

STAT 460/560 Class 8: Consistency of the bootstrap

Daniel McDonald & Ben Bloem-Reddy

Reading: Chapter 23.1-23.2, [van98].

1. Justification for the Pivotal and Percentile Bootstrap Intervals

Recall the pivotal interval. Let $\theta = T(F)$ and $\hat{\theta}_n = T(\hat{F}_n)$. Define the **pivot** $R_n = \hat{\theta}_n - \theta$. If we knew the CDF of the pivot,

$$H(r) = \mathbb{P}(R_n \leq r) , \quad (8.1)$$

then we could compute an exact $1 - \alpha$ confidence interval as $C_n^* = (a, b)$, where

$$a = \hat{\theta}_n - H^{-1}(1 - \alpha/2) \quad \text{and} \quad b = \hat{\theta}_n - H^{-1}(\alpha/2) . \quad (8.2)$$

Why does this work? Note that:

$$\begin{aligned} \mathbb{P}(\theta \in C_n^*) &= \mathbb{P}(a \leq \theta \leq b) \\ &= \mathbb{P}(a - \hat{\theta}_n \leq \theta - \hat{\theta}_n \leq b - \hat{\theta}_n) \\ &= \mathbb{P}(\hat{\theta}_n - b \leq \hat{\theta}_n - \theta \leq \hat{\theta}_n - a) \\ &= \mathbb{P}(\hat{\theta}_n - b \leq R_n \leq \hat{\theta}_n - a) \\ &= H(\hat{\theta}_n - a) - H(\hat{\theta}_n - b) \\ &= H(H^{-1}(1 - \alpha/2)) - H(H^{-1}(\alpha/2)) \\ &= 1 - \alpha/2 - \alpha/2 = 1 - \alpha . \end{aligned}$$

Then, approximating H with the bootstrap distribution of $R_n^* = \hat{\theta}_n^* - \hat{\theta}_n$ yields the pivotal bootstrap interval: $\hat{C}_n = (\hat{a}, \hat{b})$, where

$$\hat{a} = \hat{\theta}_n - r_{1-\alpha/2}^* = 2\hat{\theta}_n - \theta_{1-\alpha/2}^* \quad (8.3)$$

$$\hat{b} = \hat{\theta}_n - r_{\alpha/2}^* = 2\hat{\theta}_n - \theta_{\alpha/2}^* . \quad (8.4)$$

For the percentile interval,

$$C_n = (\theta_{\alpha/2}^*, \theta_{1-\alpha/2}^*) , \quad (8.5)$$

we need slightly more. Suppose there exists a monotone function g such that $U = g(T) \sim \mathcal{N}(\phi, c^2)$. We don't actually need to know g (or ϕ or c^2), but we do need to assume that such a transformation exists. Let $U_b^* = g(T_b^*)$ be the bootstrap version of U . Let u_β^* be the β -sample quantile of the U_b^* 's. Because monotone transformations preserve quantiles, we have that

$$\mathbb{P}(\theta_{\alpha/2}^* \leq \theta \leq \theta_{1-\alpha/2}^*) = \mathbb{P}(g(\theta_{\alpha/2}^*) \leq g(\theta) \leq g(\theta_{1-\alpha/2}^*)) \quad (8.6)$$

$$= \mathbb{P}(u_{\alpha/2}^* \leq \phi \leq u_{1-\alpha/2}^*) \quad (8.7)$$

$$= \mathbb{P}(U - cz_{\alpha/2} \leq \phi \leq U + cz_{1-\alpha/2}) \quad (8.8)$$

$$= \mathbb{P}(-z_{\alpha/2}^* \leq (U - \phi)/c \leq z_{1-\alpha/2}^*) = 1 - \alpha . \quad (8.9)$$

So, relative to the pivotal interval, the coverage guarantee of the percentile interval requires the additional assumption that there exists a monotone transformation of the statistic of interest that is asymptotically normal.

2. Setup for the bootstrap consistency results

Today we'll establish the consistency of bootstrap methods, and in particular, bootstrap confidence intervals. Extending the empirical CDF to more general settings, for n i.i.d. observations $X_1, \dots, X_n$ taking values in $\mathbb{R}^d$, $\hat{P}_n$ is the empirical measure, defined as

$$\hat{P}_n(A) = \frac{1}{n} \sum_{i=1}^n \mathbf{I}_A(X_i), \quad A \subset \mathbb{R}^d.$$

As with the empirical CDF, sampling from $\hat{P}_n$ is just sampling with replacement from $X_1, \dots, X_n$.

The asymptotics today may seem a little strange, because many of the statements will be something like “ $S_n \rightsquigarrow T$, conditionally almost surely.” What this means is that for P -almost every sequence $X_1, X_2, \dots$, the conditional distribution of S_n given $X_1, \dots, X_n$ converges to the distribution of T . Just think of this as the sequence of conditional distributions converging, but keep in mind that a conditional distribution is a random distribution (it is a function of the random variables on which it conditions), which makes the “almost surely” necessary.

