Sampling-based Approximation Algorithms for Data Driven Capacitated Stochastic Inventory Control Models

Wang Chi Cheung, David Simchi-Levi

Sampling-based Approximation Algorithms for Data Driven Capacitated Stochastic Inventory Control Models

(Authors' names blinded for peer review)

Problem Definition. We study the classical multi-period capacitated stochastic inventory control models in a data-driven setting. At the start of each period, the decision maker decides how many units of inventory to order, subject to an upper limit on the ordering quantity. After ordering, the demand for the period is observed, and it is satisfied to the maximum extent by the available inventory. The excess inventory units, if any, are carried over to the next period; otherwise the unsatisfied demand units are backlogged. A holding cost is incurred for each unit of the remaining inventory, and a backlog cost is incurred for each unit of the unsatisfied demand. The demand distributions across the periods are independent, but not necessarily identically distributed. The objective of the decision maker is to minimize the sum of expected holding and backlog costs across all the periods.

In our data-driven model, the decision maker has no access to the cumulative distribution functions (CDFs) of the demand distributions. Rather, she only has access to independent random samples of the distributions. The only assumption on the underlying demand distributions is that they have finite means, in order for the problem to be well-defined. There is no parametric assumption, and we do not assume their high moments to be finite.

Related Works. It is well known that the optimal expected cost can be achieved by a modified base stock policy. When the CDFs of the demand distributions are given, the optimal set of base stocks can be computed by a suitable dynamic programming (DP). However, in our data-driven setting, it is not even clear that the expected cost function value can be estimated within any given accuracy, which makes the classical approach not directly applicable here.

The data-driven stochastic inventory control models without capacity constraints have been studied by Levi *et al.* [2]. Levi *et al.* [2] proposes a near-optimal sampling based policy that uses polynomially many samples and return a near optimal base stock policy with any desired confidence probability. Interestingly, the policy proposed in [2] is not a Sample Average Approximation (SAA) algorithm in general.

Contributions. We analyze the number of samples needed to ensure that, with any specified confidence probability, the expected cost under modified base stock policy output by the SAA algorithm has a small relative error compared to the optimal expected cost. We show the number of samples required can be chosen to be independent of the underlying demand distributions, and polynomial in the number of periods and the accuracy parameters.

By modifying a result of Halman *et al.* [1], we show that the empirical problem is computationally intractable. More precisely, it is NP-hard to compute the optimal policy, even when the demand distributions are explicitly given and has polynomial size support. This implies that even when the SAA algorithm uses polynomially many samples, it is unlikely to be a polynomial time algorithm. Thus, we propose a polynomial time approximation scheme that still achieves a small relative error to the optimal expected cost, by introducing a sparsification procedure to the SAA algorithm. This procedure can be seen as a first order analogue to the notion of *K*-approximation proposed in [1].

Finally, similar to [2], our sampling bound has a pseudo-polynomial dependence on the unit holding and backlog costs. We show that such dependence is in fact necessary for achieving near optimality, even for the problem with only one period, namely the newsvendor problem.

Note that our results can be easily extended to the case when ordering costs and lead times are present, and we are currently trying to extend our algorithms and analyses to the serial inventory systems.

Our Approach. We analyze the performances of our proposed sampling based algorithms by a first order analysis on the dynamic program (DP) of the problem. We estimate the subgradients of the cost-to-go functions of the DP based on the empirical demand distributions constructed using the samples. By suitably choosing the number of samples drawn in each period, we ensure that the estimation errors do not grow too quickly, while making sure that only polynomially many samples are used. Finally, we use the estimated subgradients to construct a modified base stock policy, and we prove its performance guarantees by using a Lemma of [2], which relates the bounds on the estimation errors on the subgradients to performance guarantee on the objective value. Similar to [2], the challenge in the analysis is to control the estimation error bounds due to random sampling and upper-bound the error accumulated from the previous stages in the DP. However, the presence of capacity constraints makes the analysis even more complicated and intricate than in [2].

References.

[1] N. Halman, D. Klabjan, M. Mostagir, J. Orlin and D. Simchi-Levi. A Fully Polynomial Time Approximation Scheme for Single-item Inventory Control with Discrete Demand. *Mathematics of Operations Research*, 34:674-685, 2009.

[2] R. Levi, R. Roundy, and D. B. Shmoys. Provably Near-optimal Sampling-based Policies for Stochastic Inventory Control Models. *Mathematics of Operations Research*, 32:821–838, 2007.