DATA 1010
Problem Set 3
Due 28 September 2018 at 11 PM

Problem 1

Calculate, by hand, the gradient and Hessian of the function shown below. Show that the values returned by the `ForwardDiff` package are correct.

```julia
using ForwardDiff
f(x,y) = x^2 + y^2 - 2y
f(v::Vector) = f(v...) # equivalent to f(v[1],v[2])
x = [1.5,-3.25]
ForwardDiff.gradient(f,x)
ForwardDiff.hessian(f,x)
```

Problem 2

The *Rosenbrock* function is defined by

$$f(x,y) = (a-x)^2 + b(y-x^2)^2,$$

where a and b are positive constants.

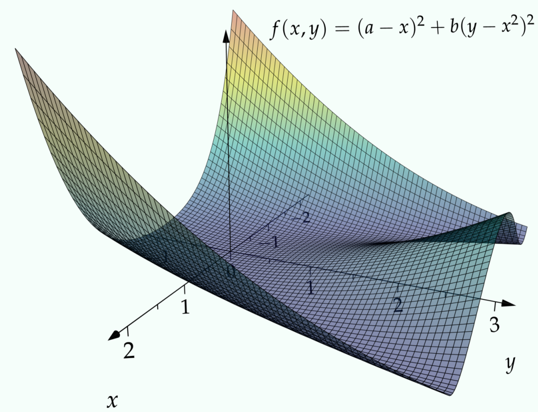

(a) Find the point (x,y) where f has a global minimum.

(b) Show that f does not have a local minimum anywhere else.

(c) Implement (from scratch) the gradient descent algorithm for this function starting from $(0,0)$, with $a = 1$ and $b = 100$, and try some different learning rates. Show that finding the minimum is challenging.

(d) Show that the Julia package `Optim` can nevertheless handle this function just fine. Hint: the code to do this can be found on the GitHub landing page for the Optim.jl package (it's 3 lines).

Notes on (c): you can calculate the gradient of f by hand if you want, and you might find that you need to use a small learning rate. Also, you probably want to use a `for` loop with a controlled number of iterations, since `while` loops might fail to terminate.

Problem 3

Suppose that five dice are rolled simultaneously. The result of the roll is identified as: all different (like 5, 6, 4, 3, 1), one pair (like 2, 3, 4, 2, 5), two pairs (6, 6, 3, 1, 3), three-of-a-kind (1, 4, 3, 4, 4), full house (5, 4, 5, 4, 4), four-of-a-kind (3, 1, 3, 3, 3), or five-of-a-kind (2, 2, 2, 2, 2).

(a) Show that the probability of five-of-a-kind is approximately 0.08%.

(b) Show that the probability of a full house is approximately 3.86%.

Problem 4

Consider the following random experiment: you flip a coin, and if it turns up heads, you roll a die. If it turns up tails, then you draw a ball from an urn containing balls labeled 1 to 11.

Define a sample space Ω for this random experiment.

Problem 5

Suppose that E is an event. Using the axioms of a probability measure (Theorem 6.2.1), show that $\mathbb{P}(E^c) = 1 - \mathbb{P}(E)$.

Problem 6

Suppose that E, F, and G are events. Come up with events $\widetilde{E}, \widetilde{F}$, and $\widetilde{G}$ which are pairwise disjoint and which satisfy

$$E = \widetilde{E}, \quad E \cup F = \widetilde{E} \cup \widetilde{F}, \quad \text{and} \quad E \cup F \cup G = \widetilde{E} \cup \widetilde{F} \cup \widetilde{G}.$$

Problem 7

Suppose that A is the event that the high temperature in Providence next Tuesday is at least 65 degrees Fahrenheit, let B be the event that the high temperature in Providence next Tuesday is at least 60 degrees Fahrenheit, and let C be the event that it rains in Providence next Tuesday. Write each of the following events using the sets A, B, and C and the operations $\cap, \cup$, and ^{c}.

(a) It will be less than 60 degrees all day and rainy next Tuesday in Providence.

(b) The high temperature in Providence next Tuesday will be at least 60 degrees but not as high as 65 degrees.

(c) In Providence next Tuesday, it will either be dry or warm (where warm is defined to mean "daily high of at least 65 degrees").

(d) The daily high temperature in Providence next Tuesday will be higher than 65 degrees and less than 60 degrees.

Problem 8

The matrix $A_1 = \begin{bmatrix} -2 & -6 & 7 \\ -2 & -2 & 2 \\ -5 & 5 & 5 \end{bmatrix}$ has exactly one eigenvalue (approximately -4.532), while $A_2 = \begin{bmatrix} -6 & -1 \\ 5 & 0 \end{bmatrix}$ has two eigenvalues (-5 and -1). Find all of the eigenvalues of the matrix

$$A = \begin{bmatrix} -2 & -6 & 7 & 0 & 0 \\ -2 & -2 & 2 & 0 & 0 \\ -5 & 5 & 5 & 0 & 0 \\ 0 & 0 & 0 & -6 & -1 \\ 0 & 0 & 0 & 5 & 0 \end{bmatrix}.$$

Problem 9

Suppose that A is a 5×5 diagonalizable matrix with eigenvalues $-1, -1, -1, +1$, and $+1$. Show that $A^2 = I$.

Problem 10

Write a Julia function which accepts a two-column array as an argument and returns the number of rows for which the first column contains a number strictly greater than 200 and the second column contains the string `"blue"`.

```julia
M = [370.512 "red"; 937.542 "blue"; 782.404 "blue"; 697.21 "blue";
     154.13 "blue"; 819.013 "red"; 568.343 "red"; 928.226 "red";
     947.238 "red"; 656.98 "blue"]
myrowcount(M) # should return 4
```