Nearly Optimal Catoni's Estimator for Infinite Variance

Sujay Bhatt, Guanhua Fang, Ping Li, & Gennady Samorodnitsky

Baidu Research & Cornell University

The 39th International Conference on Machine Learning (ICML 2022)

I.I.D r.vs $\{X_i\}_{i=1}^n$, $\mathbb{E}(X_1) = \mu$ and $\mathbb{E}|X_1 - \mu|^{1+\varepsilon} \leq \upsilon$ for $\varepsilon \in (0, 1]$.

- **Catoni's M-estimator:** $\widehat{\mu}_n$ is a solution to the equation

$$\sum_{i=1}^{n} \psi\Big(\alpha(X_i - \widehat{\mu}_n)\Big) = 0$$

  with $\psi : \mathbb{R} \to \mathbb{R}$ a non-decreasing influence function and $\alpha > 0$.
- **Key-Contribution**: For a given *confidence* $\delta \in (0, 1)$, we design the tightest possible $\varrho = \varrho(n, \delta)$ such that $\mathbb{P}\Big\{\Big|\widehat{\mu}_n - \mu\Big| > \varrho\Big\} \leq \delta$.

Motivation

Sample complexity reduction in best-arm identification for heavy-tailed bandits. Independent interest.

Literature & Key Ideas

- Catoni (2012)[1]: Near-optimal sub-Gaussian estimator for finite variance!
  - Uses novel influence functions.
  - Estimator obtained as a root of quadratic polynomial.
  - New techniques needed to relax finite variance assumptions.
- Chen et.al (2021)[2] extend the analysis to infinite variance. Key-drawbacks:
  - Obtain large coefficients owing to loose characterization of roots of polynomial of degree < 2.
  - Minimum data requirements are not decoupled from the moment bound (v) – not suitable for online learning in bandits!

[1] Catoni, O., 2012. Challenging the empirical mean. In Annales de l'IHP Probabilités et statistiques.

[2] Chen et. al, 2021. A generalized Catoni's M-estimator under finite α-th moment assumption with $\alpha \in (1, 2)$. Electronic Journal of Statistics.

Theorem: Confidence Interval

Let the minimum samples $n = O(\log(1/\delta)/\varepsilon\tau)$ for arbitrary $\tau > 0$. For a carefully chosen α,

$$\left|\widehat{\mu}_n - \mu\right| < \upsilon^{1/(1+\varepsilon)} \cdot \left(\frac{\log(2/\delta)}{n}\right)^{\frac{\varepsilon}{1+\varepsilon}} \cdot G(\varepsilon, \tau)$$

Scaling of lower bound $\rightarrow$ $G(\varepsilon, \tau)$.

- When $\varepsilon = 1$, $G(\varepsilon, \tau) \approx (2 + \gamma)^{1/2}$ for any $\gamma > 0$. Same order and asymptotic constant as in Catoni (2012).
- When $\varepsilon < 1$, the scaling $G(\varepsilon, \tau) \approx (1 + \varepsilon)^{1/2}(1 - \varepsilon)^{\frac{1}{1+\varepsilon} - \frac{1}{2}}$ as $\tau \downarrow 0$. Much sharper than Chen et. al (2021).

Other extensions compared to Chen et. al (2021) include:

- Adaptation to the case of unknown moment bound υ. Lepski's Method for moments.
- Relaxing i.i.d. Only require $\mathbb{E}[X_{t+1}|\mathcal{F}_t] = \mu$ and $\mathbb{E}[|X_{t+1} - \mu|^{1+\varepsilon}|\mathcal{F}_t] \leq \upsilon_\varepsilon$ for any filtration $\mathcal{F}_t$.

Extends the scope of the results for applications!

Applications: Bandits

SE-TEA (Yu et.al (2018)[3]). Successive Elimination with Catoni and Phase-based Elimination with Catoni ($O(\log\left(\frac{K\log(1/\Delta_i)}{\delta}\right)/\Delta_i^{1+\varepsilon/\varepsilon})$).

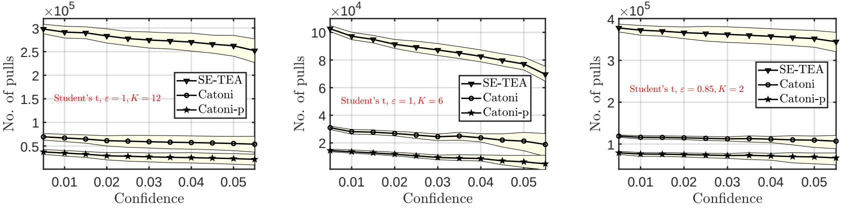

Figure: Average number of pulls for a fixed confidence over 50 iterations for different number of arms $K = 2, 6, 12$, and $\varepsilon = 1$ and $\varepsilon < 1$.

END

[3] Yu et. al., 2018. Pure exploration of multi-armed bandits with heavy-tailed payoffs. In UAI.