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

Yifei Wang[1], Tavor Baharav[1], Yanjun Han[2], Jiantao Jiao[3], David Tse[1]
[1]Department of Electrical Engineering, Stanford [2]Statistics and Data Science Center, MIT IDSS [3]Department of EECS and Department of Statistics, UC Berkeley

NEURAL INFORMATION PROCESSING SYSTEMS

Introduction

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

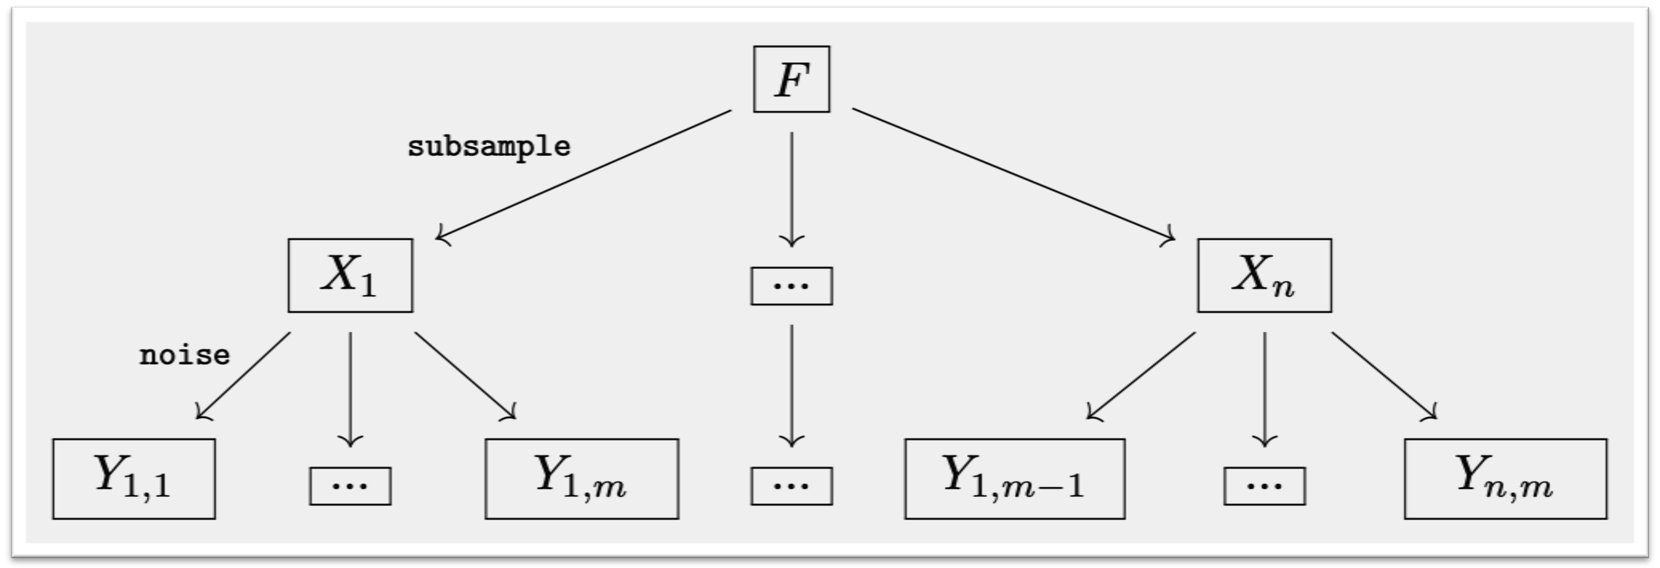

Setting

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

Action

At each time instance, the learner can:

- Either sample an arm that has been already observed in the past,
- Or draw and sample from a new arm, whose average reward is drawn from an underlying distribution $F(x)$.

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)
  $\text{Arms: } X_1, \dots, X_n, \quad \text{Observations: } Y_i = X_{a_i} + Z_i.$
- Infinite-armed bandit (Level 0 - 2)
  $\text{Arms: } X_1, \dots, X_i, \dots \sim F(x) \quad \text{Observations: } Y_i = X_{a_i} + Z_i.$
- The objective of infinite-armed bandit is minimizing simple or cumulative regret, while we are interested in estimating the distribution functional.

