High-dimensional (Group) Adversarial Training in Linear Regression

Yiling Xie, Xiaoming Huo

School of Industrial and Systems Engineering
Georgia Institute of Technology

Adversarial Training

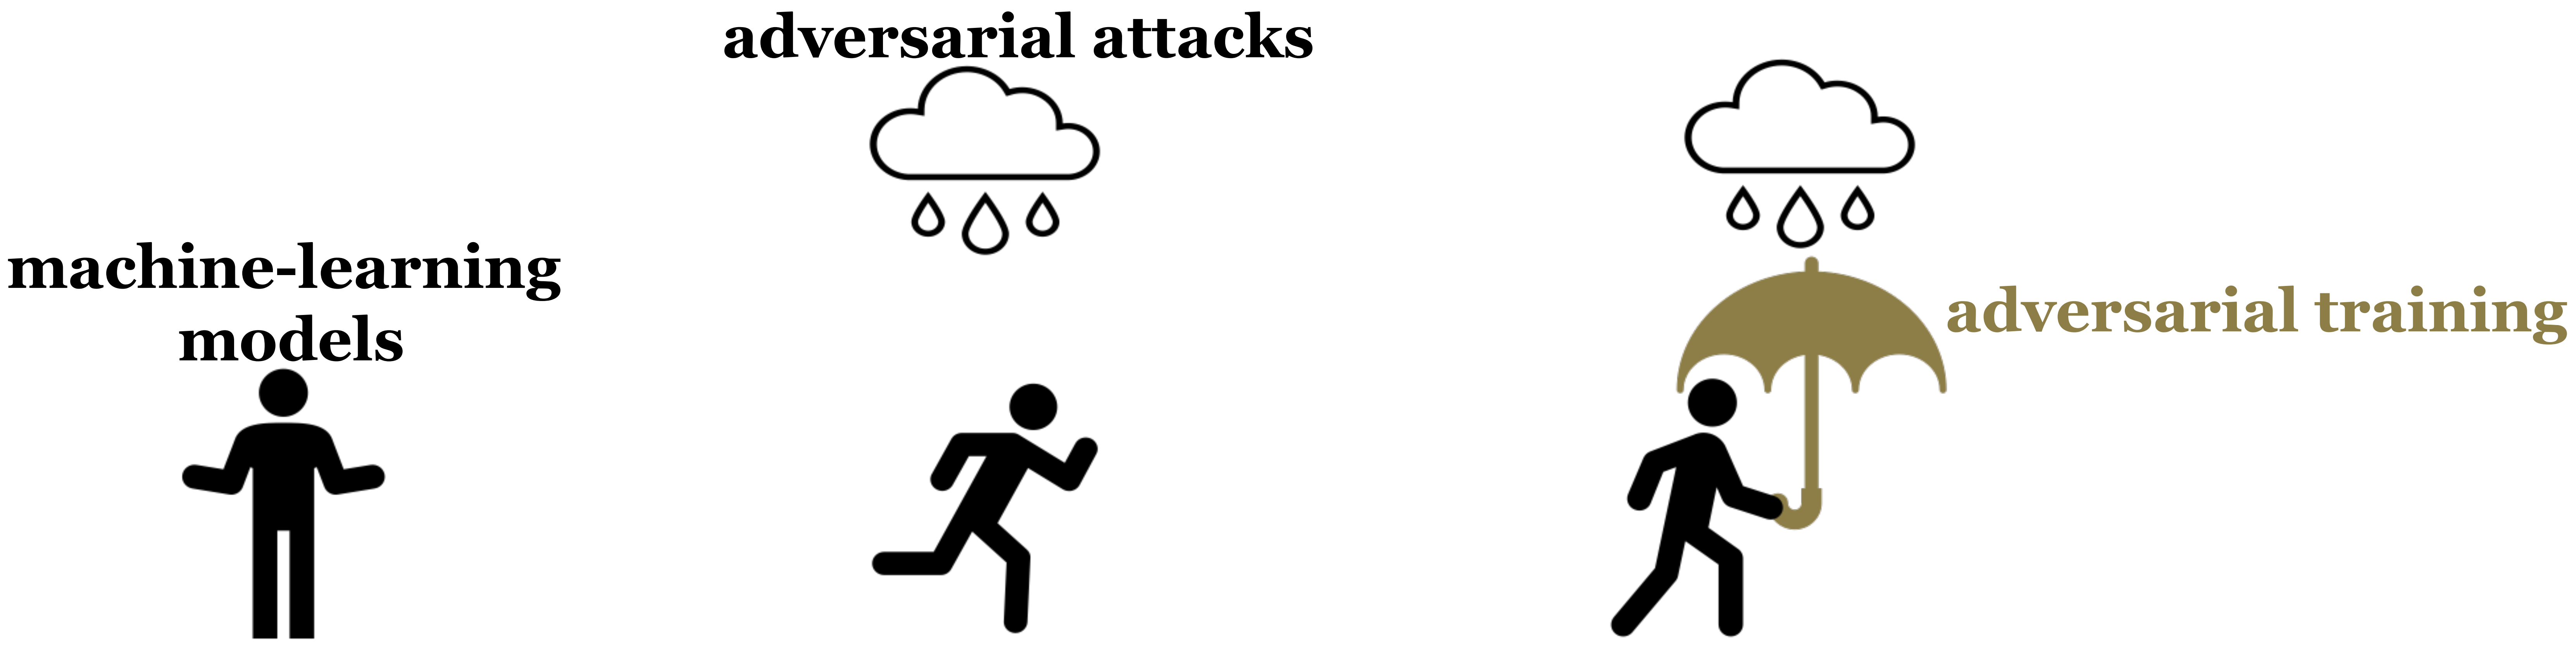

Empirical Risk Minimization

$$\min_{\beta} \frac{1}{n} \sum_{i=1}^{n} L\left(X_i, Y_i, \beta\right)$$

Adversarial Training

$$\min_{\beta} \frac{1}{n} \sum_{i=1}^{n} \underbrace{\sup_{\|\Delta\| \leq \delta} L\left(X_i + \Delta, Y_i, \beta\right)}_{\text{worst-case loss}}$$

Non-asymptotic Convergence Rate

1. **Minimax Optimality**
2. Group Adversarial Training

Adversarial Training in Linear Regression

High-dimensional adversarially-trained linear regression under ℓ_∞-perturbation

$$\beta^n \in \arg\min_{\beta} \frac{1}{n} \sum_{i=1}^{n} \sup_{\|\Delta\|_\infty \leq \delta_n} \left((X_i + \Delta)^\top \beta - Y_i\right)^2$$

square loss

ℓ_∞-perturbation

High-dimension: parameter β has p dimensions, $p > n$
Sparsity: s dimensions of the ground-truth β_* are nonzero, $p > s$

Non-asymptotic Convergence Rate

$$\beta^n \in \arg\min_{\beta} \frac{1}{n} \sum_{i=1}^{n} \sup_{\|\Delta\|_\infty \leq \delta_n} \left((X_i + \Delta)^\top \beta - Y_i \right)^2$$

Under certain conditions, then the following holds with a high probability:

$$\text{PredictionError}(\beta^n) = \mathcal{O}\left(\frac{s \log p}{n} \right)$$