Before we get to the bootstrap, we need an intermediate result, the proof of which can be found as Lemma 21.2 in [van98]. Recall that if F is a CDF, the **quantile function**, denoted F^{-1} , is the generalized inverse of F ,

$$F^{-1}(p) = \inf\{x: F(x) \geq p\}.$$

If F is continuous and strictly increasing then this is just the usual inverse. The generalized part comes in when F has jumps, flat spots, etc.

Lemma 8.1 (Lemma 21.2, [van98]). *Let F_n be any sequence of CDFs and F another CDF. $F_n^{-1}(p) \rightarrow F^{-1}(p)$ at every p where F^{-1} is continuous if and only if $F_n(x) \rightarrow F(x)$ at every x where F is continuous.*

Note that this kind of convergence is equivalent to $X_n \rightsquigarrow X$ if F_n is the CDF of X_n and F is the CDF of X . For that reason, we write $F_n \rightsquigarrow F$ (but this is a convenient overloading of notation).

3. Consistency of interval from convergence of estimator

Recall that the bootstrap conditions on an estimate $\hat{P}_n$ of P (formed from the sample), and generates a bootstrap sample $X_1^*, \dots, X_n^*$. In practice, we'll simulate the bootstrap sample many times to estimate things like the distribution of $(\hat{\theta}_n^* - \hat{\theta}_n)/\hat{\sigma}_n^*$. For the purposes of analysis, we'll assume that we have access to the actual bootstrap distribution (i.e., what we would obtain in the limit $B \rightarrow \infty$).

To form the bootstrap interval, we will estimate quantiles of $(\hat{\theta}_n - \theta)/\sigma$ by quantiles of the bootstrap quantity $(\hat{\theta}_n^* - \hat{\theta}_n)/\hat{\sigma}_n^*$. In particular, let

$$\hat{G}_n(x) := P\left(\frac{\hat{\theta}_n^* - \hat{\theta}_n}{\hat{\sigma}_n^*} \leq x \mid \hat{P}_n\right), \quad (8.10)$$

with $\hat{\xi}_{n,1-\alpha} = \inf\{x: \hat{G}_n(x) \geq 1-\alpha\}$. In [van98], $\hat{F}_n$ is used to denote this conditional CDF, but we will use $\hat{G}_n$ to avoid confusion with the empirical CDF $\hat{F}_n$. Note that this is a random number, where the randomness comes from conditioning on $\hat{P}_n$. If T is a random variable with CDF F , we say that $(\hat{\theta}_n^* - \hat{\theta}_n)/\hat{\sigma}_n^* \rightsquigarrow T$ conditionally on $X_1, X_2, \dots$ if $\hat{G}_n \rightsquigarrow F$ conditioned on $X_1, X_2, \dots$.

Using the estimated quantile, we can form an approximate interval as

$$\hat{I}_n(\alpha, \beta) = [\hat{\theta}_n - \hat{\xi}_{n,\beta} \hat{\sigma}_n, \hat{\theta}_n - \hat{\xi}_{n,1-\alpha} \hat{\sigma}_n]. \quad (8.11)$$

Here, we are letting α and β be the tail probabilities, so that the interval has approximate coverage $1 - \alpha - \beta$. We can also write $\alpha = \beta$ to get the typical symmetric interval (with coverage $1 - 2\alpha$), but this is not necessary. We'll continue to use α and β to match the notation in [van98].

A sufficient condition for consistency of the bootstrap intervals is that the bootstrap quantities converge in distribution to the same place as the non-bootstrap quantities. For the purposes of $\hat{I}_n(\alpha, \beta)$, the bootstrap quantiles are the only portion that depends on the bootstrap sample, so we only need to check that the bootstrap quantiles converge in distribution to the same place as the non-bootstrap quantiles. First, we need a lemma that shows that if the bootstrap quantiles converge in distribution to the same place as the non-bootstrap quantiles, then the coverage of the interval converges to the nominal level.

Lemma 8.2. *Suppose that $(\hat{\theta}_n - \theta)/\hat{\sigma}_n \rightsquigarrow T$, and that for almost every sequence $X_1, X_2, \dots$, $(\hat{\theta}_n^* - \hat{\theta}_n)/\hat{\sigma}_n^* \rightsquigarrow T$ conditionally on $X_1, X_2, \dots$, for a random variable T with a continuous CDF F . Then*

$$P(\theta \in \hat{I}_n(\alpha, \beta)) \rightarrow 1 - \alpha - \beta. \quad (8.12)$$

Before going through the proof, we will make a couple of assumptions that are not strictly necessary, but that simplify the proof. The first is that we assume that F is a continuous function. This rules out having to deal with CDFs that can differ on a null set of discontinuities (a headache that doesn't offer additional insight). The second is that we assume that convergence in distribution $(\hat{\theta}_n^* - \hat{\theta}_n)/\hat{\sigma}_n^* \rightsquigarrow T$ holds almost surely, conditionally on $X_1, X_2, \dots$. This can be weakened to convergence in probability by arguing along subsequences but again, that doesn't offer additional insight.

