Using R to plot the probability density function (PDF)

In probability theory, a probability density function (PDF) or density of a continuous random variable (which takes on a numerous infinite number of possible values) is defined as a function which indicates the relative likelihood for this random variable to take on a given value. That means if we are interested in knowing the probability that X falls in some interval (a, b) of a given set of values, we will need to find $P(a<X<b)$. In other words, the probability of the random variable falling within a particular range of values is given by the integral of this variable's density over that range, i.e. it is given by the area under the density function, being above the horizontal axis and between the lowest and greatest values of the range.

The probability density function $f(x)$ of a continuous random variable X with a supporting base B therefore has the following properties:

a. $f(x)$ is positive everywhere in the supporting base B, that is $f(x)>0$ for all x in B

b. The area under the curve $f(x)$ in the base support B is 1, i.e. $\int b.f(x)dx = 1$

c. If $f(x)$ is the PDF of x, then the probability that x belongs to A, where A is some interval within the range, is given by the integral of $f(x)$ over that interval, i.e.

$$P(X \in A) = \int A.f(x)dx$$

Now, let's see how we can use R language to plot a density function

Define a vector x over the domain. We can then apply the distribution's density function to x and then plot the result. The code sniper plots the standard normal distribution:

```
> x<-seq(from=-3,to=+3,length.out=100)
> plot(x,dnorm(x))
>
```

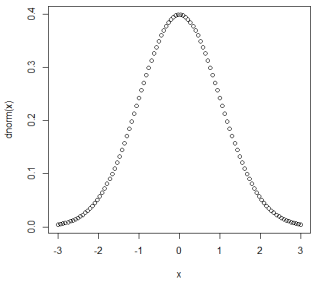

Whilst,

```
> x<-seq(from=-3,to=+3,length.out=1000)
> plot(x,dnorm(x))
>
```

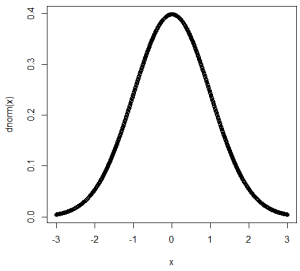

It is noted that all the R built-in probability distributions include a density function. For a particular density, the function name is "d" prepended to the density name. For example, the density function for the normal distribution is dnorm, the density for the gamma distribution is dgamma, and so forth.

If the first argument of the density function is a vector, then the function calculates the density at each point and returns the vector of densities.

The following code creates a 2 x 2 plot of four densities as shown in the figures on page 3:

```
x<-seq(from=-3,to=+3,length.out=1000)
> plot(x,dnorm(x))
>
> x<-seq(from=0,to=6,length.out=100)  # Define the density domain
> ylim<-c(0,0.6)
>
> par(mfrow=c(2,2))  # Create a 2x2 plotting area
>
> plot(x,dunif(x,min=2,max=4),main="Uniform",type="l",ylim=ylim)  #Plot a uniform density
>
> plot(x,dnorm(x,mean=3,sd=1),main="Normal",type="l",ylim=ylim)  #Plot a Normal density
>
>plot(x,dexp(x,rate=1/2),main="Exponential",type="l",ylim=ylim)  #Plot an Exponential density
>
>  plot(x,dgamma(x,shape=2,rate=1),main="Gamma",type="l",ylim=ylim)  #Plot a Gamma density
>
```

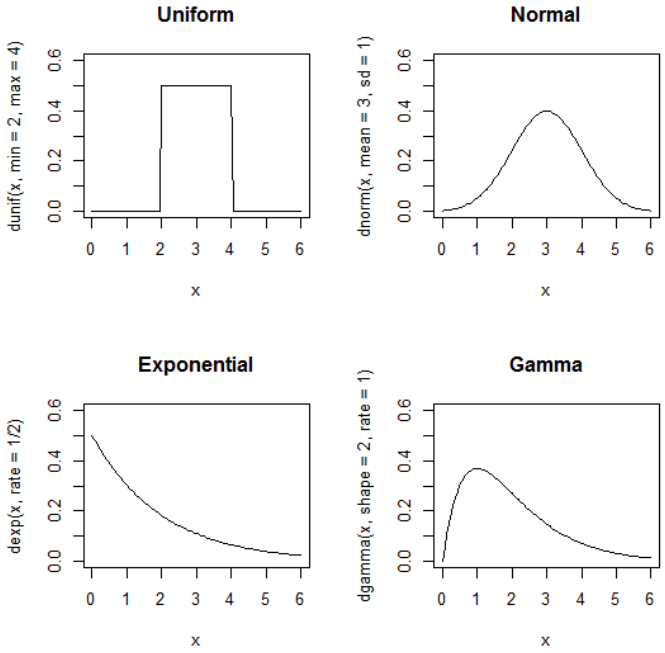
Uniform
dunif(x, min = 2, max = 4)
0.0
0.2
0.4
0.6
0 1 2 3 4 5 6
x
Normal
dnorm(x, mean = 3, sd = 1)
0.0
0.2
0.4
0.6
0 1 2 3 4 5 6
x
Exponential
dexp(x, rate = 1/2)
0.0
0.2
0.4
0.6
0 1 2 3 4 5 6
x
Gamma
dgamma(x, shape = 2, rate = 1)
0.0
0.2
0.4
0.6
0 1 2 3 4 5 6
x