

DFF: Decision-Focused Fine-Tuning for Smarter Predict-Then-Optimize with Limited Data

Jiaqi Yang^{1*}, Enming Liang^{2*}, Zicheng Su^{1†}, Zhichao Zou³, Peng Zhen³, Jiecheng Guo³,
Wanjiang Ma¹, Kun An^{1†}

¹The Key Laboratory of Road and Traffic Engineering of the Ministry of Education, Tongji University, China

²Department of Data Science, City University of Hong Kong

³Didi Chuxing, Beijing, China

2410196@tongji.edu.cn, eliang4-c@my.cityu.edu.hk, suzicheng@tongji.edu.cn, zouzhichao@didiglobal.com, zhenpeng@didiglobal.com, jasonguo@didiglobal.com, mawanjiang@tongji.edu.cn, kunan@tongji.edu.cn

Abstract

Decision-focused learning (DFL) offers an end-to-end approach to the predict-then-optimize (PO) framework by training predictive models directly on decision loss (DL), enhancing decision-making performance within PO contexts. However, the implementation of DFL poses distinct challenges. Primarily, DL can result in deviation from the physical significance of the predictions under limited data. Additionally, some predictive models are non-differentiable or black-box, which cannot be adjusted using gradient-based methods.

To tackle the above challenges, we propose a novel framework, Decision-Focused Fine-tuning (DFF), which embeds the DFL module into the PO pipeline via a novel bias correction module. DFF is formulated as a constrained optimization problem that maintains the proximity of the DL-enhanced model to the original predictive model within a defined trust region. We theoretically prove that DFF strictly confines prediction bias within a predetermined upper bound, even with limited datasets, thereby substantially reducing prediction shifts caused by DL under limited data. Furthermore, the bias correction module can be integrated into diverse predictive models, enhancing adaptability to a broad range of PO tasks. Extensive evaluations on synthetic and real-world datasets, including network flow, portfolio optimization, and resource allocation problems with different predictive models, demonstrate that DFF not only improves decision performance but also adheres to fine-tuning constraints, showcasing robust adaptability across various scenarios.

Introduction

Predict-then-Optimize (PO) is a framework that uses machine learning to address decision problems under uncertainty, as the parameters of the optimization problem are likely unknown before making the decision (Bertsimas and Kallus 2020). This framework operates in two stages: In the first stage, auxiliary features are utilized to predict the unknown parameters; in the second stage, decisions are made based on these predictions. A critical limitation of this two-stage framework is that during the training process of the

first stage, the generally used loss functions such as MSE aim to minimize fitting error. However, this may not align with the objective of the final decision task. The new framework proposed to address this issue is Decision-focused Learning (DFL), which has been demonstrated to improve the quality of the final decision by customizing the training of the predictive model based on the decision task (Kotary et al. 2021; Mandi et al. 2023; Sadana et al. 2024). However, directly replacing the two-stage PO framework with DFL still poses significant challenges.

The first challenge is the convergence issues associated with training from scratch based on decision loss¹ (DL). This arises from the highly non-convex nature of DL (e.g., regret), combined with potential discontinuities, leading to substantial computational costs in calculating the gradient of the downstream problem (Elmachtoub and Grigas 2022; Tang and Khalil 2022). A promising solution for this problem is constructing a surrogate loss for DFL. In existing works, the design of the surrogate losses is based on Fisher consistency, i.e., the surrogate losses are equivalent with DL when the data is infinite (Elmachtoub and Grigas 2022; Shah et al. 2024). However, in real-world scenarios with limited data, DFL may fail to converge or converge to a sub-optimum.

The second challenge is the significant deviation in predictions induced by directly minimizing DL through training. As the accuracy of the predictions becomes less controllable, this can lead to a loss of the inherent physical meaning present in the two-stage PO framework. For example, suppose we use a simple linear model with coefficient matrix β to fit the dataset $\{\mathbf{X}, \mathbf{y}\}$, where $\mathbf{X} \in \mathbb{R}^{N \times p}$ and $\mathbf{y} \in \mathbb{R}^{N \times d}$. First, the PO framework adopts the MSE loss, leading to a closed-form coefficient matrix:

$$\beta_{\text{PO}} = (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{y} \quad (1)$$

As for the DFL, the predictive model is trained with DL, which is implicitly defined by the decision problem. However, it has been shown that there exists an input-dependent

*These authors contributed equally.

†Corresponding authors.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹In this paper, unless otherwise specified, decision loss refers to loss functions used to measure decision quality, such as regret, surrogate loss, approximated loss, etc.




Figure 1: The distribution of predictions generated by different methods for the shortest path problem (Elmachtoub and Grigas 2022) is shown in the figure. It includes the two-stage PO (MSE) and three representative DFL methods, including SPO+ (Elmachtoub and Grigas 2022), PFYL (Berthet et al. 2020), and LSLTR (Mandi et al. 2022). The results indicate that the predictions from the MSE-based method maintain a distribution similar to the ground truth, whereas the predictions from the DFL models exhibit significant deviations.

positive semi-definite matrix \mathbf{Q}_i for $i = 1, \dots, N$, leading to a quadratic loss $L = \frac{1}{N} \sum_{i=1}^N (\mathbf{y}_i - \hat{\mathbf{y}}_i)^T \mathbf{Q}_i (\mathbf{y}_i - \hat{\mathbf{y}}_i)$, which is Fisher consistent with DL (Shah et al. 2024). Therefore, the coefficient matrix under the surrogate loss is derived from such a generalized least squares problem:

$$\beta_{\text{DL}} = \left(\sum_{i=1}^N \mathbf{x}_i \mathbf{Q}_i \mathbf{x}_i^T \right)^{-1} \sum_{i=1}^N \mathbf{x}_i \mathbf{Q}_i \mathbf{y}_i \quad (2)$$

Since \mathbf{Q}_i is not an identity matrix, DFL with limited data can introduce biases originating from the downstream decision-making objective. In other words, the predictions obtained by DFL focus on decision quality, which can result in undesired phenomena, such as multiplicative shifts (Tang and Khalil 2022) as shown in Figure 1 and fail to capture the underlying physical meanings of predictions.

