Beyond the Best: Distribution Functional Estimation in Infinite-Armed Bandits

Yifei Wang

Department of Electrical Engineering, Stanford University

Nov. 29th, Neurips 2022

Yifei Wang (Stanford)

Tavor Baharav (Stanford)

Yanjun Han (MIT)

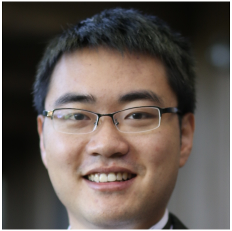

Jiantao Jiao (UC Berkeley)

David Tse (Stanford)

Setting

- Goal: estimate some functional $g(F)$ of an underlying distribution F.

Setting

- Goal: estimate some functional $g(F)$ of an underlying distribution F.
- Setting:

$$\boxed{F}$$

  - Level 0: distribution F

Setting

- Goal: estimate some functional $g(F)$ of an underlying distribution F.
- Setting:

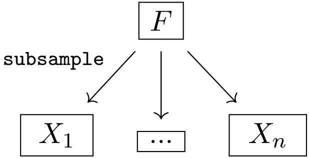

  - Level 0: distribution F
  - Level 1: unobserved samples $X_1, \ldots, X_n, \cdots \sim F$

Setting

- Goal: estimate some functional $g(F)$ of an underlying distribution F.
- Setting:

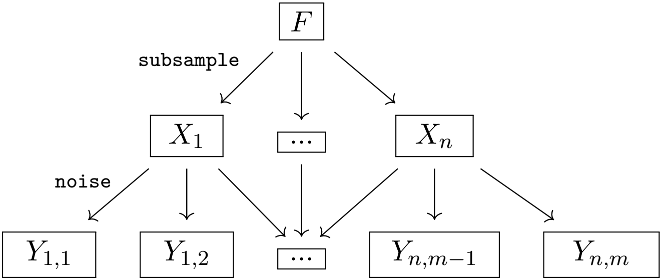

  - Level 0: distribution F
  - Level 1: unobserved samples $X_1, \ldots, X_n, \cdots \sim F$
  - Level 2: noisy observations: $Y_{i,j} \sim \mathcal{N}(X_i, 1)$

Examples

- Mean estimation in single-cell RNA-sequencing
- Benjamini-Hochberg (BH) threshold in multiple hypothesis testing
- Personalized recommendation in large-scale distributed learning

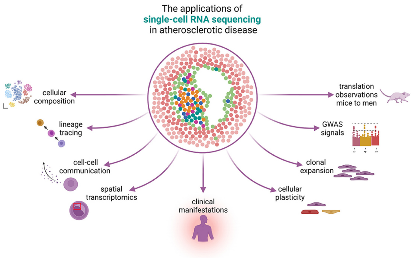

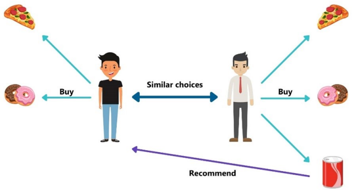

Connection to multi-armed bandits

- Classical multi-armed bandits (Level 1-2):

$$\text{Arms: } X_1, \ldots, X_n, \quad \text{Observations: } Y_i = X_{a_i} + Z_i.$$

- Infinite-armed bandit[1] (Level 0-2):

$$\text{Arms: } X_1, \ldots, X_i, \cdots \sim F(x) \quad \text{Observations: } Y_i = X_{a_i} + Z_i.$$

- The objective of infinite-armed bandit is minimizing simple or cumulative regret, while we are interested in estimating the distribution functional.

[1] Berry, D. A. et al. Bandit problems with infinitely many arms. The Annals of Statistics (1997).

Sampling algorithms

- Offline sampling algorithm: uniformly sample each arm

$$\hat{X}_i \sim \mathcal{N}(X_i, 1/m).$$

- Online sampling algorithm: requires continual interactions with the environment

$$Y_{t+1} \sim \mathcal{N}(X_{A_{t+1}}, 1), \quad A_{t+1} = \pi(Y_{1:t}, A_{1:t}).$$

- Difficulty:
  - Mixture model / density deconvolution
  - Statistical complexity of online sampling algorithms

Our contribution

- Offline and online sample complexity of
  - Mean
  - Median
  - Trimmed mean
  - Maximum
- Unified meta algorithms
- General lower bounds for both offline and online sampling

Structure of the functional

- The functional g can be represented as

$$g(F) = \mathbb{E}\left[X|F(X) \in [\alpha_1, \alpha_2]\right]$$

for $0 \leq \alpha_1 \leq \alpha_2 \leq 1$.

