Stat 710: Mathematical Statistics
Lecture 30

Jun Shao

Department of Statistics
University of Wisconsin
Madison, WI 53706, USA

Example 6.24 ($r \times c$ contingency tables)

The following $r \times c$ contingency table is a natural extension of the 2×2 contingency table considered in Example 6.12:

	A_1	A_2	$\cdots$	A_c	Total
B_1	X_{11}	X_{12}	$\cdots$	X_{1c}	n_1
B_2	X_{21}	X_{22}	$\cdots$	X_{2c}	n_2
$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$
B_r	X_{r1}	X_{r2}	$\cdots$	X_{rc}	n_r
Total	m_1	m_2	$\cdots$	m_c	n

where A_i's are disjoint events with $A_1 \cup \cdots \cup A_c = \Omega$ (the sample space of a random experiment), B_i's are disjoint events with $B_1 \cup \cdots \cup B_r = \Omega$, and X_{ij} is the observed frequency of the outcomes in $A_j \cap B_i$.

There are two important applications in this problem.

- testing independence of $\{A_j : j = 1,...,c\}$ and $\{B_i : i = 1,...,r\}$;
- testing equality of multinomial distributions.

Lecture 30: Contingency tables and Bayes tests

Example 6.24 ($r \times c$ contingency tables)

The following $r \times c$ contingency table is a natural extension of the 2×2 contingency table considered in Example 6.12:

	A_1	A_2	$\cdots$	A_c	Total
B_1	X_{11}	X_{12}	$\cdots$	X_{1c}	n_1
B_2	X_{21}	X_{22}	$\cdots$	X_{2c}	n_2
$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$	$\cdots$
B_r	X_{r1}	X_{r2}	$\cdots$	X_{rc}	n_r
Total	m_1	m_2	$\cdots$	m_c	n

where A_i's are disjoint events with $A_1 \cup \cdots \cup A_c = \Omega$ (the sample space of a random experiment), B_i's are disjoint events with $B_1 \cup \cdots \cup B_r = \Omega$, and X_{ij} is the observed frequency of the outcomes in $A_j \cap B_i$.

There are two important applications in this problem.

- testing independence of $\{A_j : j = 1, ..., c\}$ and $\{B_i : i = 1, ..., r\}$;
- testing equality of multinomial distributions.

Testing independence

Testing independence of $\{A_j : j = 1, ..., c\}$ and $\{B_i : i = 1, ..., r\}$ is equivalent to testing hypotheses

$$H_0 : p_{ij} = p_{i\cdot}p_{\cdot j} \;\text{ for all } i,j \qquad \text{versus} \qquad H_1 : p_{ij} \neq p_{i\cdot}p_{\cdot j} \;\text{ for some } i,j,$$

where $p_{ij} = P(A_j \cap B_i) = E(X_{ij})/n$, $p_{i\cdot} = P(B_i)$, and $p_{\cdot j} = P(A_j)$, $i = 1, ..., r$, $j = 1, ..., c$.

In this case, $X = (X_{ij}, i = 1, ..., r, j = 1, ..., c)$ has the multinomial distribution with parameters p_{ij}, $i = 1, ..., r$, $j = 1, ..., c$.

Under H_0, MLE's of $p_{i\cdot}$ and $p_{\cdot j}$ are $\bar{X}_{i\cdot} = n_i/n$ and $\bar{X}_{\cdot j} = m_j/n$, respectively, $i = 1, ..., r$, $j = 1, ..., c$ (exercise).

The number of free parameters is $rc - 1$.

Under H_0, the number of free parameters is $r - 1 + c - 1 = r + c - 2$.

The difference of the two is $rc - r - c + 1 = (r-1)(c-1)$.

