Why Adaptively Collected Data Have *Negative* Bias and How to Correct for It

Stanford University

Xinkun Nie, Xiaoying Tian, Jonathan Taylor, James Zou

Abstract

Collecting data adaptively is ubiquitous, from clinical trials to online A/B testing. Such adaptivity introduces complex correlations between the data and the collection procedure.

We prove that when the data collection procedure satisfies natural conditions, the sample means of the data have systematic negative biases. This would imply, for example, the naïve empirical mean estimator can *underestimate* the true effects of a drug in a clinical trial setting.

We also propose a novel debiasing algorithm based on selective inference techniques. In experiments, our method can effectively reduce bias and estimation error.

Motivating Example: Greedy

$X \sim Bernoulli(0.5)$

$Y \sim Bernoulli(0.5)$

** Break ties by sampling X*

t=1	Sample X	Observe: 1	t=1	Sample X	Observe: 0
t=2	Sample Y		t=2	Sample Y	
t=3	Sample X with p=1		t=3	Sample X with p=0.5	

Empirical mean of X: 0.75 — *Reversion to the mean*

Empirical mean of X: 0.125 — *Stuck at unlucky values*

In expectation, E[empirical mean] < true mean!

Greedy

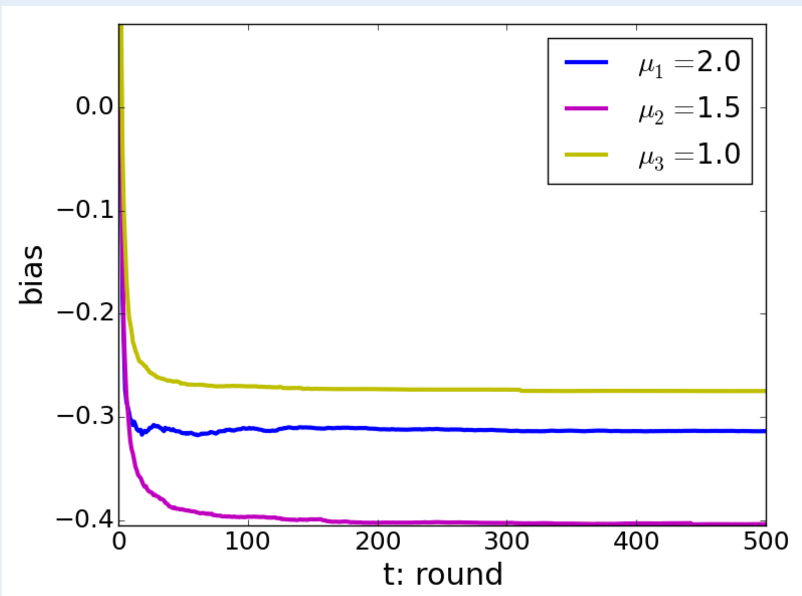

lil' UCB
(Jamieson et al., 2014)

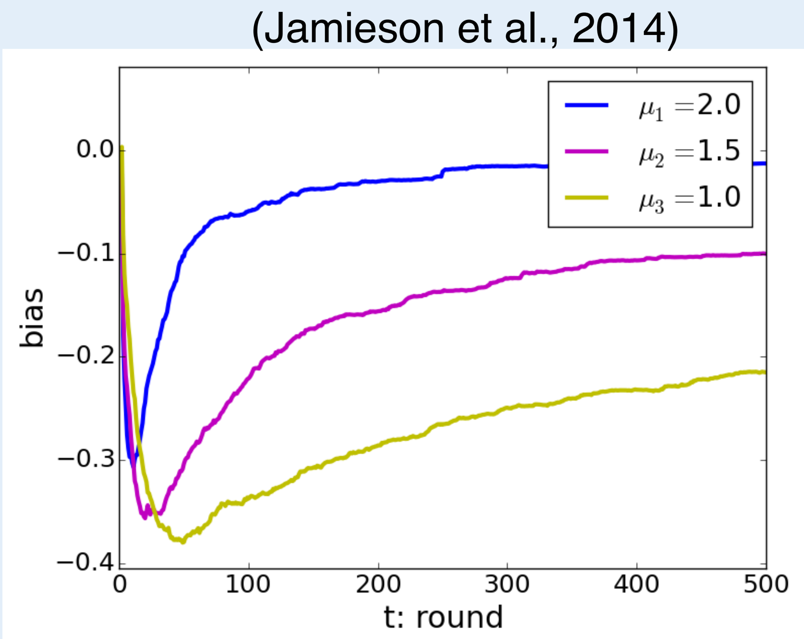

Main Theorem (informal)

