UCB) for the true expected cost of the minimizer, utilizing the other half of the sample data. Further, if there exists a decision with a lower UCB of confidence 1, we employ it as our decision, disregarding the minimizer obtained at the first step. In this way, we can be offered a finite-sample performance guarantee. As we assume nothing about $\mathbb{P}$ except boundedness of Ξ, we can straightforwardly formulate the optimization problem in the first step, unlike with DRO. Nonetheless, the proposed approach admits an interpretation as an approximate solution method for a DRO problem concerning the distributional uncertainty of the cost as a random variable. Thus, it is distributionally robust as well.

Unfortunately, however, the minimization problem to solve in the first step is non-convex and intractable in general. To obtain a solution to it, we introduce additional assumptions on $\mathcal{X}$, Ξ, and f, including the monotonicity of the cost function with respect to each random variable. Subsequently, we design a tractable algorithm for the problem, which consists of solving finitely many linear programs (LPs). Using two test problems in different application areas, we numerically show that our approach can be less conservative compared to the DRO methods [4], [5], and [7].

Our main contributions are summarized as follows:

- We propose a data-driven approach for (1) that offers a finite-sample guarantee, asymptotic consistency, and distributional robustness without using an ambiguity set.
- We identify a case where the novel approach is tractable and correspondingly design an algorithm for it.

The remainder of this paper is organized as follows. In Section II, we describe in detail the two steps of the proposed approach for (1). In Section III, we analyze the three theoretical properties of the proposed approach. In Section IV, we present a tractable algorithm for the proposed approach. In Section V, we discuss the simulation results. In Section VI, we provide concluding remarks.

II. Proposed Method

In this section, we explain the proposed data-driven stochastic optimization method for (1). It builds on the following mild assumption about $\mathcal{X}$, Ξ, and f.

Assumption 1: First, $\mathcal{X}$ and Ξ are non-empty compact sets in $\mathbb{R}^n$ and $\mathbb{R}^m$, respectively. Second, f is continuous.

Assumption 1 implies that $f(x,\xi)$ attains the maximum and minimum in terms of $\xi \in \Xi$ for any $x \in \mathcal{X}$. Further, we observe that $f(x,\xi)$ is a random variable on the same probability space as ξ for any $x \in \mathcal{X}$. The support of $f(x,\xi)$ is denoted by $\mathcal{F}(x) := [l(x), u(x)] \subset \mathbb{R}$, where

$$l(x) := \min_{\xi\in\Xi} f(x,\xi) \quad \text{and} \quad u(x) := \max_{\xi\in\Xi} f(x,\xi).$$

We also assume that $N > 1$ samples $\xi_1, \ldots, \xi_N$ of ξ are available. Let $\xi^{(1)} := (\xi_1, \ldots, \xi_M)$ and $\xi^{(2)} := (\xi_{M+1}, \ldots, \xi_{2M})$ where $M := \lfloor N/2 \rfloor$. For any $x \in \mathcal{X}$, we denote the sample means of $f(x,\xi)$ for $\xi^{(1)}$ and $\xi^{(2)}$ by $\hat{\mu}_M^{(1)}(x) := \frac{1}{M}\sum_{i=1}^{M} f(x,\xi_i)$ and $\hat{\mu}_M^{(2)}(x) := \frac{1}{M}\sum_{i=1}^{M} f(x,\xi_{M+i})$, respectively.

The proposed approach comprises two steps. The first step is to minimize an approximation to μ, which is derived using $\xi^{(1)}$. Specifically, we solve

$$\min_{x\in\mathcal{X}} \overline{\mu}_M^{(1)}(x) := \hat{\mu}_M^{(1)}(x) + \sqrt{\frac{1}{2M}\log\frac{1}{1-\gamma}}\, D(x), \quad (2)$$

where $\gamma \in (0,1)$ denotes a user-defined parameter and $D(x) := u(x) - l(x)$. The role of γ becomes clear in Section III-A. Problem (2) always has an optimal solution due to Assumption 1. Let $x_M^{(1)*}$ and $\overline{\mu}_M^{(1)*}$ denote any solution and the optimal value of (2), respectively. In the second step of our approach, we compute $\overline{\mu}_M^{(2)}(x_M^{(1)*})$, where

$$\overline{\mu}_M^{(2)}(x) := \hat{\mu}_M^{(2)}(x) + \sqrt{\frac{1}{2M}\log\frac{1}{1-\gamma}}\, D(x),$$

and solve

$$\min_{x\in\mathcal{X}} u(x). \quad (3)$$

Let $x^{\mathrm{r}*}$ and u^* denote any solution and the optimal value of (3), respectively. Then, we define our data-driven solution to (1) as

$$x_N^* := \begin{cases} x_M^{(1)*} & \text{if} \quad \overline{\mu}_M^{(2)}\left(x_M^{(1)*}\right) \le u^* \\ x^{\mathrm{r}*} & \text{otherwise} \end{cases}$$

and let

$$\overline{\mu}_N^* := \min\left\{\overline{\mu}_M^{(2)}\left(x_M^{(1)*}\right), u^*\right\}.$$

In the following section, we analyze the theoretical properties of x_N^* and $\overline{\mu}_N^*$.