By Theorem 6.9, the χ^2-test rejects H_0 when $\chi^2 > \chi^2_{(r-1)(c-1),\alpha}$, where

$$\chi^2 = \sum_{i=1}^{r}\sum_{j=1}^{c} \frac{(X_{ij} - n\bar{X}_{i\cdot}\bar{X}_{\cdot j})^2}{n\bar{X}_{i\cdot}\bar{X}_{\cdot j}}$$

Testing independence

and $\chi^2_{(r-1)(c-1),\alpha}$ is the $(1-\alpha)$th quantile of the chi-square distribution $\chi^2_{(r-1)(c-1)}$.

One can also obtain the modified χ^2-test by replacing $n\bar{X}_{i\cdot}\bar{X}_{\cdot j}$ by X_{ij} in the denominator of each term of the sum in χ^2.

Testing equality of multinomial distributions

Suppose that $(X_{1j},...,X_{rj})$, $j=1,...,c$, are c independent random vectors having the multinomial distributions with parameters $(p_{1j},...,p_{rj})$, $j=1,...,c$, respectively.

Consider the problem of testing whether c multinomial distributions are the same, i.e.,

$$H_0: p_{ij} = p_{i1} \text{ for all } i,j \qquad \text{versus} \qquad H_1: p_{ij} \neq p_{i1} \text{ for some } i,j.$$

Since $(X_{1j},...,X_{rj})$ has the multinomial distribution with size n_j and probability vector $(p_{1j},...,p_{rj})$, the MLE of p_{ij} is X_{ij}/n.

Let $Y_i = \sum_{j=1}^{c} X_{ij}$.

Testing independence

and $\chi^2_{(r-1)(c-1),\alpha}$ is the $(1-\alpha)$th quantile of the chi-square distribution $\chi^2_{(r-1)(c-1)}$.

One can also obtain the modified χ^2-test by replacing $n\bar{X}_{i\cdot}\bar{X}_{\cdot j}$ by X_{ij} in the denominator of each term of the sum in χ^2.

Testing equality of multinomial distributions

Suppose that $(X_{1j},...,X_{rj})$, $j=1,...,c$, are c independent random vectors having the multinomial distributions with parameters $(p_{1j},...,p_{rj})$, $j=1,...,c$, respectively.

Consider the problem of testing whether c multinomial distributions are the same, i.e.,

$$H_0: p_{ij} = p_{i1} \text{ for all } i,j \qquad \text{versus} \qquad H_1: p_{ij} \neq p_{i1} \text{ for some } i,j.$$

Since $(X_{1j},...,X_{rj})$ has the multinomial distribution with size n_j and probability vector $(p_{1j},...,p_{rj})$, the MLE of p_{ij} is X_{ij}/n.

Let $Y_i = \sum_{j=1}^{c} X_{ij}$.

Testing equality of multinomial distributions

Under H_0, $(Y_1,...,Y_r)$ has the multinomial distribution with size n and probability vector $(p_{11},...,p_{r1})$.
Hence, the MLE of p_{i1} under H_0 is $\bar{X}_{i\cdot} = Y_i/n$.
Note that $n_j = n\bar{X}_{\cdot j}$, $j = 1,...,c$.
Hence, under H_0, the MLE of the expected (i,j)th frequency is $n\bar{X}_{i\cdot}\bar{X}_{\cdot j}$.
The number of free parameters in this case is $c(r-1)$.
Under H_0, the number of free parameters is $r-1$.
The difference of the two is $c(r-1)-(r-1) = (r-1)(c-1)$.
Hence, by Theorem 6.9, $\chi^2 \to_d \chi^2_{(r-1)(c-1)}$ under H_0, where χ^2 is the same as that in testing independence.
The rejection region of the χ^2-test is still $\chi^2 > \chi^2_{(r-1)(c-1),\alpha}$.

LR tests

One can also obtain the LR test in this problem.
When $r = c = 2$, the LR test is equivalent to Fisher's exact test given in Example 6.12, which is a UMPU test.
When $r > 2$ or $c > 2$, however, a UMPU test does not exist.