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

- Adaptivity helps in functional estimation: offline and online sample complexity of Mean / Median / Trimmed mean / Maximum
- Achieve the upper bound by a unified meta algorithms
- General lower bounds for both offline and online sampling

Background

Structure of the functional

- The functional $g(F)$ can be represented as
  $$g(F) = \mathbb{E}\left[X | F(X) \in [\alpha_1, \alpha_2]\right]$$
  where $0 \le \alpha_1 \le \alpha_2 \le 1.$
- This essentially encompasses all "nice" 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	$\mathbb{E}[X \vert F(X) \in [\alpha_1, \alpha_2]]$	α	$1-\alpha$	$\alpha \in (0, 1/2)$

Indicator-based functionals

Sample complexity of different functionals

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	$\tilde{\Theta}(\epsilon^{-3})$	$\tilde{\Theta}(\epsilon^{-2.5})$	$g(F) = \mathbb{E}\{X \vert F(X) \in [\alpha, 1-\alpha]\}$

Offline Algorithms

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

Theorem (Offline PAC sample complexity)
An (ε, δ)- PAC offline uniform-sampling-based algorithm for estimating $g(F)$ requires $\Theta_\varepsilon(nm)$ samples where (n, m) are chosen based on (ε, δ), the functional $g(F)$, 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 Algorithms

Intuition

- Find (n, m) such that the plug-in estimator $\hat{G}_{n,m}$ will be an $(\varepsilon/2, \delta/2)$ - PAC approximation of $g(F)$.
- Adaptively compute an $(\varepsilon/2, \delta/2)$ - PAC approximation of $\hat{G}_{n,m}$ using significantly fewer samples by constructing **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 $S = [F^{-1}(\alpha_1), F^{-1}(\alpha_2)].$

Lower bound via Wasserstein distance

Mean and Maximum

Lemma For any offline algorithm π, it holds that
$$D_{\text{KL}}(p_{\pi, F_1, n} \| p_{\pi, F_2, n}) \le \frac{mn}{2}\mathcal{W}_2^2(F_1, F_2).$$
For any online algorithm π which queries T samples, it holds that
$$D_{\text{KL}}(p_{\pi, F_1} \| p_{\pi, F_2}) \le \frac{T}{2}\mathcal{W}_\infty^2(F_1, F_2).$$
Here $\mathcal{W}_2^2(P,Q) = \inf_{\gamma \in \Gamma} \mathbb{E}_{(X,Y)\sim\Gamma}[(X-Y)^2]$, $\mathcal{W}_\infty(P,Q) = \inf_{\gamma \in \Gamma} \operatorname{esssup}_{(X,Y)\sim\Gamma} |X - Y|$
and Γ is the class of all couplings between P and Q.

- The sample complexity of offline algorithms
  $$\Omega(1/\min\{\mathcal{W}_2^2(F_1, F_2) : F_1, F_2 \in \mathcal{F}, |g(F_1) - g(F_2)| \ge 2\varepsilon\})$$
- The sample complexity of online algorithms
  $$\Omega(1/\min\{\mathcal{W}_\infty^2(F_1, F_2) : F_1, F_2 \in \mathcal{F}, |g(F_1) - g(F_2)| \ge 2\varepsilon\})$$

Lower bound via thresholding phenomenon

Median and Trimmed mean

A **key** quantity to prove lower bounds for median estimation is
$$\text{KL}_\sigma(\varepsilon) := \min\{D_{\text{KL}}(F_1 * \mathcal{N}(0,\sigma^2) \| F_2 * \mathcal{N}(0,\sigma^2)) : F_1, F_2 \in \mathcal{F}, |F_1^{-1}(1/2) - F_2^{-1}(1/2)| \ge 2\varepsilon\}.$$

Lemma (offline) For $\varepsilon \in (0, 1/4)$, the following characterization of $\text{KL}_\sigma(\varepsilon)$ holds as a function of σ:
$$\text{KL}_\sigma(\varepsilon) \begin{cases} \in [c_1\varepsilon^2, c_2\varepsilon^2] & \text{if } \sigma \le c\varepsilon^{1/2}, \\ \le C(\theta, \kappa)\varepsilon^\kappa & \text{if } \sigma \ge \varepsilon^{1/2-\theta}, \end{cases}$$
where $\theta \in (0, 1/4).$

Lemma (online) Fix any $\varepsilon, \theta \in (0, 1/4)$ and $\kappa \in \mathbb{N}$. There exists two distributions $F_1, F_2 \in \mathcal{F}$ with $|F_1^{-1}(0.5) - F_2^{-1}(0.5)| \ge \varepsilon$, and
$$D_{\text{KL}}(F_1 * \mathcal{N}(0,\sigma^2) \| F_2 * \mathcal{N}(0,\sigma^2)) \begin{cases} \in [c_1\varepsilon^{3/2}, c_2\varepsilon^{3/2}] & \text{if } \sigma \le c\varepsilon^{1/2}, \\ \le C(\theta, \kappa)\varepsilon^\kappa & \text{if } \sigma \ge \varepsilon^{1/2-\theta}, \end{cases}$$

Theorem (Median adaptive lower bound)
The $(\varepsilon, 0.1)$ - PAC sample complexity for median estimation is $\Omega(\varepsilon^{-2.5})$ for any online algorithm.