III. Theoretical Properties

In this section, we prove that the proposed approach has the two desirable properties of our interest, i.e., a finite-sample guarantee and asymptotic consistency. Furthermore, we show that our approach is also distributionally robust.

A. Finite-Sample Performance Guarantee

Ideally, a data-driven approach for (1) should allow the true expected cost to be bounded above with confidence for any finite number of samples. This is because we may incur unexpectedly high costs arbitrarily often otherwise. For example, SAA is unreliable as it lacks such a guarantee and often underestimates the true expected cost [3]. In this subsection, we prove that our approach offers a finite-sample guarantee, which is a corollary to the following theorem.

Theorem 1: For any $x \in \mathcal{X}$, we have

$$\Pr\left[\mu(x) < \overline{\mu}_M^{(2)}(x)\right] \geq \gamma.$$

Proof: Fix any $x \in \mathcal{X}$. As $f(x,\xi) - \mu(x)$ can be considered as a $D(x)/2$-subgaussian random variable, it follows from [10, Corollary 5.5] that

$$\Pr\left[\mu(x) \geq \hat{\mu}_M^{(2)}(x) + \alpha\right] \leq \exp\left[-\frac{2M\alpha^2}{D(x)^2}\right] \quad \forall \alpha \geq 0.$$

We have the result by letting $\alpha = \sqrt{\frac{1}{2M}\log\frac{1}{1-\gamma}}D(x)$. ■

Corollary 1 (Finite-sample performance guarantee): The true expected cost $\mu(x_N^*)$ of x_N^* is bounded above by $\overline{\mu}_N^*$ with confidence at least γ for any $N < \infty$, i.e., $\Pr[\mu(x_N^*) < \overline{\mu}_N^*] \geq \gamma$.

Proof: We have $\mu(x^{\mathrm{r}*}) \leq u^*$ by the definitions of $x^{\mathrm{r}*}$ and u^*. By the theorem and the law of total probability, we also have

$$\Pr\left[\mu\left(x_M^{(1)*}\right) < \overline{\mu}_M^{(2)}\left(x_M^{(1)*}\right)\right] \geq \gamma.$$

Hence, the statement holds. ■

The finite-sample guarantee can also be enjoyed in theory by the DRO approaches [3]–[7] using an ambiguity set that contains $\mathbb{P}$ with confidence. However, constructing such an ambiguity set is usually a cumbersome task, which might even require the identification of $\mathbb{P}$. Thus, the confidence level associated with our approach can be tuned more straightforwardly than that associated with a general DRO method. Furthermore, [4] provides a rule to build a theoretically valid ambiguity set without making an assumption on $\mathbb{P}$ beyond the boundedness of Ξ, which is worth comparing to our approach. We numerically investigate the relationship between our approach and this DRO method in Section V.

B. Asymptotic Consistency

It is also crucial for a data-driven approach for (1) to be asymptotically consistent with solving (1). Otherwise, we cannot be assured that sample data are efficiently exploited to achieve our ultimate goal of minimizing the true expected cost. For example, a DRO method with an ambiguity set based on moment conditions [11], which does not satisfy this property, may not reach an optimum even when many samples are used [7]. The following theorem shows that our approach is asymptotically consistent with solving (1).

Theorem 2 (Asymptotic consistency): Let μ^* and $\mathcal{M}^*$ denote the optimal value and the set of solutions of (1), respectively. We have that $d(x_N^*, \mathcal{M}^*) \to 0$ a.s. as $N \to \infty$, where $d(x_N^*, \mathcal{M}^*) := \inf_{y\in\mathcal{M}^*}\|x_N^* - y\|_2$ denotes a distance between x_N^* and $\mathcal{M}^*$, and $\overline{\mu}_N^* \to \mu^*$ a.s. as $N \to \infty$.

