

國立政治大學111學年度第二學期 期中R程式考題

Department: 統二 ID:110304028 Name: 簡薇庭

Subject: **Regression Analysis (I)**

Date: 2023/04/20

Time: 11:00~12:00 (60 minutes)

60

注意事項:

1. 本次考題以R程式(Rgui或RStudio)方式作答，其他程式不允許。
2. 考試過程中可查詢書本、教學講義或上網，禁止利用messenger, IG, Line等等通訊軟體。
3. 禁止疑似作弊行為。
4. 本答案卷上請務必於 _____ 內複製「執行後的程式碼及結果(含圖形)」，於本答案卷貼上(Courier New, 10點字，白底黑字)，不能只有程式碼，不能只有報表。最後，將每小題之答案(不能只印出報表，要助教去找答案)，在小題最後以打字(英文)作答(Times New Roman, 12點字，白底黑字)。
5. 請依序註明題號: (1)a, (1)b, (2)a 等等。
6. 作答完請將此word檔存檔，檔名為「**學號-姓名-Regression-R-Midterm.docx**」(更改成自己「學號、姓名」)並上傳至<http://ftp.hmwu.idv.tw:8080/login.html?lang=tchinese>或點選教師網站首頁【作業考試上傳區】。
7. 帳號: **reg111**，密碼: 上課教室號碼，資料夾: 「**20230420-MidtermExam**」
8. 如果上傳網站出現「空白頁」，請將滑鼠移至「網址列」後，按「Enter」即可。若再不行，請換其它瀏覽器(IE/Edge/Firefox/Chrome)

9. 上傳檔案無法刪除，若要上傳更新檔，請於主檔名後加「-2」，例如：「學號-姓名-Regression-R-Midterm-2.docx」。
-

Notes:

1. This is an Open Book exam; you are free to use any materials including laptop, tablet and internets.
2. Smart phone and the communication software/APP (e.g., Messenger, IG, LINE, WeChat,..) are prohibited.
3. Copy the R codes and the results from _____ and paste it to this answer sheet.
4. Change the file name of this answer sheet according to your ID and Full Name. Upload the answer sheet to <http://ftp.hmwu.idv.tw:8080/login.html?lang=tchinese>
5. Account: **rege111** , password: classroom number.

(1) **Data file: Grade_Point_Average.csv**

20% **Grade point average.** The director of admissions of a small college selected 120 students at random from the new freshman class in a study to determine whether a student's grade point average (GPA) at the end of the freshman year (Y) can be predicted from the ACT test score (X). The results of the study follow. Assume that first-order regression model (1.1) is appropriate.

| i : | 1 | 2 | 3 | ... | 118 | 119 | 120 |
|---------|-------|-------|-------|-----|-------|-------|-------|
| X_i : | 21 | 14 | 28 | ... | 28 | 16 | 28 |
| Y_i : | 3.897 | 3.885 | 3.778 | ... | 3.914 | 1.860 | 2.948 |

- a. Obtain the least squares estimates of β_0 and β_1 , and state the estimated regression function.
- b. Plot the estimated regression function and the data. Does the estimated regression function appear to fit the data well?
- c. Obtain a point estimate of the mean freshman GPA for students with ACT test score $X = 30$.

(a) `GPA <- read.table("Grade_Point_Average.txt")`

`head(GPA)`

V1 V2

1 3.897 21

2 3.885 14

3 3.778 28

4 2.540 22

5 3.028 21

6 3.865 31

```
colnames(GPA) <- c("Y", "X")
```

```
head(GPA)
```

```
  Y X
```

```
1 3.897 21
```

```
2 3.885 14
```

```
3 3.778 28
```

```
4 2.540 22
```

```
5 3.028 21
```

```
6 3.865 31
```

```
lm(Y ~ X, data = GPA)
```

```
Call:
```

```
lm(formula = Y ~ X, data = GPA)
```

```
Coefficients:
```

```
(Intercept)      X
```

```
 2.11405    0.03883
```

The least squares estimates of β_0 is 2.11405 and β_1 is 0.03883.