Testing equality of multinomial distributions

Under H_0, $(Y_1, ..., Y_r)$ has the multinomial distribution with size n and probability vector $(p_{11}, ..., p_{r1})$.
Hence, the MLE of p_{i1} under H_0 is $\bar{X}_{i\cdot} = Y_i/n$.
Note that $n_j = n\bar{X}_{\cdot j}$, $j = 1, ..., c$.
Hence, under H_0, the MLE of the expected (i,j)th frequency is $n\bar{X}_{i\cdot}\bar{X}_{\cdot j}$.
The number of free parameters in this case is $c(r-1)$.
Under H_0, the number of free parameters is $r-1$.
The difference of the two is $c(r-1) - (r-1) = (r-1)(c-1)$.
Hence, by Theorem 6.9, $\chi^2 \to_d \chi^2_{(r-1)(c-1)}$ under H_0, where χ^2 is the same as that in testing independence.
The rejection region of the χ^2-test is still $\chi^2 > \chi^2_{(r-1)(c-1),\alpha}$.

LR tests

One can also obtain the LR test in this problem.
When $r = c = 2$, the LR test is equivalent to Fisher's exact test given in Example 6.12, which is a UMPU test.
When $r > 2$ or $c > 2$, however, a UMPU test does not exist.

Bayes tests

An LR test actually compares $\sup_{\theta\in\Theta_0}\ell(\theta)$ with $\sup_{\theta\in\Theta_1}\ell(\theta)$ for testing $H_0 : \theta \in \Theta_0$ versus $H_1 : \theta \in \Theta_1$.

Instead of comparing two maximum values, one may compare two averages such as

$$\widehat{\pi}_j = \int_{\Theta_j} \ell(\theta) d\Pi(\theta) \Big/ \int_{\Theta} \ell(\theta) d\Pi(\theta), \quad j = 0, 1,$$

where $\Pi(\theta)$ is a c.d.f. on Θ, and reject H_0 when $\widehat{\pi}_1 > \widehat{\pi}_0$.

If Π is treated as a prior c.d.f., then $\widehat{\pi}_j$ is the posterior probability of Θ_j, and this test is a particular Bayes action (see Exercise 18 in §4.6) and is called a Bayes test.

In Bayesian analysis, one often considers the *Bayes factor* defined as

$$\beta = \frac{\text{posterior odds ratio}}{\text{prior odds ratio}} = \frac{\widehat{\pi}_0/\widehat{\pi}_1}{\pi_0/\pi_1},$$

where $\pi_j = \Pi(\Theta_j)$ is the prior probability of Θ_j.

If there is a statistic sufficient for θ, then the Bayes test and Bayes factor depend only on the sufficient statistic.

Bayes tests

Consider the special case where $\Theta_0 = \{\theta_0\}$ and $\Theta_1 = \{\theta_1\}$ are simple hypotheses.

For given $X = x$,

$$\widehat{\pi}_j = \frac{\pi_j f_{\theta_j}(x)}{\pi_0 f_{\theta_0}(x) + \pi_1 f_{\theta_1}(x)}.$$

Rejecting H_0 when $\widehat{\pi}_1 > \widehat{\pi}_0$ is the same as rejecting H_0 when

$$\frac{f_{\theta_1}(x)}{f_{\theta_0}(x)} > \frac{\pi_0}{\pi_1}.$$

This is equivalent to the UMP test T_* in Theorem 6.1 with $c = \pi_0/\pi_1$ and $\gamma = 0$.

The Bayes factor in this case is

$$\beta = \frac{\widehat{\pi}_0 \pi_1}{\widehat{\pi}_1 \pi_0} = \frac{f_{\theta_0}(x)}{f_{\theta_1}(x)}.$$

Thus, the UMP test T_* in Theorem 6.1 is equivalent to the test that rejects H_0 when the Bayes factor is small.

The rejection region of the Bayes test depends on prior probabilities, whereas the Bayes factor does not.

