Gaming Helps! Learning from Strategic Interactions in Natural Dynamics

Yahav Bechavod Hebrew University
Katrina Ligett Hebrew University
Zhiwei Steven Wu Carnegie Mellon University
Juba Ziani University of Pennsylvania

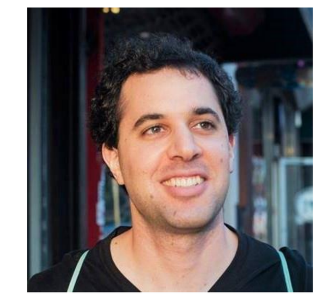
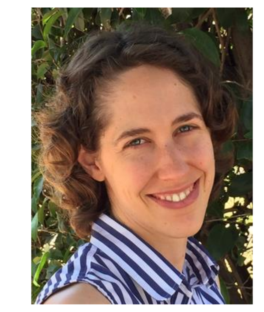

Gaming

Strategic modification of measurements, which individuals anticipate would positively affect the outcome of a decision rule.

Examples: College admissions, Credit, Insurance, Hiring, …

Machine learning algorithms are now heavily involved.

Problem: Feature modifications might make individuals appear better than they actually are.

Approaches in prior work:

1. Obfuscation of decision rule.
   - May leak over time.
   - Individuals can learn from past examples.
2. Robustness to gaming.
   - Additional burden on qualified individuals.
   - Cripples ability to recover or improve.

Our Approach

Gaming can be actually be **helpful!**

Idea: Distinguish false feature manipulation from improvement.

Manipulation Vs. Improvement

Manipulation
Obtain additional credit cards
Raise your credit limits
…

Improvement
Reduce your debt
Increase your income
…

Model

Online. Linear Regression.

Underlying model: $y = {\beta^*}^T x + \epsilon$

True Label, True regressor, Features, Noise

Meaningful Vs. Non-meaningful features:

$\beta^* \in R^d$

Meaningful: ${\beta^*}_i \neq 0$

Non-Meaningful: ${\beta^*}_i = 0$

Algorithms

Algorithm 1: Online Regression with Epoch-Based Strategic modification (Epoch size n)

Learner picks (any) initial $\hat{\beta}_0$.
for *every epoch* $E \in \mathbb{N}$ **do**
  for $t \in \{(E-1)n+1, \ldots, En\}$ **do**
    Agent t reports $\bar{x}_t \in M(\hat{\beta}_{E-1}, c_t, B_t)$.
    Learner observes $\bar{y}_t = {\beta^*}^\top \bar{x}_t + \varepsilon_t$.
  end
  Learner picks $\hat{\beta}_E \in LSE(\tau(E))$.
end

Tie-breaking matters! (Intuition for Algorithm 2)

Weight on meaningful features
➔ Incentivize improvements of said features
➔ May prevent exploration of remaining features

Solution: put weight in unexplored directions:

- Retains accuracy on directions seen so far
- Incentivizes exploration of unseen directions

Motivation

If distribution over X is not full-rank, recovery of β^* is **impossible**.

Optimizing for $\hat{\beta}$ over a rank-deficient space implies:

1. Non-zero weight on non-meaningful features -> Susceptibility to false manipulations.
2. Less weight on meaningful features -> Reduced utility.

Results

Our provided algorithm + tie-breaking scheme guarantee:

1. Recovery of the true underlying model $\hat{\beta}$.
2. Achieving recovery within the confinements of **natural dynamics**. At any point, deployed scoring rule projected to the recovered subspace is optimal.

Theorem 5.2 (Recovery Guarantee with Tie-Breaking Scheme (Algorithm 2)). *Suppose the epoch size satisfies* $n \geq \frac{\kappa d^2}{\lambda}\sqrt{2T\log(24d/\delta)}$*, and take* α *to be*

$$\alpha \geq \gamma\left(\sqrt{d} + \frac{Kd\sqrt{2T\log(8d/\delta)}}{\lambda n}\right),$$

where γ, K, κ, λ *are instance-specific constants that only depend on* σ, $\mathcal{C}$, Σ, *and* $\lambda > 0$. *If* $T \geq dn$, *we have with probability at least* $1-\delta$ *that at the end of the last epoch* T/n,

$$\left\|\hat{\beta}_{T/n} - \beta^*\right\|_2 \leq \frac{K\sqrt{2dT\log(8d/\delta)}}{\lambda n},$$

under the tie-breaking rule of Algorithm 2 .

Acknowledgements + Disclosure of Funding

Part of this work was done while the authors were visiting the Simons Institute for the Theory of Computing.The work of Yahav Bechavod and Katrina Ligett was supported in part by Israel Science Foundation (ISF) grants \#1044/16 and 2861/20, the United States Air Force and DARPA under contracts FA8750-16-C-0022 and FA8750-19-2-0222, and the Federmann Cyber Security Center in conjunction with the Israel national cyber directorate. Yahav Bechavod was also supported in part by the Apple Scholars in AI/ML PhD Fellowship. Katrina Ligett was also funded in part by in part by a grant from Georgetown University and Simons Foundation Collaboration 733792. Zhiwei Steven Wu was supported in part by the NSF FAI Award \#1939606, a Google Faculty Research Award, a J.P. Morgan Faculty Award, a Facebook Research Award, and a Mozilla Research Grant. Juba Ziani was supported in part by the Inaugural PIMCO Graduate Fellowship at Caltech, the National Science Foundation through grant CNS-1518941, as well as the Warren Center for Network and Data Sciences at the University of Pennsylvania.Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Air Force and DARPA. We thank Mohammad Fereydounian and Aaron Roth for useful discussions.