If the selection function f satisfies *Exploit* and *Independent of Irrelevant Options (IIO)*, then $\forall k$ and given horizon T,

$$E[\bar{X}_T^{(k)}] \le \mu_k.$$

Empirical Mean ↑ True Mean ↑

*Equality holds iff number of times a distribution is sampled is *independent* from observed values

1st Condition: *Exploit*

Higher empirical mean ⇒ More likely to be selected*

* Conditioning on # pulls, the realization of all other arms

***Example**: Greedy* *Break ties by sampling X*

t=1	Sample X	1	t=1	Sample X	0
t=2	Sample Y	0	t=2	Sample Y	0
t=3	Sample X ⇐		t=3	Sample X	

2nd Condition: *Independent of Irrelevant Options(IIO)*

Conditioning on not selecting distr. *K* at t

⇓

Which distr. to select past samples from distr. *k*

Proposition (informal)

Standard bandit algorithms including Greedy, -greedy, lil-UCB, and Thompson Sampling* all satisfy *Exploit* and *Independent of Irrelevant Options (IIO)*, and thus exhibit the negative bias.

*under natural prior conditions

Correcting for the Bias

Data Splitting (naïve approach)

At any given time, select some *k*

X ⇒ Used by selection alg.

X' ⇒ Used for *estimation*

Unbiased ☺
High variance ☹

We are discarding **half** of the data!

Cox, 1975, Sladek et al., 2007, Meinshausen, Meier, and Bühlmann, 2009, Wasserman and Roeder, 2009, Fithian, Sun, and Taylor, 2014 Xu, Qin, and Liu, 2013

Conditional Maximum Likelihood Estimator (**our approach**)

$$\arg\max \Pr_\theta[\text{ sample values } \mid \text{ selection history }]$$

conditional likelihood

$$\propto \Pr_\theta[\text{ sample values }] \times \prod_{t=1}^{T} \Pr[\text{ choice(t) } \mid \text{ history(t) }]$$

original likelihood — selection likelihood

Adding Gumbel noise to all samples before selection ⇒ Asymptotic consistency using conditional MLE

Tian & Taylor 2015
Panigrahi et al 2016

⇓

We can solve conditional MLE efficiently!

Our Algorithm

① At each timestep, add i.i.d. Gumbel noise to samples

② At the end of a data collection trial, use MLE conditioned on the selection history

Experiments

5 Gaussian Distributions, repeated 1000 trials

Bias Reduction

	lil' UCB		ϵ-Greedy	
	orig.	cMLE	orig.	cMLE
T=20	-0.32	**14.9%**	-0.31	**9.1%**
T=40	-0.35	**14.2%**	-0.27	**8.8%**

bias in original algorithm — % bias compared to original

"held": data splitting
"prop": propensity weighting

Mean Squared Error (MSE) Reduction

	lil' UCB			$\epsilon-$Greedy			
	orig.	held	cMLE	orig.	held	prop	cMLE
T=20	0.57	112%	**99%**	0.52	123%	401%	**94%**
T=40	0.54	104%	**52%**	0.39	135%	312%	**62%**

MSE in original algorithm — % MSE compared to original

Summary

- Rigorously proved negative bias of sample mean in adaptive data collection under two natural conditions: *Exploit* and *Independent of Irrelevant Options (IIO)*
- A wide range of applications (e.g. A/B testing, online search advertising, science experiments, etc.) manifest this negative bias
- Proposed an algorithm conditional MLE that achieves asymptotic consistency and lowers MSE empirically