Proof: Since $x^{\mathrm{r}*}$ and u^* are invariant with respect to the sample data, it is enough to show that

$$d(x_M^{(1)*}, \mathcal{M}^*) \to 0 \quad \text{a.s.} \tag{4}$$

and

$$\overline{\mu}_M^{(2)*}(x_M^{(1)*}) \to \mu^* \quad \text{a.s.} \tag{5}$$

We first prove (4). Let $\overline{D} := \max_{x\in\mathcal{X}} D(x)$, which is finite. According to [1, Proposition 8], the sequence of functions $\{\hat{\mu}_M^{(1)}\}$ uniformly converges to μ a.s. as $M \to \infty$. This implies that, for any $\varepsilon > 0$, there exists a natural number $M_1 > \frac{\overline{D}^2}{2\varepsilon^2}\log\frac{1}{1-\gamma}$ such that

$$\sup_{x\in\mathcal{X}}|\hat{\mu}_M^{(1)}(x) - \mu(x)| < \varepsilon \quad \text{a.s.}$$

for any $M > M_1$. Thus, there exists a natural number $M_2 > M_1$ such that

$$\sup_{x\in\mathcal{X}}|\overline{\mu}_M^{(1)}(x) - \mu(x)| < \varepsilon \quad \text{a.s.}$$

for any $M > M_2$. This indicates that $\overline{\mu}_M^{(1)}$ uniformly converges to μ a.s. as $M \to \infty$. The same holds for $\overline{\mu}_M^{(2)}$. Further, we observe that μ is continuous due to Assumption 1 [1, Proposition 1]. Thus, according to [1, Theorem 9], we have $\overline{\mu}_M^{(1)*} \to \mu^*$ and $d(x_M^{(1)*}, \mathcal{M}^*) \to 0$ a.s. as $M \to \infty$.

Next, we prove (5). From the proof of (4), we observe the almost sure uniform convergence of $\overline{\mu}_M^{(1)}$ and $\overline{\mu}_M^{(2)}$ to μ as well as the almost sure convergence of $\overline{\mu}_M^{(1)*}$ to μ^*. Thus, for any $\varepsilon > 0$, there exists a natural number M_3 such that

$$\sup_{x\in\mathcal{X}}\left|\overline{\mu}_M^{(2)}(x) - \mu(x)\right| + \sup_{x\in\mathcal{X}}\left|\overline{\mu}_M^{(1)}(x) - \mu(x)\right| + \left|\overline{\mu}_M^{(1)*} - \mu^*\right| < \varepsilon \quad \text{a.s.}$$

for any $M > M_3$. Since we have

$$\left|\overline{\mu}_M^{(2)}\left(x_M^{(1)*}\right) - \mu^*\right| \leq \sup_{x\in\mathcal{X}}\left|\overline{\mu}_M^{(2)}(x) - \mu(x)\right| + \sup_{x\in\mathcal{X}}\left|\overline{\mu}_M^{(1)}(x) - \mu(x)\right| + \left|\overline{\mu}_M^{(1)*} - \mu^*\right|$$

for any $(\xi^{(1)}, \xi^{(2)})$, (5) holds and the proof is complete. ■

C. Distributional Robustness

As the name suggests, DRO methods for (1) are distributionally robust, i.e., robust against misspecifications of $\mathbb{P}$ as well as changes in $\mathbb{P}$. This is another advantageous quality of a data-driven stochastic optimization method, since it can mitigate the optimizer's curse [3] and the data-generating distribution $\mathbb{P}$ might vary over time. In this subsection, we show that our approach is also distributionally robust.

To this end, we first define

$$h_{M,\theta}^{(2)}(x) := \max_{\mathbb{Q}\in\mathcal{Q}_{M,\theta}^{(2)}(x)} \mathbb{E}_{f(x,\xi)\sim\mathbb{Q}}[f(x,\xi)] \tag{6}$$

for any $x \in \mathcal{X}$. Let $\mathcal{Q}(x)$ and $\delta_{(\cdot)}$ denote the family of all probability distributions supported on $\mathcal{F}(x)$ and the

