Interval Estimation

Andrew Nobel

October, 2025

Interval Estimation

Setting: Family $\mathcal{P} = \{f(\cdot|\theta) : \theta \in \Theta\}$ of distributions on sample space $\mathcal{X}$

- For now, assume parameter space $\Theta \subseteq \mathbb{R}$

Goal: Given observation $X \sim f(\cdot|\theta) \in \mathcal{P}$ find an interval $[L(X), U(X)]$ that is

- Small
- Contains θ with high probability

Interval Estimation

Definition: An *interval estimator* is a pair (L, U) with $L, U : \mathcal{X} \to \mathbb{R} \cup \{-\infty, +\infty\}$ such that $L(x) \leq U(x)$ for every $x \in \mathcal{X}$.

Interval estimator (L, U) associates each $x \in \mathcal{X}$ with a *confidence interval* (CI)

$$[L(x), U(x)] \subseteq \mathbb{R}$$

for the parameter θ. Confidence interval is

- *Two-sided* if $L(x)$ and $U(x)$ are finite
- *One-sided* if one of $L(x)$ and $U(x)$ is infinite

Coverage Probability and Confidence Coefficients

Definition: The *coverage probability* of an interval estimator (L, U) at θ is

$$\mathbb{P}_\theta(\theta \in [L(X), U(X)]) \ = \ \mathbb{P}_\theta(L(X) \leq \theta \leq U(X))$$

Note: Observation X is random, parameter θ is fixed

Definition: The *confidence coefficient* of (L, U) is worst case coverage probability

$$\min_{\theta \in \Theta} \mathbb{P}_\theta(L(X) \leq \theta \leq U(X))$$

Terminology: If (L, U) has confidence coefficient $(1 - \alpha)$ then $[L(x), U(x)]$ is a $(1 - \alpha)$ confidence interval for θ based on data x

Examples

Ex 1. Observe $X_1, \ldots, X_n$ i.i.d. $\sim \mathcal{N}(\mu, \sigma^2)$ with σ^2 known. For fixed $c > 0$ the interval estimator

$$L(x) = \overline{x} - \frac{c\sigma}{\sqrt{n}} \qquad\qquad U(x) = \overline{x} + \frac{c\sigma}{\sqrt{n}}$$

has confidence coefficient $1 - 2\Phi(-c)$.

Ex 2. Observe $X = X_1, \ldots, X_n$ i.i.d. $\sim$ U$(0, \theta)$. Let $M(X) = \max\{X_1, \ldots, X_n\}$. For fixed $1 \leq a \leq b$ the interval estimator

$$L(x) = a\, M(x) \qquad\qquad U(x) = b\, M(x)$$

has confidence coefficient $a^{-n} - b^{-n}$.

Confidence Sets

Setting: General parameter space Θ, not necessarily a subset of $\mathbb{R}$

Definition: A *set estimator* is a map $C : \mathcal{X} \to 2^{\Theta}$ from data points to subsets of Θ.
The *confidence coefficient* of a set estimator C is

$$\min_{\theta \in \Theta} \mathbb{P}_{\theta}(\theta \in C(X))$$

Task: Find confidence sets for different statistical models. Two approaches

- Inverting hypothesis tests
- Pivots

Hypothesis Tests and Confidence Sets

Common threads:

- Hypothesis test with simple null: identify data x consistent with $\theta = \theta_0$
- Confidence set: identify parameters θ consistent with observed data x

Idea: Given data x, consider the set of parameters θ_0 for which we accept the null $\mathrm{H}_0\colon \theta = \theta_0$ based on x

Constructing Confidence Sets by Inverting Hypothesis Tests

Goal: Construct $(1-\alpha)$ confidence set for θ based on $X \sim f(\cdot|\theta) \in \mathcal{P}$

Step 1. Hypothesis tests. For each $\theta_0 \in \Theta$ let $(X, R(\theta_0))$ be a level-α test of

$$\mathrm{H}_0\colon \theta = \theta_0 \;\; \text{vs} \;\; \mathrm{H}_1 : \theta \neq \theta_0$$

Step 2. Invert the tests.

- For each $\theta_0 \in \Theta$ define the *acceptance region* $A(\theta_0) = R(\theta_0)^c$
- For $x \in \mathcal{X}$ define the confidence set $C(x) = \{\theta_0 : x \in A(\theta_0)\}$

Note: By definition, $\theta_0 \in C(x)$ if and only if $x \in A(\theta_0)$ $(*)$

Inverting Hypothesis Tests, cont.

Interpretation: Confidence set $C(x)$ is the set of parameters θ_0 for which we accept $\text{H}_0\colon \theta = \theta_0$ when we observe data x

Fact: Set estimate $C(X)$ is a $(1-\alpha)$ confidence set for θ

Fact: If $C(X)$ is a $(1-\alpha)$ confidence set for θ then for every $\theta_0 \in \Theta$

$$R(\theta_0) = \{x : \theta_0 \notin C(x)\}$$

is the rejection region of a level α test of $\text{H}_0 : \theta = \theta_0$ vs. $\text{H}_1 : \theta \neq \theta_0$.

