

Exam 2 Review

18.05 Spring 2014

Jeremy Orloff and Jonathan Bloom

- Cannot cover everything.
- Jerry will go through a list of examples.
- Jon, Peter, Ruthi, Erika will take questions.

Summary

- Data: x_1, \dots, x_n
- Basic statistics: sample mean, sample variance, sample median
- Likelihood, maximum likelihood estimate (MLE)
- Bayesian updating: prior, likelihood, posterior, predictive probability, probability intervals; prior and likelihood can be discrete or continuous
- NHST: H_0 , H_A , significance level, rejection region, power, type 1 and type 2 errors, p -values.

Basic statistics

Data: x_1, \dots, x_n .

$$\text{sample mean} = \bar{x} = \frac{x_1 + \dots + x_n}{n}$$

$$\text{sample variance} = s^2 = \frac{\sum_{i=1}^n (x_i - \bar{x})^2}{n - 1}$$

sample median = middle value

Example. Data: 1, 2, 3, 6, 8.

$$\bar{x} = 4, \quad s^2 = \frac{9+4+1+4+16}{4} = 8.5, \quad \text{median} = 3.$$

Likelihood

x = data

θ = parameter of interest or hypotheses of interest

Likelihood:

$p(x | \theta)$ (discrete distribution)

$f(x | \theta)$ (continuous distribution)

Log likelihood :

$\ln(p(x | \theta)).$

$\ln(f(x | \theta)).$

Likelihood examples

Examples. Find the likelihood function of each of the following.

1. Coin with probability of heads θ . Toss 10 times get 3 heads.
2. Wait time follows $\text{exp}(\lambda)$. In 5 trials wait 3,5,4,5,2
3. Usual 5 dice. Two rolls, 9, 5. (Likelihood given in a table)
4. $x_1, \dots, x_n \sim N(\mu, \sigma^2)$
5. $x = 6$ drawn from $\text{uniform}(0, \theta)$
6. $x \sim \text{uniform}(0, \theta)$

MLE

Methods for finding the maximum likelihood estimate (MLE).

- Discrete hypotheses: compute each likelihood
- Discrete hypotheses: maximum is obvious
- Continuous parameter: compute derivative (often use log likelihood)
- Continuous parameter: maximum is obvious

Examples. Find the MLE for each of the examples in the previous slide.

Bayesian updating: discrete prior-discrete likelihood

Jon has 1 four-side, 2 six-sided, 2 eight-sided, 2 twelve sided, and 1 twenty-sided dice. He picks one at random and rolls a 7.

- 1 For each type of die, find the posterior probability Jon chose that type.
- 2 What are the posterior odds Jon chose the 20-sided die?
- 3 Compute the prior predictive probability of rolling a 7 on roll 1.
- 4 Compute the posterior predictive probability of rolling a 8 on roll 2.

Bayesian updating: conjugate priors

1. Beta prior, binomial likelihood

Data: $x \sim \text{binomial}(n, \theta)$. θ is unknown.

Prior: $f(\theta) \sim \text{beta}(a, b)$

Posterior: $f(\theta | x) \sim \text{beta}(a + x, b + n - x)$

Example. Suppose $x \sim \text{binomial}(30, \theta)$, $x = 12$.

If we have a prior $f(\theta) \sim \text{beta}(1, 1)$ find the posterior.

2. Beta prior, geometric likelihood

Data: x

Prior: $f(\theta) \sim \text{beta}(a, b)$

Posterior: $f(\theta | x) \sim \text{beta}(a + x, b + 1)$.

Example. Suppose $x \sim \text{geometric}(\theta)$, $x = 6$.

If we have a prior $f(\theta) \sim \text{beta}(4, 2)$ find the posterior.

Normal-normal

3. Normal prior, normal likelihood:

$$a = \frac{1}{\sigma_{\text{prior}}^2} \qquad b = \frac{n}{\sigma^2}$$
$$\mu_{\text{post}} = \frac{a\mu_{\text{prior}} + b\bar{x}}{a + b}, \qquad \sigma_{\text{post}}^2 = \frac{1}{a + b}.$$

Example. In the population IQ is normally distributed:

$$\theta \sim N(100, 15^2).$$

An IQ test finds a person's 'true' IQ + random error $\sim N(0, 10^2)$.

Someone takes the test and scores 120.

Find the posterior pdf for this person's IQ:

Bayesian updating: continuous prior-continuous likelihood

Examples. Update from prior to posterior for each of the following with the given data. Graph the prior and posterior in each case.

