

Locally Convex Global Loss Network for Decision-Focused Learning

Haeun Jeon*, Hyunglip Bae*, Minsu Park, Chanyeong Kim, Woo Chang Kim†

KAIST

{haeun39, qogudflq, mspark0425, kim.chanyeong, wkim}@kaist.ac.kr

Abstract

In decision-making problems under uncertainty, predicting unknown parameters is often considered independent of the optimization part. Decision-focused learning (DFL) is a task-oriented framework that integrates prediction and optimization by adapting the predictive model to give better decisions for the corresponding task. Here, an inevitable challenge arises when computing the gradients of the optimal decision with respect to the parameters. Existing research copes with this issue by smoothly reforming surrogate optimization or constructing surrogate loss functions that mimic task loss. However, they are applied to restricted optimization domains. In this paper, we propose Locally Convex Global Loss Network (LCGLN), a global surrogate loss model that can be implemented in a general DFL paradigm. LCGLN learns task loss via a partial input convex neural network which is guaranteed to be convex for chosen inputs while keeping the non-convex global structure for the other inputs. This enables LCGLN to admit general DFL through only a single surrogate loss without any sense for choosing appropriate parametric forms. We confirm the effectiveness and flexibility of LCGLN by evaluating our proposed model with three stochastic decision-making problems.

1 Introduction

Decision-making problems under uncertainty arise in various real-world applications such as production optimization, energy planning, and asset liability management (Shina and Birge 2003; Carino et al. 1994; Fleten and Kristoffersen 2008; Delage and Ye 2010; Garlappi, Uppal, and Wang 2007). These problems involve two main tasks: *Prediction* and *Optimization*. The *Prediction* task aims to create a model for uncertainty and estimate unknown parameters from some input features. This task is often performed using machine learning (ML) techniques such as regression or deep learning. In the *Optimization* task, the corresponding optimization problem is solved using diverse off-the-shelf solvers with the estimated parameters from the prediction task. For example, in asset liability management, asset returns are estimated in the prediction stage, and the optimal

portfolio to repay the liability is obtained in the optimization stage based on these predictions.

Prediction-focused learning (PFL) framework is the most commonly used approach to solve such problems, handling the two tasks in separate and independent steps. PFL first trains the ML model on the input features to produce predictions that closely match the observed parameters. Subsequently, the decision-making problem is defined using the parameters predicted by the trained ML model and is solved to obtain an optimal decision. PFL is based on the underlying belief that accurate predictions will lead to good decisions. However, ML models frequently fail to achieve perfect accuracy, resulting in sub-optimal decisions. Consequently, in many applications, the prediction and optimization tasks are not distinct but are intricately linked and should ideally be considered jointly.

Decision-focused learning (DFL) achieves the above purpose by directly training ML models to make predictions that lead to good decisions. DFL combines prediction and optimization tasks by creating an end-to-end system that aims to minimize *task loss* directly. Task loss is defined as the quality of decisions derived from the predictive model and therefore, depends on the optimal solution of the associated optimization problem. To train an ML model in this context, differentiating through the optimization problem is required to calculate the gradient of the task loss. This presents a key challenge in integrating prediction and optimization. However, differentiation may be impossible if the decision variable is discrete or the objective function is discontinuous.

While surrogate optimization still necessitates differentiation of the optimization problem, efforts have been made to develop solver-free surrogate loss functions that can effectively obtain gradients. We refer to this approach as *Surrogate DFL*. Recent successful works in Surrogate DFL include SPO (Elmachtoub and Grigas 2022), Contrastive Loss (Mulamba et al. 2020), LTR Loss (Mandi et al. 2022), LODL (Shah et al. 2022), EGL (Shah et al. 2024), LANCER (Zharmagambetov et al. 2024), and TaskMet (Bansal et al. 2024). SPO, Contrastive loss, and LTR loss are specifically designed for linear objectives while LODL can be implemented for general DFL problems. However, LODL creates and trains a surrogate with a specific parametric form for each data instance resulting in expensive computation, with

*These authors contributed equally.

†Corresponding author.