The third challenge is that some prediction models are non-differentiable and may not be directly combined with DFL. Current training schemes for DFL primarily focus on first-order gradient-based methods. However, some popular prediction models, such as tree-based models, are not differentiable (Breiman 2001). Moreover, in cases where the confidence in the predictive model is low, white-box or semi-black-box simulation models developed based on principles of the physical world can be used for prediction and offline evaluation (Wu et al. 2022). These models can explicitly provide state transitions between variables, offering robust estimates in the presence of many unobserved confounding variables. Nonetheless, it remains unclear how to effectively apply DFL in non-differentiable backbone scenarios.

To address the aforementioned challenges and make DFL more effective, we propose Decision-Focused Fine-Tuning (DFF), which consists of a given backbone predictive model

and an additional bias correction layer, as shown in Figure 2. The contributions of this paper are as follows:

- By utilizing the bias correction layer, DFF can fine-tune the output of any backbone model to produce predictions that are aligned with the decision-making objective.
- By explicitly constraining the correction layer, DFF can prevent the fine-tuned output from shifting the inherent physical meaning of the backbone model’s predictions.
- We theoretically analyze the performance of DFF on the adherence to fine-tuning constraints and strictly limit the bias in the predictions.
- We conduct extensive simulations, including benchmark network flow and portfolio optimization problems and two real-world ride-sourcing subsidy allocation problems. The results show that DFF consistently achieves better decision quality compared to backbone models.

Related Works

Decision Focused Learning

Since the introduction of OptNet by Amos and Kolter, which integrates optimization problems as individual layers within end-to-end trainable deep networks, many works have explored how to better integrate generic optimization problems with machine learning. For convex optimization problems, the gradient can be accurately computed using implicit differentiation methods to guide training (Donti, Amos, and Kolter 2017; Wilder, Dilikina, and Tambe 2019). However, discrete optimization problems and linear optimization problems often lack gradient information that can guide training, necessitating the development of techniques to obtain approximate gradients (Mandi et al. 2023). These techniques include relaxing and smoothing the optimization problem (Wilder, Dilikina, and Tambe 2019), approximating gradients using random perturbations (Berthet et al. 2020; Dalle et al. 2022), and designing surrogate losses (Mandi and Guns 2020). The SPO+ loss function proposed by Elmachtoub, Liang, and McNellis (2020) is the first surrogate loss with theoretical guarantees, suitable for optimization problems with any linear objective function. It has been proven to be Fisher consistent with DL. Subsequent research has further explored designing surrogate loss functions for general optimization problems, leading to methods such as LODL (Shah et al. 2022), EGL (Shah et al. 2024), LANCER (Zharmagambetov et al. 2024), and TaskMet (Bansal et al. 2024). As for constrained optimization, several studies propose effective methods for addressing scenarios in which uncertain parameters are incorporated into the constraints of optimization problems (Hu, Lee, and Lee 2023, 2024). Furthermore, thanks to packages like `CvxpyLayers` (Agrawal et al. 2019) and `PyEPO` (Tang and Khalil 2022), deploying these complex loss functions now comes with lower engineering costs. Therefore, DFL has the potential to become a modular tool that allows predictive models to be enhanced for specific optimization problems.

DFL with Special Prediction Models

However, DFL faces challenges under some special upstream predictive models, e.g., tree-based model (Elmach-

toub, Liang, and McNellis 2020), semi-parametric (Wu et al. 2022; Zhao et al. 2019), and simulation-based physical model (She et al. 2024). These models are favored under specific settings; for example, tree-based models are preferred in the industry due to their high accuracy and strong interpretability. However, these models are non-differentiable, making it difficult for the gradient of the DL to be backpropagated efficiently.

Existing works have combined tree-based models directly with DFL. Elmachtoub, Liang, and McNellis (2020) propose SPO Trees and SPO Forests by modifying the splitting criteria of decision trees, which shows smaller decision regret in linear optimization problems compared to the decision tree and random forest based on MSE. Butler and Kwon (2023) combine gradient boosting with DFL and propose Dboost, which performs better in convex cone optimization problems. However, numerical experiments by Butler and Kwon (2023) reveal that the decision quality of Dboost and SPO forest is inferior to gradient boosting based on MSE in certain scenarios, which limits their practicality. To date, there is a lack of an efficient framework to incorporate DFL into general predictive models.

Fine-tuning Approach

Fine-tuning is a crucial technique in deep learning, which involves taking a pre-trained model and adjusting its parameters to better fit a particular task. This approach allows practitioners to leverage the knowledge embedded in the pre-trained model, significantly speeding up the training process and improving performance on the new task (Zhang et al. 2023; Ding et al. 2023; Fu et al. 2023; Fan et al. 2024).

