Contextual Stochastic Optimization Methods for Decision Making Under Uncertainty

Utsav Sadana*, Erick Delage†

Executive Summary: The field of decision-making under uncertainty is evolving towards data-driven approaches that incorporate side information (covariates) into decision processes. Recent advancements in machine learning and stochastic optimization have fostered a shift towards contextual optimization, utilizing covariates observed before decisions to mitigate uncertainty. Our survey paper covers three major frameworks of contextual optimization: decision rule optimization, sequential learning and optimization, and integrated learning and optimization. Practical challenges include identifying relevant covariates and developing scalable models. Despite these hurdles, contextual optimization holds significant potential for dynamic, personalized decision-making across various domains, necessitating further exploration and development.

Acknowledgments: We thank Prof. Yang Haoxiang and Prof. Georgio Consigli for providing us the opportunity to share our views on contextual stochastic optimization and suggesting improvements.

The field of decision-making under uncertainty is increasingly embracing data-driven approaches, which leverage historical data to prescribe actions. Traditionally, these methods have concentrated on quantifying uncertainty (assigning probability distributions in stochastic programming or constructing uncertainty sets for robust optimization) based only on previously observed realizations of the uncertain parameters affecting the performance of the decisions. However, a significant paradigm shift has occurred recently toward contextual optimization methods, propelled by advancements in machine learning and stochastic optimization, and the need for effective high-stakes decision-making in various domains, e.g., energy grid management, portfolio optimization, supply chain management, etc. This shift involves incorporating side information (covariates), which is revealed just before an action is taken, into the decision-making pipeline to mitigate the uncertainty more effectively . The side information could include weather conditions, which affect road congestion, in a last-mile delivery problem, or social media posts reflecting on publicly traded companies' recent achievements in a portfolio management problem.

Specifically, in a contextual optimization problem, a decision maker observes a set of covariates that are correlated with the uncertain parameters before taking an action. One further assumes that the decision-maker has access to a set of historical observations of covariates and corresponding uncertain parameters' values sampled from the true, yet unknown, underlying distribution. The objective of the decision-maker is to find a policy, which suggests an action as a function of the covariates, that minimizes an expected cost that depends on the joint distribution of the covariates and uncertain parameters. [9] identified three data-driven frameworks that together cover most of the approaches proposed in the literature on contextual optimization. In this column, we summarize these three frameworks and point to recent advancements to stimulate further interest in this topic. While it's impractical to delve into all relevant studies in this

*Department of Computer Science and Operations Research, University of Montreal, utsav.sadana@umontreal.ca

†GERAD and Department of Decision Sciences, HEC Montreal, erick.delage@hec.ca

column, we have selected a representative set of papers based on our subjective judgment.

The first approach, which we refer to as decision rule optimization, involves explicitly defining the policy as a parametric function of observed covariates and solving the resulting empirical risk minimization problem. This method gained prominence in the operations research community through the seminal paper [2] "Big Data Newsvendor: Practical Insights from Machine Learning", where a regularized newsvendor problem is solved with an order quantity that depends linearly in the covariates. Unfortunately, linear decision rules can be too rigid to achieve the best possible performance even with an arbitrarily large dataset. To enhance performance, recent studies have explored projecting the covariates to a reproducing kernel Hilbert space and using linear decision rules in the lifted space. This approach retains some of the computational advantages of linear methods while achieving greater adaptability. Alternatively, the adoption of nonlinear policies through deep neural networks, decision trees and distributionally robust decision rules has further expanded the capability of decision models to handle more complex, nonlinear relationships between covariates and decision outcomes.

A second approach proposes learning a conditional probability distribution that is then exploited by a stochastic program to produce a policy. This can address the asymptotic optimality issue of decision rules when the distribution model is well specified given that with enough data, the true model could be recovered. We refer to this framework as sequential learning and optimization (SLO, also known as estimate then optimize) where the first step consists in learning a prediction of the conditional distribution and the second step involves solving a stochastic program with this model conditioned on the new observed realization of the covariates is observed. An interesting and widely studied special case arises when the cost function is linear in the parameters, e.g., in shortest-path problems. By linearity of expectation, SLO reduces to predicting the conditional mean of parameters, served by training using a least square error loss, followed by a deterministic optimization involving the expected cost vector. For non-linear cost functions, a popular method employs a discrete conditional distribution model, e.g., residuals-based distribution, k-nearest neighbour, kernel density estimation, random forests, etc. We refer interested readers to [5] for a benchmarking study of such SLO and decision rule methods.

The effectiveness of sequential learning and optimization heavily depends on the accuracy of the prediction model. When there are prediction errors, which is usually the case with limited data, the decisions could be suboptimal. By separating the prediction and optimization stages, we may fail to capture the potential feedback loop where decision-quality could be used to tune the prediction model. This limitation of sequential learning and optimization is addressed with an Integrated Learning and Optimization framework (see Figure 1) that aims to learn a prediction model to improve the decision-quality, rather than reducing the prediction/statistical error [3]. For instance, in a health supply chain management context, [6] observe that an integrated learning and optimization model can discriminate between the poorly stocked and well-stocked facilities based on the inventory optimization problem, and thereby, shifts its predictive precision towards the facilities where ensuring adequate supplies is crucial. This targeted prediction significantly reduces the risk of understocking in regions with high demand and lower prior inventory levels.

