Towards Practical Mean Bounds for Small Samples

My Phan, Philip Thomas, Erik Learned-Miller

August 9, 2021

University of Massachusetts - Amherst

Non-parametric Upper-Confidence Bound

Assumption.

- The distribution has support $\mathcal{D} \subseteq [0, 1]$.

Definition.

Given a confidence level $1 - \alpha \in (0, 1)$, an upper confidence bound $\mu_{\mathtt{upper}}^{1-\alpha}$ has *guaranteed coverage* if, for all sample sizes $1 \leq n \leq \infty$ and for all distributions F with support on $[0, 1]$, it satisfies

$$Prob_F[\mu \leq \mu_{\mathtt{upper}}^{1-\alpha}(X_1, X_2, ..., X_n)] \geq 1 - \alpha, \tag{1}$$

where μ is the mean of the unknown distribution F.

Existing confidence bounds for the mean of distributions with support $\mathcal{D} \subseteq [0, 1]$.

- With guarantees, but bad performance for small samples (Maurer-Pontil, Hoeffding and Anderson)
- No guarantees (assume Gaussianity), good performance for small samples (Student-t)

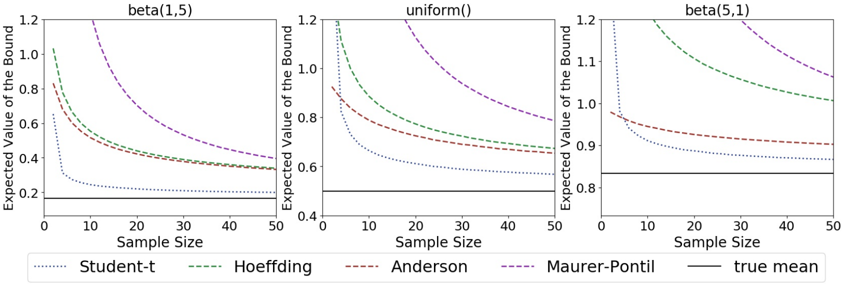

Figure 1: The expectation of the confidence bounds for small samples. .

Existing Results

Existing confidence bounds for the mean of distributions with support $\mathcal{D} \subseteq [0, 1]$.

- With guarantees, but bad performance for small samples (Maurer-Pontil, Hoeffding and Anderson)
- No guarantees (assume Gaussianity), good performance for small samples (Student-t).

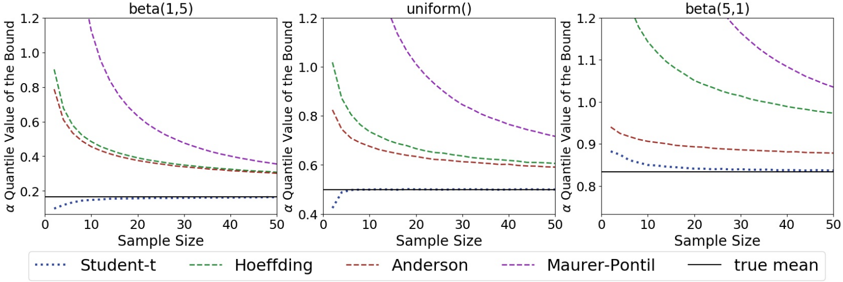

Figure 2: The α-quantile of the confidence bounds for small samples. Student-t does not have guarantee for small sample size.

Motivations

Consider the problem of finding mean bounds for **small** sample sizes. Examples:

- Phase 1 Clinical Trial.
- Safe Reinforcement Learning: importance-weighted estimators are used to estimate the policy's return.

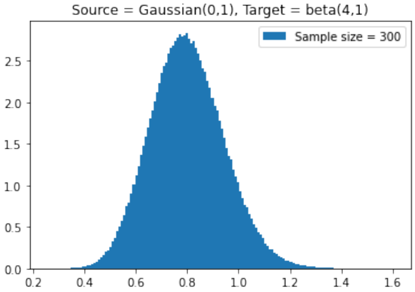

Figure 3: The sample mean of importance-weighted estimator is skewed even for large sample size.

Problem Statement

Find an upper confidence bound:

- with guaranteed coverage for distributions on $[0, 1]$
- with good performance for small sample size

Our Bound

- Input:
  - A sample $\mathbf{x} = (x_1, \cdots, x_n)$ of size n.
  - A confidence level α.
  - A function $T : [0, 1]^n \to R$.
    - Example: $T(\mathbf{x})$ is the sample mean of $\mathbf{x}$.
- Output: upper-confidence bound $b_T^\alpha(\mathbf{x})$.
- Our bound is parameterized by a function T. Function T defines the order of the bound: If $T(\mathbf{x}) \leq T(\mathbf{y})$ then $b_T^\alpha(\mathbf{x}) \leq b_T^\alpha(\mathbf{y})$.
  - Note: Hoeffding's bound and Student-t's bound order the samples by the sample mean.
  - We produce a guaranteed bound with **any** function T.

