Solutions to the Exercises

1. Newton's Method is a discrete dynamical system of the form $x_{i+1} = F(x_i)$. Let $g(x)$ be the function for which we're determining the root (using Newton's Method).

   (a) What is F (in terms of g)?
   SOLUTION: From the formula for Newton's Method,
   $$x_{i+1} = x_i - \frac{g(x_i)}{g'(x_i)} = F(x_i)$$

   (b) Show that the fixed point for F corresponds to the root(s) of g.
   SOLUTION: The fixed points are found by solving $x = F(x)$ for x:
   $$x = x - \frac{g(x)}{g'(x)} \quad \Rightarrow \quad \frac{g(x)}{g'(x)} = 0$$
   Assuming the derivative is not zero (so the expression is defined), $g(x) = 0$ is the solution.

2. Find all fixed points for $x_{i+1} = F(x_i)$ and classify them as attracting, repelling or neutral:

   (a) $F(x) = \sqrt{x}$
   SOLUTION: The fixed points are $x = \sqrt{x}$, or $x^2 - x = 0$ so that $x = 1$ or $x = 0$. The derivative is:
   $$F'(x) = \frac{1}{2\sqrt{x}}.$$
   The function is not differentiable at $x = 0$, but graphical analysis shows it is repelling. Further, $|F'(1)| < 1$, so $x = 1$ is an attracting fixed point.

   (b) $F(x) = x(1 - x)$
   SOLUTION: The fixed points are the solutions to
   $$x = x - x^2 \quad \Rightarrow \quad x = 0$$
   Now, $F'(x) = 1 - 2x$, so at $x = 0$, $F'(0) = 1$. Thus, the fixed point is *neutral.*

   (c) $F(x) = 1/x^2$
   SOLUTION: The fixed points are the solutions to
   $$x = \frac{1}{x^2} \quad \Rightarrow \quad x^3 = 1$$
   The only real solution is $x = 1$. Computing the derivative,
   $$F'(x) = \left.\frac{-2}{x^3}\right|_{x=1} = -2$$
   Therefore, $|F'(1)| > 1$, and the fixed point is repelling.

(d) $F(x) = (2 - x)/10$
SOLUTION: The fixed point is the solution to

$$x = \frac{1}{5} - \frac{1}{10}x \quad \Rightarrow \quad \frac{11}{10}x = \frac{1}{5} \quad \Rightarrow \quad x = \frac{2}{11}$$

And $|F'(x)| = \frac{1}{10}$, so this is an attracting fixed point.

3. Use graphical analysis to describe the fate of all orbits, if $x_{i+1} = F(x_i)$, and

(a) $F(x) = -2x + 1$: All orbits (except the fixed point) will diverge.

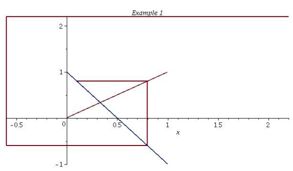

(b) $F(x) = \sin(x)$: All orbits go to zero.

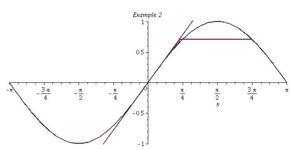

4. Determine the general solution:

(a) $x_{n+1} = \frac{3}{2}x_n + 1$
SOLUTION: Recall that, if $x_{n+1} = ax_n + b$, then

$$\begin{aligned} x_1 &= ax_0 + b \\ x_2 &= a(ax_o + b) + b = a^2x_0 + ab + b \\ x_3 &= a^3x_0 + b(1 + a + a^2) \\ &\vdots \\ x_n &= a^nx_0 + b(1 + a + a^2 + \cdots + a^{n-1}) \end{aligned}$$

You should be able to derive the formula we need for that finite geometric series:

$$x_n = a^nx_0 + b\frac{1 - a^n}{1 - a}$$

In our case, $a = \frac{3}{2}$, $b = 1$, so we get the closed form:

$$x_n = \frac{3}{2}^n x_0 - 2\left(1 - \frac{3}{2}^n\right)$$

(b) $x_{n+1} = 5x_n + n^2$
SOLUTION: For the equation $x_{n+1} = ax_n$, the solution is $x_n = a^n x_0$, or more generally, $a^n C$ for some constant C. For the n^2 part, we'll guess that the solution is of the form

$$An^2 + Bn + C$$

and put it into the equation to see if we can solve for A, B, C:

