Robust Losses for Decision-Focused Learning
Encore Abstract

Noah Schutte[1][0000−0001−8499−6578], Krzysztof Postek[2][0000−0002−4028−3725], and Neil Yorke-Smith[1][0000−0002−1814−3515]

[1] Delft University of Technology, The Netherlands
{n.j.schutte,n.yorke-smith}@tudelft.nl
[2] Independent Researcher krzysztof.postek@gmail.com

Abstract. Novel robust regret-based loss-functions are presented for decision-focused learning (end-to-end predict-then-optimize) to improve performance when training data is limited or there is a noisy relationship between the input-output variables of the predictive problem. This is an encore abstract of work with the same title published at IJCAI 2024 [6].

Keywords: decision-focused learning · robustness · predict and optimize

1 Motivation

Real-world optimization (decision) problems, such as shortest path or scheduling problems, are mostly non-deterministic. However, their uncertain parts are often correlated with contextual information. This makes it natural to use prediction. In a typical prediction (regression) problem, the goal is to maximize predictive accuracy. In a setting where the prediction is used to solve an uncertain optimization problem, performance should be evaluated on the decision quality. This is the main premise of *decision-focused learning* (DFL) [1, 4, 5].

More formally, the goal of the decision maker is to solve the following stochastic optimization problem with linear objective, observing some contextual information z as feature values: $\min_{x\in X} \mathbb{E}_{c\sim\mathcal{C}_z}[c^T x|z]$.

Since $\mathcal{C}_z$ is unknown, training a parametric predictor $\hat{c} := f_\theta(z)$ can assist in picking good decisions $x^*(\hat{c}) = \arg\min_{x\in X} \hat{c}^T x$. Instead of training the predictor by minimizing a predictive error loss, DFL aims at minimizing decision error. In this setting, *regret* [5] is natural and most used: $l_{\text{emp}}(\hat{c}, c) = c^T x^*(\hat{c}) - c^T x^*(c)$. While regret is determined based on empirically (emp) observed samples (z, c), given our problem definition we would like to find a good decision w.r.t. underlying distribution $\mathcal{C}_z$ for each z. This means that in l_{emp}, c is used as an estimator for $\mathbb{E}_{c\sim\mathcal{C}_z}[c|z]$ (using linearity). This is not different from a typical training pipeline, but due to the potential combinatorial nature of the optimization problem it is easier for the predictive model to become biased towards empirically observed optimal decisions. This leads to our main idea and contribution:

We introduce three regret-based loss-functions that have (1) a lower variance estimator, and/or (2) consider decisions that are robust against the estimator's estimation error as optimal.

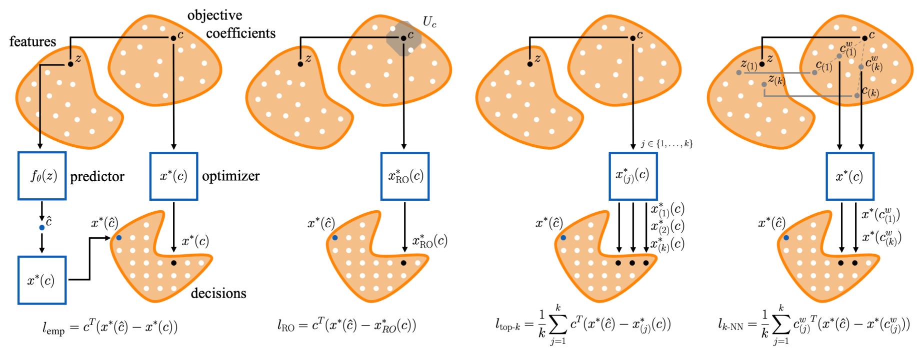

Fig. 1: Visualization of the (empirical) regret loss in DFL (left) and the robust losses in comparison (predictive pipeline is equal). From left to right: empirical regret (l_{emp}), RO loss (l_{RO}), top-k loss ($l_{\text{top-}k}$), k-NN loss ($l_{k\text{-NN}}$).

2 Contribution

Figure 1 shows the full formulations of the introduced losses. Based on the aforementioned idea, we introduce a k-NN based loss, where the k-NN estimator given feature values z is defined as the arithmetic mean of $c_{(j)}$, where $(z_{(j)}, c_{(j)}) = \arg\min_{(z',c')\in D\setminus\bigcup_{i=1}^{j-1}\{(z_{(i)},c_{(i)})\}} ||z' - z||$ is the j-th closest data point in the feature space, $||\cdot||$ is some norm (we use Euclidean) and D the set of (training) data points. Given an indefinitely increasing number of data points, the k-NN estimator converges to the true expectation [3] and a variance of $\text{Var}_{c\sim\mathcal{C}_z}(c)/k$. Since we know we have limited data, we adjust the values of the neighbours using some interpolation weight $w \in [0,1]$: $c^w_{(j)} = wc_{(j)} + (1-w)c$.

For the other two losses, no new estimator is used. Instead, the empirically optimal decision $x^*(c)$ is replaced by a decision that is robust against the estimation error of estimator c. The Robust Optimization (RO) loss considers a RO-formulation for optimal decisions: $x^*_{\text{RO}}(c) = \min_{x\in X}\max_{c\in U_c} c^T x$. The top-$k$ loss considers the best k decisions as (equally) optimal: $x^*_{(j)}(c) := \arg\min_{x\in X\setminus\{x^*_{(1)}(c),\ldots,x^*_{(j-1)}(c)\}}\{c^T x\}$.

3 Results

We perform experiments on shortest path, travelling salesperson, and energy-cost aware scheduling problems. The introduced losses are compared with empirical regret based on two state-of-the-art DFL gradient-approximation approaches, Smart "Predict, then Optimize" [5] and the Perturbed Fenchel–Young Loss [2]. In these experiments we use different training data sizes and relationships with varying noise between the variables of the regression problem. We see that especially when there is little training data and/or a more noisy relationship, the robust losses perform significantly better. Especially the k-NN loss is consistently better across different problems. Complete results are in the published paper [6].

References

1. Bengio, Y.: Using a financial training criterion rather than a prediction criterion. International Journal of Neural Systems **8**, 433–443 (1997)
2. Berthet, Q., Blondel, M., Teboul, O., Cuturi, M., Vert, J.P., Bach, F.: Learning with differentiable perturbed optimizers. Advances in Neural Information Processing Systems **33**, 9508–9519 (2020)
3. Cheng, P.E.: Strong consistency of nearest neighbor regression function estimators. Journal of Multivariate Analysis **15**, 63–72 (1984)
4. Donti, P., Amos, B., Kolter, J.Z.: Task-based end-to-end model learning in stochastic optimization. Advances in Neural Information Processing Systems **30** (2017)
5. Elmachtoub, A.N., Grigas, P.: Smart "predict, then optimize". Management Science **68**, 9–26 (2022)
6. Schutte, N., Postek, K., Yorke-Smith, N.: Robust losses for decision-focused learning. In: Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI-24). pp. 4868–4875 (2024)