

ETC 2420/5242 Lab 9 2017

SOLUTION

Week 9

Purpose

This lab is related to Bayesian thinking, to compute conditional probabilities and posterior probabilities.

Warmup

Watch the video “A visual guide to Bayesian thinking”. And also the snippet from “Big Bang Theory”.

Exercise 1 (4pts)

A situation where Bayes theorem is routinely used is the spam filter in your mail server. The message is scrutinized for the appearance of key words which make it likely that the message is spam. Let us describe how one of these filters might work. We imagine that the evidence for spam is that the subject message of the mail contains the word “easy”. We define events *spam* (the message is spam) and *easy* (the subject line contains this word).

From previous examination of my mail I have learned that 70% of emails are spam, 2% of spam email have “easy” in the subject line, and .1% of non-spam emails have this word in the subject line.

Use Bayes Theorem to tackle to find $P(\text{spam}|\text{free})$.

Let A=mail is spam, B=mail has “easy” in subject. What is

- $P(A)$? 0.70
- $P(B|A)$? 0.02
- $P(B)$? 0.0143
- $P(A|B)$? $0.7 \times 0.02 / 0.0143 = 0.4903$

Exercise 2 (8pts)

For the dice problem explained in the lecture, write down the problem in tabular format, and specify $P(D)$.

H	P(H)	P(D H)	P(H)P(D H)	P(H D)
4	1/5	if $D < 5$ 1/4 else 0	if $D < 5$ 1/20 else 0	if $D < 5$ 0.5 else 0
6	1/5	if $D < 7$ 1/6 else 0	if $D < 7$ 1/30 else 0	if $D < 5$ 0.333, else if $5 \leq D < 7$ 0.417, else 0

H	P(H)	P(D H)	P(H)P(D H)	P(H D)
8	1/5	if D<9 1/8 else 0	if D<9 1/40 else 0	if D<5 0.25, else if 5<=D<7 0.3125, else if 7<=D<9 0.417 else 0
12	1/5	if D<13 1/12 else 0	if D<13 1/60 else 0	if D<5 0.167, else if 5<=D<7 0.208, else if 7<=D<9 0.278, else if 9<=D<13 0.333, else 0
20	1/5	1/20	1/100	if D<5 0.100, else if 5<=D<7 0.125, else if 7<=D<9 0.167, else if 9<=D<13 0.25, else 0.5

P(D) has different probabilities depending on the frequency of the numbers, 0.1, 0.1, 0.1, 0.1, 0.08, 0.08, 0.06, 0.06, 0.04, 0.04, 0.04, 0.04, 0.04, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02, 0.02.

Exercise 3 (3pts)

Suppose there were 3 20-side dice. Re-do the calculations for P(H|D) if a 6 is thrown. What is P(D)? Prior? Give an explanation for why the probability of the 20-sided dice increases.

P(D) = 0.0777778, 0.0777778, 0.0777778, 0.0777778, 0.0666667, 0.0666667, 0.0555556, 0.0555556, 0.0444444, 0.0444444, 0.0444444, 0.0444444, 0.0333333, 0.0333333, 0.0333333, 0.0333333, 0.0333333, 0.0333333, 0.0333333, 0.0333333, 0.0333333.

prior is now $P(4) = 1/7, P(6) = 1/7, P(8) = 1/7, P(12) = 1/7, P(20) = 3/7$.

```

dice <- c(4, 6, 8, 12, 20)
prior <- c(rep(1/7, 4), 3/7)
prob_D <- c(rep(7/90, 4), rep(6/90, 2), rep(5/90, 2),
            rep(4/90, 4), rep(3/90, 8))

```

```

dice_likelihood <- function(x, dice) {
  if (x>dice)
    l <- 0
  else
    l <- 1/dice
  return(l)
}
dice_posterior <- function(x, dice, prior, prob_D) {
  post <- NULL
  pD <- prod(prob_D[x])
  for (i in 1:length(dice)) {
    post <- c(post, prior[i]*dice_likelihood(x, dice[i])/pD)
  }
  df <- data.frame(dice, post=post/sum(post))
  return(df)
}
dice_posterior(6, dice, prior, prob_D)
#   dice      post
# 1     4 0.0000000
# 2     6 0.3174603
# 3     8 0.2380952
# 4    12 0.1587302
# 5    20 0.2857143

```

There are more 20-sided dice, so it is understandable that even though 6 is small, there's a higher chance that it came from one of the 20-sided dice.