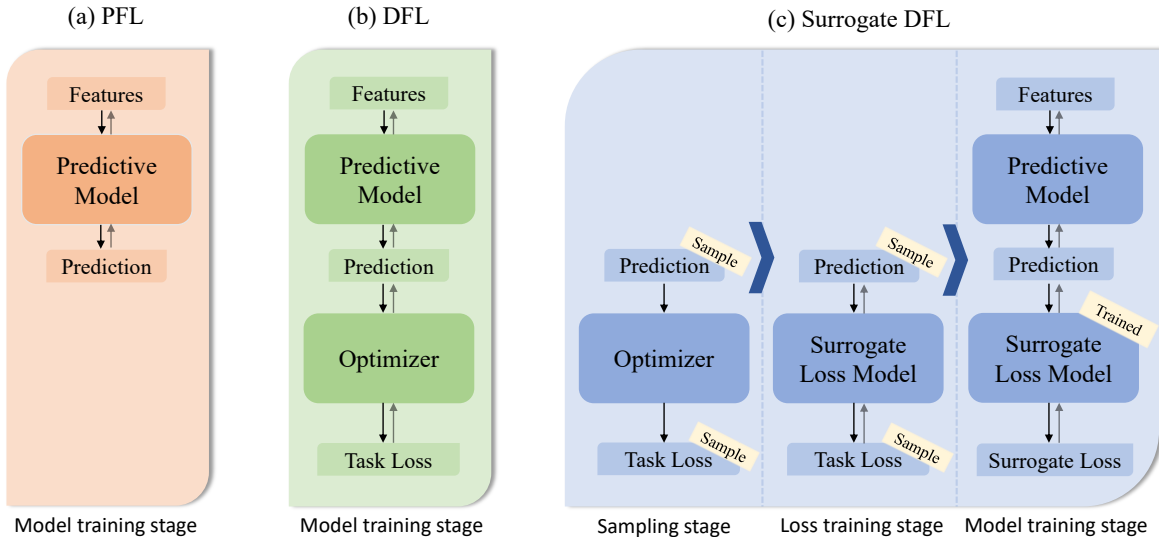


Figure 1: A model training pipeline for PFL, DFL, and surrogate DFL. PFL trains the predictive model by minimizing the prediction loss. DFL directly delivers gradients minimizing the task loss. Surrogate DFL first learns a surrogate loss model that follows the true task loss by sampling predictions and its task losses. Then, it trains the predictive model to convey useful gradients derived from the trained surrogate loss model in an end-to-end manner.

decision quality heavily dependent on the chosen parametric surrogate form. While EGL and LANCER have partially addressed the computational cost issue by extending a local model to a global model, the challenge of selecting an appropriate parametric surrogate form still remains. Alongside the surrogates, TaskMet minimized the prediction error while maintaining the decision quality. We illustrated the model training pipeline for PFL, DFL, and Surrogate DFL in Figure 1.

Notably, there have been attempts to construct surrogate loss functions using neural networks (Shah et al. 2022). However, these approaches generally performed poorly in experiments because naive MLPs fail to capture the local characteristics of the true underlying loss.

In this paper, we propose Locally Convex Global Loss Network (LCGLN), which is a global and general surrogate loss for DFL. LCGLN adopts partial input convex neural network (PICNN) (Amos, Xu, and Kolter 2017) as a parametric surrogate form to approximate task loss. With PICNN, we guarantee the surrogate loss function to be convex for the chosen inputs while maintaining the general structure for the others. Furthermore, users are not required to possess an artistic sense for selecting the suitable parametric function form for the task loss. Therefore, LCGLN can handle the general DFL problems using only one surrogate loss, regardless of the number of observed data. In the experiment section, we demonstrate the capability of LCGLN on three stochastic decision-making problems, namely inventory stock problem, budget allocation problem, and portfolio optimization problem. We show that LCGLN learns task loss well with a single surrogate loss.

2 Related Works

Various methodologies based on gradient-based learning in DFL have been developed (Mandi et al. 2023). Some approaches directly differentiate the constrained optimization problem to update the model parameters. For instance, stochastic optimization problems were tackled by directly differentiating convex QPs using KKT optimality conditions and employing Optnet (Amos and Kolter 2017) for efficient differentiation (Donti, Amos, and Kolter 2017). For general convex optimization problems, the differentiable solver Cvxpylayers (Agrawal et al. 2019) were developed.

Alternatively, some works introduce regularization terms to smooth the optimization mapping. For example, the Euclidean norm of the decision variables was added to use quadratic programming methods (Wilder, Dilkina, and Tambe 2019), while the logarithmic barrier term was added to differentiate linear programming problems (Mandi and Guns 2010). Furthermore, entropy terms were incorporated to solve multi-label problems (Martins and Kreutzer 2017; Amos, Koltun, and Kolter 2019), while constrained optimization problems were addressed using the perturb-and-MAP (Papandreou and Yuille 2011) framework, which added regularization through perturbations (Niepert, Minervini, and Franceschi 2021; Berthet et al. 2020).

