Optimization Under Epistemic Uncertainty Using Prediction

Noah Schutte
Delft University of Technology
n.j.schutte@tudelft.nl

Abstract

Due to the complexity of randomness, optimization problems are often modeled to be deterministic to be solvable. Specifically epistemic uncertainty, i.e., uncertainty that is caused due to a lack of knowledge, is not easy to model, let alone easy to subsequently solve. Despite this, taking uncertainty into account is often required for optimization models to produce robust decisions that perform well in practice. We analyze effective existing frameworks, aiming to improve robustness without increasing complexity. Specifically we focus on robustness in decision-focused learning, which is a framework aimed at making context-based predictions for an optimization problem's uncertain parameters that minimize decision error.

1 Introduction

Optimization under uncertainty has always been a research area of interest, as real-life problems are almost never fully deterministic. To solve these problems, several types of modelling frameworks have been proposed.

One type is exact modelling frameworks, i.e., modelling frameworks that can guarantee a solution is optimal given certain assumptions. Stochastic and robust optimization are two effective and well-studied modelling frameworks [Bertsimas *et al.*, 2018], that model uncertain variables with some probability distribution, or some set of possible values. A more recent modelling framework and generalization of these two is distributionally robust optimization. Here, uncertain variables are modeled to have a set of potential distributions named an *ambiguity set* [Rahimian and Mehrotra, 2019] and the goal is to find the best expected objective given the worst-case distribution. This makes this framework suitable for the distinction between *epistemic* and *aleatoric* uncertainty, as aleatoric uncertainty is inherent to the uncertain process and captured by a distribution, while the epistemic uncertainty is defined as uncertainty due to lack of knowledge and modeled by the uncertainty around the distribution.

Additionally, search algorithms or non-exact modelling frameworks have also shown to be effective in practice, as they aim specifically to find good solutions in reasonable time. Specifically metaheuristics, i.e., general search methods that utilize heuristics have shown to be effective, even for optimization problems under uncertainty [Bianchi *et al.*, 2009; Juan *et al.*, 2023]. However, most research in this area is aimed at deterministic problems and not tailored towards efficiently solving stochastic problems, which is what our first contribution tackles.

Finally, data-driven approaches have become more relevant due to both better data-availability and the rapid improvements in machine learning. This caused more interest in *contextual optimization*, where contextual variables are correlated with the optimization problem's uncertain variables. This framework in which prediction can improve the decision result is also been called *predict-then-optimize*, where *smart predict-then-optimize* or *decision-focused learning* (DFL) is a framework where the predictive model is trained by minimizing decision error instead of maximizing predictive accuracy [Elmachtoub and Grigas, 2022]. In our second contribution, we were able to identify potential generalization issues with the commonly used training loss in DFL and proposed robust losses to mitigate this.

2 Main Research

In the introduction we present several modelling frameworks for optimization under uncertainty. The main research question we pose is if methods for optimization under uncertainty can be improved by getting more robust solutions without increasing computational complexity. This started with analysis on exact modelling frameworks like distributionally robust optimization, followed by analysis on non-exact modelling frameworks. This led to our first contribution at CPAIOR 2024 [Schutte *et al.*, 2024].

Improving Metaheuristic Efficiency for Stochastic Optimization by Sequential Predictive Sampling. Metaheuristics are known to be effective in finding good solutions in combinatorial optimization, but solving stochastic problems is costly due to the need for evaluation of multiple scenarios. We propose a general method to reduce the number of scenario evaluations per solution and thus improve metaheuristic efficiency. We use a sequential sampling procedure exploiting estimates of the solutions' expected objective values. These values are obtained with a predictive model, which is founded on an estimated discrete probability distribution linearly related to all solutions' objective distributions; the probability distribution is continuously refined based on incoming solution evaluation. The proposed method is tested using simulated annealing, but in general applicable to single solution metaheuristics. The method's performance is compared to descriptive sampling and an adaptation of a sequential sampling method assuming noisy evaluations. Experimental results on three problems indicate the proposed method is robust overall, and performs better on average than the baselines on two of the problems.

Furthermore, we investigated DFL. DFL has shown to be effective, but in cases with a lot of uncertainty around the relationship between features and outcome variables or when there is limited training data available, DFL trained predictors can run into issues. We identified this could be solved adjusting the commonly used loss function to a more robust version. This led to our second contribution at IJCAI 2024 [Schutte *et al.*, 2023].