Dirac distribution, respectively. In (6), $\mathcal{Q}^{(2)}_{M,\theta}(x) \subseteq \mathcal{Q}(x)$ represents the 1-Wasserstein ball of radius $\theta > 0$ centered at the empirical distribution $\mathbb{Q}^{(2)}_M(x) := \frac{1}{M}\sum_{i=1}^M \delta_{f(x,\xi_{M+i})}$ of $f(x,\xi)$ for $\xi^{(2)}$, defined with the 1-norm as the metric on $\mathcal{F}(x)$. That is, we let

$$\mathcal{Q}^{(2)}_{M,\theta}(x) := \left\{\mathbb{Q} \in \mathcal{Q}(x) : d^{\mathrm{w}}\left(\mathbb{Q}, \mathbb{Q}^{(2)}_M(x)\right) \leq \theta\right\},$$

where d^{w} denotes the Wasserstein metric of order 1 defined with the 1-norm for two probability distributions supported on $\mathcal{F}(x)$, i.e., for any $\mathbb{Q}_1, \mathbb{Q}_2 \in \mathcal{Q}(x)$,

$$d^{\mathrm{w}}(\mathbb{Q}_1, \mathbb{Q}_2) := \inf_{\pi \in \Pi(\mathbb{Q}_1, \mathbb{Q}_2)} \int_{\mathcal{F}(x)\times\mathcal{F}(x)} |y_1 - y_2| \pi(dy_1, dy_2).$$

Here, $\Pi(\mathbb{Q}_1, \mathbb{Q}_2)$ denotes the set of joint distributions supported on $\mathcal{F}(x) \times \mathcal{F}(x)$ with marginals $\mathbb{Q}_1$ and $\mathbb{Q}_2$. Thus, $h^{(2)}_{M,\theta}(x)$ represents the worst-case expected cost of decision x over the decision-dependent Wasserstein ball $\mathcal{Q}^{(2)}_{M,\theta}(x)$.

Subsequently, we present the following theorem, which implies a relationship between the proposed approach and the minimization of $h^{(2)}_{M,\theta}$ for a specific choice of θ.

Theorem 3: For any $x \in \mathcal{X}$, we have

$$h^{(2)}_{M,\overline{\theta}_M(x)}(x) = \min\left\{\overline{\mu}^{(2)}_M(x), u(x)\right\},$$

where $\overline{\theta}_M(x) := \sqrt{\frac{1}{2M}\log\frac{1}{1-\gamma}}D(x)$.

Proof: For any $x \in \mathcal{X}$, $h^{(2)}_{M,\theta}(x)$ is the expectation of an affine function of the random variable $f(x,\xi)$ with the bounded support $\mathcal{F}(x)$ with respect to worst-case distributions in $\mathcal{Q}^{(2)}_{M,\theta}(x)$. Thus, according to [3, Theorem 4.4], $h^{(2)}_{M,\theta}(x)$ is equal to the optimal value of the problem

$$\begin{aligned}\max_{q_i\in\mathbb{R}} \quad & \hat{\mu}^{(2)}_M(x) + \frac{1}{M}\sum_{i=1}^M q_i \\ \text{s.t.} \quad & \frac{1}{M}\sum_{i=1}^M |q_i| \leq \theta \\ & l(x) \leq f(x,\xi_i) + q_i \leq u(x) \quad \forall i \leq M.\end{aligned}$$

By imposing non-negativeness on q_i for each $i = 1, \ldots, M$ without loss of optimality, we can rewrite this problem as

$$\begin{aligned}\max_{q_i\in\mathbb{R}_+} \quad & \hat{\mu}^{(2)}_M(x) + \frac{1}{M}\sum_{i=1}^M q_i \\ \text{s.t.} \quad & \frac{1}{N}\sum_{i=1}^M q_i \leq \theta \\ & q_i \leq u(x) - f(x,\xi_i) \quad \forall i = 1, \ldots, M,\end{aligned}$$

where $\mathbb{R}_+$ denotes the set of non-negative real numbers. Thus, we have

$$h^{(2)}_{M,\theta}(x) = \min\left\{\hat{\mu}^{(2)}_M(x) + \theta, u(x)\right\}$$

for any $\theta > 0$, and the proof is complete. ∎

By Theorem 3, we can consider the proposed approach as an approximate solution method for the Wasserstein DRO problem

$$\min_{x\in\mathcal{X}} h^{(2)}_{M,\overline{\theta}_M(x)}(x),$$

whose optimal value is equal to the minimum of $\overline{\mu}^{(2)*}_M$ and u^*. The following corollary describes more specifically the distributional robustness inherent in our approach.