$$A(n^2 + 2n + 1) + B(n + 1) + C = 5An^2 + 5Bn + 5C + n^2$$

Equate the n^2 terms: $A = 5A + 1$, so we have $A = -1/4$
Equate the n terms: $2A + B = 5B$, so $B = -1/8$
Equate the constant terms: $A + B + C = 5C$, so $C = -3/32$.
Therefore, the general solution is

$$x_n = 5^n C - \frac{1}{4}n^2 - \frac{1}{8}n - \frac{3}{32}$$

from which we see that $C = x_0 + \frac{3}{32}$, and so we could express this as

$$x_n = 5^n \left(x_0 + \frac{3}{32} \right) - \frac{1}{4}n^2 - \frac{1}{8}n - \frac{3}{32}$$

(c) $x_{n+1} = \frac{1}{2}x_n + 4n^2 + 2n + 1$
We solve this in much the same way as the previous problem. If we assume the particular part of the solution is $An^2 + Bn + C$, we solve the three equations:
Equate the n^2 terms: $A = A/2 + 4$, so we have $A = 8$
Equate the n terms: $2A + B = B/2 + 2$, so $B = -28$
Equate the constant terms: $A + B + C = C/2 + 1$, so $C = 42$.
We can write the general solution as:

$$x_n = \frac{1}{2}^n C + 8n^2 - 28n + 42$$

And, solving for C in terms of x_0, we get $C = x_0 - 42$:

$$x_n = \frac{1}{2}^n (x_0 - 42) + 8n^2 - 28n + 42$$

5. Let $F(x) = |x - 2|$. Use graphical analysis to display a variety of orbits under F. If we were to ask about points of period 2, what might you say?

SOLUTION: The point $x = 1$ is fixed, an all other points are eventually period 2 points. For example, starting at $x_0 = 4$,

$$4, 2, 0, 2, 0, 2, 0, 2, \cdots$$

or if $x_0 = \sqrt{2}$, then the orbit is:

$$2 - \sqrt{2},\ \sqrt{2},\ 2 - \sqrt{2},\ \sqrt{2}, \cdots$$

6. (Calculator) You currently have \$5000 in a savings account that pays 6% interest per year. Interest is compounded monthly, and you add \$200 per month to your account. How much money do you have after 5 years? What is the total interest earned?

   SOLUTION: We'll count by months, so that 5 years is 60 months. Let's see what kind of system we have. Let S_0 be the amount at the beginning, and S_n be the amount after n months. Then the amount we have at month $n+1$ will be:

$$S_{n+1} = S_n + \frac{0.06}{12} S_n + 200 = \left(1 + \frac{0.06}{12}\right) S_n + 200 = 1.005 S_n + 200$$

   which is of the form

$$x_{n+1} = a x_n + b$$

   Therefore, we know the solution:

$$S_n = (1.005)^n S_0 + 200 \frac{1 - 1.005^n}{1 - 1.005}$$

   Using a calculator, we see that $S_{60} \approx 20698.26$. To find the interest, we can subtract the amount we've invested:

$$5000 + (200)(60) = 17000$$

   so over the five years, we made \$3698.26

7. (Calculator) Mary receives \$5000 as a graduation gift. She deposits the money is a savings account at an annual interest rate of 3%, and interest is compounded monthly.

   Mary plans on working for three years and during this time, she will be making equal monthly deposits.

   After the three years, Mary will withdraw \$1200 per month during the fourth year, \$1300 per month for her fifth, \$1400 per month for her sixth, and finally \$1500 per month for her seventh year.

   What must the monthly payment be so that the savings account balance is zero at the end of the seventh year? What is the total interest earned over the seven year period?

   SOLUTION: The best way to handle this problem is in pieces. For the first three years, Mary is depositing equal amounts. If we let b be that value, then after 3 years, Mary has

$$S_{36} = 1.0025^{36}(5000) + b \frac{1 - 1.0025^{36}}{1 - 1.0025} \approx 5470.257 + 37.6206 b = T_0$$

   Now, calling this T_0 and running through the next year, we have:

$$T_{12} = 1.0025^{12} T_0 - 1200 \frac{1 - 1.0025^{12}}{1 - 1.0025} \approx 1.03042 T_0 - 14599.659$$

