Homework # 9, Stat355, Spring 2021

Please use R to do all of the following problems. Submit your homework by a WORD or PDF document with main results and the corresponding code.

1. Confidence intervals are often used in practice. To understand how to implement the method, compute confidence intervals in the following data.

   (a) Data: $X_1, \cdots, X_n \sim^{iid} N(\mu, \sigma^2)$. The z-confidence interval for μ is $\bar{x} \pm z_{\alpha/2} s/\sqrt{n}$. The t-confidence interval for μ is $\bar{x} \pm t_{\alpha/2,n-1} s/\sqrt{n}$. Suppose we observe data

$$5.691, 5.107, 4.644, 4.970, 5.218, 4.418, 5.021, 5.239, 4.752, 6.071.$$

   Compute the 95% z and t confidence intervals for μ, respectively, and compare their lengths. You need to include your mathematical details as well as your R code and output.

   Solution; Based on the data, we have $\bar{x} = 5.113$, $s^2 = 0.2379 = 0.4878^2$, $n = 10$, $z_{0.025} = 1.96$, and $z_{0.025,9} = 2.262$. The 95% z-confidence interval for μ is

$$\bar{x} \pm z_{0.025} \frac{s}{\sqrt{n}} = 5.113 \pm 1.96 \frac{0.4878}{\sqrt{10}} = [4.811, 5.415].$$

   The 95% t-confidence interval for μ is

$$\bar{x} \pm t_{0.025,9} \frac{s}{\sqrt{n}} = 5.114 \pm 2.262 \frac{0.4878}{\sqrt{10}} = [4.764, 5.462].$$

   Their lengths are 0.604 and 0.698, respectively. The t-confidence interval is longer.

   R code and output

```
> x <- c(5.691,5.107,4.644,4.970,5.218,4.418,5.021,5.239,4.752,6.071)
> n <- length(x)
> x.mean <- mean(x)
> x.var <- var(x)
> x.mean
[1] 5.1131
> x.var
[1] 0.2378894
> qnorm(0.975)
[1] 1.959964
> qt(0.975,n-1)
[1] 2.262157
> c(x.mean-qnorm(0.975)*sqrt(x.var/n),x.mean+qnorm(0.975)*sqrt(x.var/n))
[1] 4.810802 5.415398
> c(x.mean-qt(0.975,n-1)*sqrt(x.var/n),x.mean+qt(0.975,n-1)*sqrt(x.var/n))
[1] 4.764192 5.462008
```

(b) Data: $X \sim Bin(50, \theta)$. The binomial z confidence interval is $\hat{\theta} \pm z_{\alpha/2}\sqrt{\hat{\theta}(1-\hat{\theta})/n}$, where $\hat{\theta} = X/n$. Suppose we observe $x = 15$. Compute the 95% z confidence interval for θ. You need to include your mathematical details as well as your R code and output.

Solution: Based on the data, we have $\hat{\theta} = 15/50 = 0.3$. Thus, the 95% z confidence interval for θ is

$$\hat{\theta} \pm 1.96\sqrt{\frac{\hat{\theta}(1-\hat{\theta})}{n}} = 0.3 \pm 1.96\sqrt{\frac{0.3(1-0.3)}{50}} = [0.1730, 0.4270].$$

R code and output

```
> n <- 50
> x <- 15
> theta.hat <- x/n
> theta.hat
[1] 0.3
> c(theta.hat-1.96*sqrt(theta.hat*(1-theta.hat)/n),
+ theta.hat+1.96*sqrt(theta.hat*(1-theta.hat)/n))
[1] 0.1729775 0.4270225
```

(c) Data; $X_1, \cdots, X_n \sim Poisson(\theta)$. The Poisson z confidence interval is $\bar{x} \pm z_{\alpha/2}\sqrt{\bar{x}/n}$. Suppose we observe data

$$21, 23, 25, 12, 19.$$

Compute the 95% z confidence interval for θ. You need to include your mathematical details as well as your R code and output.

Solution: Based on the data, we have $\bar{x} = 20$. Thus, the 95% z-confidence interval for θ is

$$\bar{x} \pm 1.96\sqrt{\bar{x}/n} = 20 \pm 1.96\sqrt{20/n} = [16.08, 23.92].$$

```
> x <- c(21,23,25,12,19)
> x.mean <- mean(x)
> x.mean
[1] 20
> c(x.mean-1.96*sqrt(x.mean/5),x.mean+1.96*sqrt(x.mean/5))
[1] 16.08 23.92
```

2. z and t confidence intervals are both popular. In particular, let $X_1, \cdots, X_n \sim^{iid} N(\mu, \sigma^2)$, where both μ and σ^2 are unknown. Then, the $1-\alpha$ level z confidence interval for μ is $\bar{x} \pm z_{\alpha/2}s/\sqrt{n}$, and the $1-\alpha$ level t confidence intevral for μ is $\bar{x} \pm t_{\alpha/2,n-1}s/\sqrt{n}$. Use simulation with 10^4 replications to evaluate the coverage probability and the length of the 95% z and t confidence intervals for μ when $n = 5, 10, 20, 50, 100$ with $\mu = 0, 1, 2, 3, 4, 5$ and fixed $\sigma^2 = 1$, respectively. You need to provide a simple description of your simulation, a summary of your findings, and your R code and output.