Proof. By assumption, $\hat{G}_n \rightsquigarrow F$, where F is the CDF of T . By Lemma 8.1, $\hat{G}_n^{-1}(p) \rightarrow F^{-1}(p)$ at every continuity point of F^{-1} . Choose α so that $1 - \alpha$ is such a continuity point. Then $\hat{\xi}_{n,1-\alpha} = \hat{G}_n^{-1}(1 - \alpha)$ converges almost surely to $F^{-1}(1 - \alpha)$. Hence, by Slutsky's lemma,

$$(\hat{\theta}_n - \theta)/\hat{\sigma}_n - \hat{\xi}_{n,1-\alpha} \rightsquigarrow T - F^{-1}(1 - \alpha).$$

Therefore,

$$P(\theta \geq \hat{\theta}_n - \hat{\sigma}_n \hat{\xi}_{n,1-\alpha}) = P\left(\frac{\hat{\theta}_n - \theta}{\hat{\sigma}_n} \leq \hat{\xi}_{n,1-\alpha} \mid P\right) \rightarrow P(T \leq F^{-1}(1 - \alpha)) = 1 - \alpha.$$

This must hold for all but at most countably many values of α . Both the LHS and the RHS of the previous math display are monotone functions of α and the RHS is continuous in α , the convergence must hold for every α .

The same argument applies to β . □

With this in hand, showing that bootstrap confidence intervals are consistent amounts to showing that our statistic of interest converges in distribution, and that the bootstrap version of our statistic of interest converges in distribution to the same place.

Theorem 8.3. *Let $X_1, X_2, \dots$ be i.i.d. random variables in $\mathbb{R}$ with mean μ and variance σ^2 . Then conditionally on $X_1, X_2, \dots$, for almost every sequence $X_1, X_2, \dots$,*

$$\sqrt{n}(\bar{X}_n^* - \bar{X}_n) \rightsquigarrow \mathcal{N}(0, \sigma^2).$$

Exercise 8.1. (on Assignment 4) Show that

$$\mathbb{E}[X_i^* | \hat{P}_n] = \bar{X}_n, \quad \text{and} \quad \mathbb{E}[(X_i^* - \bar{X}_n)^2 | \hat{P}_n] = \frac{1}{n} \sum_{i=1}^n X_i^2 - \bar{X}_n^2.$$

By the (strong) law of large numbers, these converge almost surely to μ and to σ^2 , respectively. In other words, for P -almost every sequence $X_1, X_2, \dots$, $\bar{X}_n \rightarrow \mu$ and similarly for the variance.

Showing the asymptotic normality of $\bar{X}_n^*$ requires a strong version of the CLT that is suitable for the case in which the X_i^* 's are sampled from a different distribution for each n . In particular, the Lindeberg–Feller CLT applies to *triangular arrays*, and although van der Vaart [van98, Prop. 2.27] calls it “the simplest extension of the classical central limit theorem,” it is extremely useful. An infinite triangular array is formed by taking a vector Y_n of length k_n , putting it in the n -th row, and appending infinitely many zeros. In the theorem, the entries in each row are assumed to be independent, but there may be dependence between rows. It can also be extended to the situation in which each entry of the array is actually a vector of fixed length [see van98, Prop. 2.27].

Theorem 8.4 (Lindeberg–Feller CLT). *For each $n \geq 1$, let $(Y_{n,1}, \dots, Y_{n,k_n})$ be a vector of independent random variables with finite variances and for every $\epsilon > 0$,*

$$\sum_{i=1}^{k_n} \mathbb{E}[\|Y_{n,i}\|^2 \mathbf{I}\{\|Y_{n,i}\| > \epsilon\}] \rightarrow 0 \quad \text{as } n \rightarrow \infty. \quad (8.13)$$

Moreover, assume that $\sum_{i=1}^{k_n} \text{Var}(Y_{n,i}) \rightarrow \sigma^2$. Then as $n \rightarrow \infty$,

$$\sum_{i=1}^{k_n} (Y_{n,i} - \mathbb{E}[Y_{n,i}]) \rightsquigarrow \mathcal{N}(0, \sigma^2). \quad (8.14)$$

Note that the typical use case for the Lindeberg–Feller CLT is when $k_n = n$, but it is not necessary. Furthermore, the sum of the variances converging is more general than assuming that each $Y_{n,i}$ has the same variance and looking at sample averages, which is the case in the classical CLT. In the classical CLT compared with (8.14), averaging is used to ensure that the sum of the variances converges to a finite value.