Or, substituting our value for T_0:

$$T_{12} \approx -8963.019 + 38.765b$$

Now, for the next 12 months, Mary withdraws \$1300 per month:

$$T_{24} = 1.0025^{12} T_{12} - 1300 \frac{1 - 1.0025^{12}}{1 - 1.0025} \approx 1.0304 T_{12} - 15816.2976$$

Which simplifies to:

$$T_{24} \approx -25051.936 + 39.944b$$

And for the next two years:

$$T_{36} = 1.0025^{12} T_{24} - 1400 \frac{1 - 1.0025^{12}}{1 - 1.0025} \approx 1.0304 T_{24} - 18249.574 = -44063.489 + 41.159b$$

$$T_{48} = 1.0025^{12} T_{36} - 1500 \frac{1 - 1.0025^{12}}{1 - 1.0025} \approx -62399.652 + 42.411b$$

So setting this to zero, $b \approx 1471.32$. How much interest was paid over the whole 7 years? To answer this, we can compute how much Mary deposited and how much (in total) that Mary got back- The difference would be the interest paid.

Mary deposited a total of $5000 + 1471.32 \times 36 \approx 57976.46$.

Mary withdrew a total of $(1200 + 1300 + 1400 + 1500)12 = 64800$ over the four year period.

Over the 7 years, Mary was paid about \$6832.55 in interest.

8. Short Answer: Say (in words) what the following Matlab commands do:

   (a) `x=8:-2:1`

   This assigns a vector of decreasing integers from 8 to 1 (by 2 integers) to the value x. That is, $x = [8, 6, 4, 2]$

   (b) `exp([-1, 0, 1, 2])`

   We get a vector whose elements are e^{-1}, e^{0}, e^{1} and e^{2}.

   (c) `probs=exp(P)./sum(exp(P))` (for some vector P that was previously defined, and could have any real numbers as entries).

   We get a vector that can be used in the softmax algorithm. The resulting values in `prob` form a probability measure (they are non-negative and sum to 1).

   (d) `z=linspace(1,8,150);`

   This creates a (row) vector z that has 150 evenly spaced points between 1 and 8. This is typically used to create a plot or quickly create a vector.

(e) `2.^[-1,0,1,2]`

Like the exponential before, this exponentiates each term in the vector, giving a vector whose elements are $1/2, 1, 2, 4$.

(f) `Q=[3 1 2 2 1]; idx=find(Q==max(Q));`

This piece of code returns the indices of Q where the maximum element resides. In this case, `idx` is just the number 1.

Note: `Q==max(Q)` is the vector `[1,0,0,0,0]`

9. In Matlab, if X is a matrix, write the Matlab command that gives you:

   (a) The entire 3rd row. SOLUTION: $X(3, :)$

   (b) The entire 4th column. SOLUTION: $X(:, 4)$

   (c) The odd rows, with columns 1 and 2. SOLUTION: `X(1:2:end, 1:2)`

10. Name several differences between a *script* and a *function* in Matlab.

    SOLUTION: Lots of different possible answers. Here are a few salient points:

    A *script* is a text document that has a series of Matlab commands. They are executed in order as if you were typing them into the command window. To run a script file, in the command window, we just type the name of the file (without the ".m").

    A *function* is a text document whose first line starts with `function`, and defines input and output variables, then the algorithm. We can use functions in scripts or live at the keyboard, but we have to define the inputs/outputs while using it.

11. What was the $n-$armed bandit problem? In particular, what were the two competing goals, and why were they "competing"?

    SOLUTION: In the $n-$armed bandit problem, we have a slot machine with n arms (or n one-armed bandits). We don't know anything about the machines, but we'd like to play in such a way as to maximize our rewards!

    The two competing goals are: (1) We have to "explore" the payoffs that the machines are making in order to estimate the underlying actual payoffs. (2) To maximize our reward, we should "exploit" our estimates and always play the machine with the best reward.

12. For the $n-$armed bandit problem, what is *the greedy algorithm*? How is it modified to get *the $\epsilon-$greedy algorithm*? How did the curve that you produced depend on ϵ (and was that reasonable)?

    SOLUTION: The greedy algorithm is when we always play the machine with the largest current estimated payoff. In the $\epsilon-$greedy algorithm, we (with probability ϵ) choose a machine at random (otherwise, play greedy). In our figures, we saw that increasing ϵ resulted in a larger and larger average payout.

13. What is the "softmax" action selection? In particular, how did we change a set of payouts Q_i to a set of probabilities, P_i?

