Horizon-Free Minimax Optimal Online Linear Regression

Alan Malek   Peter Bartlett

Model: Online Linear Regression

Given: covariate set $\mathcal{X} \in \mathbb{R}^d$, label set $\mathcal{Y} \in \mathbb{R}$.
On each round $t = 1, 2, \dots,$

- Environment reveals $\boldsymbol{x}_t \in \mathcal{X}(\boldsymbol{x}_1^{t-1})$
- We play $\hat{y}_t \in \mathbb{R}$
- Environment reveals true label $y_t \in \mathcal{Y}(\boldsymbol{x}_1^t)$
- We incur loss $(y_t - \hat{y}_t)^2$
- Environment may continue or end the game

Environment may be adversarial; we cannot control loss, but we can control the *Regret*,

$$\mathcal{R} := \underbrace{\sum_{t \geq 1} (\hat{y}_t - y_t)^2}_{\text{Our Loss}} - \underbrace{\min_{\theta \in \mathbb{R}^d} \sum_{t \geq 1} \left(\theta^\top \boldsymbol{x}_t - y_t\right)^2}_{\text{Best Linear Predictor}}$$

In 1-D

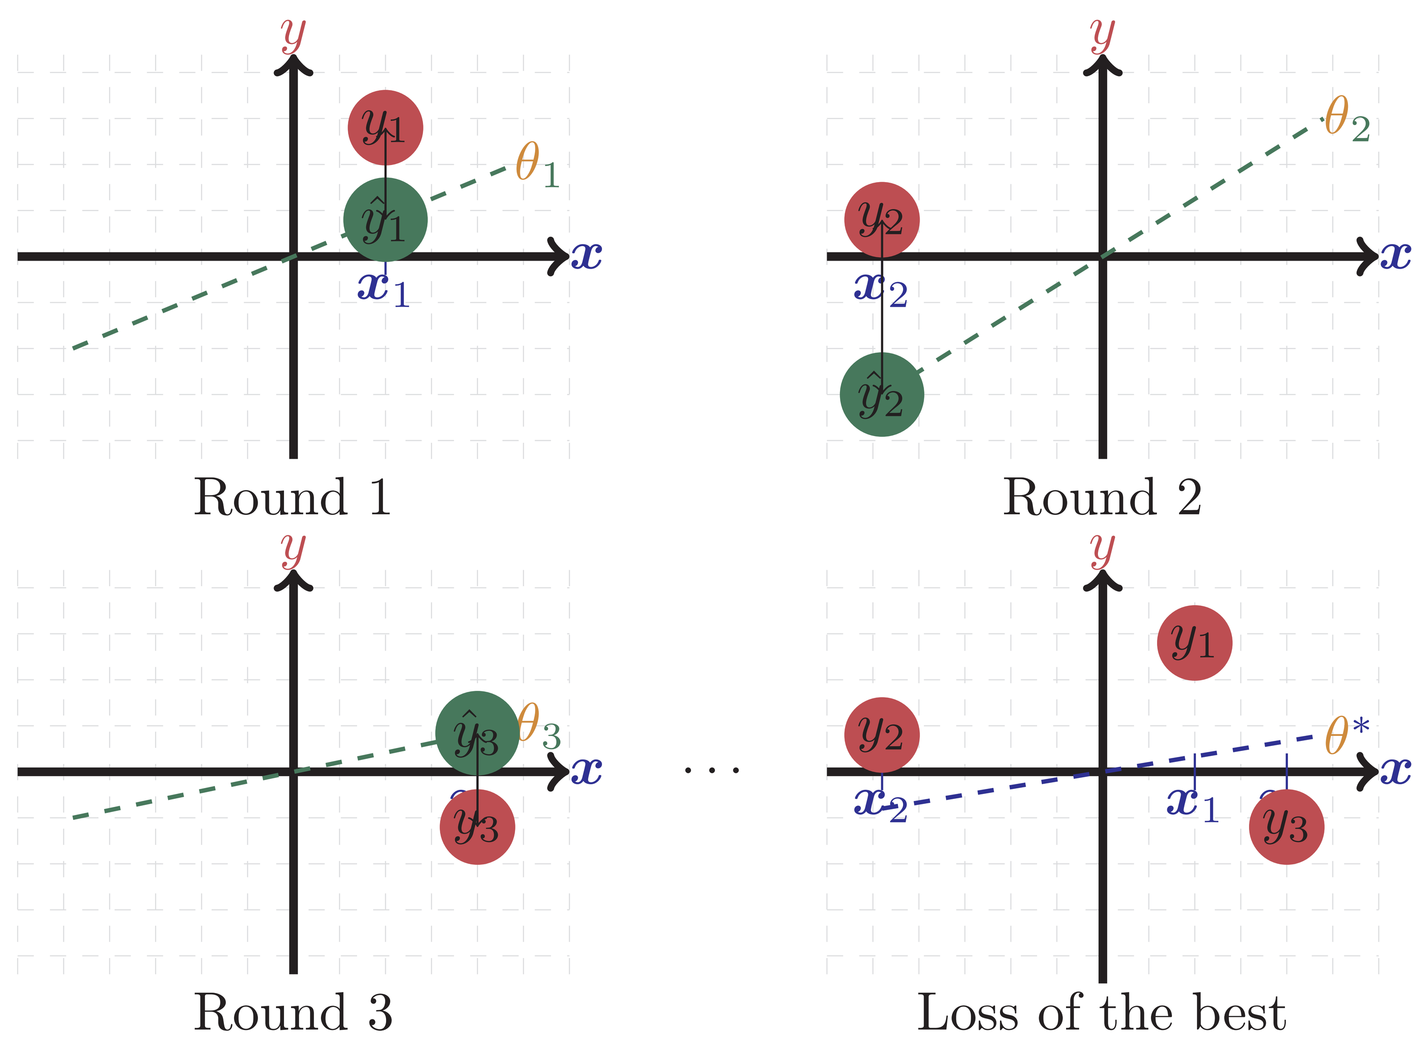

Minimax Regret

The best we can do against the worst case adversary is

