

國立政治大學 111 學年度第二學期

迴歸分析(一)期末 R 程式加分考題

Department: 風管三 ID: 109308028 Name: 蔡念澄

24

Subject: **Regression Analysis (I)**

Date: 2023/06/15, Time: 11:00~12:00 (60 minutes)

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 **R Console** 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: 「**20230615-FinalExam**」
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 **R Console** 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: **reg111** , password: classroom number.

(1)

30%

Data file: CDI.csv

Refer to the **CDI** data set in Appendix C.2. The number of active physicians (Y) is to be regressed against total population (X_1), total personal income (X_2), and geographic region (X_3, X_4, X_5).

- Fit a first-order regression model. Let $X_3 = 1$ if NE and 0 otherwise, $X_4 = 1$ if NC and 0 otherwise, and $X_5 = 1$ if S and 0 otherwise.
- Examine whether the effect for the northeastern region on number of active physicians differs from the effect for the north central region by constructing an appropriate 90 percent confidence interval. Interpret your interval estimate.
- Test whether any geographic effects are present; use $\alpha = .10$. State the alternatives, decision rule, and conclusion. What is the P -value of the test?

Data Set C.2 CDI

This data set provides selected county demographic information (CDI) for 440 of the most populous counties in the United States. Each line of the data set has an identification number with a county name and state abbreviation and provides information on 14 variables for a single county. Counties with missing data were deleted from the data set. The information generally pertains to the years 1990 and 1992. The 17 variables are:

Variable Number	Variable Name	Description
1	Identification number	1–440
2	County	County name
3	State	Two-letter state abbreviation
4	Land area	Land area (square miles)
5	Total population	Estimated 1990 population
6	Percent of population aged 18–34	Percent of 1990 CDI population aged 18–34
7	Percent of population 65 or older	Percent of 1990 CDI population aged 65 years old or older
8	Number of active physicians	Number of professionally active nonfederal physicians during 1990
9	Number of hospital beds	Total number of beds, cribs, and bassinets during 1990
10	Total serious crimes	Total number of serious crimes in 1990, including murder, rape, robbery, aggravated assault, burglary, larceny-theft, and motor vehicle theft, as reported by law enforcement agencies
11	Percent high school graduates	Percent of adult population (persons 25 years old or older) who completed 12 or more years of school
12	Percent bachelor's degrees	Percent of adult population (persons 25 years old or older) with bachelor's degree
13	Percent below poverty level	Percent of 1990 CDI population with income below poverty level
14	Percent unemployment	Percent of 1990 CDI labor force that is unemployed
15	Per capita income	Per capita income of 1990 CDI population (dollars)
16	Total personal income	Total personal income of 1990 CDI population (in millions of dollars)
17	Geographic region	Geographic region classification is that used by the U.S. Bureau of the Census, where: 1 = NE, 2 = NC, 3 = S, 4 = W

1	2	3	4	5	6	7	8	9	10
1	Los_Angeles	CA	4060	8863164	32.1	9.7	23677	27700	688936
2	Cook	IL	946	5105067	29.2	12.4	15153	21550	436936
3	Harris	TX	1729	2818199	31.3	7.1	7553	12449	253526
...
438	Montgomery	TN	539	100498	35.7	7.9	87	188	6537
439	Maui	HI	1159	100374	26.2	11.3	192	182	7130
440	Morgan	AL	582	100043	26.3	11.7	122	464	4693

11	12	13	14	15	16	17
70.0	22.3	11.6	8.0	20786	184230	4
73.4	22.8	11.1	7.2	21729	110928	2
74.9	25.4	12.5	5.7	19517	55003	3
...
77.9	16.5	10.8	8.0	13169	1323	3
77.0	17.8	5.7	3.2	18504	1857	4
69.4	15.5	9.4	7.1	16458	1647	3

(2)
30%

Data file: Kidney_Function_Data.csv

Kidney function. Creatinine clearance (Y) is an important measure of kidney function, but is difficult to obtain in a clinical office setting because it requires 24-hour urine collection. To determine whether this measure can be predicted from some data that are easily available, a kidney specialist obtained the data that follow for 33 male subjects. The predictor variables are serum creatinine concentration (X_1), age (X_2), and weight (X_3).