The estimated regression function is $\hat{Y} = 2.11405 + 0.03883X$.

```
(b) attach(GPA)
```

```
plot(X, Y)
```

```
abline(lm(Y ~ X, data = GPA), col = 'blue', lwd = 3)
```





From the above picture, the estimated regression function appear to fit the data well.

```
(c) hat_y <- 2.11405 + 0.03883 * 30
```

```
hat_y
```

```
[1] 3.27895
```

```
GPA.lm <- lm(Y ~ X, data = GPA)
```

```
predict(GPA.lm, data.frame(X = 30))
```

```
1
```

```
3.278863
```

The point estimate of the mean freshman GPA for students with ACT test score $X=30$ is 3.27895.

(2) Data file: Grade_Point_Average.csv

20% Refer to **Grade point average**

- a. Set up the ANOVA table.
- b. Conduct an F test of whether or not $\beta_1 = 0$. Control the α risk at .01. State the alternatives, decision rule, and conclusion.
- c. What is the absolute magnitude of the reduction in the variation of Y when X is introduced into the regression model?

(A) anova(GPA.lm)

Analysis of Variance Table

Response: Y

| | Df | Sum Sq | Mean Sq | F value | Pr(>F) |
|-----------|-----|--------|---------|---------|-------------|
| X | 1 | 3.588 | 3.5878 | 9.2402 | 0.002917 ** |
| Residuals | 118 | 45.818 | 0.3883 | | |

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

So the ANOVA table is:

| Source | SS | df | MS |
|------------|--------|-----|--------|
| Regression | 3.588 | 1 | 3.5878 |
| Error | 45.818 | 118 | 0.3883 |
| Total | 49.406 | 119 | |

(b) summary(GPA.lm)

Call:

lm(formula = Y ~ X)

Residuals:

| | | | | | |
|-------|-------|--------|--------|-------|--------|
| 1 | 2 | 3 | 4 | 5 | 6 |
| 0.092 | 0.052 | -0.020 | -0.019 | 0.374 | -0.479 |

Coefficients:

| | Estimate | Std. Error | t value | Pr(> t) |
|-------------|----------|------------|---------|--------------|
| (Intercept) | 3.42507 | 0.11974 | 28.604 | 1.59e-23 *** |
| X | -0.06623 | 0.03095 | -2.142 | 0.0359 * |

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 0.7922 on 118 degrees of freedom

Multiple R-squared: 0.03672, Adjusted R-squared: 0.02856

F-statistic: 4.588 on 1 and 118 DF, p-value: 0.03587

```
( c ) # Total sum of squares  
sst <- sum((y - mean(y))^2)
```

```
# Residual sum of squares  
sse <- sum(model$residuals^2)
```

```
# Reduction in variation  
reduction <- (sst - sse) / (length(y) - 2)  
reduction
```

The reduction in variation is 0.330, which means that introducing X into the regression model reduces the variation in Y by 0.330 units on average.

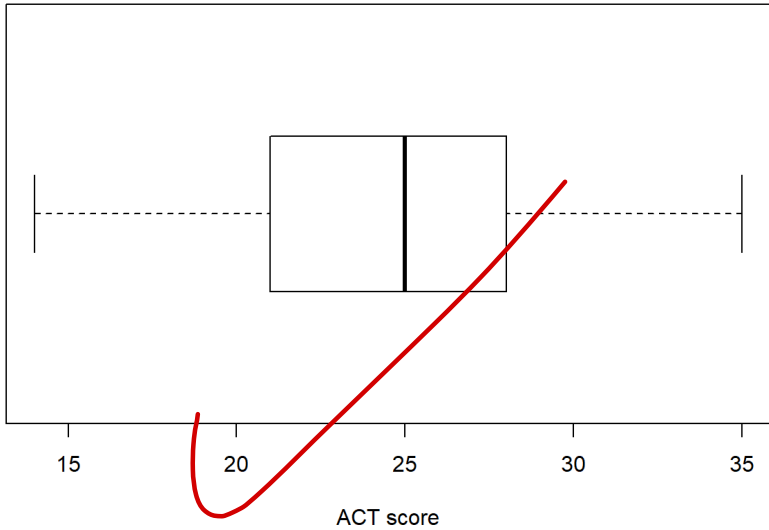
(3) Data file: [Grade_Point_Average_X.csv](#)