$$\left\langle \max_{\boldsymbol{x}_t} \min_{\hat{y}_t} \max_{y_t} \right\rangle_{t=1}^{T} \sum_{t=1}^T (\hat{y}_t - y_t)^2 - \min_{\theta \in \mathbb{R}^d} \sum_{t=1}^T (\theta^\top \boldsymbol{x}_t - y_t)^2,$$

calculated by the backward induction (for known T)

$$V_T(\boldsymbol{x}_1^T, y_1^T) := -\min_{\theta} \sum_{t=1}^T (\theta^\top \boldsymbol{x}_t - y_t)^2$$

$$V_{t-1}(\boldsymbol{x}_1^{t-1}, y_1^{t-1}) := \max_{\boldsymbol{x}_t} \min_{\hat{y}_t} \max_{y_t} (y_t - \hat{y}_t)^2 + V_t(\boldsymbol{x}_1^t, y_1^t)$$

Fixed Design Case [previous work]

Theorem 1 *Assume that $\boldsymbol{x}_1^T$ is fixed and let B_1^T be a sequence of label bounds. If $\boldsymbol{x}_1^T \in \mathcal{B}(B_1^T) := \left\{ \boldsymbol{x}_1^T : B_t \geq \sum_{s=1}^{t-1} \left|\boldsymbol{x}_t^\top \boldsymbol{P}_t \boldsymbol{x}_s\right| B_s \text{ for } 2 \leq t \right\}$ and $y_1^T \in \mathcal{L}(B_1^T) = \{y_t : |y_t| \leq B_t\}$, then the minimax strategy calculates states $\boldsymbol{s}_t = \sum_{s=1}^t y_s \boldsymbol{x}_s$ and $\Pi_t = \sum_{s=1}^t \boldsymbol{x}_s \boldsymbol{x}_s^\top$ and plays*