Exercise 4 (5pts)


In the lecture we used a power law prior to dampen down the estimate of the number of locomotives. Instead fit a $N(\mu, 100)$ ($\sigma = 100$) prior. Plot the posterior distributions.

```

loco_likelihood <- function(x, loco){
  if (x>loco)
    l <- 0
  else
    l <- 1/loco
}
l <- NULL
for (i in 1:1000)
  l <- c(l, loco_likelihood(60, i))
df <- data.frame(hyp=1:1000, l, post=l/sum(l))
alpha=0.9
prior <- (1/1:1000)^alpha
df <- data.frame(hyp=1:1000, l, unif=l/sum(l))
df$power <- l*prior
df$power <- df$power/sum(df$power)
prior2 <- exp(-(1:1000/100)^2)/sqrt(2*pi)
df$norm <- l*prior2
df$norm <- df$norm/sum(df$norm)
df_m <- df %>% gather(prior, posterior, unif, power, norm)
ggplot(df_m, aes(x=hyp, y=posterior, color=prior)) + geom_line() + xlab("Number of trains") +
  ylab("Probability") +

```

```
scale_color_brewer(palette="Dark2")
```



Exercise 5 (12pts)

It helps to read the lecture notes to find the equations that you need for this question.

a. Suppose we have a sample $X_1, \dots, X_n \sim N(\theta, 3)$, $\sigma = 3$.

- (2pts) Write down the likelihood function.

$$l(x_1, \dots, x_n | \theta) = \prod_{i=1}^n \exp \left\{ -\frac{1}{2} \left(\frac{x_i - \theta}{3} \right)^2 \right\} = \exp \left\{ -\frac{n}{2} \left(\frac{\bar{x} - \theta}{3} \right)^2 \right\}$$

- (2pts) What is the MLE estimate $\hat{\theta}_{\text{MLE}}$, theoretically? (If we actually had a sample, we could use the function `fitdistr` from the `MASS` library to estimate it.)

$$\bar{X}$$

- (2pts) If we have reason to believe that θ is not a fixed value, but can vary according to $\theta \sim N(\mu, \tau)$, what is the posterior distribution, $\pi(\theta | x_1, \dots, x_n)$?

Plug in 9 instead of σ^2 in the lecture notes

$$\begin{aligned} \pi(\theta | x_1, \dots, x_n) &= \pi(\theta) l(x_1, \dots, x_n | \theta) \\ &= \exp \left\{ -\frac{1}{2} \left(\frac{n}{9} + \frac{1}{\tau^2} \right)^{-1} \left(\theta - \frac{n\bar{x}/9 + \mu/\tau^2}{n/9 + 1/\tau^2} \right)^2 \right\} \end{aligned}$$

- (2pts) What is the posterior mean $E[\theta|x_1, \dots, x_n]$?

$$\frac{n\bar{x}/9 + \mu/\tau^2}{n/9 + 1/\tau^2}$$


- b. Suppose the “true” value is $\theta = 2$ (at one instance of time to allow for simulation). Consider (1) $\mu = 5$ and $\tau = 1$, and (2) $\mu = 2$ and $\tau = 2$. For $n \in \{1, 10, 20, 50, 100, 10000\}$. Use Nat’s shiny app for this.
 - Simulate a data set consisting of n observations
 - Plot on the same graphic $\pi(\theta)$, $\pi(\theta|x_1, \dots, x_n)$ and $\hat{\theta}_{MLE}$.
- c. (4 pts) Discuss the behavior of $\pi(\theta|x_1, \dots, x_n)$ as n increases and the impact of the prior distribution.

As n increases the prior has less influence on the parameter estimates. In the first example the prior is very different from the “true” value, and for small sample sizes pulls the estimates away from this true value. In the second example, the prior is quite close to the “true” value, and this is reflected in the posterior distributions being centered around the “true” value regardless of the sample size.

Exercise 6 (8pts)

Suppose your prior distribution for θ , the proportion of Australians who will vote yes on the Marriage Equality Postal Survey, is beta with mean 0.5 and standard deviation 0.4.

- a. (2pts) What are the parameters α and β of your prior distribution. Plot the prior density function.



- b. (2pts) Why might it be reasonable to consider θ in a Bayesian way?

People might change their viewpoint depending on information available when they are asked, or when they turn their vote in.

- c. (4pts) A random sample of 1000 Australians is taken, and 65% say they will vote “yes”. What are the posterior mean and variance for θ ? Plot the posterior density function.


Mean and variance of a beta distribution (**Beta(a, b)**), are given by $\frac{a}{a+b}$ and $\frac{ab}{(a+b)^2(a+b+1)}$.

Posterior is given by **Beta(a+650, b+350)**

```

a <- 0.5+650
b <- 0.4+1000-650
a/(a+b)
# [1] 0.6499151
a*b/((a+b)^2*(a+b+1))
# [1] 0.000227094
betabin_posterior <- function(x, n, a, b) {
  p <- seq(0, 1, 0.01)
  post <- dbeta(p, a+x, b+n-x)
  df <- data.frame(p, post)
  return(df)
}
df <- betabin_posterior(650, 1000, 0.5, 0.4)
ggplot(df, aes(x=p, y=post)) + geom_line()

```



TURN IN

- Your .Rmd file
- Your html file that results from knitting the Rmd.