Under the context of DFL, fine-tuning refers to further adjusting the model’s predictions to reduce the loss associated with the DL. Since the training process of DFL typically requires repeatedly solving the optimization problem, it can reduce the overall training time and improve convergence efficiency. For example, a checkpoint with good prediction performance is obtained by training on MSE, and then further fine-tuning is conducted with DL until convergence (Mandi et al. 2020; Kotary et al. 2022).

Beichter et al. (2024) propose a retraining method that fine-tunes a pre-trained predictive model with weighted MSE and DL and applies it to the dispatchable feeder optimization problem, which resulted in significant performance improvements.

Nevertheless, existing fine-tuning/retraining can not be applied to general prediction models and there is no performance guarantee of the tuned model when considering the shift under limited data.

Problem Definition

Predict-then-Optimize

In the two-stage predict-then-optimize framework, we first train a predictive model M using the MSE loss, with input features x that produce predictions $\hat{c} = M(x)$. These predictions serve as the input parameters for the second-stage optimization problem and determine the optimal decision $w^*(\hat{c})$:

$$w^*(\hat{c}) = \arg \min_w f(w, \hat{c}) \quad (3)$$

$$\text{s.t. } g_j(w) \leq 0, \quad \text{for } j \in \{1, 2, \dots, J\}. \quad (4)$$

where $f(\cdot)$ represents the objective function of the downstream optimization task, and $g_j(w)$ denotes the constraints on the decision variable w . Here, $\hat{c} \in \mathbb{R}^d$ represents the unknown parameters, with dimension d typically greater than 1. For each decision problem, our goal is to minimize the decision regret, introduced by Elmachtoub and Grigas (2022):

$$DR(c, \hat{c}) = f(w^*(\hat{c}), c) - f(w^*(c), c) \quad (5)$$

Given a dataset with N samples, denoted as $\mathcal{D} = \{(x_1, c_1), (x_2, c_2), \dots, (x_N, c_N)\}$, the predictive model M is trained to minimize the average decision regret as follows:

$$\overline{DR} = \frac{1}{N} \sum_{i=1}^N DR(c_i, M(x_i)) \quad (6)$$

Open Issues for DFL

As shown in previous research, the DFL method tends to outperform two-stage PO when the predictive model is misspecified or when there are limitations to further improving prediction accuracy (Hu, Kallus, and Mao 2022; Elmachtoub et al. 2023). However, the application of DFL in such scenarios is also hindered by several limitations.

Biased Prediction under Limited Data DFL tends to deliver biased predictions, especially under limited data conditions. The reasons are as follows. Firstly, as we mentioned earlier, models trained using DL inherently introduce bias to better align with the decision-making objective. Secondly, existing surrogate loss functions may not satisfy Fisher’s consistency under limited data, leading to potential bias compared to DL (Elmachtoub and Grigas 2022). Lastly, the multiplicative shifts induced by DL can cause predictions to lose their inherent physical meaning (Tang and Khalil 2022), which is critical in helping analyze downstream decision tasks.

Non-differentiable Predictive Model There are scenarios where designing a suitable predictive model and training it to achieve sufficient accuracy can be challenging. In such cases, an alternative is to utilize simulation-based models. These models can incorporate explicit state transitions based on principles of the physical world and merge prior knowledge of the task. However, it is unclear how to directly apply DFL to these non-differentiable models, and replacing simulation models with DFL-based predictive models is unrealistic. Moreover, since these simulation-based models have been validated as reliable in numerous cases, DFL is expected to enhance decision-making by building on the simulation results rather than simply overturning them. This motivates us to explore how to leverage DFL in these problems while maintaining the existing solution framework and making only minor adjustments to the outputs.




Figure 2: Illustration of the Decision-Focused Fine-tuning Framework.

Decision-Focused Fine-tuning

In this section, we introduce the proposed framework, Decision-Focused Fine-Tuning (DFF), as illustrated in Figure 2. We consider a scenario where an upstream predictive model M has already been established, and its predictions are denoted as \hat{c} . We then design a neural network F_θ with a constrained output range, by which \hat{c} is explicitly represented through a residual connection. All parameters θ of F_θ are optimized to perform a linear transformation on \hat{c} that minimizes the DL, with the output denoted as \tilde{c} .

Constrained Fine-tuning

To preserve the advantages of the two-stage framework when applying DFL, our objective is to tune the existing predictive model M to achieve better decision quality, while ensuring that the adjusted predictions \tilde{c} remain within a reasonable range from the predictions of the original model, i.e. \hat{c} . This can be formulated as a constrained optimization (CO) problem as follows:

$$\min_{F_\theta} \mathbb{E}_{x,c} [DR(c, F_\theta(x))] \quad (7)$$

$$\text{s.t. } \mathbb{D}(F_\theta(x), M(x)) \leq \epsilon \quad (8)$$

where \mathbb{D} is a distance metric, e.g., Euclidean distance, in a specific scenario, and ϵ defines the acceptable range for adjusting the upstream predictive model. This formulation can be interpreted as improving the decision performance by fine-tuning the upstream predictive model M via F_θ within the trust region, particularly under limited data.

We make the following remarks on the effects of the parameter ϵ :

- When $\epsilon = 0$, no adjustments are made to the upstream predictive model, leading to $F_\theta = M$.
- When $\epsilon \rightarrow \infty$, it indicates that F_θ is trained entirely based on the DL without any restrictions, leading to $F_\theta = \arg \min_{F_\theta} \mathbb{E}_{x,c} [DR(c, \tilde{c})]$.