$$\hat{y}_{t+1} = \boldsymbol{x}_{t+1}^\top \boldsymbol{P}_{t+1} \boldsymbol{s}_t \qquad \text{(MMS)}$$

simple linear predictor

with coefficient matrices defined by

$$\boldsymbol{P}_T = \Pi_T^\dagger \text{ and recursion } \boldsymbol{P}_t = \boldsymbol{P}_{t+1} + \boldsymbol{P}_{t+1} \boldsymbol{x}_{t+1} \boldsymbol{x}_{t+1}^\top \boldsymbol{P}_{t+1}.$$

- Can compute $V_t(y_1^t)$ by backwards induction efficiently
- Applying Sherman-Morrison yields

$$\boldsymbol{P}_t^\dagger = \Pi_t + \sum_{s=t+1}^T \frac{\boldsymbol{x}_s^\top \boldsymbol{P}_s \boldsymbol{x}_s}{1 + \boldsymbol{x}_s^\top \boldsymbol{P}_s \boldsymbol{x}_s} \boldsymbol{x}_s \boldsymbol{x}_s^\top \succeq \Pi_t$$

Forward Recursion

- How can we generalize to adversarial design? Looks hard:
  - For fixed-design, every $\boldsymbol{P}_t$ depends on all $\boldsymbol{x}_1^T$
  - Full backwards induction with $\max_{\boldsymbol{x}_t}$ is intractable
- Key observation: we can invert the recursion. Given base case $\boldsymbol{\Sigma}_0 = \boldsymbol{P}_0^\dagger$, define

$$\boldsymbol{P}_t := \boldsymbol{P}_{t-1} - \frac{a_t}{b_t^2} \boldsymbol{P}_{t-1} \boldsymbol{x}_t \boldsymbol{x}_t^\top \boldsymbol{P}_{t-1} \qquad (1)$$

for $b_t^2 := \boldsymbol{x}_t^\top \boldsymbol{P}_{t-1} \boldsymbol{x}_t$, $a_t := \left(\sqrt{4b_t^2 + 1} - 1\right)\left(\sqrt{4b_t^2 + 1} + 1\right)^{-1}$

- Treat $\boldsymbol{\Sigma}$ as a *covariate budget* for the adversary:

$$\mathcal{A}(\boldsymbol{\Sigma}) := \left\{ \boldsymbol{x}_1^T : \text{for } \boldsymbol{P}_0, \dots, \boldsymbol{P}_T \text{ defined by (1), } \boldsymbol{P}_0^\dagger \preceq \boldsymbol{\Sigma} \right\}$$
$$= \left\{ \boldsymbol{x}_1^T : \boldsymbol{P}_0^\dagger = \boldsymbol{\Sigma} \text{ and } \boldsymbol{P}_t^\dagger \succeq \Pi_t \ \forall t \geq 1 \right\}$$

- If the adversary plays $\boldsymbol{x}_1^T$ with $\boldsymbol{P}_0(\boldsymbol{x}_1^T)^\dagger = \boldsymbol{\Sigma}$, then we responded optimally!

Hold Your Horses

Lemma 1 *Fix any $\boldsymbol{\Sigma}$ and any $\{B_t\}$ with $B_t \geq b > 0$ for all t. Then, for any $M > 0$, there exists $\boldsymbol{x}_1^T \in \mathcal{A}(\boldsymbol{\Sigma}) \cap \mathcal{B}(B_t, \boldsymbol{\Sigma})$ and $y_1^T \in \mathcal{L}(B_t)$ such that the minimax regret is larger than M.*

$\mathcal{B}$ using $\boldsymbol{P}_t$ defined by (1)

- In general, constraining $\boldsymbol{x}_1^T \in \mathcal{A}(\boldsymbol{\Sigma})$ is *not* sufficient
- For a finite sequence, the regret of $\boldsymbol{x}_1^T$ can be $O(\log(T))$