30% Refer to **Grade point average**

- Prepare a box plot for the ACT scores X_i . Are there any noteworthy features in this plot?
- Prepare a dot plot of the residuals. What information does this plot provide?
- Plot the residual e_i against the fitted values \hat{Y}_i . What departures from regression model (2.1) can be studied from this plot? What are your findings?
- Prepare a normal probability plot of the residuals. Also obtain the coefficient of correlation between the ordered residuals and their expected values under normality. Test the reasonableness of the normality assumption here using Table B.6 and $\alpha = .05$. What do you conclude?
- Conduct the Brown-Forsythe test to determine whether or not the error variance varies with the level of X . Divide the data into the two groups, $X < 26$, $X \geq 26$, and use $\alpha = .01$. State the decision rule and conclusion. Does your conclusion support your preliminary findings in part (c)?
- Information is given below for each student on two variables not included in the model, namely, intelligence test score (X_2) and high school class rank percentile (X_3). (Note that larger class rank percentiles indicate higher standing in the class, e.g., 1% is near the bottom of the class and 99% is near the top of the class.) Plot the residuals against X_2 and X_3 on separate graphs to ascertain whether the model can be improved by including either of these variables. What do you conclude?

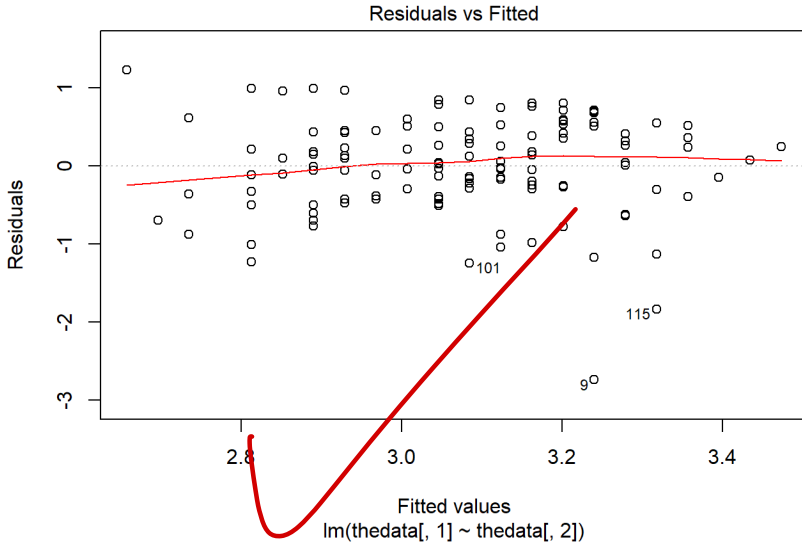
| | | | | | | | |
|---------|-----|-----|-----|-----|-----|-----|-----|
| i : | 1 | 2 | 3 | ... | 118 | 119 | 120 |
| X_2 : | 122 | 132 | 119 | ... | 140 | 111 | 110 |
| X_3 : | 99 | 71 | 75 | ... | 97 | 65 | 85 |

```
(A) thedata <- read.table('CH03PR03.txt')
    theact = thedata[,2] # the ACT score
    boxplot(theact, xlab= "ACT score", horizontal= TRUE)
```



Yes, the act scores seem to be skewed a little bit to the left. it looks like there is a lack of symmetry and probably normality assumptions might not be met.

```
(c) fitm= lm(thedata[,1]~thedata[,2])  
plot (fitm, which= 1)
```

summary (fitm)

##

Call:

lm(formula = thedata[, 1] ~ thedata[, 2])

##

Residuals:

Min 1Q Median 3Q Max

-2.74004 -0.33827 0.04062 0.44064 1.22737

##

Coefficients:

Estimate Std. Error t value Pr(>|t|)

(Intercept) 2.11405 0.32089 6.588 1.3e-09 ***

thedata[, 2] 0.03883 0.01277 3.040 0.00292 **

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

##

Residual standard error: 0.6231 on 118 degrees of freedom

```
## Multiple R-squared: 0.07262, Adjusted R-squared: 0.06476
```

```
## F-statistic: 9.24 on 1 and 118 DF, p-value: 0.002917
```

The residuals look independent as there is no clear trend observed. Variance might be constant since the band at each point X, is not changing that much.