- This essentially encompasses all "nice" functionals.

Table: Indicator-based functionals.

Functional	$g(F)$	α_1	α_2	Comment
Mean	$\mathbb{E}[X]$	0	1	
Quantile	$F^{-1}(\alpha)$	α	α	$\alpha \in (0,1)$, e.g. $\alpha = 1/2$ for median
Maximum	$F^{-1}(1)$	1	1	$\alpha_1 = \alpha_2 = 0$ for minimum
Trimmed mean	$\mathbb{E}[X\|F(X) \in [\alpha_1, \alpha_2]]$	α	$1-\alpha$	$\alpha \in (0, 1/2)$

Sample complexity

Table: Sample complexity of estimating different functionals $g(F)$, where F is the cumulative distribution function (CDF) of the distribution to estimate.

Functional	Offline complexity	Online complexity	Comments
Mean	$\Theta(\epsilon^{-2})$	$\Theta(\epsilon^{-2})$	No gain from online sampling
Median	$\Theta(\epsilon^{-3})$	$\tilde{\Theta}(\epsilon^{-2.5})$	Holds for quantile too
Maximum	$\Theta(\epsilon^{-(2+\beta)})$	$\tilde{\Theta}(\epsilon^{-\max(\beta,2)})$	Depends on the tail regularity β
Trimmed mean	$\tilde{\Theta}(\epsilon^{-3})$	$\tilde{\Theta}(\epsilon^{-2.5})$	$g(F) = \mathbb{E}\{X \mid F(X) \in [\alpha, 1-\alpha]\}$

Offline algorithm

- For $\epsilon > 0$ and $\delta \in (0,1)$, we call an estimator $\hat{G}$ an (ϵ, δ)-PAC approximation of $g(F)$ if $P(|\hat{G} - g(F)| > \epsilon) \leq \delta$.

Theorem (Offline PAC sample complexity)

An (ϵ, δ)-PAC offline uniform-sampling-based algorithm for estimating $g(F)$ requires $\Theta_\epsilon(nm)$ samples where n, m are chosen based on ϵ, δ, the functional g, and information about F.

Functional	number of samples per point m	number of points n
Mean	$\Theta(1)$	$\Theta(\epsilon^{-2})$
Median	$\Theta(\epsilon^{-1})$	$\Theta(\epsilon^{-2})$
Maximum	$\Theta(\epsilon^{-2})$	$\Theta(\epsilon^{-\beta})$
Trimmed mean	$\Theta\left(\epsilon^{-1}\log\left(\epsilon^{-1}\right)\right)$	$\Theta(\epsilon^{-2})$

Online algorithm

1. Find n, m such that the plug-in estimator $G_{n,m}$ will be an $(\epsilon/2, \delta/2)$-PAC approximation of $g(F)$.
2. Adaptively compute an $(\epsilon/2, \delta/2)$-PAC approximation of $G_{n,m}$ using significantly fewer samples based on confidence intervals.

Theorem (Meta algorithm)

Our online meta algorithm provides an (ϵ, δ)-PAC estimate of $g(F)$ with M samples when given the requisite inputs. The expected sample complexity is bounded by

$$\mathbb{E}[M] = O\left(n\mathbb{E}\left[\min\left(m, \frac{\log(n/\delta)}{[\text{dist}(X, \partial S)]^2}\right)\right]\right),$$

where $S = [F^{-1}(\alpha_1), F^{-1}(\alpha_2)]$.

Conclusion

- Offline and online algorithms for estimating functionals of distributions
- Unified algorithms for estimating
  - Mean
  - Median
  - Trimmed mean
  - Maximum
- Information theoretic lower bounds
  - mean and maximum: different Wasserstein distances
  - median and trimmed mean: an interesting thresholding phenomenon of the noise level