Expanded conditions

- Define label budget $\gamma_0 > 0$ with update $\gamma_t = \gamma_{t-1} - B_t^2 \boldsymbol{x}_t^\top \boldsymbol{P}_t \boldsymbol{x}_t$
- Define the *continuation conditions*

$$\mathcal{C}(\boldsymbol{\Sigma}, \gamma_0) := \left\{ \boldsymbol{x}_1^T : \boldsymbol{s}_t^\top \left(\Pi_t^\dagger - \boldsymbol{P}_t\right) \boldsymbol{s}_t \leq \gamma_t \ \forall t \geq 0, \forall \boldsymbol{s}_t \right\},$$

where $\boldsymbol{s}_t$ ranges over all $\boldsymbol{x}_1^T \in \mathcal{A}(\boldsymbol{\Sigma}) \cap \mathcal{B}(B_t, \boldsymbol{\Sigma})$ and $y_1^T \in \mathcal{L}(B_t)$

Theorem 2 *For any $\{B_t\} > 0$, $\boldsymbol{\Sigma} \succ 0$ and $\gamma_0 \geq 0$, (MMS) with $\boldsymbol{P}_t$ defined by (1) has minimax regret γ_0 and is horizon-independent minimax optimal; that is,*

over all strategies

$$\sup_{T} \left(\sup_{\boldsymbol{x}_1^T, y_1^T} R_T((\text{MMS}), \boldsymbol{x}_1^T, y_1^T) - \min_{s} \sup_{\boldsymbol{x}_1^T, y_1^T} R_T(s, \boldsymbol{x}_1^T, y_1^T) \right) = 0,$$

over $\boldsymbol{x}_1^T \in \mathcal{A}(\boldsymbol{\Sigma}) \cap \mathcal{B}(B_t) \cap \mathcal{C}(\boldsymbol{\Sigma}, \gamma_0)$ and $y_1^T \in \mathcal{L}(B_t)$.

- Removing $\mathcal{A}$, $\mathcal{B}$, or $\mathcal{C}$ lets the adversary cause infinite regret
- We can compete with strategies that know T and γ_0
- The game-length measure T is replaced with the more natural measure $\boldsymbol{\Sigma}$

Proof Intuition

- The proof defines an "early stopping game", where $\boldsymbol{x}_1^T$ are fixed but the adversary can stop any round
- We calculate the difference in regret between stopping at t and continuing until T
- Shorter games may cause more regret because (MMS) is over-regularizing
- The $\mathcal{C}$ condition holds when the adversary always causes more regret by continuing
- Under $\mathcal{C}$, the adversary wants to play out the budget, which is the setting where (MMS) is optimal

FTRL

(MMS) is *Follow the Regularized Leader* that plays

$$\hat{y}_t = \hat{\theta}_t^\top \boldsymbol{x}_t \text{ where } \hat{\theta}_t := \min_{\theta} \sum_{s=1}^{t-1} (\theta^\top \boldsymbol{x}_s - y_s)^2 + \theta^\top \boldsymbol{R}_t \theta$$

with $\boldsymbol{R}_0 := \boldsymbol{\Sigma}^{-1}$ and $\boldsymbol{R}_t := \boldsymbol{R}_{t-1} + \frac{1}{1 + \boldsymbol{x}_t^\top \boldsymbol{P}_t \boldsymbol{x}_t} \boldsymbol{x}_t \boldsymbol{x}_t^\top - \boldsymbol{x}_{t-1} \boldsymbol{x}_{t-1}^\top$

Comparison with other methods:

1. Ridge regression: $\boldsymbol{R}_t = \lambda_t \boldsymbol{I}$
2. Last-step-minimax: $\boldsymbol{R}_t = \boldsymbol{x}_t \boldsymbol{x}_t^\top$
3. OLS: $\boldsymbol{R}_t = 0$