Robust Losses for Decision-Focused Learning. Optimization models used to make discrete decisions often contain uncertain parameters that are context-dependent and estimated through prediction. To account for the quality of the decision made based on the prediction, decision-focused learning (end-to-end predict-then-optimize) aims at training the predictive model to minimize regret, i.e., the loss incurred by making a suboptimal decision. Despite the challenge of the gradient of this loss w.r.t. the predictive model parameters being zero almost everywhere for optimization problems with a linear objective, effective gradient-based learning approaches have been proposed to minimize the expected loss, using the empirical loss as a surrogate. However, empirical regret can be an ineffective surrogate because empirical optimal decisions can vary substantially from expected optimal decisions. To understand the impact of this deficiency, we evaluate the effect of aleatoric and epistemic uncertainty on the accuracy of empirical regret as a surrogate. Next, we propose three novel loss functions that approximate expected regret more robustly. Experimental results show that training two state-of-the-art decision-focused learning approaches using robust regret losses improves test–sample empirical regret in general while keeping computational time equivalent relative to the number of training epochs.

3 Future Research

The work on robust losses in DFL left us with many open research questions. One of them is on the use of point predictors in DFL. Even though we know we cannot predict the value of an uncertain variable with certainty, we do not try to estimate this uncertain variable with some distribution or other uncertainty-quantifying measure. Effectively we are trying to predict the expectation. Given that in a lot of problems in literature the uncertain variables are linear in the objective function, the expectation is the only measure that is relevant to us. However, we do know that problems under uncertainty become specifically challenging when we cannot simply change the order of the expectation and the objective function. We are currently investigating when point predictors are sub-optimal, and how this would affect DFL.

A second research direction we are currently pursuing is on generalization in DFL. Out-of-distribution generalization is an extensive field in machine learning, particularly in computer vision. Currently there are no works in DFL that try to understand generalization of DFL predictors. In some ways DFL predictors are trained to be specific to a certain optimization problem, however generalization to for example a slightly adjusted version of the problem is desirable. Does DFL still perform better than traditional predict-then-optimize in a setting like this? Can we make DFL predictors robust to generalize better?

Answering these questions can lead to multiple contributions, which could lead to the dissertation being more focused on DFL than the original idea of improving on multiple modelling frameworks.

Acknowledgements

I would like to thank my supervisors Krzysztof Postek and Neil Yorke-Smith for their continued support throughout this project. This work was partially supported by Epistemic AI and by TAILOR, both funded by EU Horizon 2020 under grants 964505 and 952215 respectively.

References

[Bertsimas *et al.*, 2018] Dimitris Bertsimas, Vishal Gupta, and Nathan Kallus. Data-driven robust optimization. *Mathematical Programming*, 167:235–292, 2 2018.

[Bianchi *et al.*, 2009] Leonora Bianchi, Marco Dorigo, Luca Maria Gambardella, and Walter J. Gutjahr. A survey on metaheuristics for stochastic combinatorial optimization. *Natural Computing*, 8:239–287, 2009.

[Elmachtoub and Grigas, 2022] Adam N. Elmachtoub and Paul Grigas. Smart "predict, then optimize". *Management Science*, 68:9–26, 2022.

[Juan *et al.*, 2023] Angel A. Juan, Peter Keenan, Rafael Martí, Seán McGarraghy, Javier Panadero, Paula Carroll, and Diego Oliva. A review of the role of heuristics in stochastic optimisation: from metaheuristics to learnheuristics. *Annals of Operations Research*, 320:831–861, 2023.

[Rahimian and Mehrotra, 2019] Hamed Rahimian and Sanjay Mehrotra. Distributionally robust optimization: A review. *arXiv preprint arXiv:1908.05659*, 2019.

[Schutte *et al.*, 2023] Noah Schutte, Krzysztof Postek, and Neil Yorke-Smith. Robust losses for decision-focused learning. *arXiv preprint arXiv:2310.04328*, 2023.

[Schutte *et al.*, 2024] Noah Schutte, Krzysztof Postek, and Neil Yorke-Smith. Improving metaheuristic efficiency for stochastic optimization by sequential predictive sampling. *Integration of Constraint Programming, Artificial Intelligence, and Operations Research*, pages 158–175, 2024.