1. With pen and paper or LaTeX, derive the maximum likelihood estimator for the mean μ of the normal distribution, $N(\mu, \sigma^2 = 4)$. Show your work.

$$f(x|\mu, \sigma) = (2\pi\sigma^2)^{-1/2} \exp\left(\frac{-(x-\mu)^2}{2\sigma^2}\right).$$

Solution. The simplified log-likelihood is proportional to ($\propto$) the following expression, since nothing else depends on μ.

$$l(\mu|\mathbf{x}) \propto \sum_{n=1}^{N}(x_n - \mu)^2$$

Next take the derivative with respect to μ of l and set it equal to zero.

$$\sum_{N=1}^{N} x_n = N * \mu$$

Solving for μ gives

$$\hat{\mu} = N^{-1} \sum_{n=1}^{n} x_n = \bar{x}.$$

2. With R, generate data from $N(\mu = 5, \sigma^2 = 4)$, and then use `optim` to estimate μ and σ via the log-likelihood function for the normal distribution. Show your work.

```
## Set up simplified log-likelihood,
## recognizing that we now want both mu and sigma.
## Use a proportional to argument again.
## Don't forget R does minimization
## and we want maximiation.

ll_normal <- function(theta, X) {
    mu <- theta[1]
    sigma <- theta[2]
    N <- length(X)

    N*log(sigma) + sum((x-mu)^2)/(2*sigma^2)
}
## assume population known, generate data
x <- rnorm(1001, mean=5, sd=2)
optim(c(10, pi), ll_normal,
      method="L-BFGS", lower=c(-1e6, 1e-6), X=x)$par

## [1] 5.026873 2.027332
```