Minimax Fixed-Design Linear Regression

Peter L. Bartlett, Wouter M. Koolen, **Alan Malek**,
Eiji Takimoto, Manfred Warmuth

Berkeley UNIVERSITY OF CALIFORNIA
QUT Queensland University of Technology
KYUSHU UNIVERSITY
UNIVERSITY OF CALIFORNIA SANTA CRUZ

Conference on Learning Theory
Paris, France
July 5th, 2015

Context: Linear regression

- We have data $(\boldsymbol{x}_1, y_1), \ldots, (\boldsymbol{x}_T, y_T)$
- Offline linear regression: predict $\hat{y} = \theta^\intercal \boldsymbol{x}$, where

$$\theta = (X^\intercal X)^{-1} X^\intercal Y.$$

Context: Linear regression

- We have data $(\boldsymbol{x}_1, y_1), \ldots, (\boldsymbol{x}_T, y_T)$
- Offline linear regression: predict $\hat{y} = \theta^\intercal \boldsymbol{x}$, where

$$\theta = (X^\intercal X)^{-1} X^\intercal Y.$$

- Online fixed-design linear regression:
  1. Covariates $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_T$ are fixed at the start
  2. Need to predict $\hat{y}_t$ before seeing y_t

Protocol

Given: $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_T \in \mathbb{R}^d$
For $t = 1, 2, \ldots, T$:

- Learner predicts $\hat{y}_t \in \mathbb{R}$,
- Adversary reveals $y_t \in \mathbb{R}$,
- Learner incurs loss $(\hat{y}_t - y_t)^2$.

Figure: Fixed-design protocol

Minimax

Our goal is to find a strategy that achieves the minimax regret:

$$\min_{\hat{y}_1} \max_{y_1} \cdots \min_{\hat{y}_T} \max_{y_T} \sum_{t=1}^{T} (\hat{y}_t - y_t)^2 - \min_{\theta \in \mathbb{R}^d} \sum_{t=1}^{T} (\theta^\intercal \boldsymbol{x}_t - y_t)^2$$

Minimax

Our goal is to find a strategy that achieves the minimax regret:

$$\min_{\hat{y}_1} \max_{y_1} \cdots \min_{\hat{y}_T} \max_{y_T} \underbrace{\sum_{t=1}^{T} (\hat{y}_t - y_t)^2}_{\text{algorithm}} - \min_{\theta \in \mathbb{R}^d} \sum_{t=1}^{T} (\theta^\intercal \boldsymbol{x}_t - y_t)^2$$

Minimax

Our goal is to find a strategy that achieves the minimax regret:

$$\min_{\hat{y}_1} \max_{y_1} \cdots \min_{\hat{y}_T} \max_{y_T} \underbrace{\sum_{t=1}^{T} (\hat{y}_t - y_t)^2}_{\text{algorithm}} - \underbrace{\min_{\theta \in \mathbb{R}^d} \sum_{t=1}^{T} (\theta^\intercal \boldsymbol{x}_t - y_t)^2}_{\text{best linear predictor}}$$

The Minimax Strategy

- Is linear

$$\hat{y}_t = \boldsymbol{s}_{t-1}^\intercal \boldsymbol{P}_t \boldsymbol{x}_t \qquad \text{where} \qquad \boldsymbol{s}_t = \sum_{q=1}^{t} \boldsymbol{x}_q y_q,$$

- with coefficients:

$$\boldsymbol{P}_t^{-1} = \sum_{q=1}^{t} \boldsymbol{x}_q \boldsymbol{x}_q^\intercal + \sum_{q=t+1}^{T} \frac{\boldsymbol{x}_q^\intercal \boldsymbol{P}_q \boldsymbol{x}_q}{1 + \boldsymbol{x}_q^\intercal \boldsymbol{P}_q \boldsymbol{x}_q} \boldsymbol{x}_q \boldsymbol{x}_q^\intercal .$$

- Cheap recursive calculation, can be done before seeing y_ts.
- Minimax under alignment condition and $|y_t| \leq B$

The Minimax Strategy

- Is linear

$$\hat{y}_t \;=\; \boldsymbol{s}_{t-1}^\intercal \boldsymbol{P}_t \boldsymbol{x}_t \qquad \text{where} \qquad \boldsymbol{s}_t \;=\; \sum_{q=1}^{t} \boldsymbol{x}_q y_q,$$

- with coefficients:

$$\boldsymbol{P}_t^{-1} \;=\; \underbrace{\sum_{q=1}^{t} \boldsymbol{x}_q \boldsymbol{x}_q^\intercal}_{\text{least squares}} + \underbrace{\sum_{q=t+1}^{T} \frac{\boldsymbol{x}_q^\intercal \boldsymbol{P}_q \boldsymbol{x}_q}{1 + \boldsymbol{x}_q^\intercal \boldsymbol{P}_q \boldsymbol{x}_q} \boldsymbol{x}_q \boldsymbol{x}_q^\intercal}_{\text{re-weighted future instances}}.$$

- Cheap recursive calculation, can be done before seeing y_ts.
- Minimax under alignment condition and $|y_t| \leq B$

Guarantees

- If the adversary plays y_t with

$$\sum_{t=1}^{T} y_t^2 \boldsymbol{x}_t^\intercal \boldsymbol{P}_t \boldsymbol{x}_t \;=\; R,$$

  we are minimax against all y_ts in this set
- Explains re-weighting:

$$\boldsymbol{P}_t^{-1} \;=\; \sum_{q=1}^{t} \boldsymbol{x}_q \boldsymbol{x}_q^\intercal + \sum_{q=t+1}^{T} \underbrace{\frac{\boldsymbol{x}_q^\intercal \boldsymbol{P}_q \boldsymbol{x}_q}{1 + \boldsymbol{x}_q^\intercal \boldsymbol{P}_q \boldsymbol{x}_q}}_{\text{future regret potential}} \boldsymbol{x}_q \boldsymbol{x}_q^\intercal$$

- Minimax strategy does not depend on R
- We achieve regret exactly $R = O(\log T)$

Guarantees

- If the adversary plays y_t with

$$\sum_{t=1}^{T} y_t^2 \boldsymbol{x}_t^\intercal \boldsymbol{P}_t \boldsymbol{x}_t \;=\; R,$$

  we are minimax against all y_ts in this set
- Explains re-weighting:

$$\boldsymbol{P}_t^{-1} \;=\; \sum_{q=1}^{t} \boldsymbol{x}_q \boldsymbol{x}_q^\intercal + \sum_{q=t+1}^{T} \underbrace{\frac{\boldsymbol{x}_q^\intercal \boldsymbol{P}_q \boldsymbol{x}_q}{1 + \boldsymbol{x}_q^\intercal \boldsymbol{P}_q \boldsymbol{x}_q}}_{\text{future regret potential}} \; \boldsymbol{x}_q \boldsymbol{x}_q^\intercal$$

- Minimax strategy does not depend on R
- We achieve regret exactly $R = O(\log T)$
- Thanks!