Lecture Summary

7.5 Maximum Likelihood Estimation

Likelihood

- When the joint $\mathrm{pdf}/\mathrm{pf}\, f(\mathrm{x} \mid \theta)$ is regarded as a function of θ for given observations $x_1, \ldots, x_n$ it is called the likelihood function.

Maximum likelihood estimator

(MLE): For any given observations x we pick the $\theta \in \Omega$ that maximizes $f(\mathrm{x} \mid \theta)$.

Frame Title

- Given $X = x$, the maximum likelihood estimate (MLE) will be a function of x. Notation: $\hat{\theta} = \delta(X)$
- Potentially confusing notation: Sometimes $\hat{\theta}$ is used for both the estimator and the estimate.
- Note: The MLE is required to be in the parameter space Ω.
- Often it is easier to maximize the log-likelihood $L(\theta) = \log(f(x \mid \theta)$

MLE

- We pick the parameter that makes the observed data most likely.
- But: The likelihood is not a pdf/pf: If the likelihood of θ_1 is larger than the likelihood of θ_1, i.e. $f(x \mid \theta_2) > f(x \mid \theta_1)$ it does NOT mean that θ_2 is more likely.
- Remember: θ is not random here

Examples

- Let $X \sim \text{Binomial}(\theta)$. Find the maximum likelihood estimator of θ. Say we observe $X = 3$, what is the maximum likelihood estimate of θ?
- Let $X_1, \ldots, X_n$ be i.i.d. $N\left(\mu, \sigma^2\right)$.
- Find the MLE of μ when σ^2 is known.
- Find the MLE of μ and σ^2 (both unknown).
- Let $X_1, \ldots, X_n$ be i.i.d. Uniform $[0, \theta]$, where $\theta > 0$. Find $\hat{\theta}$
- Let $X_1, \ldots, X_n$ be i.i.d. Uniform $[\theta, \theta + 1]$. Find $\hat{\theta}$

Limitations of the MLE

- Does not always exist
- Is not always appropriate
- Is not always unique

Practice Exercises

7.5 2, 4, 6