- When $\epsilon \in (0, \infty)$, F_θ is trained using DL within the trust region of the trained predictive model. The trade-off between these two losses can be problem-specific and defined by the user through the parameter ϵ .

We can note that the CO formulation bridges the naive weighted sum of MSE and DL by applying the distance metric \mathbb{D} as a constraint. A similar approach can be found in Beichter et al. (2024), which incorporates the MSE loss into the objective function (7) as a Lagrangian relaxation. This can be considered as a special case of our CO formulation. The superiority of the CO formulation lies in its applicability to various types of distance metrics \mathbb{D} , allowing for a more efficient adjustment of the trade-off through the parameter ϵ .

Such a trust-region optimization formulation has been successfully applied in various areas such as reinforcement learning (Schulman et al. 2015), online learning (Wu et al. 2017), and fine-tuning (Kurutach et al. 2018). It ensures that the fine-tuned results do not deviate significantly from the upstream predictions, enabling the model to effectively handle decision tasks where data is limited and stability is critical (Queeney, Paschalidis, and Cassandras 2021), while further enhancing decision quality. However, the challenges of constrained fine-tuning are twofold: some predictive models are non-differentiable, and enforcing the constraint can be difficult (Mandi et al. 2020; Kotary et al. 2022; Beichter et al. 2024). We address these challenges by proposing a bias correction layer, as introduced in the following section.

Bias Correction Module

To tackle the above challenges, we propose a bias correction module F_θ , motivated by the hyper-network design (Ha, Dai, and Le 2016; She et al. 2024). We set F_θ in the following semi-parametric form as:

$$F_\theta(x) = \phi(x) \odot M(x) + b(x) \quad (9)$$

where \odot is the Hadamard product. $\phi(x)$ and $b(x)$ represent the input-dependent weights and bias, respectively. Together, they fine-tune the predictive model M to reach

an improved decision performance. Notably, F_θ does not directly tune the parameters of the predictive model M ; instead, it learns an input-dependent linear transformation. This approach allows it to apply the DL to general non-differentiable predictive models. Moreover, the distance constraint outlined in Equation (8) can be easily enforced by designing the weight and bias modules accordingly.

Specifically, we define the distance metric in the trust-region constraint (8) as the percentage error:

$$\mathbb{D}_i(\tilde{c}_i, \hat{c}_i) = \left| \frac{\tilde{c}_i - \hat{c}_i}{\hat{c}_i} \right| \leq \epsilon, \quad i \in \{1, 2, \dots, d\} \quad (10)$$

Here, ϵ represents a pointwise adjustable percentage. The distance measurement of percentage error is advantageous as it normalizes the predictions into unitless values, thereby avoiding the additional costs of parameter adjustment that may arise due to the differences in the magnitudes of the predictions.

Given the constraint shown in Equation (10), we set $\mathbf{b}(\mathbf{x}) = 0$ and apply an offset-scaled Sigmoid transformation to ensure that the output of $F_\theta(\mathbf{x})$ strictly falls within the specified range and exhibits center-symmetry, i.e. $\frac{\tilde{c}_i}{\hat{c}_i} \in [1 - \epsilon, 1 + \epsilon]$. The outputs of $\phi(\mathbf{x})$ can be represented as:

$$\phi(\mathbf{x}) = [(1 - \epsilon) + 2\epsilon \cdot \sigma(\mathbf{h}(\mathbf{x}))] \quad (11)$$

Here, $\mathbf{h}(\mathbf{x})$ denotes the output from the penultimate layer, and $\sigma(\mathbf{x}) = \frac{1}{1+e^{-x}}$ denotes the standard Sigmoid function.

It is important to clarify that the percentage error shown in Equation (8) is not the only instance of the distance metric, but rather an intuitive and effective example. Further discussion on different types of distance metrics can be found in Appendix A.

We now present the upper bound on the increment of the RMSE induced by the bias correction layer, and the maximum angular gap between the original predictions \hat{c} and the corrected predictions \tilde{c} .

Theorem 1. *For a given ϵ , the increment in RMSE of \tilde{c} obtained in Equation (11) in the main text has an upper bound such that:*

$$RMSE(\tilde{c}, c) - RMSE(\hat{c}, c) \leq \frac{\epsilon}{\sqrt{d}} \|\hat{c}\|_2 \quad (12)$$

where $\|\cdot\|_2$ denotes the L_2 norm and d represents the dimension of the vector \hat{c} .

Moreover, the cosine similarity between \tilde{c} and \hat{c} can be calculated as:

$$\cos\langle \tilde{c}, \hat{c} \rangle = \frac{\tilde{c} \cdot \hat{c}}{\|\tilde{c}\|_2 \|\hat{c}\|_2} \quad (13)$$

where $\cos\langle \cdot, \cdot \rangle \in [-1, 1]$ denotes the similarity of directions between two vectors, with a value of 1 indicating the same direction. The lower bound of $\cos\langle \tilde{c}, \hat{c} \rangle$ is as follows:

$$\cos\langle \tilde{c}, \hat{c} \rangle \geq \sqrt{1 - \epsilon^2} \quad (14)$$

Proof. See Appendix A.

Theorem 1 demonstrates that the point-wise distance metric defined in Equation (10) can impose a strict bound on the vector-wise distance metrics. These bounds on auxiliary metrics enhance the reliability of predictions derived from DFF.

