

Question 1: [15 Marks] For some function f (which we can compute), we wish to find a solution to $f(x) = 0$, with an absolute accuracy of 10^{-10} or better. (That is, we wish to find an x such that $|x - x^*| \leq 10^{-10}$, where x^* is some true solution to $f(x) = 0$.)

We know that f is continuous, and that $f(10) = 12$ and $f(110) = -2$. We plan to use bisection to solve $f(x) = 0$, starting with the interval from 10 to 110. How many iterations of bisection (each of which evaluates f once, not counting the values at the initial endpoints) will be needed to guarantee that we will obtain a solution with the required accuracy? Explain your answer.

Note: $2^{10} = 1024 \approx 10^3$.

Bisection operates by narrowing the interval in which the solution must be by a factor of two each iteration.

We need to get the interval down to the size 2×10^{-10} , at which point the mid-point of the interval must be within 10^{-10} of the solution.

The initial interval is of size $110 - 10 = 100$.

So we need to shrink the interval by the factor $100 / (2 \times 10^{-10}) = 10^{12}/2$. This is approximately $2^{40}/2 = 2^{39}$.

So 39 iterations will be needed.

Question 2: [70 Marks Total] Suppose we model positive real data values, y_1, \dots, y_n , as being independently generated from a gamma distribution, whose density function is

$$f(y) = \frac{b^a}{\Gamma(a)} y^{a-1} \exp(-by)$$

where a and b are positive real model parameters. We wish to find the maximum likelihood estimates for a and b based on the data y_1, \dots, y_n .

Note that $\Gamma(a)$ is the “gamma function”. The log of the gamma function is computed by the R function `lgamma`. The derivative of the log of the gamma function is computed by R’s `digamma` function, and the second derivative by `trigamma`. All these functions may take a vector as their argument, and return the vector of corresponding function values.

- a) [10 Marks] Write an R function below that computes the log likelihood for parameter values a and b given a data vector y . The beginning of the function definition is already shown below, to which you need to add the body of the function. Do not use R’s built-in `dgamma` function.

```
log_lik <- function (y,a,b)
  sum (a*log(b) - lgamma(a) + (a-1)*log(y) - b*y)
```

- b) [6 Marks] Write an R function below that computes the derivative of the log likelihood with respect to the a parameter:

```
log_lik_deriv_a <- function (y,a,b)
  sum (log(b) - digamma(a) + log(y))
```

- c) [6 Marks] Write an R function below that computes the derivative of the log likelihood with respect to the b parameter:

```
log_lik_deriv_b <- function (y,a,b)
  sum (a/b-y)
```

- d) [6 Marks] Write an R function below that computes the second derivative of the log likelihood with respect to the a parameter:

```
log_lik_deriv2_a <- function (y,a,b)
  length(y) * -trigamma(a)
```

- e) [6 Marks] Write an R function below that computes the second derivative of the log likelihood with respect to the b parameter:

```
log_lik_deriv2_b <- function (y,a,b)
  length(y) * -a/b^2
```

Suppose we decide to find the maximum likelihood estimates for a and b by using an alternating maximization procedure (also known as non-linear Gauss-Seidel iteration). That is, we alternately maximize the log likelihood with respect to a , with b fixed, then with respect to b , with a fixed, then with respect to a again, etc. We decide to use univariate Newton iteration to maximize with respect to a or with respect to b .

- f) [15 Marks] Write an R function below that tries to find a zero of a univariate function $f1$, whose derivative is the function $f2$, starting from the initial point x , using m iterations of Newton iteration. The beginning of the function is already present; you should write the body of the function. Your function should use only basic R facilities, not R's `nlm` function.

```
newton_iteration <- function (f1,f2,x,m)
{
  for (i in 1:m)
    x <- x - f1(x)/f2(x)
  x
}
```

- g) [21 Marks] Write an R function that uses the `newton_iteration` function (with m set to 10) and the functions for computing the log likelihood and its derivatives to find the maximum likelihood estimate for a and b given a data vector y . Your function should use alternating maximization as described above. The value returned should be a vector of estimates for a and b . The beginning of the functions is already present, with the argument n being the number of iterations of alternating maximization to do, and the arguments a and b being the initial values to use to start the iterations.

```

mle <- function (y,a,b,n)
{
  for (i in 1:n) {
    a <- newton_update (function (a) log_lik_deriv_a(y,a,b),
                       function (a) log_lik_deriv2_a(y,a,b), a, 10)
    b <- newton_update (function (b) log_lik_deriv_b(y,a,b),
                       function (b) log_lik_deriv2_b(y,a,b), b, 10)
  }

  c(a,b)
}

```

Question 3: [15 Marks] We wish to compute an approximation to the integral of some function, f , of d variables, over the range $(0,1)^d$. The function f is continuous and infinitely differentiable.

For example, if $d = 3$, we wish to compute an approximation to

$$I = \int_0^1 \int_0^1 \int_0^1 f(x, y, z) dz dy dx$$

We approximate I by nested estimates of univariate integrals. So for the $d = 3$ example, we define

$$h(x, y) = \int_0^1 f(x, y, z) dz$$

and

$$g(x) = \int_0^1 h(x, y) dy$$

so that

$$I = \int_0^1 g(x) dx$$

Suppose we approximate all the integrals we need to compute using the Trapezoid Rule. For example, we approximate the integral of $g(x)$ above that gives I using the Trapezoid Rule. Similarly, for every evaluation of $g(x)$ at some point x , we approximate the integral of $h(x, y)$ above that gives $g(x)$ using the Trapezoidal Rule, and similarly for the integral of $f(x, y, z)$ that gives $h(x, y)$. Suppose that for all uses of the Trapezoidal Rule we divide the interval $(0,1)$ into the same number of sub-intervals (and hence the number of points where we evaluate the integrand is this number plus one).

Find the rate at which the error declines as the total number, n , of points at which f is evaluated increases, explaining how you obtained your answer, and how the rate depends on d .

If we evaluate the integrand at m points for each use of the Trapezoid Rule in the scheme above, then we will evaluate the function f at m^d points. So we need to use $m = n^{1/d}$.

The univariate Trapezoid Rule with m points has error proportional to m^{-2} (for large m), which is $n^{-2/d}$.

The errors at each stage of integration will tend to add (we can't count on being so lucky that they cancel out). So the error in I will decline in proportion to $n^{-2/d}$.

Suppose that instead of the Trapezoid Rule we use Simpson's Rule for all integrations. Then how fast will the error declines with n ?

The error for the univariate Simpson's Rule with m points declines in proportion to m^{-4} , so the same argument as above leads to the error in I declining in proportion to $n^{-4/d}$.