Recent works on DFL have focused on constructing differentiable surrogate loss models. Smart "Predict, Then Optimize" (SPO) (Elmachtoub and Grigas 2022) proposed a convex SPO+ loss where the loss upper bounds the task loss. To update the predictive model, they obtained a subgradient of their proposed surrogate loss. NCE (Mulamba et al. 2020) used a noise contrastive approach (Gutmann and Hyvärinen 2010) to build a family of surrogate loss for linear objec-




Figure 2: A simple example of a knapsack problem. There are two items valued \$40, \$30 each, marked with a yellow star. We predict the value of items and choose the higher one. Blue dots and red crosses are predicted values representing good and bad decisions respectively. PFL gives the same prediction loss for every prediction while DFL gives \$10 loss in red cross and 0 in blue.

tives. LTR (Mandi et al. 2022) applied this approach to ranking problems. SO-EBM (Kong et al. 2022) used an energy-based differentiable layer to model the conditional probability of the decision. The energy-based layer acts as a surrogate of the optimization problem. LODL (Shah et al. 2022) used supervised machine learning to locally construct surrogate loss models to represent task loss. They first sampled predictions near true instances for each instance and built a surrogate loss model respectively. They proposed three families for the model (e.g. WeightedMSE, Quadratic) to design parametric surrogate loss models. EGL (Shah et al. 2024) extends LODL to a global surrogate model and LANCER (Zharmagambetov et al. 2024) learns such handcraft global surrogate loss via actor-critic framework.

3 Preliminaries

In this section, we motivate the necessity of DFL approach from a simple knapsack example. We also briefly summarize surrogate loss models and input convex neural networks.

Comparison on PFL and DFL Prediction-focused learning (PFL) and decision-focused learning (DFL) are two big learning pipelines for decision-making problems under uncertainty as in Figure 1. While PFL learns the predictive model focusing on the prediction from the predictive model, DFL focuses on the objective of the downstream optimization problem, commonly referred to *task loss*. Starting with

PFL, the model training step is divided into two stages, prediction and optimization. In the prediction stage, PFL learns a predictive model by minimizing the prediction loss such as MSE. In the optimization stage, it solves an optimization problem using the prediction of the predictive model as parameters. By contrast, the predictive model in DFL is trained in an end-to-end manner. It prioritizes learning the model to make good decisions (or actions) by minimizing the task loss, rather than optimizing the prediction loss.

We introduce a motivating example with a simple knapsack problem in Figure 2 to give motivation in using DFL. Let’s consider two items: item 1 with a value of \$40 and item 2 with \$30. Two axes in Figure 2 represent the value of each item. The true optimal value (40, 30) is marked by a yellow star. Assume we can only select one item. Our objective is to predict the value of items by only observing the item features and choosing one with a higher value. If we choose item 1 (the optimal decision), we would be satisfied, earning \$40. Conversely, choosing item 2 would be a sub-optimal decision, resulting in a relative loss of \$10. Suppose we predict the values as blue dots and red crosses. From a PFL point of view, the prediction loss is equal to \$20. Any prediction on the gray circle results in the same prediction loss, which is not informative as the blue dots are correct decisions, whereas the red crosses are not. From a DFL standpoint, blue dots return zero task loss while red crosses give task loss of \$10. This example shows that minimizing the prediction loss *may not* lead to lower task loss.

Surrogate DFL DFL computes task loss gradients with respect to the input parameters by, for example, directly differentiating the optimization problem (or optimizer) (Donti, Amos, and Kolter 2017). However, the optimizer cannot be differentiated easily in most cases. To tackle this, methods using surrogate loss functions were proposed (Elmachtoub and Grigas 2022). Surrogate loss models are differentiable and it approximates the mapping between the prediction and the task loss. Consequently, it can be used to calculate the gradient to update the predictive model efficiently.

Given dataset $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$ with N instances, our goal is to minimize the regret \mathcal{R} for the optimization problem:

$$\begin{aligned} \min_{\theta} \quad & \mathcal{R}(\hat{y}, y) := \mathcal{L}_{task}(a^*(\hat{y}), y) - \mathcal{L}_{task}(a^*(y), y) \\ \text{s.t.} \quad & \hat{y} = \mathcal{M}_{\theta}(x) \\ & a^*(\hat{y}) = \arg \min_{a \in \mathcal{A}} \mathcal{L}_{task}(a, \hat{y}) \end{aligned} \tag{1}$$

where \mathcal{L}_{task} is the task loss to be optimized, \hat{y} is the prediction from the predictive model \mathcal{M} parameterized by θ and $a^* \in \mathcal{A}$ is an optimal action (or decision) in a feasible region \mathcal{A} derived by any off-the-shelf solver. Note that the second term of the objective is nothing but a constant optimal loss, and therefore our objective is equivalent to minimizing $\mathcal{L}_{task}(a^*(\hat{y}), y)$.