Loss Function and Training

After determining the constraint on F_θ , the training process is transformed into a standard DFL training process, which takes the following loss function:

$$L = \frac{1}{N} \sum_{i=1}^N DR(c_i, F_\theta(x_i)) \quad (15)$$

We need to solve the following core gradient optimization problem:

$$\frac{\partial L}{\partial \theta} = \frac{\partial L}{\partial \mathbf{w}^*} \frac{\partial \mathbf{w}^*}{\partial \mathbf{c}} \frac{\partial \mathbf{c}}{\partial \theta} \quad (16)$$

where the first two terms involve the gradient of the optimization problem, and the third term is the gradient optimization of the predictive model.

It is worth noting that our framework is general; we only need to select an appropriate gradient calculation method based on the characteristics of the optimization task to compute $\frac{\partial L}{\partial \theta}$. In this paper, we employ a direct approximation method for calculating $\frac{\partial L}{\partial \mathbf{c}}$ using the SPO+ method because of its superior performance in linear objective function scenarios (Elmachtoub and Grigas 2022). This method constructs a convex upper bound surrogate loss function $L = DR(\mathbf{c}, \tilde{\mathbf{c}})$:

$$L_{SPO+} = \min_{\mathbf{w}} \{(2\tilde{\mathbf{c}} - \mathbf{c})^T \mathbf{w}\} + 2\tilde{\mathbf{c}}^T \mathbf{w}^*(\mathbf{c}) - f^*(\mathbf{c}) \quad (17)$$

For this loss function, the gradient can be computed using the following formula:

$$\frac{\partial L}{\partial \mathbf{c}} \approx \frac{\partial L_{SPO+}}{\partial \mathbf{c}} = \mathbf{w}^*(2\tilde{\mathbf{c}} - \mathbf{c}) - \mathbf{w}^*(\mathbf{c}) \quad (18)$$

Case Study

We evaluate the effectiveness of the proposed DFF on three two-staged PO problems using various datasets and different predictive models. Firstly, the DFF is tested on two well-established benchmarks using synthetic data. We then validate the DFF on the resource allocation problem with real-world data from the ride-hailing platform DiDi Chuxing. Lastly, we employ the DFF to adjust the predictions from a non-differentiable simulation model. Note that all experiments are run 10 times with different random seeds and the average results are reported.

Benchmarks with Synthetic Data

Problem Settings and Dataset In line with Elmachtoub and Grigas (2022) and Butler and Kwon (2023), the proposed DFF framework is first evaluated on two well-established optimization problems: the network flow problem and the portfolio optimization problem. The formulations for these benchmarks are illustrated in Appendix B. Specifically, both problems aim to minimize a linear objective function, while the portfolio optimization problem involves a nonlinear constraint. To conduct a fair comparison among different predictive models, we generate two distinct datasets using different mechanisms, including polynomial functions, periodic functions, and piecewise functions with different coefficients. The details of the synthetic data are presented in Appendix B.

Algorithm	Network Flow Problem				Portfolio Optimization Problem			
	Dataset 1		Dataset 2		Dataset 1		Dataset 2	
	NDR	MSE	NDR	MSE	NDR	MSE	NDR	MSE
Random Forest	6.24%	4.43×10^{-1}	8.70%	5.75×10^{-1}	23.21%	4.60×10^{-1}	35.69%	6.92×10^{-1}
SPO Forest	5.99% (✓)	5.19×10^{-1}	8.49% (✓)	6.44×10^{-1}	22.50% (✓)	4.76×10^{-1}	34.50% (✓)	7.03×10^{-1}
Boost	5.26%	3.69×10^{-1}	5.75%	3.85×10^{-1}	16.44%	3.26×10^{-1}	16.38%	3.47×10^{-1}
Dboost	5.86%	5.78×10^{-1}	8.32%	6.83×10^{-1}	16.56%	3.88×10^{-1}	17.87%	4.78×10^{-1}
NN-MSE	5.06%	3.37×10^{-1}	9.44%	6.08×10^{-1}	14.46%	2.87×10^{-1}	21.98%	4.49×10^{-1}
NN-SPO+	4.63% (✓)	1.92	9.43% (✓)	1.80	14.74%	4.98×10^{-1}	23.55%	7.08×10^{-1}
2-fold Boost	4.49%	3.03×10^{-1}	4.88%	3.23×10^{-1}	16.08%	3.18×10^{-1}	15.92%	3.39×10^{-1}
DFF (ours)	4.39% (✓)	3.11×10^{-1}	4.81% (✓)	3.26×10^{-1}	15.97% (✓)	3.36×10^{-1}	15.91% (✓)	3.41×10^{-1}

Table 1: NDR and MSE of benchmarking methods on two well-established PO problems

Baseline In this paper, we compare the DFF framework with three state-of-the-art DFL methods, as well as their original predictive models trained with the MSE loss. The following are brief introductions to the baseline models:

1. **Random Forest**: Random forest trained with MSE.
2. **SPO Forest** (Elmachtoub, Liang, and McNellis 2020): Random forest trained with the decision regret shown in Equation (5) using zero-order gradient.
3. **Boost**: Gradient boosting tree trained with MSE.
4. **Dboost** (Butler and Kwon 2023): Gradient boosting tree trained with the decision regret in Equation (5) using the fixed-point argmin differentiation.
5. **NN-MSE**: Neural network trained with MSE.
6. **NN-SPO+** (Elmachtoub and Grigas 2022): Neural network trained with the SPO+ loss shown in Equation (17).
7. **2-fold Boost**: Gradient boosting tree based on MSE with 2-fold cross-validation.
8. **DFF (ours)**: Fine-tuning the **2-fold Boost** with the SPO+ loss shown in Equation (17).

