Contribution We propose a novel differentiable LP layer, based on interior point solving of LPs. This allows us to perform end-to-end training of predict+optimize LP problems, more specifically in this paper the LP relaxation of (mixed) Integer Linear Programming problems. The key challenge is finding the gradients of the model parameters across the argmin of the LP problem. We propose two novelties, namely the use of the standard log barrier in LPs instead of adding a quadratic penalty and Quadratic Programming gradients, and to differentiate the Homogeneous Self-Dual formulation of the LP instead of the KKT conditions (of the QP). We show that this approach is able to do better than the QPTL approach, as well as the SPO sub-gradient approach.

Runtime Our implementation of the HSD, which is written in plain python, requires longer time for training. For example for a large Energy-cost aware scheduling problem, per epoch runtime of SPO, QPTL and our approach are approximately 7, 900 and 1600 seconds respectively. We want to point out that SPO uses an industry-grade (I)LP solver and the subgradient is straightforward to compute, and that QPTL uses carefully engineered matrix operations of OptNet. There is potential for improving matrix operations and reuse of decompositions in our approach too. An industry-grade interior point solver could also be used in our approach if it allows early stopping and exposes the full solving state.

MILP We want to reiterate that our proposed approach does solve an LP problem and finds gradient of the same. Our motivation and application domain however is predict+optimize of (M)ILP problems, where previous works (Wilder et al., 2019; Mandi et al., 2019) has shown that they can be effectively solved by considering the continuous relaxation of it. Tightening the LP relaxation before the predict+optimize is something that can benefit all approaches equally Ferber et al. (2020). This motivates us to use benchmarks of MILP problems, which are abound in real-life and have wide applications. We write MILP, as the proposed approach can equally deal with mixed and non-mixed integer problems, given that we start from the LP relaxation, but the experiments indeed involve only ILP problems.

Warmstart Warmstarting the IP is a research topic, but it is not a limiting factor in this research especially not as we use the HSD. We start from a fixed initial point and perform the predictor-corrector method of Mehrota (Mehrotra, 1992), a standard practice in interior point. We acknowledge the idea of accelerating the solving process by means of warm-staring across the mini-batches, as potential for future research.

Previous Work Vlastelica et al. (2019) proposed to derive the gradients for such problem by interpolating two locally constant regions in a piecewise-linear manner. The gradient they propose is of the form: $\left(x^* \left(\hat{c} + c_{perturbed}\right) - x^* \left(\hat{c}\right)\right)$; where $x^* \left(\hat{c}\right)$ is the solution for the predicted $\hat{c}$. For SPO, the proposed gradient is $\left(x^* \left(c\right) - x^* \left(2\hat{c} - c\right)\right)$. In our benchmarks we included SPO, we did not include this; but this work should be mentioned in related work.

Finally, the approach proposed by Agrawal et al. (2019) is similar to OptNet, which is used by the QPTL method, but generalized to Convex cone programs. They compute the gradients from the adjoint matrix and considering 0 in place of the dual differentials similar to QPTL. Their paper unfortunately only includes a runtime comparison to QPTL, so we don't know if it also improves quantitatively on QPTL. We should discuss it in our related work.

Fig1 Scatterplot Fig 1 aims to demonstrate how the MSE of our approach is high yet the regret is low. We want to visualise that because of the scale invariance property of the linear objective, IntOpt predictions are typically shifted by several magnitudes from the groundtruth. If looked at carefully it can be indeed be seen to predict extreme cases wrongly, which we can only assume will have little effect on the optimisation problem as attains lower regret.

Other Remarks

- t in Eq. 4 and Eq. 8 are same, in Eq. 4 it was introduced as a slack variable, whereas in Eq. 8 it is derived by the log-barrier formulation, suggesting why log-barrier formulation is more intrinsic to the LP formulation, instead of adding an external quadratically penalized term. We will make this explicit in the paper.
- The purpose of introducing the Tikhonov damping is to overcome the issue of the matrix not being positive definite for some instances during implementation. This is different from adding a quadratic term as it does not change the intrinsic objective function.
- In the real estate investment experiment, our approach is somewhat outdone by the twostage, because the prediction problem seems easier and the predictive model generates highly accurate predictions. In some other knapsack instances, our approach outperforms both the SPO and the two-stage. We would be glad to add them in the appendix.
- We did not have space for the learning curves and the ILP formulations; the learning curves, as one would expect, are the most natural ones; we can add them in the appendix for completeness.
- Finally, we are thankful to the reviewers for their comments and their interest.