matches **minimax lower bound**
of the prediction error in linear regression

minimax optimal in sparse high-dimensional linear regression

Non-asymptotic Convergence Rate

1. Minimax Optimality
2. Group Adversarial Training

Group Adversarial Training

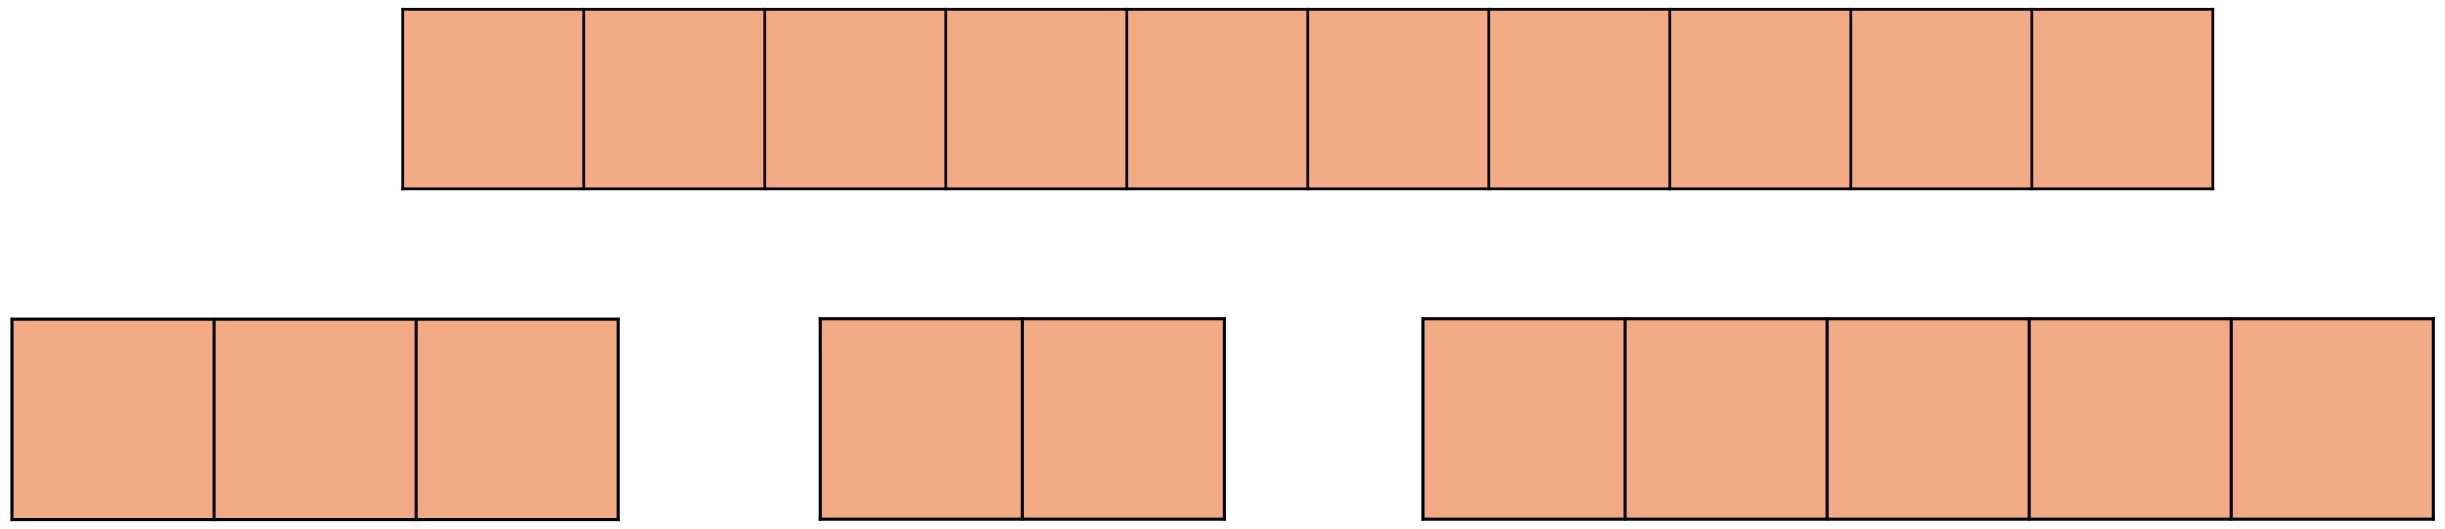

$$\min_{\beta} \frac{1}{n} \sum_{i=1}^{n} \sup_{\|\Delta_{\boldsymbol{\omega}}\|_{2,\infty} \leq \delta_n} \left((X_i + \Delta)^\top \beta - Y_i \right)^2$$

$$\Delta = (\Delta^1, ..., \Delta^L)$$

$$\boldsymbol{\omega} = (\omega^1, ..., \omega^L)$$

$$\|\Delta_{\boldsymbol{\omega}}\|_{2,\infty} = \max_{1 \leq l \leq L} \|\omega_l \Delta^l\|_2$$

Group Adversarial Training Improvement

$$\beta^n \in \arg\min_{\beta} \frac{1}{n} \sum_{i=1}^{n} \sup_{\|\Delta\|_\infty \le \delta_n} \left((X_i + \Delta)^\top \beta - Y_i\right)^2$$

$$\widehat{\beta}^n \in \arg\min_{\beta} \frac{1}{n} \sum_{i=1}^{n} \sup_{\|\Delta_{\boldsymbol{\omega}}\|_{2,\infty} \le \delta_n} \left((X_i + \Delta)^\top \beta - Y_i\right)^2$$

Under certain conditions and group assumption, then the following holds with a high probability:

$$\text{PredictionError}(\beta^n) > \text{PredictionError}(\widehat{\beta}^n)$$

Georgia Institute of Technology

Thanks for your attention!