To update the predictive model via gradient descent w.r.t. θ , one must find the gradient $\partial \mathcal{R} / \partial \theta$. Using the chain rule,

the gradient can be decomposed into:

$$\frac{\partial \mathcal{R}(\hat{y}, y)}{\partial \theta} = \frac{\partial \mathcal{R}(\hat{y}, y)}{\partial \hat{y}} \cdot \frac{\partial \hat{y}}{\partial \theta} \quad (2)$$

$$\approx \frac{\partial \mathcal{L}_\psi(\hat{y}, y)}{\partial \hat{y}} \cdot \frac{\partial \hat{y}}{\partial \theta} \quad (3)$$

where $\mathcal{L}_\psi(\hat{y}, y)$ is the parametric surrogate loss with parameter ψ that is trained to approximate $\mathcal{R}(\hat{y}, y)$. The global surrogate loss can be trained to richly approximate the true task loss as:

$$\psi^* = \arg \min_{\psi} \mathbb{E}_{\hat{y}, y} [|\mathcal{L}_\psi(\hat{y}, y) - \mathcal{R}(\hat{y}, y)|] \quad (4)$$

Training the predictive model \mathcal{M}_θ can be easily done with backpropagation using the gradient $\partial \mathcal{L}_\psi / \partial \theta$.

Partial Input Convex Neural Network Amos et al. (Amos, Xu, and Kolter 2017) proposed an input convex neural network (PICNN) that ensures convexity with respect to the chosen inputs. PICNN is constructed by introducing additional weights to connect the input layer to each hidden layer. Given this structure, non-decreasing convex activation functions such as softplus and non-negativity of weights connecting between hidden layers are required to guarantee the convexity of PICNN. LCGLN uses PICNN to richly represent the non-convex nature of the true loss mapping, and simultaneously ensure local convexity.

4 Locally Convex Global Loss Network

In various optimization problems, each has its own specific objectives known as *task losses*. Our goal is to devise a surrogate loss function capable of accurately representing the true task loss across a range of optimization problems. In this paper, we introduce LCGLN, where a single loss representation can replace the true task loss. While DFL methods suffer from differentiating through the optimization solver when the optimization problem is not smooth, our LCGLN is easily differentiable and therefore can be readily back-propagated when updating the predictive model via gradients.

We now introduce the end-to-end training procedure for LCGLN. The goal is to learn a regret mapping $(\hat{y}, y) \rightarrow \mathcal{R}(\hat{y}, y)$ to obtain gradient $\partial \mathcal{R} / \partial \theta$ for updating the predictive model \mathcal{M}_θ . The training consists of three main steps: generating samples, learning a global surrogate loss, and training a predictive model.

Generating Samples To train the LCGLN in a supervised learning manner, we generate K samples for each N instance, i.e. for each instance y_i , we sample $\tilde{y}_i^{(1)}, \tilde{y}_i^{(2)}, \dots, \tilde{y}_i^{(K)}$.

Some previous research assumed that predictions \tilde{y} would closely approximate the true labels y , typically generating sample predictions by simply adding Gaussian noise. However, the Gaussian sampling method may be challenging to apply for two reasons. First, in some cases, the true label y is unknown. For example, in the inventory stock problem used in our experiments, the true y represents the probability vector of each demand occurring. Yet, the observed data only

Algorithm 1: Training Predictive Model with LCGLN \mathcal{L}_ψ

Dataset: $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$
1. Generate Samples:
 Sample set $\mathcal{S} \leftarrow \emptyset$
for $i = 1, \dots, N$ **do**
 $\mathcal{S} \leftarrow \mathcal{S} \cup (y_i, y_i, 0)$
end for
 Initialize sampling model \mathcal{M}_ξ .
for $k = 1, \dots, K - 1$ **do**
 for $i = 1, \dots, N$ **do**
 $\tilde{y}_i^{(k)} = \mathcal{M}_\xi(x_i)$
 $\xi \leftarrow \text{Update}(\xi, \nabla_\xi \|\tilde{y}_i^{(k)} - y_i\|)$
 Solve $a_i^{*(k)} = \arg \min_a \mathcal{L}_{task}(a(\tilde{y}_i^{(k)}), y_i)$.
 $\mathcal{S} \leftarrow \mathcal{S} \cup (\tilde{y}_i^{(k)}, y_i, \mathcal{R}(\tilde{y}_i^{(k)}, y_i))$
 end for
end for
2. Learn LCGLN \mathcal{L}_ψ :
 Initialize surrogate loss model \mathcal{L}_ψ .
for $(\hat{y}, y, \mathcal{R}(\hat{y}, y))$ in \mathcal{S} **do**
 $\psi \leftarrow \text{Update}(\psi, \nabla_\psi \|\mathcal{L}_\psi(\hat{y}, y) - \mathcal{R}(\hat{y}, y)\|)$
end for
3. Train Predictive Model \mathcal{M}_θ :
 Initialize predictive model \mathcal{M}_θ .
 $\theta \leftarrow \text{Update}(\theta, \nabla_\theta \hat{y} \cdot \nabla_{\hat{y}} \mathcal{L}_\psi(\hat{y}, y))$

provides a specific demand value realized according to that probability. In such cases, since the true y is unknown, there is no target to which Gaussian noise can be added. Second, even if the true y is known, determining the appropriate standard deviation for sampling around y can be difficult, often requiring repeated tuning.