With the advancements made in automatic differentiation for end-to-end training of differentiable composite functions, there is a surge of interest in identifying convex loss functions and differentiable proxies for non-differentiable components in the integrated learning and optimization framework. One of the computational challenges is that when the cost function is bilinear in the action and uncertain parameter, the solution of the linear program is piecewise constant with respect to the uncertain parameter's distribution, hence offering no useful information regarding the direction of improvement. To circumvent this issue, the expected cost estimator can be learned using a convex approxi-

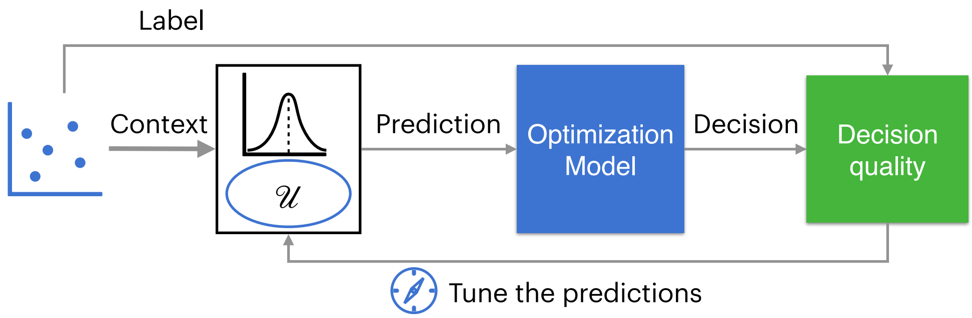

Figure 1: Integrated learning and optimization

mation of the decision-based loss function [8] or by randomly perturbing the expected cost predictions [4]. For quadratic program (or more generally strongly convex decision models), the prescribed decision is linked to the distribution model through the Karush–Kuhn–Tucker (KKT) conditions. The latter form equations that can be differentiated through using the implicit function theorem [1, 7]. Several other methods have been proposed to improve computability [Section 5 of 9] which has led to the development of numerous packages e.g., TorchOpt, JAXopt, Theseus, PyEPO and benchmarks for integrated learning and optimization.

For practical application of integrated learning and optimization, several theoretical and applied challenges must be overcome. One of the problems is the difficulty to recognize, within the limited real-world data, covariates that are strongly correlated with the uncertain parameters. Most current works rely on synthetic environments or on real-world environments (e.g., stock market data) but with a very limited number of covariates. To empirically validate the efficacy of these methods in large-scale systems, it is essential to develop surrogate models and loss functions that reduce the computational burden of training. Furthermore, identifying the appropriate hypothesis class of prediction models that will generalize to new data remains a complex challenge. On the theoretical side, there are a few results on the finite sample guarantees on their performance. To achieve robust and risk-aware prescriptions, contextual optimization methods will need to incorporate uncertainty in the constraints and its possible decision-dependence, two aspects that are overlooked by existing methods. In light of mandates such as the European Union's General Data Protection Regulation, which enforces the "right to explanation," one can expect that real-world implementation of contextual optimization methods will be subjected to fairness, privacy, explainability, and interpretabality requirements. To close, we wish to share our view that contextual optimization is a rich field of research that has the potential to transform the way we employ stochastic programming in a world where decisions need to be reactive and personalized, and are expected to make best use of a constantly growing collection of data. We encourage interested readers to pursue their investigation of contextual stochastic optimization by referring to our recently published survey [9].

References

[1] Brandon Amos and J Zico Kolter. OptNet: Differentiable optimization as a layer in neural networks. In *International Conference on Machine Learning*, volume 70, pages 136–145. PMLR, 2017.

[2] Gah-Yi Ban and Cynthia Rudin. The Big Data Newsvendor: Practical Insights from Machine Learning. *Operations Research*, 67(1):90–108, 2019.

[3] Y. Bengio. Using a financial training criterion rather than a prediction criterion. *International Journal of Neural Systems*, 8(4):433–443, 1997.

[4] Quentin Berthet, Mathieu Blondel, Olivier Teboul, Marco Cuturi, Jean-Philippe Vert, and Francis Bach. Learning with differentiable perturbed optimizers. In *Advances in Neural Information Processing Systems*, volume 33, pages 9508–9519. Curran Associates, Inc., 2020.

[5] Simone Buttler, Andreas Philippi, Nikolai Stein, and Richard Pibernik. A meta analysis of data-driven newsvendor approaches. In *ICLR 2022 Workshop on Setting up ML Evaluation Standards to Accelerate Progress*, 2022.

[6] Tsai-Hsuan Chung, Vahid Rostami, Hamsa Bastani, and Osbert Bastani. Decision-aware learning for optimizing health supply chains. *arXiv preprint arXiv:2211.08507*, 2022.

[7] Priya Donti, Brandon Amos, and J Zico Kolter. Task-based end-to-end model learning in stochastic optimization. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.

[8] Adam N. Elmachtoub and Paul Grigas. Smart "Predict, then Optimize". *Management Science*, 68(1):9–26, 2022.

[9] Utsav Sadana, Abhilash Chenreddy, Erick Delage, Alexandre Forel, Emma Frejinger, and Thibaut Vidal. A survey of contextual optimization methods for decision-making under uncertainty. *European Journal of Operational Research*, 2024.