Solution: For each combination of μ and n, I generate data 10^4 times. In each replication of the generation, I compute the z and t-confidence intervals for μ. I say the confidence interval is correct and output 1 if it contains μ use the definition. I also calculate their lengths. I use the average of the correctness of the confidence interval as the approximation of the coverage probability. I use the average of the length as the length of the confidence interval. Al those are derive based on R. I find that the coverage probabilities of z-confidence interval is much lower that 0.95 when n is small, indicating that it is inappropriate for small n (e.g., $n = 5$). The coverage probabilities of t-confidence interval is always around 0.95, indicating that it is always appropriate. The length of z-confidence interval is much lower that that of t-confidence interval for small n (e.g, $n = 5$). Their lengths are close to each other when n is not small (e.g., $n = 100$).

```
> sigma <- 1
> n.o <- c(5,10,20,50,100)
> mu.o <- c(0,1,2,3,4,5)
> run <- 1e4
> COVERAGE <- matrix(0,length(n.o)*length(mu.o),4)
> LENGTH <- matrix(0,length(n.o)*length(mu.o),4)
>
> result <- matrix(0,run,4)
> for(i in 1:length(n.o)){
+  for(j in 1:length(mu.o)){
+   n <- n.o[i]
+   mu <- mu.o[j]
+   for(k in 1:run){
+    x <- rnorm(n,mean=mu,sd=sigma)
+    x.mean <- mean(x)
+    x.var <- var(x)
+    CI.z <- c(x.mean-1.96*sqrt(x.var/n),x.mean+1.96*sqrt(x.var/n))
+    CI.t <- c(x.mean-qt(0.975,n-1)*sqrt(x.var/n),x.mean+qt(0.975,n-1)*sqrt(x.var/n))
+    result[k,] <- c(CI.z,CI.t)
+   }
+   COVERAGE[length(mu.o)*(i-1)+j,] <- c(n,mu,mean((result[,1]<=mu)
+  *(result[,2]>=mu))
+  ,mean((result[,3]<=mu)*(result[,4]>=mu)))
+   LENGTH[length(mu.o)*(i-1)+j,] <- c(n,mu,mean(result[,2]-
+  result[,1]),mean(result[,4]-result[,3]))
+  }
+ }
> COVERAGE
      [,1] [,2]   [,3]   [,4]
 [1,]    5    0 0.8757 0.9465
 [2,]    5    1 0.8772 0.9466
 [3,]    5    2 0.8764 0.9507
 [4,]    5    3 0.8765 0.9495
 [5,]    5    4 0.8772 0.9516
```

```
 [6,]    5    5 0.8769 0.9495
 [7,]   10    0 0.9156 0.9450
 [8,]   10    1 0.9108 0.9455
 [9,]   10    2 0.9150 0.9475
[10,]   10    3 0.9133 0.9467
[11,]   10    4 0.9183 0.9473
[12,]   10    5 0.9183 0.9505
[13,]   20    0 0.9322 0.9495
[14,]   20    1 0.9377 0.9500
[15,]   20    2 0.9383 0.9517
[16,]   20    3 0.9331 0.9484
[17,]   20    4 0.9344 0.9484
[18,]   20    5 0.9334 0.9505
[19,]   50    0 0.9416 0.9473
[20,]   50    1 0.9488 0.9543
[21,]   50    2 0.9458 0.9522
[22,]   50    3 0.9459 0.9523
[23,]   50    4 0.9466 0.9523
[24,]   50    5 0.9424 0.9485
[25,]  100    0 0.9463 0.9490
[26,]  100    1 0.9444 0.9470
[27,]  100    2 0.9509 0.9538
[28,]  100    3 0.9430 0.9468
[29,]  100    4 0.9447 0.9475
[30,]  100    5 0.9470 0.9502
> LENGTH
      [,1] [,2]      [,3]      [,4]
 [1,]    5    0 1.6450092 2.3302438
 [2,]    5    1 1.6426055 2.3268388
 [3,]    5    2 1.6383204 2.3207687
 [4,]    5    3 1.6475883 2.3338972
 [5,]    5    4 1.6428533 2.3271898
 [6,]    5    5 1.6466438 2.3325592
 [7,]   10    0 1.2061673 1.3921122
 [8,]   10    1 1.2051346 1.3909204
 [9,]   10    2 1.2011886 1.3863660
[10,]   10    3 1.2057224 1.3915988
[11,]   10    4 1.2071103 1.3932006
[12,]   10    5 1.2036975 1.3892616
[13,]   20    0 0.8661867 0.9249743
[14,]   20    1 0.8637670 0.9223904
[15,]   20    2 0.8672357 0.9260945
[16,]   20    3 0.8647519 0.9234421
```

```
[17,]   20    4 0.8632113 0.9217969
[18,]   20    5 0.8679184 0.9268235
[19,]   50    0 0.5520224 0.5659849
[20,]   50    1 0.5527567 0.5667379
[21,]   50    2 0.5518809 0.5658399
[22,]   50    3 0.5519585 0.5659195
[23,]   50    4 0.5522228 0.5661904
[24,]   50    5 0.5510085 0.5649454
[25,]  100    0 0.3916174 0.3964561
[26,]  100    1 0.3911581 0.3959911
[27,]  100    2 0.3912275 0.3960613
[28,]  100    3 0.3912050 0.3960386
[29,]  100    4 0.3911971 0.3960306
[30,]  100    5 0.3910814 0.3959134
```