On account of this, we adopt the model-based sampling (MBS) approach (Shah et al. 2024). MBS involves constructing a sampling model \mathcal{M}_ξ that mirrors the architecture of the predictive model and training within a PFL paradigm using MSE. During the training, we take inferences from the intermediate model and generate samples. Unlike the Gaussian sampling, neither the true y nor the standard deviation for the noise is required for MBS.

Learning Global Surrogate Loss LCGLN Our objective is to learn a mapping $(\hat{y}, y) \rightarrow \mathcal{R}$ for conveying informative gradients of task loss. We propose leverage of partial input convex neural network (PICNN) (Amos, Xu, and Kolter 2017) as our Locally Convex Global Loss Network (LCGLN). Our motivation for using PICNN as a surrogate model for task loss is fourfold:

- *Expressiveness:* A good surrogate loss model should have a sufficient number of parameters to accurately approximate the true task loss since a lack of expressiveness may lead to under-performance. Amos, Xu, and Kolter (2017) proved that a k -layer PICNN can represent any k -layer feedforward network. Since feedforward neural networks are known as universal approximators (Cybenko 1989; Funahashi 1989; Hornik 1991), this ensures the PICNN can accurately model task loss.
- *Easily Differentiable:* We need to differentiate our loss model with respect to \hat{y} for gradients training the pre-

PARADIGM	METHODS	PROBLEM		
		INVENTORY	BUDGET (500 FAKES)	PORTFOLIO
2 STAGE	PFL	0.242 ± 0.005	0.513 ± 0.016	0.189 ± 0.002
EXACT DIFF	DFL	0.228 ± 0.002	0.532 ± 0.020	0.187 ± 0.002
LOCAL SURROGATE	LODL-DQ	0.378 ± 0.007	0.503 ± 0.020	0.193 ± 0.002
GLOBAL SURROGATE	LANCER	0.182 ± 0.004	0.490 ± 0.010	0.246 ± 0.008
	EGL-WMSE	0.371 ± 0.002	0.510 ± 0.013	0.187 ± 0.001
	EGL-DQ	0.369 ± 0.007	0.492 ± 0.005	0.256 ± 0.002
	LCGLN	0.174 ± 0.002	0.468 ± 0.009	0.185 ± 0.000

Table 1: The table contains normalized test regret $\mathcal{R}_{test}/\mathcal{R}_{worst}$ with standard error mean (SEM) tested on three stochastic optimization problems. The metric is *lower the better* with an optimal regret of zero. The best-performing results for each problem are bold-lettered. We evaluate PFL, DFL, local, and global surrogate loss models. We use 32 samples for the surrogate loss models. The global surrogate loss LCGLN outperformed the baselines across all three problems.

dictive model. Discontinuous or non-differentiable loss models face significant challenges in such tasks. In contrast, PICNN can be easily differentiated by using built-in backward functions.

- *Locally Convex, but not Globally:* We want the model to capture the highly non-convex structure of the true underlying task loss mapping. Task loss can be expressed as $f(x^*(\hat{y}), y) - f(x^*(y), y)$, where it achieves its minimum regret zero when $\hat{y} = y$. While task loss may not be convex in \hat{y} for a given y in general, we use PICNN to drive \hat{y} towards the true value y by eliminating local minima except y . This property enables the model to provide informative gradients for training the predictive model with a small sample size. Furthermore, we induce the local minima by adding $\{(y_i, y_i, 0)\}_{i=1}^N$ to the dataset for every instance.
- *Generality:* The adoption of PICNN helps generalize the function approximators and allows the approximation of differentiable optimization. Using PICNN, we can only focus on the overall network architecture and perform a simple hyperparameter search instead of requiring well-trained experts’ efforts in choosing the right parametric function forms for specific problems.

Training Predictive Model The predictive model M_θ learns a mapping $x \rightarrow y$. To obtain gradients for the predictive model training, we first generated samples using the MBS approach. For each generated sample, we derived the true task loss in regret form and added the sample, instance, and regret pair to the dataset. Additionally, we included $\{(y_i, y_i, 0)\}_{i=1}^N$ into the dataset to induce local minima for every instance y_i . Using this dataset, we trained the global surrogate loss model structured with PICNN.

Building on these steps, we now train the predictive model M_θ via gradient descent utilizing the gradients provided by the global surrogate loss \mathcal{L}_ψ . The entire procedure from generating the dataset to training the predictive model is summarized in Algorithm 1.

5 Experiments and Results

We validate our methodology with three stochastic optimization problems.

5.1 Experimental Settings

In this section, we explain the experimental details and the baselines used to evaluate our method. Each problem is elaborated in two stages: the parameter prediction stage and the optimization stage.