Bayes tests

When either Θ_0 or Θ_1 is not simple, however, Bayes factors also depend on the prior Π.

If Π is an improper prior, the Bayes test is still defined as long as the posterior probabilities $\widehat{\pi}_j$ are finite.

However, the Bayes factor may not be well defined when Π is improper.

Example 6.25

Let $X_1, ..., X_n$ be i.i.d. from $N(\mu, \sigma^2)$ with an unknown $\mu \in \mathscr{R}$ and a known $\sigma^2 > 0$.

Let the prior of μ be $N(\xi, \tau^2)$.

Then the posterior of μ is $N(\mu_*(x), c^2)$, where

$$\mu_*(x) = \frac{\sigma^2}{n\tau^2+\sigma^2}\xi + \frac{n\tau^2}{n\tau^2+\sigma^2}\bar{x} \qquad \text{and} \qquad c^2 = \frac{\tau^2\sigma^2}{n\tau^2+\sigma^2}$$

(see Example 2.25).

Consider first the problem of testing $H_0 : \mu \le \mu_0$ versus $H_1 : \mu > \mu_0$.

Let Φ be the c.d.f. of the standard normal.

Bayes tests

When either Θ_0 or Θ_1 is not simple, however, Bayes factors also depend on the prior Π.

If Π is an improper prior, the Bayes test is still defined as long as the posterior probabilities $\widehat{\pi}_j$ are finite.

However, the Bayes factor may not be well defined when Π is improper.

Example 6.25

Let $X_1,...,X_n$ be i.i.d. from $N(\mu,\sigma^2)$ with an unknown $\mu \in \mathscr{R}$ and a known $\sigma^2 > 0$.

Let the prior of μ be $N(\xi,\tau^2)$.

Then the posterior of μ is $N(\mu_*(x),c^2)$, where

$$\mu_*(x) = \frac{\sigma^2}{n\tau^2+\sigma^2}\xi + \frac{n\tau^2}{n\tau^2+\sigma^2}\bar{x} \qquad \text{and} \qquad c^2 = \frac{\tau^2\sigma^2}{n\tau^2+\sigma^2}$$

(see Example 2.25).

Consider first the problem of testing $H_0 : \mu \le \mu_0$ versus $H_1 : \mu > \mu_0$.

Let Φ be the c.d.f. of the standard normal.

Example 6.25 (continued)

Then the posterior probability of Θ_0 and the Bayes factor are, respectively,

$$\widehat{\pi}_0 = \Phi\left(\frac{\mu_0 - \mu_*(x)}{c}\right) \qquad \text{and} \qquad \beta = \frac{\Phi\left(\frac{\mu_0 - \mu_*(x)}{c}\right)\Phi\left(\frac{\xi - \mu_0}{\tau}\right)}{\Phi\left(\frac{\mu_*(x) - \mu_0}{c}\right)\Phi\left(\frac{\mu_0 - \xi}{\tau}\right)}.$$

It is interesting to see that if we let $\tau \to \infty$, which is the same as considering the improper prior $\Pi =$ the Lebesgue measure on $\mathscr{R}$, then

$$\widehat{\pi}_0 \to \Phi\left(\frac{\mu_0 - \bar{x}}{\sigma/\sqrt{n}}\right),$$

which is exactly the p-value $\widehat{\alpha}(x)$ derived in Example 2.29.

Consider next the problem of testing $H_0 : \mu = \mu_0$ versus $H_1 : \mu \neq \mu_0$.
In this case the prior c.d.f. cannot be continuous at μ_0.
We consider

$$\Pi(\mu) = \pi_0 I_{[\mu_0,\infty)}(\mu) + (1 - \pi_0)\Phi\left(\frac{\mu - \xi}{\tau}\right).$$

Let $\ell(\mu)$ be the likelihood function based on $\bar{x}$.

