

Computational Statistics. Chapter 4: Classical simulation. Solution of exercises

Thierry Denoeux

9/20/2021

Exercise 1

Question 1a

We first write a function that computes the likelihood for an i.i.d. sample from the Poisson distribution:


```
lik <- function(lambda,x) return(prod(sapply(x,dpois,lambda=lambda)))
```

We then compute functions q and e for the range $\lambda \in [1, 10]$:

```
x<- c(8,3,4,3,1,7,2,6,2,7)
meanlog<-log(4)
sdlog<-0.5
lambda<-seq(from=1,to=10,by=0.01)
L<-sapply(lambda,lik,x=x) # Likelihood
prior<-sapply(lambda,dlnorm,meanlog = meanlog, sdlog = sdlog) # Prior
q<-L*prior
e<-lik(mean(x),x)*prior
```

We plot the two curves:

```
plot(lambda,q,type="l",ylim=range(q,e),col="red",lwd=2,xlab=expression(lambda))
lines(lambda,e,lty=2,col="blue",lwd=2)
legend("topright",legend=c("q","e"),col=c("red","blue"),lty=c(1,2),lwd=c(2,2))
```



Question 1b


This is an implementation of the rejection sampling algorithm:

```
n<-1000 # desired sample size
Lambda<-vector(n,mode="numeric")
i<-0
likh<-lik(mean(x),x) # maximum likelihood
while(i<n){
  l<-rlnorm(1,meanlog = meanlog, sdlog = 0.5)
  u <- runif(1)
  if(u<lik(l,x)/likh){
    i<-i+1
    Lambda[i]<-l
  }
}
```

We plot the histogram of generated values together with the prior density:

```
h<-hist(Lambda,breaks=20,plot=FALSE)
ylim=range(0,prior,h$density)
hist(Lambda,freq=FALSE,breaks=20,ylim=ylim,xlim=range(lambda),xlab=expression(lambda),
     main="RS")
lines(lambda,prior,lty=2)
```

RS



Question 1c

```
Linf=mean(Lambda)-sd(Lambda)*qnorm(0.975)/sqrt(n)
Lsup=mean(Lambda)+sd(Lambda)*qnorm(0.975)/sqrt(n)
print(c(Linf,Lsup))
```

```
## [1] 4.238677 4.315563
```

Question 1d

R implementation of the SIR algorithm:

```
m<-50000
n<-1000
Lambda0<-rlnorm(m,meanlog = meanlog, sdlog = sdlog) # sample from prior
L<-sapply(Lambda0,lik,x=x)
w<-L/sum(L) # standardized importance weights
Lambda1<-sample(Lambda0,size=n,replace=TRUE,prob=w) # resampling
```


Plot of the histogram together with the prior density:

```
h<-hist(Lambda1,freq=FALSE,breaks=20,plot=FALSE)
```

```
## Warning in hist.default(Lambda1, freq = FALSE, breaks = 20, plot = FALSE):
## argument 'freq' is not made use of
```

```
ylim=range(0,prior,h$density)
hist(Lambda1,freq=FALSE,breaks=20,ylim=ylim,xlim=range(lambda),xlab=expression(lambda),main="SIR")
lines(lambda,prior,lty=2)
```

SIR



We can see that the obtained distribution is very similar to the one found using the rejection sampling algorithm. The confidence interval is close to the previous one:


```
Linf1=mean(Lambda1)-sd(Lambda1)*qnorm(0.975)/sqrt(n)
Lsup1=mean(Lambda1)+sd(Lambda1)*qnorm(0.975)/sqrt(n)
print(c(Linf1,Lsup1))
```

```
## [1] 4.250918 4.329941
```

Question 1d

We can compare the two distributions graphically using a QQ-plot:

```
qqplot(Lambda,Lambda1,xlab="rejection sampling",ylab="SIR")
abline(0,1,col="red",lwd=2)
```



We can also plot boxplots of the two distributions side by side:

```
boxplot(Lambda,Lambda1,ylab=expression(lambda),names=c("RS","SIR"))
```