The Lindeberg condition (8.13) is a technical condition that ensures that no single term in the sum dominates the others. Another way for it to fail is if some terms have rare but large values. It is satisfied if the $Y_{n,i}$'s are uniformly bounded, or if they are identically distributed with finite variance.

Back to our bootstrap analysis, $\sqrt{n}\bar{X}_n^* = \sum_{i=1}^n \frac{1}{\sqrt{n}} X_{n,i}^* = \sum_{i=1}^n Y_{n,i}$, where we have defined $Y_{n,i} = \frac{1}{\sqrt{n}} X_{n,i}^*$, and $X_{n,i}^*$ is the i -th bootstrap sample from $\hat{P}_n$. We just need to check the conditions of the Lindeberg–Feller CLT. To that end, fix arbitrary $\epsilon > 0$, and consider

$$\begin{aligned} \sum_{i=1}^n \mathbb{E} \left[\|Y_{n,i}\|^2 \mathbf{I}\{\|Y_{n,i}\| > \epsilon\} \mid \hat{P}_n \right] &= \sum_{i=1}^n \mathbb{E} \left[\frac{1}{n} \|X_{n,i}^*\|^2 \mathbf{I}\{\|X_{n,i}^*\| > \sqrt{n}\epsilon\} \mid \hat{P}_n \right] \\ &= \frac{n}{n} \frac{1}{n} \sum_{i=1}^n \|X_{n,i}\|^2 \mathbf{I}\{\|X_{n,i}\| > \sqrt{n}\epsilon\} \\ &\leq \frac{1}{n} \sum_{i=1}^n \|X_{n,i}\|^2 \mathbf{I}\{\|X_{n,i}\| > M\} \xrightarrow{\text{a.s.}} \mathbb{E}[\|X_{n,i}\|^2 \mathbf{I}\{\|X_{n,i}\| > M\}], \end{aligned}$$

for $M \leq \sqrt{n}\epsilon$, and where the convergence follows from the strong law of large numbers.

Now, we assumed that $\text{Var}[X_i] = \sigma^2$, which implies that

$$\lim_{M \rightarrow \infty} \mathbb{E}[\|X_{n,i}\|^2 \mathbf{I}\{\|X_{n,i}\| > M\}] = 0. \quad (8.15)$$

(See the optional Appendix A on integrability below for details.) Hence, for each $\eta > 0$, there is some M_η such that

$$\mathbb{E}[\|X_{n,i}\|^2 \mathbf{I}\{\|X_{n,i}\| > M_\eta\}] < \eta,$$

which implies that for n large enough,

$$\sum_{i=1}^n \mathbb{E} \left[\|Y_{n,i}\|^2 \mathbf{I}\{\|Y_{n,i}\| > \epsilon\} \mid \hat{P}_n \right] < \eta ,$$

for P -almost every sequence $X_1, X_2, \dots$. That shows that the first condition of the Lindeberg–Feller CLT is satisfied for P -almost every sequence.

Finally, using the activity above and the strong law of large numbers,

$$\begin{aligned} \sum_{i=1}^{k_n} \text{Var}(Y_{n,i}) &= \frac{1}{n} \sum_{i=1}^n \mathbb{E}[(X_i^* - \bar{X}_n)^2 \mid \hat{P}_n] \\ &= \frac{1}{n} \sum_{i=1}^n X_i^2 - \bar{X}_n^2 \xrightarrow{\text{a.s.}} \sigma^2 . \end{aligned}$$

Hence, the conclusion of the Lindeberg–Feller CLT applied here is that

$$\sum_{i=1}^{k_n} (Y_{n,i} - \mathbb{E}[Y_{n,i} \mid \hat{P}_n]) = \sqrt{n}(\bar{X}_n^* - \bar{X}_n) \rightsquigarrow \mathcal{N}(0, \sigma^2)$$

holds for P -almost every sequence $X_1, X_2, \dots$.

Theorem 23.5 in [van98] extends this substantially by proving the validity of the delta method for the bootstrap.

References

- [Çin11] E. Çinlar. *Probability and Stochastics*. Springer New York, 2011.
- [van98] A. W. van der Vaart. *Asymptotic Statistics*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 1998.

A Integrability (optional)

A random variable X is *integrable* if $\mathbb{E}[\|X\|] < \infty$. An equivalent condition [see, e.g., Çin11, Lemma 3.10] is that

$$\lim_{b \rightarrow \infty} \mathbb{E}[\|X\| \mathbf{I}\{\|X\| > b\}] = 0 .$$

Applied to the situation above, (8.15) holds because $\mathbb{E}[X_i^2] = \mu^2 + \sigma^2 < \infty$ (i.e., it is integrable).