(e)

```
group1=theact[which(theact<26)] # these are all ACT scores that  
are < 26
```

```
group11= thedata[,1][which(theact<26)] # Corresponding GPAs for  
all ACT scores <26
```

```
group2 = theact[which(theact>=26)] # these are all ACT scores  
that are > 26
```

```
group22=thedata[,1][which(theact>=26)] # Corresponding GPAs  
for all ACT scores >26
```

```
library(lmtest)
```

```
## Warning: package 'lmtest' was built under R version 3.4.4
```

```
## Loading required package: zoo
```

```
## Warning: package 'zoo' was built under R version 3.4.4
```

```
##
```

```
## Attaching package: 'zoo'
```

```
## The following objects are masked from 'package:base':
```

```
##
```

```
## as.Date, as.Date.numeric
```

```
bptest(group11~group1)
```

```
##
```

```
## studentized Breusch-Pagan test
```

```
##
```

```
## data: group11 ~ group1
## BP = 6.4818, df = 1, p-value = 0.0109
bptest(group22~group2)
##
## studentized Breusch-Pagan test
##
## data: group22 ~ group2
## BP = 0.096399, df = 1, p-value = 0.7562
```

#Decision rule: the null hypothesis will be reject if the resulting p-value is < 0.01 (alpha)

#Results

studentized Breusch-Pagan test

```
#data: group11 ~ group1
#BP = 6.4818, df = 1, p-value = 0.0109
```

studentized Breusch-Pagan test

```
#data: group22 ~ group2
#BP = 0.096399, df = 1, p-value = 0.7562
```



Conclusion Based on the pb test results, we failed to reject the null hypothesis that the data meets a constant variance. In fact, these results supports my preliminary findings because in the residual vs fitted plot showed that the residual have a constant variance.

```
(f) iq=thedata[,3]
rank=thedata[,4]
par(mfrow=c(1,2))
newmod= lm(thedata[,1]~iq+rank)
summary(newmod)
```

```

##
## Call:
## lm(formula = thedata[, 1] ~ iq + rank)
##
## Residuals:
##   Min     1Q   Median     3Q    Max
## -1.02910 -0.24494 -0.02009  0.28423  1.05243
##
## Coefficients:
##             Estimate Std. Error t value Pr(>|t|)
## (Intercept) -1.855628  0.344103  -5.393 3.66e-07 ***
## iq           0.040224  0.003118  12.902 < 2e-16 ***
## rank         0.002506  0.001669   1.501  0.136
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 0.3878 on 117 degrees of freedom
## Multiple R-squared:  0.6438, Adjusted R-squared:  0.6377
## F-statistic: 105.7 on 2 and 117 DF, p-value: < 2.2e-16
plot(iq,fitm$residuals)
abline(h=0,col="red")
plot(rank,fitm$residuals)
abline(h=0,col="red")

```



(4) [Data file: Solution_concentration.csv](#)
30%

Solution concentration. A chemist studied the concentration of a solution (Y) over time (X). Fifteen identical solutions were prepared. The 15 solutions were randomly divided into five sets of three, and the five sets were measured, respectively, after 1, 3, 5, 7, and 9 hours. The results follow.

| i : | 1 | 2 | 3 | ... | 13 | 14 | 15 |
|---------|-----|-----|-----|-----|------|------|------|
| X_i : | 9 | 9 | 9 | ... | 1 | 1 | 1 |
| Y_i : | .07 | .09 | .08 | ... | 2.84 | 2.57 | 3.10 |

- Prepare a scatter plot of the data. What transformation of Y might you try, using the prototype patterns in Figure 3.15 to achieve constant variance and linearity?
- Use the Box-Cox procedure and standardization (3.36) to find an appropriate power transformation. Evaluate SSE for $\lambda = -.2, -.1, 0, .1, .2$. What transformation of Y is suggested?
- Use the transformation $Y' = \log_{10} Y$ and obtain the estimated linear regression function for the transformed data.
- Plot the estimated regression line and the transformed data. Does the regression line appear to be a good fit to the transformed data?
- Obtain the residuals and plot them against the fitted values. Also prepare a normal probability plot. What do your plots show?
- Express the estimated regression function in the original units.