Theoretical Results

A new bound with guarantee coverage and good performance:

- For any sample size, for any α, for any sample $\mathbf{x}$, our bound is smaller than or equal to Anderson's [Anderson, 1969], one of the best bound for small sample sizes.
  - For any sample size, for any α, for any sample $\mathbf{x}$, Anderson's bound is smaller than or equal to Hoeffding [Hoeffding, 1963].
- Therefore for any sample size, for any α, for any sample $\mathbf{x}$, our bound is smaller than or equal to Hoeffding's.

Computation

- We use Monte Carlo simulation to compute the bound.
  - The Monte Carlo output is proven to converge to the theoretical value of the bound as the number of Monte Carlo samples increases.
  - Given any error threshold $\epsilon > 0$, it is possible to compute the number of Monte Carlo samples required such that with probability at least $1 - \alpha$:

$$\mu \leq \widehat{b_T^\alpha(\mathbf{x})} \leq b_T^\alpha(\mathbf{x}) + \epsilon \tag{2}$$

    where $b_T^\alpha(\mathbf{x})$ is the theoretical value of the bound and $\widehat{b_T^\alpha(\mathbf{x})}$ is the Monte Carlo output.
  - Computing the bound for one time takes seconds.
- The bound outputs the same value for samples $\mathbf{x}$ with the same $T(\mathbf{x})$. Therefore we can pre-compute a table mapping $T(\mathbf{x})$ to the value of the bound.

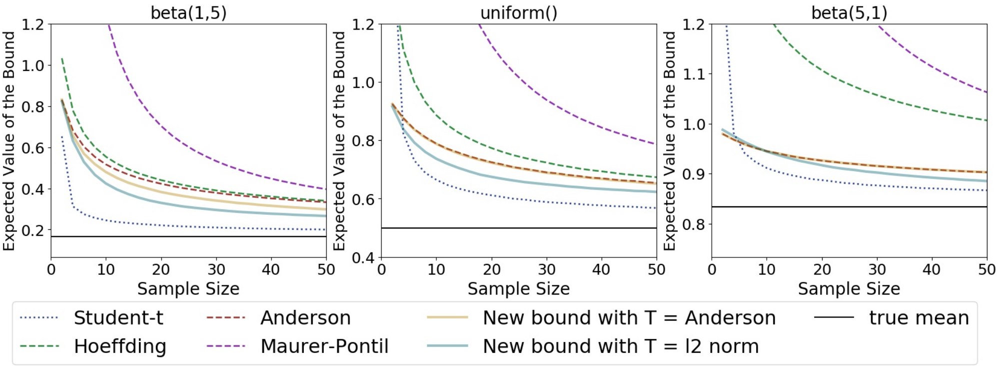

Figure 4: Expected values of the bounds versus sample size n. For each n, we sample $\mathbf{X} \in \mathcal{R}^n$ 10,000 times, and take the average of the bound. Our new bound has better performance than Anderson's, Hoeffding's and Maurer and Pontil's.

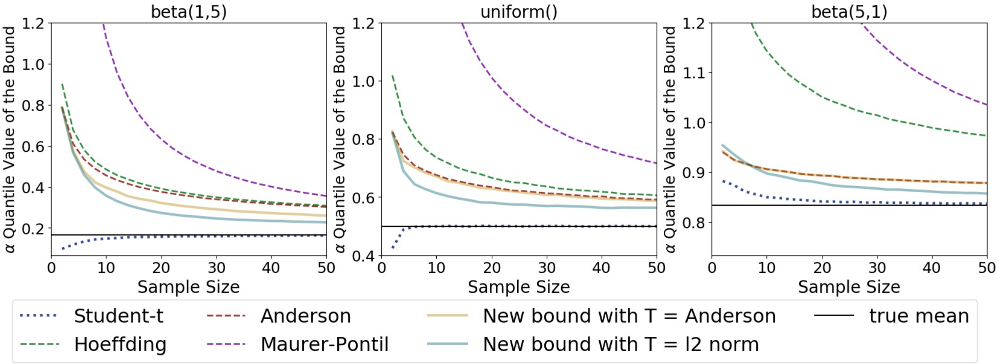

Figure 5: The α-quantiles of bound distributions. If α portion of the samples' bound is below the true mean, the bound does not have guarantee. For the $\mathtt{uniform}(0,1)$ and $\mathtt{beta}(1,5)$ distribution, when the sample size is small, Student-t does not have guarantee.

Thank You!

My Phan
myphan@cs.umass.edu

References

T. W. Anderson. Confidence limits for the value of an arbitrary bounded random variable with a continuous distribution function. *Technical Report Number 1, Department of Statistics, Stanford University*, 1969.

W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.