The baselines are evaluated with the normalized decision regret (NDR) (Tang and Khalil 2022), defined as follows:

$$NDR = \frac{\sum_{i=1}^N f(\mathbf{w}^*(\hat{\mathbf{c}}_i), \mathbf{c}_i) - f(\mathbf{w}^*(\mathbf{c}_i), \mathbf{c}_i)}{\sum_{i=1}^N |f(\mathbf{w}^*(\mathbf{c}_i), \mathbf{c}_i)|} \quad (19)$$

Hyperparameters and Training Process In this paper, all tree-based models are trained with a maximum depth of 2 and no more than 100 trees. In particular, the Random Forest model samples the training data with a 50% sampling rate. As for the neural network, it consists of 3 layers with 32 neurons on each layer, and the ReLU activation function is used. The parameter of constrained distance ϵ is set to 0.5 for the DFF. Since the Boost model produces the lowest MSE error among the prediction models, we use it as the backbone model to fine-tune its predictions with the SPO+ loss, as shown in Equation (17). To avoid self-fitting and make full use of data, we split the training set into two disjoint datasets, which is in line with the 2-fold cross-fitting method proposed by Chernozhukov et al. (2018). In Appendix C, we demonstrate that the proposed DFF module also enhances the decision performance of other predictive models.

Results The normalized decision regret and MSE of different methods are summarized in Table reftab:case1. The checkmark ‘✓’ represents that the predictive model trained using the DL can improve the NDR of its counterpart trained with MSE. Firstly, we can note that the proposed DFF can improve its base model in both problems across all the datasets. This demonstrates that DFF can effectively enhance the downstream optimization task by adjusting the backbone predictive model within the trust region, even in the presence of nonlinear constraints. Moreover, the DFF method achieves the lowest normalized regrets in all settings, except for the portfolio optimization problem with Datasets 1. This is because the base model, 2-fold Boost, is less efficient than NN-MSE in this scenario, and the fine-tuning based on DFF is constrained by the trust region.

Resource Allocation Problem with Real-world Data

We further validate the DFF framework with the resource allocation problem for the ride-hailing platform DiDi Chuxing. The key task is to allocate the subsidy budget between different cities based on the predicted subsidy conversion rate, thereby maximizing the platform’s revenue. The detailed formulation is presented in Appendix D.

In practice, the platform employs Extreme Gradient Boosting (XGBoost) as the predictive model, as it offers sufficient predictions, high robustness, and ease of implementation. However, it is challenging to apply DFL to XGBoost for two reasons. On the one hand, it requires calculating the second-order derivative of the loss function to train XGBoost, for which there is currently no corresponding method in existing DFL frameworks. On the other hand, there is limited data available for training such a model, as only the historical market data within three months can be used to make a timely prediction.

We fine-tune the XGBoost model using DFF with market data from 102 consecutive days from Didi Chuxing and subsequently allocate the subsidy budget across 105 cities. The results are summarized in the first part of Table 2. The proposed DFF framework outperforms the original prediction backbone XGBoost and the NN-SPO+ method in terms of normalized regret. Notably, the MSE of the DFF framework is significantly lower than that of the NN-SPO+ method, comparable to the XGBoost model which is specifically

Method	NDR	MSE
Fine-tuning the XGBoost model		
XGBoost	2.54%	0.45
NN-SPO+	2.40%	2.04
DFF(ours)	2.39%	0.49
Fine-tuning the simulation model		
Average allocation	3.45%	/
Simulation model	2.79%	2.59×10^{-2}
DFF(ours)	2.11%	2.61×10^{-2}

Table 2: Normalized regret and MSE across different methods in resource allocation problem.




Figure 3: Distribution of predictions with different methods

trained to minimize the MSE loss. Moreover, Figure 3 depicts the distribution of predicted values from different models, which represents the subsidy conversion rate that has a crucial physical meaning. Notably, the ground truth exhibits a significant pattern of bimodal distribution. However, the predicted values from the NN-SPO+ method show a considerable deviation, presenting a unimodal distribution. This is consistent with the findings in Tang and Khalil (2022) that direct training with DL can induce a multiplicative shift in the predicted values. In contrast, the predictions from the DFF method closely align with the ground truth in both magnitude and distribution shape.

Fine-tuning Non-differentiable Simulation Model

In the context of the resource allocation problem, the ride-hailing platform DiDi Chuxing has developed a simulation model to derive the subsidy conversion rate, enabling the analysis of domain knowledge in the ride-hailing market (Gao et al. 2024). However, the downstream DL cannot be directly backpropagated to adjust the simulation model due to its non-differentiable nature. To address this, we apply the proposed DFF module to fine-tune the simulation model based on the DL. Due to the absence of DFL methods for adjusting the simulation model, we incorporate a rule-based average allocation strategy as a baseline.