1. Romeo is late:

likelihood: $x \sim U(0, \theta)$, prior: $U(0, 1)$.

data: 0.3, 0.4, 0.4

2. Waiting times:

likelihood: $x \sim \exp(\lambda)$, prior: $\lambda \sim \exp(2)$.

data: 1, 2

3. Waiting times:

likelihood: $x \sim \exp(\lambda)$, prior: $\lambda \sim \exp(2)$.

data: x_1, x_2, \dots, x_n

NHST: Steps

- 1 Specify H_0 and H_A .
- 2 Choose a significance level α .
- 3 Choose a test statistic and determine the null distribution.
- 4 Determine how to compute a p -value and/or the rejection region.
- 5 Collect data.
- 6 Compute p -value or check if test statistic is in the rejection region.
- 7 Reject or fail to reject H_0 .

NHST: probability tables

Show tables if needed.

NHST: One-sample t -test

- Data: we assume normal data with both μ and σ unknown:

$$x_1, x_2, \dots, x_n \sim N(\mu, \sigma^2).$$

- Null hypothesis: $\mu = \mu_0$ for some specific value μ_0 .
- Test statistic:

$$t = \frac{\bar{x} - \mu_0}{s/\sqrt{n}}$$

where

$$s^2 = \frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2.$$

- Null distribution: $t(n-1)$

Example: z and one-sample t -test

For both problems use significance level $\alpha = .05$.

Assume the data 2, 4, 4, 10 is drawn from a $N(\mu, \sigma^2)$.

Take $H_0: \mu = 0$; $H_A: \mu \neq 0$.

1. Assume $\sigma^2 = 16$ is known and test H_0 against H_A .
2. Now assume σ^2 is unknown and test H_0 against H_A .

Two-sample t -test: equal variances

Data: we assume normal data with μ_x, μ_y and (same) σ unknown:

$$x_1, \dots, x_n \sim \mathbf{N}(\mu_x, \sigma^2), \quad y_1, \dots, y_m \sim \mathbf{N}(\mu_y, \sigma^2)$$

Null hypothesis H_0 : $\mu_x = \mu_y$.

Pooled variance:
$$s_p^2 = \frac{(n-1)s_x^2 + (m-1)s_y^2}{n+m-2} \left(\frac{1}{n} + \frac{1}{m} \right).$$

Test statistic:
$$t = \frac{\bar{x} - \bar{y}}{s_p}$$

Null distribution: $f(t | H_0)$ is the pdf of $T \sim t(n+m-2)$

Example: two-sample t -test

We have data from 1408 women admitted to a maternity hospital for (i) medical reasons or through (ii) unbooked emergency admission. The duration of pregnancy is measured in complete weeks from the beginning of the last menstrual period.

(i) Medical: 775 obs. with $\bar{x} = 39.08$ and $s^2 = 7.77$.

(ii) Emergency: 633 obs. with $\bar{x} = 39.60$ and $s^2 = 4.95$

1. Set up and run a two-sample t -test to investigate whether the duration differs for the two groups.
2. What assumptions did you make?

Chi-square test for goodness of fit

Three treatments for a disease are compared in a clinical trial, yielding the following data:

	Treatment 1	Treatment 2	Treatment 3
Cured	50	30	12
Not cured	100	80	18

Use a chi-square test to compare the cure rates for the three treatments

F-test = one-way ANOVA

Like t -test but for n groups of data with m data points each.

$$y_{i,j} \sim N(\mu_i, \sigma^2), \quad y_{i,j} = j^{\text{th}} \text{ point in } i^{\text{th}} \text{ group}$$

Assumptions: data for each group is an independent normal sample with (possibly) different means but the same variance.

Null-hypothesis is that means are all equal: $\mu_1 = \dots = \mu_n$

Test statistic is $\frac{MS_B}{MS_W}$ where:

$$MS_B = \text{between group variance} = \frac{m}{n-1} \sum (\bar{y}_i - \bar{y})^2$$

$$MS_W = \text{within group variance} = \text{sample mean of } s_1^2, \dots, s_n^2$$

Idea: If μ_i are equal, this ratio should be near 1.

Null distribution is F-statistic with $n-1$ and $n(m-1)$ d.o.f.:

$$\frac{MS_B}{MS_W} \sim F_{n-1, n(m-1)}$$

ANOVA example

The table shows recovery time in days for three medical treatments.

1. Set up and run an F-test.
2. Based on the test, what might you conclude about the treatments?

T_1	T_2	T_3
6	8	13
8	12	9
4	9	11
5	11	8
3	6	7
4	8	12

For $\alpha = .05$, the critical value of $F_{2,15}$ is 3.68.

NHST: some key points

1. α is not the probability of being wrong overall. It's the probability of being wrong if the null hypothesis is true.
2. Likewise, power is not a probability of being right. It's the probability of being right if a particular alternate hypothesis is true.

MIT OpenCourseWare
<http://ocw.mit.edu>

18.05 Introduction to Probability and Statistics

Spring 2014

For information about citing these materials or our Terms of Use, visit: <http://ocw.mit.edu/terms>.