Corollary 2 (Distributional robustness): The expectation of $f(x^*_N, \xi)$ as a random variable with respect to any distribution in $\mathcal{Q}^{(2)}_{M,\theta}(x^*_N)$ for $\theta = \overline{\theta}_M(x^*_N)$ is bounded above by $\overline{\mu}^*_N$, i.e.,

$$\mathbb{E}_{f(x^*_N,\xi)\sim\mathbb{Q}}[f(x^*_N,\xi)] \leq \overline{\mu}^*_N \quad \forall \mathbb{Q} \in \mathcal{Q}^{(2)}_{M,\overline{\theta}_M(x^*_N)}(x^*_N).$$

Proof: The result is immediate from the theorem and the definition (6) of $h^{(2)}_{M,\theta}$. ∎

Corollary 2 suggests that the proposed approach is robust against the imperfect representation $\mathbb{Q}^{(2)}_M(x^*_N)$ of the true distribution of $f(x^*_N, \xi)$. Moreover, even when $\mathbb{P}$ is not fixed, if $\mathbb{P}$ varies in such a way that the corresponding distribution of $f(x^*_N, \xi)$ belongs to $\mathcal{Q}^{(2)}_{M,\theta}(x^*_N)$ for $\theta = \overline{\theta}_M(x^*_N)$, then the true expected cost $\mu(x^*_N)$ is no greater than $\overline{\mu}^*_N$. In this sense, our approach is distributionally robust.

Despite its favorable theoretical properties, however, our approach may not be implementable in practice. This is mainly because (2) is generally non-convex. We discuss the implementability of our approach in the following section.

IV. Solution Method

To use our approach, we have to solve (2) and (3). Unlike (3), (2) is usually intractable as $D(x)$ is non-convex even when, for example, $f(x,\xi)$ is jointly convex. In this section, we develop a tractable algorithm for (2) under extra assumptions on $\mathcal{X}$, Ξ, and f. Problem (3) can be easily rewritten as an LP under the conditions. We introduce the following assumption, which implies Assumption 1.

Assumption 2: First, $\mathcal{X} \subset \mathbb{R}^n$ and $\Xi \subset \mathbb{R}^m$ are a compact convex polytope and a compact m-dimensional interval, respectively. Second, f can be represented as a pointwise maximum of $J < \infty$ affine functions over Ξ whose coefficient vectors and constant terms are affine functions over $\mathcal{X}$, i.e., $f(x,\xi) = \max_{j\leq J} a_j^\top(x)\xi + b_j(x)$, where $a_j : \mathcal{X} \to \mathbb{R}^m$ and $b_j : \mathcal{X} \to \mathbb{R}$ are affine for any $j = 1, \ldots, J$. Finally, for any $x \in \mathcal{X}$, the J affine functions of $\xi \in \Xi$ that constitute $f(x,\xi)$ are either simultaneously increasing or simultaneously decreasing with any entry of ξ, i.e., $a_j(x) \circ a_{j'}(x') \in \mathbb{R}^m_+$ for any $(j, j', x, x') \in \mathcal{J}^2 \times \mathcal{X}^2$ where $\circ$ denotes the entrywise product of two vectors.

With Assumption 2 holding for the rest of this section, we first explain how to solve (2). Problem (2) is rewritten as

$$\begin{aligned}\min_{\substack{x\in\mathcal{X},\\ \eta_i,\overline{\eta},\underline{\eta}\in\mathbb{R}}} \quad & \frac{1}{M}\sum_{i=1}^M \eta_i + \sqrt{\frac{1}{2M}\log\frac{1}{1-\gamma}}\left(\overline{\eta} + \underline{\eta}\right) \\ \text{s.t.} \quad & \eta_i \geq a_j^\top(x)\xi_i + b_j(x) \quad \forall i \leq M, \forall j \leq J \\ & \overline{\eta} = \max_{\xi\in\Xi, j\leq J} a_j^\top(x)\xi + b_j(x) \\ & \underline{\eta} = -\min_{\xi\in\Xi}\max_{j\leq J} a_j^\top(x)\xi + b_j(x),\end{aligned} \tag{7}$$

where η_i, $\overline{\eta}$, and $\underline{\eta}$ represent $f(x,\xi_i)$, $u(x)$, and $-l(x)$, respectively. Since the sign of each entry in $a_j(x)$ is invariant

Algorithm 1 LP method for (2) under Assumption 2

Input: Samples $\xi_1, \ldots, \xi_M$ and confidence level $\gamma \in (0,1)$
Output: Solution $x_M^{(1)*}$
  for $k \leftarrow 1$ to J **do**
    Solve (9) and obtain $(x^{(1)k}, \overline{\mu}^{(1)k})$ ▷ LP
  end for
  $k^* \leftarrow \arg\min_{k \leq J} \overline{\mu}^{(1)k}$ and $x_M^{(1)*} \leftarrow x^{(1)k^*}$

