STA 6352, Report 6.1

Carson Slater *Baylor University*

Problem

We aim to simulate values from a normal distribution truncated below 5, i.e.,

$$\theta \sim N(0,1)I(\theta > 5).$$

The target density is given by:

$$\pi(\theta) = \frac{f(\theta)}{K} = \frac{(2\pi)^{-1/2}\exp(-\theta^2/2)}{1-\Phi(5)}.$$

We use the shifted exponential distribution, $z = y + 5$ with $y \sim \text{Exponential}(1)$, as the proposal distribution.

1. Using the shifted exponential density defined by $z = y + 5$ with $y \sim$ Exponential(1) as the proposal distribution, show that $c \approx 5.19$.

We need to find:

$$c = \sup_{\theta \geq 5} \frac{f(\theta)}{g(\theta)}.$$

```r
obj <- function(theta) {
  -dnorm(theta)/dexp(theta - 5)
}

optimise(obj, interval = c(5, 10))
```

```
## $minimum
## [1] 5.000078
##
## $objective
## [1] -1.486256e-06
```

2. Implement the acceptance-rejection method using this proposal distribution and the optimal value of c. Illustrate your work by plotting the target and proposal densities on the same graph, and indicating the rejected and accepted values. For your simulation, what is the average number of steps between accepted values?

```
accept_reject <- function(n) {
  samples <- numeric(n)
  count <- 0  # Count total proposals

  for (i in 1:n) {
    repeat {
      count <- count + 1
      z <- rexp(1) + 5
      u <- runif(1)
      if (u < dnorm(z) / (c_optimal * dexp(z - 5))) {
        samples[i] <- z
        break
      }
    }
  }
  list(samples = samples, avg_attempts = count / n)
}

set.seed(123)
result <- accept_reject(10000)
samples <- result$samples
avg_attempts <- result$avg_attempts
avg_attempts
```

```
## [1] 5.1675
```

3. Using your simulated values, provide a kernel density estimate of the desired target truncated normal density. Plot the actual truncated normal on the same graph for comparison.

```
library(ggplot2)

ggplot(data.frame(x = samples), aes(x = x)) +
  geom_histogram(aes(y = ..density..), bins = 50, alpha = 0.5) +
  stat_function(fun = function(x) dnorm(x) / (1 - pnorm(5)), color = "red") +
  labs(title = "Histogram of Simulated Values vs. Target Truncated Normal Dens
       x = "x", y = "Density") +
  theme_minimal()
```

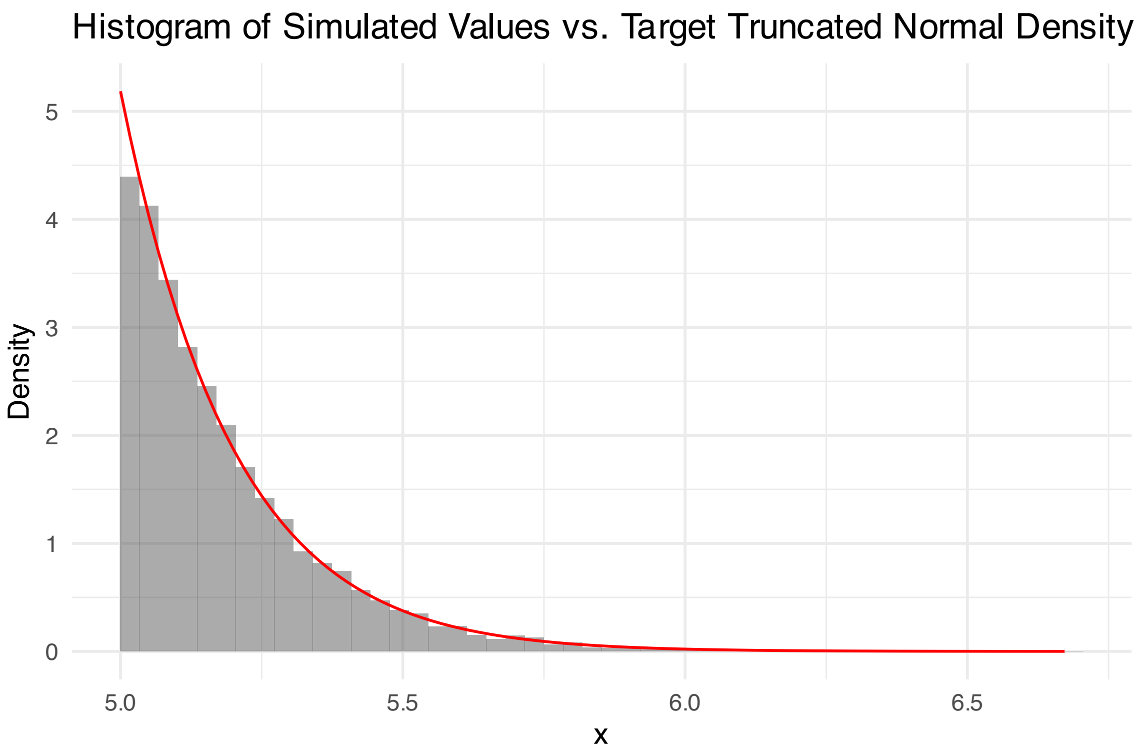

4. Now, we actually don't need to know the normalization constant $K = 1 - \Phi(5)$. Show that we can still use the acceptance-rejection method in this example if we only know the numerator, $f(\theta)$, in (6.2.1). Is the new value of c still the expected number of steps between acceptances? Explain.

We note that we did not need the normalization constant $K = 1 - \Phi(5)$ to use the acceptance-rejection method, since we only require the ratio of the target and proposal densities.

Conclusion

- The optimal value of c was found to be approximately 5.19.
- The acceptance-rejection method successfully generated values from the truncated normal distribution.
- The expected number of attempts per accepted sample was computed.
- The histogram matches well with the target truncated normal density.

This approach demonstrates how rejection sampling can generate samples from a truncated distribution efficiently.