Bootstrap Confidence Intervals for Nonlinear Least Squares Estimation

Yang Pei*

May 25, 2025

Model Setup

Suppose we observe $n = 100$ pairs of data (x_i, y_i), and we are interested in estimating a parameter vector $\beta \in \mathbb{R}^k$ from the nonlinear model:

$$y_i = f(x_i, \beta) + \varepsilon_i,$$

where $f(\cdot, \beta)$ is a known nonlinear function and ε_i is an unobserved error term.

We estimate β via nonlinear least squares (NLS), by solving:

$$\hat{\beta} = \arg\min_{\beta} \sum_{i=1}^{n} \left[y_i - f(x_i, \beta)\right]^2 .$$

Objective

We want to construct a 95% confidence interval for each element of β using the **bootstrap method**.

Step-by-Step Procedure

Step 1: Estimate $\hat{\beta}$ from the original sample.

- Using the full original dataset $\{(x_i, y_i)\}_{i=1}^n$,
- Solve the nonlinear least squares problem to get $\hat{\beta}_{\text{original}}$.
- This is your baseline estimate.

Step 2: Generate bootstrap samples.

- Repeat the following B times (e.g., $B = 1000$):
  - Randomly draw n observations *with replacement* from the original dataset.

*Yang Pei, University of Houston, `ypei5@uh.edu`

- This creates a new bootstrap sample: $\{(x_i^*, y_i^*)\}_{i=1}^n$.
- Some observations may be repeated; others may be omitted.

Step 3: Re-estimate β for each bootstrap sample.

- For each bootstrap sample $b = 1, \ldots, B$:
  - Solve the same nonlinear least squares problem:

$$\hat{\beta}^{(b)} = \arg\min_{\beta} \sum_{i=1}^{n} \left[y_i^* - f(x_i^*, \beta)\right]^2.$$

  - Store the resulting estimate $\hat{\beta}^{(b)}$.
- After B iterations, you have B estimates:

$$\hat{\beta}^{(1)}, \hat{\beta}^{(2)}, \ldots, \hat{\beta}^{(B)}.$$

Step 4: Compute the 95% bootstrap confidence intervals.

- For each parameter component β_j:
  - Extract the j-th element from each $\hat{\beta}^{(b)}$, denoted $\hat{\beta}_j^{(b)}$.
  - Form the set $\{\hat{\beta}_j^{(1)}, \ldots, \hat{\beta}_j^{(B)}\}$.
  - Sort this set from smallest to largest.
  - Compute the 2.5% and 97.5% quantiles.
- The 95% percentile confidence interval for β_j is:

$$\mathrm{CI}_{95\%}(\beta_j) = \left[\hat{\beta}_j^{(B\cdot 0.025)},\ \hat{\beta}_j^{(B\cdot 0.975)}\right]$$

- These are the lower and upper bounds such that 95% of the bootstrap estimates fall inside.

Why Use Bootstrap?

- It does not rely on analytical standard errors or asymptotic normality.
- It works well even when the model is nonlinear or the sample is small.
- It directly approximates the sampling distribution of your estimator.

Conclusion

The bootstrap is a powerful and flexible tool for constructing confidence intervals for estimated parameters. By repeatedly simulating the estimation process using resampled data, we approximate the uncertainty in our original estimate $\hat{\beta}$.