with $x \in \mathcal{X}$ and Ξ is a box, (7) is rewritten as

$$\begin{array}{ll} \min\limits_{\substack{x\in\mathcal{X},\\ \eta_i,\overline{\eta},\underline{\eta}\in\mathbb{R}}} & \dfrac{1}{M}\sum\limits_{i=1}^{M}\eta_i + \sqrt{\dfrac{1}{2M}\log\dfrac{1}{1-\gamma}}\left(\overline{\eta}+\underline{\eta}\right) \\ \text{s.t.} & \eta_i \geq a_j^\top(x)\,\xi_i + b_j(x) \quad \forall i \leq M, \forall j \leq J \\ & \overline{\eta} \geq a_j^\top(x)\,\xi^+ + b_j(x) \quad \forall j \leq J \\ & \underline{\eta} \geq \min\limits_{j\leq J} -a_j^\top(x)\,\xi^- - b_j(x) \end{array} \tag{8}$$

with $\xi^+ := \arg\max_{\xi\in\Xi} \sigma^\top \xi$ and $\xi^- := \arg\min_{\xi\in\Xi} \sigma^\top \xi$. Here, $\sigma \in \{1,-1\}^m$ denotes a vector such that $\sigma \circ a_j(x) \in \mathbb{R}^m_+$ for any $j = 1, \ldots, J$ and $x \in \mathcal{X}$, whose existence is due to Assumption 2. We address (8) by solving the LP

$$\begin{array}{ll} \min\limits_{\substack{x\in\mathcal{X},\\ \eta_i,\overline{\eta},\underline{\eta}\in\mathbb{R}}} & \dfrac{1}{M}\sum\limits_{i=1}^{M}\eta_i + \sqrt{\dfrac{1}{2M}\log\dfrac{1}{1-\gamma}}\left(\overline{\eta}+\underline{\eta}\right) \\ \text{s.t.} & \eta_i \geq a_j^\top(x)\,\xi_i + b_j(x) \quad \forall i \leq M, \forall j \leq J \\ & \overline{\eta} \geq a_j^\top(x)\,\xi^+ + b_j(x) \quad \forall j \leq J \\ & \underline{\eta} \geq -a_k^\top(x)\,\xi^- - b_k(x) \end{array} \tag{9}$$

for each $k = 1, \ldots, J$. Let $x^{(1)k}$ and $\overline{\mu}^{(1)k}$ denote a solution corresponding to x and the optimal value of (9), respectively. By definition, $x^{(1)k^*}$ for any $k^* \in \arg\min_{k\leq J} \overline{\mu}^{(1)k}$ is a solution to (8). As a solution can be obtained by solving J LPs, (2) is tractable. We provide a pseudocode for (2) in Algorithm 1. Meanwhile, (3) is rewritten as the LP

$$\min_{x\in\mathcal{X}} \quad \overline{\eta} \quad \text{s.t.} \quad \overline{\eta} \geq a_j(x)\,\xi^+ + b_j(x) \quad \forall j \leq J. \tag{10}$$

Thus, our approach can be employed by solving $(J+1)$ LPs.

In the following section, we numerically show that the proposed approach can be less conservative than existing DRO methods for a broad range of confidence levels.

V. Numerical Examples

In this section, we examine the performances of the proposed approach using numerical experiments. Specifically, we compare our approach against two Wasserstein DRO methods with ambiguity sets whose radii are set using the explicit formulae in [4] (DROa) and [5, Proposition 24] (DROb), respectively, as well as a robust SAA method (DROc) in terms of the UCBs associated with the finite-sample guarantees and the true expected costs. The GoF test for DROc is based on linear-convex ordering [7, Theorem 16]. We test the four methods on two decision-making problems that meet Assumption 2: the newsvendor problem and the farm management problem. We explain the simulation scheme in the following subsection.

A. Simulation Scheme

For each problem, we first randomly generate $N = 10, 100, 1000$ sample data according to the true distribution. After obtaining a solution and its associated UCB of each method for confidence $\gamma = 0.05, 0.1 \ldots, 0.95$, we estimate the true expected cost using SAA with another 100,000 randomly generated samples. For statistical robustness, we repeat this process 1,000 times with independent sample datasets. We then compare the average UCBs and the average true expected costs of the methods. For reference, we also estimate the minimum true expected cost, using SAA with another 1,000,000 randomly generated samples. Further, we compute the average true expected costs incurred by SAA. The source code of the simulations is available online.[2] We describe the details of each problem in the following subsection.