Problem Description We conducted experiments on three different stochastic optimization problems, each presenting unique challenges as explored in previous research (Donti, Amos, and Kolter 2017; Wilder, Dilkina, and Tambe 2019; Shah et al. 2022). Our experiments were built on top of the public codes from previous research (Donti, Amos, and Kolter 2017; Shah et al. 2022). Here we briefly describe each problem. For further problem descriptions, please refer to Appendix B.

- **Inventory Stock** (Donti, Amos, and Kolter 2017): We decide the order quantity to minimize the cost over the stochastic demand. To simplify the problem, we assume the demands are discrete.

Prediction: We predict the discrete probability distribution of the demand.

Optimization: With the predicted demand distribution, we decide the order quantity that minimizes the cost.

- **Budget Allocation** (Wilder, Dilkina, and Tambe 2019): We choose websites to advertise based on click-through rates (CTRs) of users among websites. As a variant, we conduct four different experiment settings by concatenating $\{0, 5, 50, 500\}$ size of random CTRs (fake targets) to the original CTRs and adjusting the problem difficulty.

Prediction: Given the website features, we predict the CTRs.

Optimization: With the predicted CTRs, we choose the websites to advertise that maximize the expected number of users who click on the advertisement at least once.




Figure 3: A line plot showing the normalized test regret $\mathcal{R}_{test}/\mathcal{R}_{worst}$ for each methodology and problem setting. The metric is *lower the better*, with 0 representing the optimal value. We tested sample size of $\{2, 4, 8, 16, 32\}$ for each problem. The standard error mean for each experiment is detailed in Appendix C.2. Our global surrogate loss LCGLN represented by the red straight line outperforms when 32 samples are used.

- **Portfolio Optimization** (Shah et al. 2022): We allocate weights for invested stocks to maximize the expected risk-adjusted return.

Prediction: Given the historical daily stock return data, we predict the future stock price.

Optimization: Using the predicted stock price, we assign portfolio weights that maximize the expected risk-adjusted return, given the covariance matrix.

Baselines We compare our global surrogate loss network LCGLN with the following baselines from previous research. When available, we utilized the hyperparameters specified in previous research for the baselines. In cases where these were not provided, we endeavored to fine-tune them to perform their best.

- **PFL:** A standard approach that trains a predictive model on the input features to produce predictions that closely match the observed parameters. We use the negative log-likelihood (NLL) loss for the inventory stock problem and the mean squared error (MSE) loss for the other problems.
- **DFL:** A direct approach that differentiates through the solver to derive gradients for training the predictive model. For the inventory stock and the portfolio optimization problems, we use the differentiable QP solver (Donti, Amos, and Kolter 2017). We use multilinear relaxation (Wilder, Dilkina, and Tambe 2019) for the budget allocation problem.
- **LODL:** The local surrogate model (Shah et al. 2022) that learns a surrogate loss locally for each instance, creating total N surrogate loss models. To learn a loss model, LODL generates K number of samples for each N instance and solves for the input of the model. We use directed-quadratic surrogate loss (LODL-DQ) from their paper, the most promising result among the proposed 4 different loss models.
- **LANCER:** The global surrogate model (Zharmagambevtov et al. 2024) that uses alternating optimization to learn

the predictive model and the global surrogate loss model. The prediction samples are sampled from the predictive model and are used to train the surrogate loss. The predictive model is then trained using the gradients from the updated surrogate loss.

- **EGL:** The global surrogate model (Shah et al. 2024) that learns a global surrogate loss with feature-based parameterization and model-based sampling. We use directed-quadratic (EGL-DQ) and weighted-MSE (EGL-WMSE) for the baseline, which showed the best results on the budget allocation and the portfolio optimization problems respectively.

Evaluation Metric We use the normalized test regret $\mathcal{R}_{test}/\mathcal{R}_{worst}$ as the evaluation metric, where \mathcal{R}_{test} is a test regret with regret defined in Equation 1. For the maximization tasks, such as the budget allocation and portfolio optimization problem, regrets are calculated by multiplying the negative sign on the corresponding objective. We calculate the worst case regret \mathcal{R}_{worst} for each problem and derive the normalized test regret. For the inventory stock problem, we assumed the worst-case scenario when the company ordered no stock. In the budget allocation problem, the worst-case scenario was assumed when the advertisements were allocated to the lowest predicted CTRs. For the portfolio optimization problem, we considered the worst-case scenario to be when the entire investment was allocated to the stock with the lowest predicted return. Our metric is 0 when optimal and 1 when worst.