The summarized results are presented in the lower part of Table 2. Notably, by fine-tuning the simulation model with the DFF module, the normalized decision regret is reduced




Figure 4: Sensitivity analysis of parameter ϵ on the decision loss and MSE loss

by 24.37% compared to the original simulation model, while the MSE loss remains at a similar level. This again demonstrates the effectiveness of the constrained fine-tuning design. Furthermore, we conduct a sensitivity analysis on the parameter ϵ that confines the region in which the predictive model can be adjusted. As shown in Table 4, we observe that when $\epsilon > 0.2$, the normalized decision regret stabilizes at a low level. In contrast, the MSE loss increases with a larger ϵ . Therefore, we can conclude that even by restricting the fine-tuning module within a tight region with a small value of ϵ , the DL can still be improved while maintaining the MSE loss to prevent prediction shift.

Conclusion

In scenarios with limited data for PO problems, the reliability of prediction results and the quality of final decisions are crucial. Existing DFL approaches can effectively improve the decision quality, but they may introduce significant bias in predictions, which increases the risk of decision-making from a different perspective. To address this issue, this paper introduces DFF, which employs a unique training mechanism that maximizes decision quality while ensuring that fine-tuning constraints are satisfied. The method is applicable to any upstream prediction model, including non-differentiable models. Empirical evaluations across diverse scenarios including synthetic data and real-world resource allocation problems demonstrate that DFF outperforms traditional models and existing DFL approaches.

Due to the decoupling between DFF and the upstream model, DFF has advantages in multi-task learning and multi-objective optimization scenarios. Furthermore, its constraint fine-tuning design allows it to better handle downstream problems with uncertain parameters in constraints. These are the key focus areas of our future work.

Acknowledgments

Financial supports from the National Natural Science Foundation of China (No. 52302411, No. 52472349, No. 72361137005, No. 52131204) and CCF-DiDi GAIA Collaborative Research Funds (No. 202310) are gratefully acknowledged.

References

- Agrawal, A.; Amos, B.; Barratt, S.; Boyd, S.; Diamond, S.; and Kolter, J. Z. 2019. Differentiable convex optimization layers. *Advances in Neural Information Processing Systems*, 32.
- Amos, B.; and Kolter, J. Z. 2017. Optnet: Differentiable optimization as a layer in neural networks. In *International Conference on Machine Learning*, 136–145.
- Bansal, D.; Chen, R. T.; Mukadam, M.; and Amos, B. 2024. Taskmet: Task-driven metric learning for model learning. *Advances in Neural Information Processing Systems*, 36.
- Beichter, M.; Werling, D.; Heidrich, B.; Phipps, K.; Neumann, O.; Friederich, N.; Mikut, R.; and Hagemeyer, V. 2024. Decision-Focused Retraining of Forecast Models for Optimization Problems in Smart Energy Systems. In *Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems*, 170–181.
- Berthet, Q.; Blondel, M.; Teboul, O.; Cuturi, M.; Vert, J.-P.; and Bach, F. 2020. Learning with differentiable perturbed optimizers. *Advances in Neural Information Processing Systems*, 33: 9508–9519.
- Bertsimas, D.; and Kallus, N. 2020. From predictive to prescriptive analytics. *Management Science*, 66(3): 1025–1044.
- Breiman, L. 2001. Random forests. *Machine learning*, 45: 5–32.
- Butler, A.; and Kwon, R. H. 2023. Gradient boosting for convex cone predict and optimize problems. *Operations Research Letters*, 51(1): 79–83.
- Chernozhukov, V.; Chetverikov, D.; Demirer, M.; Duflo, E.; Hansen, C.; Newey, W.; and Robins, J. 2018. Double/debiased machine learning for treatment and structural parameters. *Econometrics Journal*, 21: C1–C68.
- Dalle, G.; Baty, L.; Bouvier, L.; and Parmentier, A. 2022. Learning with combinatorial optimization layers: a probabilistic approach. *arXiv preprint arXiv:2207.13513*.
- Ding, N.; Qin, Y.; Yang, G.; Wei, F.; Yang, Z.; Su, Y.; Hu, S.; Chen, Y.; Chan, C.-M.; Chen, W.; et al. 2023. Parameter-efficient fine-tuning of large-scale pre-trained language models. *Nature Machine Intelligence*, 5(3): 220–235.
- Donti, P.; Amos, B.; and Kolter, J. Z. 2017. Task-based end-to-end model learning in stochastic optimization. *Advances in Neural Information Processing Systems*, 30.
- Elmachtoub, A. N.; and Grigas, P. 2022. Smart “Predict, then Optimize”. *Management Science*, 68: 9–26.
- Elmachtoub, A. N.; Lam, H.; Zhang, H.; and Zhao, Y. 2023. Estimate-then-optimize versus integrated-estimation optimization: A stochastic dominance perspective. *arXiv preprint arXiv:2304.06833*.
- Elmachtoub, A. N.; Liang, J. C. N.; and McNellis, R. 2020. Decision trees for decision-making under the predict-then-optimize framework. In *International Conference on Machine Learning*, 2858–2867.
- Fan, Y.; Watkins, O.; Du, Y.; Liu, H.; Ryu, M.; Boutilier, C.; Abbeel, P.; Ghavamzadeh, M.; Lee, K.; and Lee, K. 2024. Reinforcement learning for fine-tuning text-to-image diffusion models. *Advances in Neural Information Processing Systems*, 36.
- Fu, Z.; Yang, H.; So, A. M.-C.; Lam, W.; Bing, L.; and Collier, N. 2023. On the effectiveness of parameter-efficient fine-tuning. *Proceedings of the AAAI conference on artificial intelligence*, 37(11): 12799–12807.
- Gao, S.; Ran, Q.; Su, Z.; Wang, L.; Ma, W.; and Hao, R. 2024. Evaluation system for urban traffic intelligence based on travel experiences: A sentiment analysis approach. *Transportation Research Part A: Policy and Practice*, 187: 104170.
- Ha, D.; Dai, A.; and Le, Q. V. 2016. Hypernetworks. *arXiv preprint arXiv:1609.09106*.
- Hu, X.; Lee, J.; and Lee, J. 2024. Two-Stage Predict+ Optimize for MILPs with Unknown Parameters in Constraints. *Advances in Neural Information Processing Systems*, 36.
- Hu, X.; Lee, J. C.; and Lee, J. H. 2023. Predict+ Optimize for packing and covering LPs with unknown parameters in constraints. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 3987–3995.
- Hu, Y.; Kallus, N.; and Mao, X. 2022. Fast rates for contextual linear optimization. *Management Science*, 68(6): 4236–4245.
- Kotary, J.; Fioretto, F.; Van Hentenryck, P.; and Wilder, B. 2021. End-to-end constrained optimization learning: A survey. *arXiv preprint arXiv:2103.16378*.
- Kotary, J.; Fioretto, F.; Van Hentenryck, P.; and Zhu, Z. 2022. End-to-end learning for fair ranking systems. In *Proceedings of the ACM Web Conference 2022*, 3520–3530.
- Kurutach, T.; Clavera, I.; Duan, Y.; Tamar, A.; and Abbeel, P. 2018. Model-ensemble trust-region policy optimization. *arXiv preprint arXiv:1802.10592*.
- Mandi, J.; Bucarey, V.; Tchomba, M. M. K.; and Guns, T. 2022. Decision-focused learning: Through the lens of learning to rank. In *International Conference on Machine Learning*, 14935–14947.
- Mandi, J.; and Guns, T. 2020. Interior point solving for lp-based prediction+ optimisation. *Advances in Neural Information Processing Systems*, 33: 7272–7282.
- Mandi, J.; Kotary, J.; Berden, S.; Mulamba, M.; Bucarey, V.; Guns, T.; and Fioretto, F. 2023. Decision-focused learning: Foundations, state of the art, benchmark and future opportunities. *arXiv preprint arXiv:2307.13565*.
- Mandi, J.; Stuckey, P. J.; Guns, T.; et al. 2020. Smart predict-and-optimize for hard combinatorial optimization problems. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 1603–1610.
- Queoney, J.; Paschalidis, I. C.; and Cassandra, C. G. 2021. Uncertainty-aware policy optimization: A robust, adaptive trust region approach. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 9377–9385.
- Sadana, U.; Chenreddy, A.; Delage, E.; Forel, A.; Frejinger, E.; and Vidal, T. 2024. A survey of contextual optimization methods for decision-making under uncertainty. *European Journal of Operational Research*.