B. Problem Description

1) Newsvendor Problem: We first consider a single-period newsvendor problem [12] involving two continuous items with a combined stock capacity of 176. For each $i = 1, 2$, let x_i and ξ_i denote the inventory level and uncertain demand of item i, respectively. We formulate the problem as

$$\begin{array}{ll} \min\limits_{x_1,x_2\in\mathbb{R}_+} & \sum\limits_{i=1}^{2} c_i x_i + \mathbb{E}\left[\max\left\{-p_i\xi_i + q_i(x_i-\xi_i), -p_i x_i\right\}\right] \\ \text{s.t.} & x_1 + x_2 \leq 176, \end{array}$$

where c_i, p_i, and q_i denote the unit cost, price, and disposal cost of item i, respectively. We set $c_1 = 3$, $p_1 = 5$, $q_1 = 2$, $c_2 = 6$, $p_2 = 10$, and $q_2 = 6$. Assuming that ξ_1 and ξ_2 are independent random variables, we model the true distribution of ξ_1 (ξ_2) as the gamma distribution with shape 1.5 (3) and scale 20 (40) truncated over $[0, 100]$ ($[0, 120]$).

2) Farm Management Problem: The second illustrative problem is a variant of the farmer's problem depicted in [12], where a farmer who owns 500 acres of land determines how much area to allocate to three types of crops, wheat, corn, and sugar beet, with planting costs of \$150, \$230, and \$260 per acre, respectively. The farmer may need specific amounts of wheat and corn after harvesting, represented by independent random variables ξ_1 and ξ_2, respectively. Let x_1, x_2, and x_3 denote the areas of land allocated to wheat, corn, and sugar beet, respectively. We formulate the problem as

$$\begin{array}{ll} \min\limits_{x_1,x_2,x_3\in\mathbb{R}_+} & 150x_1 + 230x_2 + 260x_3 + \mathbb{E}\left[f(x,\xi_1,\xi_2)\right] \\ \text{s.t.} & x_1 + x_2 + x_3 \leq 500, \end{array}$$

where $f(x,\xi_1,\xi_2)$ is defined as the optimal value of the LP

$$\begin{array}{ll} \min\limits_{y_1,y_2,w_1,w_2,w_3,w_4\in\mathbb{R}_+} & 238y_1 - 170w_1 + 210y_2 \\ & \quad - 150w_2 - 36w_3 - 10w_4 \\ \text{s.t.} & 2.5x_1 + y_1 - w_1 \geq \xi_1 \\ & 3x_2 + y_2 - w_2 \geq \xi_2 \\ & w_3 + w_4 \leq 20x_3 \\ & w_3 \leq 6000. \end{array} \tag{11}$$

[2] https://github.com/CORE-SNU/UCB

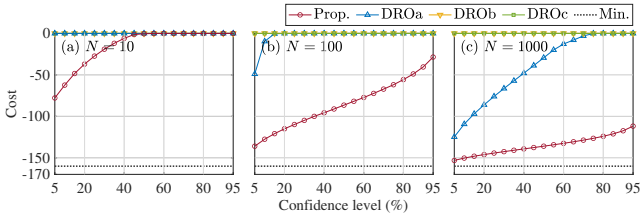

Fig. 1. The average UCB for the newsvendor problem.

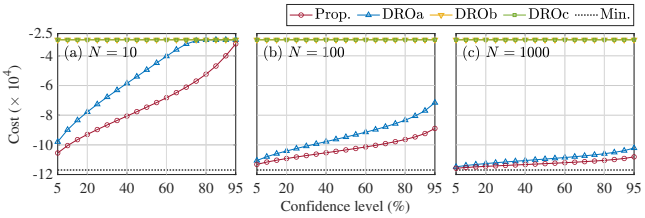

Fig. 2. The average UCB for the farm management problem.

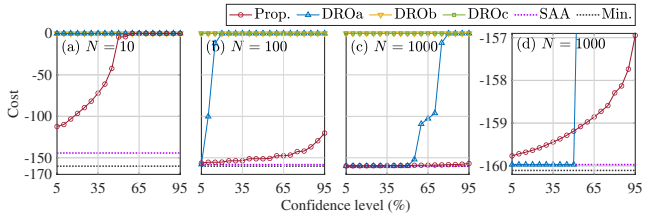

Fig. 3. The average true expected cost for the newsvendor problem. Subfigure (d) is a magnified view of (c).

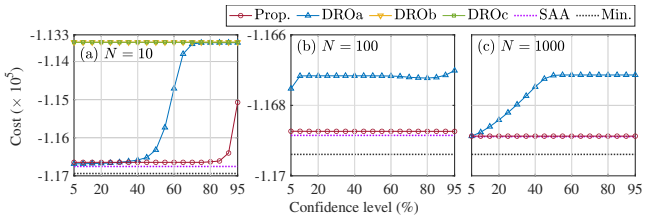

Fig. 4. The average true expected cost for the farm management problem. The results of DROb and DROc are invariant with N and γ.