We used a predictive model \mathcal{M}_θ as one hidden layer MLP and a learning rate of 0.001. 500 hidden nodes were used for the portfolio optimization and 10 for the other problems. We mirrored the predictive model exactly for the sampling model \mathcal{M}_ξ . For LCGLN, we employed a single hidden layer PICNN with two nodes per layer, a learning rate of 0.001, and a softplus activation function. A figure illustrating the LCGLN can be found in Appendix A.1. For each global surrogate loss model employing model-based sampling, we selected the learning rate that demonstrated the best perfor-

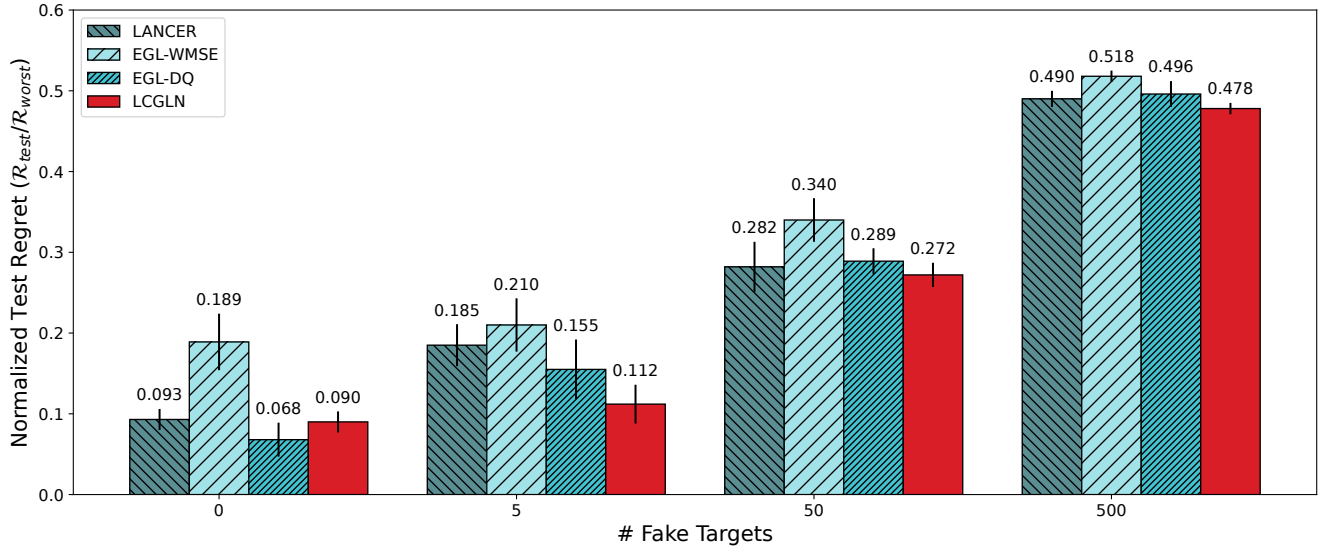


Figure 4: A histogram presenting normalized test regret $\mathcal{R}_{test}/\mathcal{R}_{worst}$ with standard error mean (SEM) for global surrogate loss models in budget allocation with varying number of fake targets. We use 16 samples for learning loss. The metric is *lower the better* and 0 when optimal. We test with $\{0, 5, 50, 500\}$ fake targets, noting that the problem becomes more challenging as the number of fake targets increases. Our LCGLN shown in red bars outperforms most settings.

mance from $\{0.01, 0.05, 0.1, 0.5, 1\}$. We run 10 experiments for each setting to ensure statistical significance. The experiments were performed on a Ryzen 7 5800X CPU and an RTX 3060 GPU with 64GB of RAM.

5.2 Results

Table 1 shows the normalized test regret $\mathcal{R}_{test}/\mathcal{R}_{worst}$ with standard error mean (SEM) for the inventory stock, budget allocation with 500 fake targets and portfolio optimization problems. For the surrogate loss models, we used 32 samples to learn the loss. We categorize the methods into four major training paradigms: two-stage, exact differentiation, local surrogate loss, and global surrogate loss.

In the inventory stock problem, LODL-DQ, EGL-WMSE, and EGL-DQ suffer from high regret, indicating that they do not provide informative gradients. LANCER and LCGLN perform well compared to others due to their *expressiveness*, as both use neural networks to learn the loss. For the budget allocation problem, we conduct four experiments with $\{0, 5, 50, 500\}$ fake targets. Detailed results with different fake targets are available in Appendix C.1. For the hardest problem with 500 fake targets, the surrogate models tend to show better results than PFL or DFL methods. In the portfolio optimization problem, good prediction in parameters showed better decisions. This is evidenced by PFL and DFL showing better results compared to surrogate models. EGL-WMSE and LCGLN also demonstrated strong performance in our settings.

To reduce the computational cost, it is crucial to train the surrogate loss with small sample sizes. We conducted experiments with varying sample sizes, as shown in the line plot in Figure 3. Note that the normalized test regret is better when

lower. Since PFL, DFL, and LANCER do not vary with changes in a number of sample predictions, they show consistent results across all sample sizes. The surrogate models showed decreased regret as the sample size increased. At a sample size of 32, LCGLN outperformed the provided baselines. For detailed experimental results for each sample size and problem, please refer to Appendix C.2.

