Optimizing the Performative Risk under Weak Convexity Assumptions

Yulai Zhao

Princeton University

Setup

- **Performative Risk** is introduced when prediction causes a change in the distribution of the target variable, i.e.,

$$\mathbf{PR(\theta) = E_{z \sim D(\theta)} l(z; \theta)}$$

  - Our ultimate goal is to find $\boldsymbol{\theta_{po}} = \boldsymbol{argmin_{\theta}\, PR(\theta)}$, performative optima
  - However, past work mainly focused on finding $\boldsymbol{\theta_{ps}} = \mathbf{argmin_{\theta}}\, \mathbf{E}_{\boldsymbol{z \sim D(\theta_{ps})}} \boldsymbol{l(z; \theta)}$, performative stable points

- Example: predicting credit default risk. A bank might estimate that a loan applicant has an elevated risk of default if he applied for a loan, and will act on it by assigning a high interest rate.

Problem

How and Under what conditions could we optimize performative risks?

We aim to answer this question through two perspectives

1. Validate several conditions in first-order optimization that could guarantee a linear convergence rate. Once recognizing such conditions, there are plenty implementations in literature.
2. Connecting the target performative optima points with stable points to take advantage of previous works

Results

- Define $\mathbf{DPR(\theta_1, \theta_2) = E_{z\sim D(\theta_1)} l(z; \theta_2)}$ for *decoupled performative risk.*

- We show: when DPR is WSC (weakly strong convex), PR is WC (weakly convex) to θ_{po}, namely

$$\mathbf{PR(\theta_{po}) \geq PR(\theta) + \langle \nabla PR(\theta), \theta_{po} - \theta \rangle}$$

Assumption 6 *For performative optimum θ_{PO} and its induced distribution $\mathcal{D} = \mathcal{D}(\theta_{\mathrm{PO}})$, suppose the optimal solution for minimizing $\mathrm{DPR}(\theta_{\mathrm{PO}}, \cdot)$ is θ^*. We say $\mathrm{DPR}(\theta_{\mathrm{PO}}, \cdot)$ satisfies $\mu-$WSC, if for any $\theta \in \Theta$ it holds that*

$$\mathrm{DPR}(\theta_{\mathrm{PO}}, \theta^*) \geq \mathrm{DPR}(\theta_{\mathrm{PO}}, \theta) + \nabla_\theta \mathrm{DPR}(\theta_{\mathrm{PO}}, \theta)^\top (\theta^* - \theta) + \frac{\mu}{2}\|\theta^* - \theta\|^2. \qquad (6)$$

Results

- We show: when DPR is RSI (Restricted Secant Inequality), PR is RSI, namely

$$\langle \nabla PR(\theta), \theta - \theta_{po} \rangle \geq \mu' |\theta_{po} - \theta|^2$$

Assumption 7 *For performative optimum θ_{PO} and its induced distribution $\mathcal{D}$, suppose the optimal solution for minimizing $\mathrm{DPR}(\theta_{\mathrm{PO}}, \cdot)$ is unique and is denoted as θ^*. We say $\mathrm{DPR}(\theta_{\mathrm{PO}}, \cdot)$ satisfies $\mu-$RSI, if for any θ it holds*

$$\langle \nabla_\theta \mathrm{DPR}(\theta_{\mathrm{PO}}, \theta), \theta - \theta^* \rangle \geq \mu \|\theta^* - \theta\|^2 \tag{8}$$

When are local properties sufficient for stability and optimality?

We could quantify gap between any θ, θ'

$$\begin{aligned}
\mathrm{PR}(\theta') &= \mathrm{DPR}(\theta', \theta') \\
&\geq \mathrm{DPR}(\theta, \theta') - LW(\mathcal{D}(\theta), \mathcal{D}(\theta')) \\
&= \mathrm{PR}(\theta) + [\mathrm{DPR}(\theta, \theta') - \mathrm{DPR}(\theta, \theta)] - LW(\mathcal{D}(\theta), \mathcal{D}(\theta')).
\end{aligned}$$

Proposition 7 *Under Assumption 3, if*

$$\Delta_\theta(\theta') \geq LW(\mathcal{D}(\theta), \mathcal{D}(\theta')),$$

it holds that

$$\mathrm{PR}(\theta) \leq \mathrm{PR}(\theta').$$

We give examples on relating PO with PS

Example 1 *Assume performative shifts are bounded by an absolute value, i.e.,*

$$W(\mathcal{D}(\theta), \mathcal{D}(\theta')) \leq B.$$

We have the following bound characterizing the suboptimality of θ

$$\mathrm{PR}(\theta) - \mathrm{PR}^* \leq LB.$$

Example 2 *Assume (1) performative shifts are bounded by an absolute value* B *and (2)* $\Delta_\theta(\theta')$ *satisfies quadratic growth, i.e.,* $\Delta_\theta(\theta') \geq \gamma\|\theta - \theta'\|^2$, *we have that performative optimal point* θ_{PO} *satisfies*

$$\|\theta_{\mathrm{PO}} - \theta\| \leq \sqrt{\frac{LB}{\gamma}}.$$

Example 3 *Under Assumption 4, suppose* $\Delta_\theta(\theta')$ *satisfies quadratic growth, the performative optimal point* θ_{PO} *satisfies*

$$\|\theta_{\mathrm{PO}} - \theta\| \leq \frac{L\epsilon}{\gamma}.$$

Open Problems

- Showing PL condition for PR.
- Understanding when and how (e.g., some structural properties of loss function or a natural set of distributions), it holds that

$$W\big(D(\theta), D(\theta')\big) \leq C \,||\, \nabla_{\theta'} DPR(\theta, \theta')||^2$$

Such a condition characterizes local properties of DPR near performative stable points, it could be more common.

- What is the impact of data pre-processing steps on the implications of performative shifts?

Contributions

I. Studied several first-order conditions (Weak Strong Convexity, Restricted Secant Inequality) for performative prediction and what structural assumptions are needed.

II. We investigate relations between stable points and optimal points.

III. We raise interesting open problems in the area of performative prediction.