One- and Two-Sided CIs via Inversion

Inversion of hypothesis test requires $\mathbb{P}_{\theta_0}(X \in R(\theta_0)) \leq \alpha$ for every $\theta_0 \in \Theta$

Rejection regions $R(\theta)$ typically tailored to one- or two-sided alternatives,

$$\mathrm{H}_1\colon \theta < \theta_0, \quad \mathrm{H}_1\colon \theta > \theta_0, \quad \mathrm{H}_1\colon \theta \neq \theta_0$$

Choice of alternative depends on problem, desired form of confidence interval

- One-sided alternative gives one-sided CI
- Two-sided alternative gives two-sided CI

Example

Exponential: $X_1, \dots, X_n$ i.i.d. $\sim$ Exp(θ) with $\theta > 0$.

- Goal: $(1-\alpha)$ CI for θ
- Inverting two-sided LRTs for $\mathrm{H}_0 : \theta = \theta_0$ vs. $\mathrm{H}_1 : \theta \neq \theta_0$ yields interval estimator

$$L(x) = b^{-1} \sum_{i=1}^{n} x_i \qquad U(x) = a^{-1} \sum_{i=1}^{n} x_i$$

for appropriate constants $0 < a < b$.

Example

Normal mean: $X_1, \ldots, X_n$ i.i.d. $\sim \mathcal{N}(\mu, \sigma^2)$ with σ^2 unknown.

- Goal: $(1-\alpha)$ upper CI for μ.
- Inverting one-sided t-tests for $\mathrm{H}_0 : \mu = \mu_0$ vs. $\mathrm{H}_1 : \mu < \mu_0$ yields interval estimator

$$L(x) = -\infty \qquad U(x) = \overline{x} + \frac{t_{n-1,\alpha} \cdot s}{\sqrt{n}}$$

Pivots

Given: Family $\mathcal{P} = \{f(\cdot|\theta) : \theta \in \Theta\}$ of distributions on sample space $\mathcal{X}$

Definition: A function $Q : \mathcal{X} \times \Theta \to \mathbb{R}$ is a *pivot* for $\mathcal{P}$ if the distribution of $Q(X, \theta)$ when $X \sim f(\cdot|\theta)$ is independent of θ

Note: If $Q(x, \theta)$ is a pivot for $\mathcal{P}$ then for all $a < b$ the probabilities

$$\mathbb{P}_\theta(a \leq Q(X, \theta) \leq b)$$

are independent of $\theta \in \Theta$. Idea: use this to define confidence sets

Examples of Pivots

Example: Observations $X_1, \ldots, X_n$ iid with $X_i \sim f(\cdot|\theta) \in \mathcal{P}_0$

- Location Family: $X_i \stackrel{d}{=} V_i + \mu$ with $V_1, \ldots, V_n$ iid G and $\mu \in \mathbb{R}$

  Pivot $Q(x, \mu) = \overline{x} - \mu$

- Scale Family: $X_i \stackrel{d}{=} \sigma V_i$ with $V_1, \ldots, V_n$ iid G and $\sigma > 0$

  Pivot $Q(x, \sigma) = \overline{x}/\sigma$

- Location-Scale Family: $X_i \stackrel{d}{=} \sigma V_i + \mu$ with $V_1, \ldots, V_n$ iid G and $\mu \in \mathbb{R}, \sigma > 0$

  Pivot $Q(x, (\mu, \sigma)) = (\overline{x} - \mu)/\sigma$

Confidence Sets from Pivots

Step 1. Suppose $\mathcal{P}$ has pivot $Q(x, \theta) \in \mathbb{R}$. Given $0 < \alpha < 1$ identify $a < b$ such that

$$\mathbb{P}_\theta(a \leq Q(X, \theta) \leq b) \ = \ 1 - \alpha \ \text{ for all } \ \theta \in \Theta$$

Step 2. For each $\theta_0 \in \Theta$ the set $A(\theta_0) = \{x : a \leq Q(x, \theta_0) \leq b\}$ is the acceptance region of a level α test of $\mathrm{H}_0 \colon \theta = \theta_0$

Step 3. Inverting acceptance regions $A(\theta_0)$ gives $(1 - \alpha)$ confidence set

$$C(x) = \{\theta : a \leq Q(x, \theta) \leq b\}$$

- If $Q(x, \cdot)$ is non-decreasing then $C(x)$ has form $[L(x, a), U(x, b)]$
- If $Q(x, \cdot)$ is non-increasing then $C(x)$ has form $[L(x, b), U(x, a)]$

Example

Observations: $X_1, \ldots, X_n$ iid $\mathcal{N}(\mu, \sigma^2)$

Case 1: Variance σ^2 known. Natural pivot $Q(x, \mu) = \frac{\overline{x} - \mu}{\sigma / \sqrt{n}}$ yields $(1 - \alpha)$ CI

$$C(x) \;=\; \left\{ \overline{x} - \frac{z_{\alpha/2}\, \sigma}{\sqrt{n}} \;\leq\; \mu \;\leq\; \overline{x} - \frac{z_{\alpha/2}\, \sigma}{\sqrt{n}} \right\}$$

Case 2: Variance σ^2 unknown. Natural pivot $Q(x, \mu) = \frac{\overline{x} - \mu}{s / \sqrt{n}}$ yields $(1 - \alpha)$ CI

$$C(x) \;=\; \left\{ \overline{x} - \frac{s\, t_{n-1,\alpha/2}}{\sqrt{n}} \;\leq\; \mu \;\leq\; \overline{x} + \frac{s\, t_{n-1,\alpha/2}}{\sqrt{n}} \right\}$$