3. z confidence interval is often used in binomial data for success/failure experiments. In particular, suppose the success/failure experiment is repeated n times, such that the total number of successes $X \sim Bin(n, \theta)$, where θ is the probability of success in each experiment. Then, the $1-\alpha$ level z confidence interval for θ is $\hat{\theta} \pm z_{\alpha/2}\sqrt{\hat{\theta}(1-\hat{\theta})/n}$, where $\hat{\theta} = X/n$. Use simulation with 10^4 replications to evaluate the coverage probability of the 95% z confidence interval for θ when θ is close 0, where you choose $\theta = 0.01, 0.02, 0.05, 0.1$ with $n = 10, 20, 50, 100, 200, 500, 1000$, respectively. You need to provide a simple description of your simulation, a summary of your findings, and your R code and output.

   Solution: For each combination of n and θ with $n = 10, 20, 50, 100, 200, 500, 1000$ and $\theta = 0.01, 0.02, 0.05, 0.1$, I generate data from $Bin(n, \theta)$ with 10^4 replications. In each replication, I calculate the 95% z-confidence interval for θ. I denote 1 if the confidence interval contains the true value of θ or 0 otherwise. The coverage probability is approximated by the proportion of 1. The result shows that the coverage probability is much lower than 0.95 for small n (e.g., $n = 10, 20, 50, 100$) when θ is very close to 0 (e.g., $\theta = 0.01, 0.02$). The coverage probability increases as either n or θ increases. I reach 0.95 when n is large (e.g. $n = 1000$) in all the cases of θ we studied.

```
> n.o <- c(10,20,50,100,200,500,1000)
> theta.o <- c(0.01,0.02,0.05,0.1)
> run <- 1e4
> COVERAGE <- matrix(0,length(n.o)*length(theta.o),3)
> for(i in 1:length(n.o)){
+  for(j in 1:length(theta.o)){
+   n <- n.o[i]
+   theta <- theta.o[j]
+   x <- rbinom(run,n,theta)
+   theta.hat <- x/n
+   lower <- theta.hat-1.96*sqrt(theta.hat*(1-theta.hat)/n)
```

```
+   upper <- theta.hat+1.96*sqrt(theta.hat*(1-theta.hat)/n)
+   COVERAGE[(i-1)*length(theta.o)+j,] <- c(n,theta,mean((lower<=theta)*(upper>=theta)))
+  }
+ }
> dimnames(COVERAGE)[[2]] <- c("n","theta","coverage")
> COVERAGE
         n theta coverage
 [1,]   10  0.01   0.0964
 [2,]   10  0.02   0.1807
 [3,]   10  0.05   0.3994
 [4,]   10  0.10   0.6461
 [5,]   20  0.01   0.1825
 [6,]   20  0.02   0.3363
 [7,]   20  0.05   0.6333
 [8,]   20  0.10   0.8756
 [9,]   50  0.01   0.3969
[10,]   50  0.02   0.6374
[11,]   50  0.05   0.9229
[12,]   50  0.10   0.8790
[13,]  100  0.01   0.6443
[14,]  100  0.02   0.8658
[15,]  100  0.05   0.8823
[16,]  100  0.10   0.9301
[17,]  200  0.01   0.8654
[18,]  200  0.02   0.9025
[19,]  200  0.05   0.9264
[20,]  200  0.10   0.9282
[21,]  500  0.01   0.8724
[22,]  500  0.02   0.9274
[23,]  500  0.05   0.9310
[24,]  500  0.10   0.9397
[25,] 1000  0.01   0.9263
[26,] 1000  0.02   0.9504
[27,] 1000  0.05   0.9468
[28,] 1000  0.10   0.9537
```