We also tested different numbers of fake targets, varying in $\{0, 5, 50, 500\}$, to show how models perform on harder problems. We used 16 samples to train the loss model. Figure 4 presents the normalized test regret with SEM for global models across each experimental setting. LCGLN consistently outperformed in most settings.

6 Conclusion

In this paper, we propose Locally Convex Global Loss Network (LCGLN), a global and general surrogate loss for DFL. LCGLN utilizes PICNN as a parametric surrogate to approximate task loss. We use PICNN to guarantee the surrogate to be convex near instances while maintaining a general non-convex structure globally. With LCGLN, users do not need to manually select the appropriate parametric function form for the task loss. Consequently, LCGLN can address general DFL problems with a single surrogate loss, regardless of the amount of observed data. We evaluated our method on three stochastic optimization problems, achieving better decision quality with fewer training samples compared to the state-of-the-art baselines. However, despite the expressive power of LCGLN, a limitation remains: achieving high decision quality requires careful selection of the samples. Thus, our future research will focus on identifying sampling strategies to improve the decision quality for surrogate loss models.

Acknowledgements

This research was supported by the National Research Foundation of Korea (NRF) grant funded by the Korean government (MSIT) (NRF-2022M3J6A1063021 and NRF-RS-2023-00208980).

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. PMLR.
- Amos, B.; Koltun, V.; and Kolter, J. Z. 2019. The limited multi-label projection layer. *arXiv preprint arXiv:1906.08707*.
- Amos, B.; Xu, L.; and Kolter, J. Z. 2017. Input convex neural networks. In *International Conference on Machine Learning*, 146–155. PMLR.
- 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.
- 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.
- Carino, D. R.; Kent, T.; Myers, D. H.; Stacy, C.; Sylvanus, M.; Turner, A. L.; Watanabe, K.; and Ziemba, W. T. 1994. The Russell-Yasuda Kasai model: An asset/liability model for a Japanese insurance company using multistage stochastic programming. *Interfaces*, 24(1): 29–49.
- Cybenko, G. 1989. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4): 303–314.
- Delage, E.; and Ye, Y. 2010. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations research*, 58(3): 595–612.
- 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(1): 9–26.
- Fleten, S.-E.; and Kristoffersen, T. K. 2008. Short-term hydropower production planning by stochastic programming. *Computers & Operations Research*, 35(8): 2656–2671.
- Funahashi, K.-I. 1989. On the approximate realization of continuous mappings by neural networks. *Neural networks*, 2(3): 183–192.
- Garlappi, L.; Uppal, R.; and Wang, T. 2007. Portfolio selection with parameter and model uncertainty: A multi-prior approach. *The Review of Financial Studies*, 20(1): 41–81.
- Gutmann, M.; and Hyvärinen, A. 2010. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, 297–304. JMLR Workshop and Conference Proceedings.
- Hornik, K. 1991. Approximation capabilities of multilayer feedforward networks. *Neural networks*, 4(2): 251–257.
- Kong, L.; Cui, J.; Zhuang, Y.; Feng, R.; Prakash, B. A.; and Zhang, C. 2022. End-to-end stochastic optimization with energy-based model. *Advances in Neural Information Processing Systems*, 35: 11341–11354.
- 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. PMLR.
- Mandi, J.; and Guns, T. 2010. Interior point solving for lp-based prediction+ optimisation, 2020. URL <http://arxiv.org/abs>.
- 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*.
- Martins, A. F.; and Kreutzer, J. 2017. Learning what’s easy: Fully differentiable neural easy-first taggers. In *Proceedings of the 2017 conference on empirical methods in natural language processing*, 349–362.
- Mulamba, M.; Mandi, J.; Diligenti, M.; Lombardi, M.; Bucarey, V.; and Guns, T. 2020. Contrastive losses and solution caching for predict-and-optimize. *arXiv preprint arXiv:2011.05354*.
- Niepert, M.; Minervini, P.; and Franceschi, L. 2021. Implicit MLE: backpropagating through discrete exponential family distributions. *Advances in Neural Information Processing Systems*, 34: 14567–14579.
- Papandreou, G.; and Yuille, A. L. 2011. Perturb-and-map random fields: Using discrete optimization to learn and sample from energy models. In *2011 International Conference on Computer Vision*, 193–200. IEEE.
- Shah, S.; Perrault, A.; Wilder, B.; and Tambe, M. 2024. Leaving the Nest: Going Beyond Local Loss Functions for Predict-Then-Optimize. arXiv:2305.16830.
- 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.
- Shiina, T.; and Birge, J. R. 2003. Multistage stochastic programming model for electric power capacity expansion problem. *Japan journal of industrial and applied mathematics*, 20: 379–397.
- 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.
- 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.