Example 6.25 (continued)

Then the posterior probability of Θ_0 and the Bayes factor are, respectively,

$$\widehat{\pi}_0 = \Phi\left(\frac{\mu_0 - \mu_*(x)}{c}\right) \qquad \text{and} \qquad \beta = \frac{\Phi\left(\frac{\mu_0 - \mu_*(x)}{c}\right)\Phi\left(\frac{\xi - \mu_0}{\tau}\right)}{\Phi\left(\frac{\mu_*(x) - \mu_0}{c}\right)\Phi\left(\frac{\mu_0 - \xi}{\tau}\right)}.$$

It is interesting to see that if we let $\tau \to \infty$, which is the same as considering the improper prior $\Pi =$ the Lebesgue measure on $\mathscr{R}$, then

$$\widehat{\pi}_0 \to \Phi\left(\frac{\mu_0 - \bar{x}}{\sigma/\sqrt{n}}\right),$$

which is exactly the p-value $\widehat{\alpha}(x)$ derived in Example 2.29.

Consider next the problem of testing $H_0 : \mu = \mu_0$ versus $H_1 : \mu \neq \mu_0$.
In this case the prior c.d.f. cannot be continuous at μ_0.
We consider

$$\Pi(\mu) = \pi_0 I_{[\mu_0,\infty)}(\mu) + (1 - \pi_0)\Phi\left(\frac{\mu - \xi}{\tau}\right).$$

Let $\ell(\mu)$ be the likelihood function based on $\bar{x}$.

Example 6.25 (continued)

Then

$$m_1(x) = \int_{\mu \neq \mu_0} \ell(\mu) d\Phi\left(\frac{\mu-\xi}{\tau}\right) = \frac{1}{\sqrt{\tau^2+\sigma^2/n}} \Phi'\left(\frac{\bar{x}-\xi}{\sqrt{\tau^2+\sigma^2/n}}\right),$$

where $\Phi'(t)$ is the p.d.f. of the standard normal distribution, and

$$\hat{\pi}_0 = \frac{\pi_0 \ell(\mu_0)}{\pi_0 \ell(\mu_0) + (1-\pi_0) m_1(x)} = \left(1 + \frac{1-\pi_0}{\pi_0 \beta}\right)^{-1},$$

where

$$\beta = \frac{\ell(\mu_0)}{m_1(x)} = \frac{\sqrt{n\tau^2+\sigma^2}\Phi'\left(\frac{\bar{x}-\mu_0}{\sigma/\sqrt{n}}\right)}{\sigma\Phi'\left(\frac{\bar{x}-\xi}{\sqrt{\tau^2+\sigma^2/n}}\right)}$$

is the Bayes factor.

More discussions about Bayesian hypothesis tests can be found in Berger (1985, §4.3.3).

Example 6.25 (continued)

Then

$$m_1(x) = \int_{\mu \neq \mu_0} \ell(\mu) d\Phi\left(\frac{\mu-\xi}{\tau}\right) = \frac{1}{\sqrt{\tau^2+\sigma^2/n}} \Phi'\left(\frac{\bar{x}-\xi}{\sqrt{\tau^2+\sigma^2/n}}\right),$$

where $\Phi'(t)$ is the p.d.f. of the standard normal distribution, and

$$\widehat{\pi}_0 = \frac{\pi_0 \ell(\mu_0)}{\pi_0 \ell(\mu_0) + (1-\pi_0) m_1(x)} = \left(1 + \frac{1-\pi_0}{\pi_0 \beta}\right)^{-1},$$

where

$$\beta = \frac{\ell(\mu_0)}{m_1(x)} = \frac{\sqrt{n\tau^2+\sigma^2}\,\Phi'\left(\frac{\bar{x}-\mu_0}{\sigma/\sqrt{n}}\right)}{\sigma \Phi'\left(\frac{\bar{x}-\xi}{\sqrt{\tau^2+\sigma^2/n}}\right)}$$

is the Bayes factor.

More discussions about Bayesian hypothesis tests can be found in Berger (1985, §4.3.3).