Lecture Summary

Example

- Let $X_1, \ldots, X_n$ be a random sample from $Bern(\theta)$.
- $f_n(x_1, \ldots, x_n; \theta) = \Pi_{i=1}^{n} \theta^{x_i}(1-\theta)^{1-x_i}$
- $\mathcal{LL}(x_1, \ldots, x_n; \theta) =$
- Set $d\mathcal{LL}/d\theta = 0$

Computation

- In many practical situations the maximization we need is not available analytically or too cumbersome.
- There exist many numerical optimization methods, Newton's Method (see definition 7.6.2) is one example.

Example: $\mathit{Gamma}(\alpha, 1)$

Assume you have samples $x_1, \ldots, x_n$ from a $\mathit{Gamma}(\alpha, 1)$ distribution. Find $\hat{\alpha}$

Method of Moments

Method of Moments (MOM)

- Let $X_1, \ldots, X_n$ be i.i.d. from $f(\boldsymbol{x} \mid \boldsymbol{\theta})$ where θ is k dimensional.
- The j^{th} sample moment is defined as $m_j = \frac{1}{n}\sum_{i=1}^{n} X_i^j$

Method of moments (MOM) estimator: match the theoretical moments and the sample moments and solve for parameters:

$$m_1 = E\left(X_1 \mid \theta\right), m_2 = E\left(X_1^2 \mid \theta\right), \ldots, m_k = E\left(X_1^k \mid \theta\right)$$

- Example: - Let $X_1, \ldots, X_n$ be i.i.d. Gamma(α, β). Then

$$E(X) = \frac{\alpha}{\beta} \quad \text{and} \quad E\left(X_2\right) = \frac{\alpha(\alpha+1)}{\beta^2}$$

Find the MOM estimator of α and β.
Find the MOM estimator for $Uniform([\theta, \theta+1])$

Properties of Estimators

Reminder

An estimator $\hat{\theta}_n = g(X_1, \ldots, X_n)$ is a function of random variables $X_1, \ldots, X_n$ and therefore has a distribution. The distribution of $\hat{\theta}_n$ is called sampling distribution.

Unbiased estimator

An estimator is unbiased if $E(\hat{\theta}_n) = \theta$.
$E(\hat{\theta}_n) - \theta$ is called the *bias* of the estimator.

Consistent estimator

An estimator is consistent if $\theta_n \xrightarrow{p} \theta$.

Example

Bernoulli MLE $\hat{\theta}_{MLE} = \sum_{i=1}^{n} x_i$. Is it unbiased? Is it consistent?

Properties of Estimators (II)

Asymptotically normal

An estimator is asymptotically normal if

$$\hat{\theta}_n \xrightarrow{d}\xrightarrow{d} \mathcal{N}(\theta, Var(\hat{\theta}_n))$$

Mean squared error of an estimator

The mean squared error of an estimator is

$$E[(\hat{\theta}_n - \theta)^2] = bias^2(\hat{\theta}_n) + Var(\hat{\theta}_n)$$

Example

What is the mean squared error of the Bernoulli MLE estimator?

Properties of MLE estimators

Under some regularity conditions, the MLE estimators are

- Consistent
- Asymptotically Normal

The MLE estimators also have another property:

Invariance

If $\hat{\theta}_n$ is the MLE for θ, $g(\hat{\theta}_n)$ is the MLE for $g(\theta)$.

Properties of MOM estimators

Under some regularity conditions, the MOM estimators are

- Consistent
- Asymptotically Normal

Example

Find the following estimators for a Bernoulli distribution:

- MLE estimator $\hat{\theta}_{MLE}$
- MOM estimator $\hat{\theta}_{MOM}$
- Bayesian estimator that minimizes squared error loss $\hat{\theta}_{Bayes}$
- Are they unbiased?
- What is the mean squared error for each of the estimators?
- Are they consistent?