Distributionally Robust Learning for Smoothed Online Optimization via Generative Ambiguity Modeling

CALVIN GLISSON, California State University, San Bernardino, USA
PENGFEI LI, University of California, Riverside, USA
QIUXIAO CHEN, California State University, San Bernardino, USA
JIANYI YANG*, University of Houston, USA

We study the Distributionally Robust Learning (DRL) algorithm in the context of smoothed online optimization. Traditional DRL approaches typically model the ambiguity set using a simple Wasserstein ball, which often fails to accurately capture the underlying data distribution. To enhance the robustness of ML under Out-Of-Distribution (OOD) testing, we propose a novel DRL method that constructs adversarial distributions using a diffusion model. This approach effectively balances distributional fidelity with adversarial strength.

ACM Reference Format:
Calvin Glisson, Pengfei Li, Qiuxiao Chen, and Jianyi Yang. 2018. Distributionally Robust Learning for Smoothed Online Optimization via Generative Ambiguity Modeling. *J. ACM* 37, 4, Article 111 (August 2018), 2 pages. https://doi.org/XXXXXXX.XXXXXXX

1 OVERVIEW

We consider a smoothed online optimization problem where at each round t, the agent receives a context y_t and decides an action x_t. The objective is to minimize the combination of the hitting cost $f(x_t, y_t)$ and the switching cost $c(x_t, x_{t-1})$, i.e. $\min_{x_{1:T}} cost(x_{1:T}, y_{1:T}) \coloneqq \sum_{t=1}^{T} f(x_t, y_t) + \gamma c(x_t, x_{t-1})$, where $\gamma > 0$ is a weight to balance the hitting and switching costs. This problem has been widely applied in data center workload scheduling, energy systems, Electrical Vehicle (EV) charging, etc. ML models have been designed to solve the smoothed online optimization problems [1]. We define the ML model as h_w, with w being the model's weight vector. And the action for each round is predicted with h_w based on the historical information. In this learning problem, we train the ML model as $\hat{w} = \min_w \mathbb{E}_{y_{1:T} \sim \hat{P}}[cost(h_w, y_{1:T})]$ given a training distribution $\hat{P}$.

Despite the powerful prediction ability of ML models, they suffer from Out-Of-Distribution (OOD) testing problems. For instance, the ML model developed for workload scheduling based on several training data center scenarios may have degraded performance when the ML model is applied for a data center with very different service types or locations. In addition, the ML training datasets may be imperfect or even corrupted, resulting in a large training-test discrepancy.

Distributionally Robust Learning (DRL) algorithms have been designed to to address the OOD testing issues. Instead of solely optimizing the ML model on the training dataset, DRL solves a

*Corresponding author.

Authors' addresses: Calvin Glisson, California State University, San Bernardino, USA, 008140867@coyote.csusb.edu; Pengfei Li, University of California, Riverside, USA, pli081@ucr.edu; Qiuxiao Chen, California State University, San Bernardino, USA, Qiuxiao.Chen@csusb.edu; Jianyi Yang, University of Houston, USA, jyang66@uh.edu.

ACM 0004-5411/2018/8-ART111
https://doi.org/XXXXXXX.XXXXXXX

distributionally robust optimization as below to train the ML model.

$$\hat{w}_{DRL} = \min_{w} \max_{Q \in \mathcal{B}(\hat{P}, \epsilon)} \mathbb{E}_{y_{1:T} \sim Q}[cost(h_w, y_{1:T})], \tag{1}$$

where $\mathcal{B}(\hat{P}, \epsilon)$ is an ambiguity set containing distributions close enough to training distribution $\hat{P}$ under a distribution discrepancy ϵ. Conventional DRL algorithms usually model the ambiguity set as a Wasserstein ball $\mathcal{B}(\hat{P}, \epsilon) = \{Q \mid d_W(Q, \hat{P}) \leq \epsilon\}$. However, it is hard to get a good generalization performance using the simple ambiguity models by Wasserstein ball: Small ϵ may fail to capture the large testing-training distribution shifts. Conversely, simply increasing ϵ does not guarantee improved generalization because the Wasserstein ball can include a lot of false distributions that are not consistent with the learning task. Given the critical role of ambiguity set modeling, we propose a new DRL algorithm by modeling the ambiguity set using generative models.

2 METHOD

Our objective is to train a distributionally robust ML model $h_{\hat{w}}$ for smoothed online optimization by constructing adversarial testing distributions P_θ. This involves maximizing the log-likelihood of P_θ on training data, encouraging distributional validity, and maximizing the expected cost of $h_{\hat{w}}$ under P_θ, capturing the adversarial nature. The overall objective function is:

$$\max_{\theta} \sum_{i=1}^{N} \log P_\theta(y_{1:T,i}) + \lambda \cdot \mathbb{E}_{y_{1:T} \sim P_\theta}[cost(h_{\hat{w}}, y_{1:T})]. \tag{2}$$

The weight $\lambda > 0$ is to balance the underlining data distribution learning and adversarial cost maximization. A larger λ means a higher adversarial level corresponding to a larger ϵ in (1). Once we get an adversarial distribution P_θ, we update the ML model weight by the generated samples of P_θ to minimize the adversarial expected cost, i.e. $\hat{w} = \min_w \mathbb{E}_{y_{1:T} \sim P_\theta}[cost(h_w, y_{1:T})]$. We repeat the updates of generative model and the ML model until convergence.

3 PRELIMINARY RESULTS

To validate our method, we explore its application to datacenter demand response [1]. We build the experiment experiment using renewable integration data of different regions of California from California ISO Open Access Same-time Information System (OASIS). We train the DRL model using data from Southern California, and then evaluate its OOD testing performance on Northern California data. The Wasserstein distance between the training and the testing of Southern California data is 0.04, while the distance between the Southern California training data and the Northern California testing data is 0.19, highlighting the distributional discrepancy between southern and northern California data. The testing costs are provided in Table 1. We can find that the DRL model

	Non-robust	$\lambda = 0.1$	$\lambda = 0.5$	$\lambda = 1$	$\lambda = 1.5$
Mean cost	0.0205	0.0192	0.0181	0.0157	0.0199

Table 1. OOD performance at different robustness levels

trained by (2) with different $\lambda > 0$ achieves better performance than non-robust ML model ($\lambda = 0$). The testing costs of (2) vary with different choices of λ. A low adversarial level such as $\lambda = 0.1$ results in less pronounced benefits due to insufficient robustness guarantee while a larger value such as $\lambda = 1.5$ is overly conservative, causing the cost to rise again. $\lambda = 1$ gives the best testing performance in our setting.

REFERENCES

[1] Pengfei Li, Jianyi Yang, and Shaolei Ren. Expert-calibrated learning for online optimization with switching costs. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 6(2):1–35, 2022.