Schulman, J.; Levine, S.; Abbeel, P.; Jordan, M.; and Moritz, P. 2015. Trust region policy optimization. In *International Conference on Machine Learning*, 1889–1897.

Shah, S.; Wang, K.; Wilder, B.; Perrault, A.; and Tambe, M. 2022. Decision-focused learning without decision-making: Learning locally optimized decision losses. *Advances in Neural Information Processing Systems*, 35: 1320–1332.

Shah, S.; Wilder, B.; Perrault, A.; and Tambe, M. 2024. Leaving the nest: Going beyond local loss functions for predict-then-optimize. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 14902–14909.

She, Y.; Atzberger, C.; Blake, A.; Gualandi, A.; and Keshav, S. 2024. MAGIC: Modular Auto-encoder for Generalisable Model Inversion with Bias Corrections. *arXiv preprint arXiv:2405.18953*.

Tang, B.; and Khalil, E. B. 2022. Pyepo: A pytorch-based end-to-end predict-then-optimize library for linear and integer programming. *arXiv preprint arXiv:2206.14234*.

Wilder, B.; Dilkina, B.; and Tambe, M. 2019. Melding the data-decisions pipeline: Decision-focused learning for combinatorial optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 1658–1665.

Wu, Y.; Mansimov, E.; Grosse, R. B.; Liao, S.; and Ba, J. 2017. Scalable trust-region method for deep reinforcement learning using kronecker-factored approximation. *Advances in Neural Information Processing Systems*, 30.

Wu, Z.; Wang, L.; Huang, F.; Zhou, L.; Song, Y.; Ye, C.; Nie, P.; Ren, H.; Hao, J.; He, R.; et al. 2022. A framework for multi-stage bonus allocation in meal delivery platform. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 4195–4203.

Zhang, R.; Han, J.; Liu, C.; Gao, P.; Zhou, A.; Hu, X.; Yan, S.; Lu, P.; Li, H.; and Qiao, Y. 2023. Llama-adapter: Efficient fine-tuning of language models with zero-init attention. *arXiv preprint arXiv:2303.16199*.

Zhao, K.; Hua, J.; Yan, L.; Zhang, Q.; Xu, H.; and Yang, C. 2019. A unified framework for marketing budget allocation. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 1820–1830.

Zharmagambetov, A.; Amos, B.; Ferber, A.; Huang, T.; Dilkina, B.; and Tian, Y. 2024. Landscape surrogate: Learning decision losses for mathematical optimization under partial information. *Advances in Neural Information Processing Systems*, 36.