

Towards Scalable Decision-Focused Learning for Combinatorial Optimization Problems

Dissertation submitted in fulfillment of the requirements for the award of the degree of

Doctor of Business Economics

by

Jayanta Mandi

Supervisors:

Prof. dr. Tias Guns

Prof. dr. Sam Verboven



Vrije Universiteit Brussel—April 2023

Abstract

The aspiration to use artificial intelligence (AI) in decision-making has been growing rapidly in recent years. In most real-world applications, the optimal decisions not only optimize certain objective functions, but they must also strictly satisfy some prespecified constraints. Purely data-driven machine learning (ML) models fail to provide any guarantee of satisfying the required constraints. An integrated framework combining ML and combinatorial optimization (CO) proves to be an effective approach to tackle such applications. In such an integration, the parameters (or a subset of the parameters) of the CO problem would be specific to each problem instance and unknown at the time of solving the CO problem. Hence, the role of the ML model is to predict these parameters from the auxiliary data, known as features. Examples include recommending a path between two locations minimizing travel time, one hour in advance. The travel time, a parameter of the CO problem, is unknown at the time of recommendation and must be predicted using an ML model.

The subject matter of this thesis is: **decision-focused learning** (DFL), where the ML model is trained to predict parameters of the CO that lead to good quality decisions. In DFL, the *decision-focused loss function* of the ML model is dependent on the output of the subsequent CO problem. So, the CO problem is in a sense embedded as a building block of the ML model. In this integration of

ABSTRACT

ML and CO, the main challenge is to formulate and solve the complex problem of empirical risk-minimization based on the output of the CO problem. In gradient-based training of neural networks, this involves **computing the derivative of the CO solution with respect to the predicted parameters**. Another major challenge in implementing DFL on real-life applications is that DFL trains slowly. This is because most CO problems are NP-hard and in DFL the **CO problem is solved for every training instance in every epoch**. This thesis develops frameworks to particularly address these two challenges.

Our first contribution is related to linear programs (LPs). We compute the derivative of the LP solution with respect to the parameter by differentiating the homogeneous self-dual embedding of the LP at the neighborhood of the optimal point. Then we focus on improving the computational cost of DFL. To this end, we investigate possible strategies to accelerate the ‘Smart Predict, then Optimize’ (SPO) framework. These include warmstarting from earlier solutions, and considering continuous relaxations in case of integer LPs. Next, we propose a new class of surrogate loss functions for DFL, where we consider non-optimal feasible solutions as negative samples and train the ML model to maximize the separation of objective values between the true optimal solution of the CO problem and these negative samples. We further extend this work by proposing a more generic surrogate loss function, where we formulate DFL as the task to learn the partial ordering of the feasible solutions as induced by the objective function of the CO problem. The computation and differentiation of these proposed surrogate loss functions require the presence of a pool of feasible points. We solve the CO problem only to form the pool of feasible points. But the **computation and back-propagation of the surrogate loss function are solver-free**. This results in a drastic decrease in training time and significantly less computational cost. We empirically validate that the proposed frameworks do *not* compromise the quality of the decisions. Furthermore, to address the computational cost of model training in the DFL framework, we propose to use the best performing feasible solution within the cache as a proxy of the true optimal solution. We validate this in DFL frameworks, which treat the CO solvers as blackbox oracles.

Overall, we explore several directions to address the computational cost of DFL. The surrogate loss functions proposed in the final two chapters demonstrate that formulation of solver-free surrogate loss functions is efficient to reduce the computational cost and training time of DFL.