Subject	i	X_{i1}	X_{i2}	X_{i3}	Y_i
1	1	.71	38	71	132
2	2	1.48	78	69	53
3	3	2.21	69	85	50
...
31	31	1.53	70	75	52
32	32	1.58	63	62	73
33	33	1.37	68	52	57

a. Using first-order and second-order terms for each of the three predictor variables (centered around the mean) in the pool of potential X variables (including cross products of the first-order terms), find the three best hierarchical subset regression models according to the AIC_p criterion.

b. Is there much difference in AIC_p for the three best subset models?

> ###2

+10 > # (a)

```
> Kidney <- read.csv("data/Kidney_Function_Data.csv",header = F)
> colnames(Kidney) <- c("Y", "X1", "X2", "X3")
> attach(Kidney)
```

The following objects are masked from Kidney (pos = 3):

X1, X2, X3, Y

The following objects are masked from Kidney (pos = 9):

X1, X2, X3, Y

The following objects are masked from Kidney (pos = 10):

X1, X2, X3, Y

The following objects are masked from Kidney (pos = 11):

X1, X2, X3, Y

The following objects are masked from Kidney (pos = 12):

X1, X2, X3, Y

```
> Kidney$X1m <- X1-mean(X1)
> Kidney$X2m <- X2-mean(X2)
> Kidney$X3m <- X3-mean(X3)
>
> Kidney$X1m2 <- Kidney$X1m^2
> Kidney$X2m2 <- Kidney$X2m^2
> Kidney$X3m2 <- Kidney$X3m^2
>
> Kidney$X1mX2m <- Kidney$X1m * Kidney$X2m
> Kidney$X2mX3m <- Kidney$X2m * Kidney$X3m
> Kidney$X1mX3m <- Kidney$X1m * Kidney$X3m
>
> library(ALSM)
> BS <- BestSub(Kidney[,5:13], Kidney[,1])
> BS
```

	p	1	2	3	4	5	6	7	8	9	SSEp	r2	r2.adj	Cp
1	2	1	0	0	0	0	0	0	0	0	11068.478	0.6429007	0.6313813	48.509656
1	2	0	1	0	0	0	0	0	0	0	17169.856	0.4460535	0.4281843	91.236011
2	3	1	1	0	0	0	0	0	0	0	7666.102	0.7526706	0.7361820	26.683707
2	3	1	0	1	0	0	0	0	0	0	8713.158	0.7188897	0.7001490	34.015966
3	4	1	1	1	0	0	0	0	0	0	4499.974	0.8548186	0.8397998	6.512138
3	4	1	0	1	0	1	0	0	0	0	5908.643	0.8093710	0.7896508	16.376681
4	5	1	1	1	0	0	0	1	0	0	3755.990	0.8788215	0.8615103	3.302215
4	5	1	1	1	1	0	0	0	0	0	3946.962	0.8726602	0.8544688	4.639546
5	6	1	1	1	0	0	1	1	0	0	3482.208	0.8876545	0.8668497	3.384990
5	6	1	1	1	0	1	0	1	0	0	3634.003	0.8827571	0.8610455	4.447976
6	7	1	1	1	0	1	1	1	0	0	3393.871	0.8905044	0.8652362	4.766392
6	7	1	1	1	0	0	1	1	0	1	3429.411	0.8893578	0.8638250	5.015266
7	8	1	1	1	1	1	1	0	1	0	3333.576	0.8924497	0.8623357	6.344160
7	8	1	1	1	0	1	1	1	1	0	3335.374	0.8923917	0.8622614	6.356750
8	9	1	1	1	1	1	1	1	1	0	3284.793	0.8940236	0.8586982	8.002543
8	9	1	1	1	1	1	1	0	1	1	3331.886	0.8925043	0.8566723	8.332326
9	10	1	1	1	1	1	1	1	1	1	3284.429	0.8940353	0.8525709	10.000000
	AICp		SBCp											
1	195.9065		198.8995								12402.603			
1	210.3953		213.3883								19676.109			
2	185.7858		190.2754								9348.907			
2	190.0107		194.5002								10326.342			
3	170.2055		176.1916								5963.853			
3	179.1929		185.1789								7940.411			
4	166.2418		173.7243								4994.071			
4	167.8784		175.3609								5276.847			
5	165.7442		174.7232								4947.502			
5	167.1522		176.1313								5680.157			
6	166.8962		177.3718								5683.431			
6	167.2400		177.7155								5031.620			
7	168.3047		180.2767								6234.421			
7	168.3225		180.2945								6220.230			
8	169.8182		183.2868								6325.155			
8	170.2879		183.7565								6620.457			
9	171.8145		186.7796								6782.970			