The specifics of the decision variables, constraints, and objective function in (11) can be easily inferred from [12]. Using duality in LPs, we can rewrite (11) in the form described in Assumption 2 for $J = 8$. We model the true distribution of ξ_1 (ξ_2) as the normal distribution of mean 200 (240) and variance 60 (72) truncated over $[0, 400]$ ($[0, 480]$).

We discuss the simulation results in the following subsection.

C. Results

The simulation results are shown in Fig. 1–4. From Fig. 1 and 2, we observe that the proposed approach yields the lowest UCB on average for any sample size N and confidence level γ regardless of the problem. Fig. 3 and 4 reveal that it also leads to the lowest true expected cost on average for any N and $\gamma \geq 0.55$ on both problems. Thus, we conclude that our approach can be much less conservative than existing DRO methods under Assumption 2 for a wide range of commonly used confidence levels.

Meanwhile, our approach reveals a significant gap between the UCB and the true expected cost, especially when the confidence level is relatively high. Thus, we may need a technique to decrease the conservativeness of our approach, which remains future work. Furthermore, SAA demonstrates the lower true expected cost on average compared to our approach for any case. Although SAA does not provide an effective finite-sample guarantee by itself, it would also be intriguing to explore additional tactics for our approach to surpass SAA in terms of the true expected cost.

VI. Conclusions

We proposed a data-driven approach for stochastic optimization that offers not only asymptotic consistency but also a finite-sample guarantee. Without relying on an ambiguity set, our approach can be implemented more straightforwardly than DRO methods that provide both performance guarantees. The simulation results demonstrated that our approach can be less conservative than existing DRO methods in terms of the finite-sample guarantee as well as the true expected cost. However, our approach is tractable only for a limited class of problems. In addition, we found numerically that our approach can result in a large disparity between the UCB and the true expected cost, with the true expected cost being higher than that incurred by SAA on average. Thus, future research directions include improving its tractability and reducing its intrinsic conservativeness.

References

[1] S. Kim, R. Pasupathy, and S. G. Henderson, "A guide to sample average approximation," in *Handbook of Simulation Optimization*, M. C. Fu, Ed. Springer, 2015, pp. 207–243.

[2] A. J. Kleywegt, A. Shapiro, and T. Homem-de Mello, "The sample average approximation method for stochastic discrete optimization," *SIAM J. Optim.*, vol. 12, no. 2, pp. 479–502, 2002.

[3] P. M. Esfahani and D. Kuhn, "Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations," *Math. Program.*, vol. 171, no. 1, pp. 115–166, 2018.

[4] C. Zhao and Y. Guan, "Data-driven risk-averse two-stage stochastic program with ζ-structure probability metrics," *Optimization Online*, 2015. [Online]. Available: https://optimization-online.org/2015/07/5014/

[5] D. Boskos, J. Cortés, and S. Martínez, "High-confidence data-driven ambiguity sets for time-varying linear systems," *arXiv preprint arXiv:2102.01142*, 2021.

[6] L. M. Chaouach, D. Boskos, and T. Oomen, "Uncertain uncertainty in data-driven stochastic optimization: towards structured ambiguity sets," in *2022 IEEE 61st Conference on Decision and Control (CDC)*. IEEE, 2022, pp. 4776–4781.

[7] D. Bertsimas, V. Gupta, and N. Kallus, "Robust sample average approximation," *Math. Program.*, vol. 171, no. 1, pp. 217–282, 2018.

[8] H. Rahimian and S. Mehrotra, "Distributionally robust optimization: A review," *arXiv preprint arXiv:1908.05659*, 2019.

[9] C. Duan, W. Fang, L. Jiang, L. Yao, and J. Liu, "Distributionally robust chance-constrained approximate AC-OPF with Wasserstein metric," *IEEE Trans. Power Syst.*, vol. 33, no. 5, pp. 4924–4936, 2018.

[10] T. Lattimore and C. Szepesvári, *Bandit Algorithms*. Cambridge University Press, 2020.

[11] E. Delage and Y. Ye, "Distributionally robust optimization under moment uncertainty with application to data-driven problems," *Oper. Res.*, vol. 58, no. 3, pp. 595–612, 2010.

[12] J. R. Birge and F. Louveaux, *Introduction to Stochastic Programming*. Springer Science & Business Media, 2011.