Conclusion?

(3)

40%

Data file: Performance_Ability_Data.csv

Performance ability. A psychologist conducted a study to examine the nature of the relation, if any, between an employee's emotional stability (X) and the employee's ability to perform in a task group (Y). Emotional stability was measured by a written test for which the higher the score, the greater is the emotional stability. Ability to perform in a task group ($Y = 1$ if able, $Y = 0$ if unable) was evaluated by the supervisor. The results for 27 employees were:

$i:$	1	2	3	...	25	26	27
$X_i:$	474	432	453	...	562	506	600
$Y_i:$	0	0	0	...	1	0	1

Logistic regression model (14.20) is assumed to be appropriate.

- Find the maximum likelihood estimates of β_0 and β_1 . State the fitted response function.
- Obtain a scatter plot of the data with both the fitted logistic response function from part (a) and a lowess smooth superimposed. Does the fitted logistic response function appear to fit well?
- Obtain $\exp(b_1)$ and interpret this number.
- What is the estimated probability that employees with an emotional stability test score of 550 will be able to perform in a task group?
- Estimate the emotional stability test score for which 70 percent of the employees with this test score are expected to be able to perform in a task group.

```
> ###3
> # (a)
> PA <- read.csv("data/Performance_Ability_Data.csv", header = F)
> colnames(PA) <- c("Y", "X")
>
> PA.lm <- glm(Y~X, data=PA, family = "binomial")
> summary(PA.lm)
```

Call:

```
glm(formula = Y ~ X, family = "binomial", data = PA)
```

Deviance Residuals:

```
      Min       1Q   Median       3Q      Max
-1.7845  -0.8350   0.5065   0.8371   1.7145
```

Coefficients:

```
              Estimate Std. Error z value Pr(>|z|)
(Intercept) -10.308925   4.376997  -2.355   0.0185 *
X              0.018920   0.007877   2.402   0.0163 *
```


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

(Dispersion parameter for binomial family taken to be 1)

Null deviance: 37.393 on 26 degrees of freedom
Residual deviance: 29.242 on 25 degrees of freedom
AIC: 33.242

Number of Fisher Scoring iterations: 4

```
>  
> # (b)  
> plot(PA$X, PA$Y)  
> lines(lowess(PA$X, PA$Y), col=2, lwd=2, lty=2)  
> new.X <- data.frame(X=seq(min(PA$X),max(PA$X),len=100))  
> predicted.Y <- predict(PA.lm, new.X, type="response")  
> new.X <- c(new.X)$X  
> lines(new.X, predicted.Y, lwd=2, col="black")  
>
```



```
> # (c)  
> exp(PA.lm$coefficients[2])
```

```
x
1.0191 ✓
>
+8 > # (d)
> predict(PA.lm, data.frame(X=550), type="response")
1
0.5242263 ✓
>
> # (e)
+8 > (log((0.7/(1-0.7)))-PA.lm$coefficients[1])/PA.lm$coefficients[2]
(Intercept)
589.6577 ✓
(a)logit(pi-hat) = -10.31 + 0.019X (b) (c) 1.0191 (d) 0.5242 (e) 589.6577
```