SOLUTION: The softmax selection is an algorithm designed to convert payoffs (both positive and negative) into percentages. Given a set of payoffs $Q_1, \ldots, Q_n$, then

$$P_i = \mathrm{e}^{Q_i/\tau} / \sum_{j=1}^{n} \mathrm{e}^{Q_j/\tau}$$

where τ is used to move from greedy (when τ is close to zero) to random with equal probability (when τ is large).

14. Suppose $Q = [-0.5, 0, 0.5, 1.0]$. Use the softmax selection technique with $\tau = 0.1$ to compute the probabilities.

SOLUTION: We get (approximately): $0, 0, 0.0067, 0.9933$

15. If $Q_1 < Q_2 < Q_3 < Q_4$ for 4 machines, how do the probabilities change (under softmax) as $\tau \to 0$? As $\tau \to 1$?

SOLUTION: TYPO: The upper limit should be ∞, not 1. Let's compute the probabilities (we've done this before for two numbers).

$$P_1 = \lim_{\tau \to 0} \frac{\mathrm{e}^{Q_1/\tau}}{\mathrm{e}^{Q_1/\tau} + \mathrm{e}^{Q_2/\tau} + \mathrm{e}^{Q_3/\tau} + \mathrm{e}^{Q_4/\tau}}$$

Dividing by $\mathrm{e}^{Q_1/\tau}$, this biggest probability becomes:

$$P_1 = \lim_{\tau \to 0} \frac{1}{1 + \mathrm{e}^{(Q_2-Q_1)/\tau} + \mathrm{e}^{(Q_3-Q_1)/\tau} + \mathrm{e}^{(Q_4-Q_1)/\tau}} = \frac{1}{1+0+0+0} = 1$$

(These expressions all go to zero since $Q_i - Q_1$ is negative). For the others, we do a similar computation and divide by the same amount. For example, the second probability becomes:

$$P_2 = \lim_{\tau \to 0} \frac{\mathrm{e}^{(Q_2-Q_1)/\tau}}{1 + + \mathrm{e}^{(Q_2-Q_1)/\tau} + \mathrm{e}^{(Q_3-Q_1)/\tau} + \mathrm{e}^{(Q_4-Q_1)/\tau}} = \frac{0}{1+0+0+0} = 0$$

Now going in the other direction, $\tau \to \infty$, we note that $Q_i/\tau \to 0$, so that

$$P_1 = \lim_{\tau \infty 0} \frac{\mathrm{e}^{Q_1/\tau}}{\mathrm{e}^{Q_1/\tau} + \mathrm{e}^{Q_2/\tau} + \mathrm{e}^{Q_3/\tau} + \mathrm{e}^{Q_4/\tau}} = \frac{1}{1+1+1+1} = \frac{1}{4}$$

And the others are calculated the same way.

16. What is the win-stay, lose-shift (or pursuit) strategy? What are the update rules?

SOLUTION: The idea behind the win-stay, lose-shift strategy is that, if we "win" with the currently selected machine, then the probability of choosing that machine should increase (and the rest decrease). For example, if we won with machine a, then

$$(P_a)_{\text{new}} = P_a + \beta(1 - P_a)$$

and all the remaining probabilities are decreased by the same amount:

$$(P_j)_{\text{new}} = P_j + \beta(0 - P_j)$$

17. Suppose we play with three machines, and machine 3 is chosen and gives a big payout (enough to make $Q_t(3)$ the maximum). Update the probabilities for win-stay, lose-shift, if they are: $P_1 = 0.3, P_2 = 0.5, P_3 = 0.2$ and $\beta = 0.3$.

    SOLUTION: P_3 increases, the rest decrease:

$$(P_3)_{\text{new}} = 0.2 + 0.3(1 - 0.2) = 0.44$$

$$(P_1)_{\text{new}} = 0.3 + 0.3(0 - 0.3) = 0.21$$
$$(P_2)_{\text{new}} = 0.5 + 0.3(0 - 0.5) = 0.35$$

    (Note that they still sum to 1).

18. Suppose we have a genetic algorithm with 4 chromosomes, and current fitness values $[-1, 1, 2, 3]$. We want to construct the probability of choosing these chromosomes for mating. Calculate the probabilities (if possible) using the current ordering, if we use:

    (a) Normalized fitness values: SOLUTION: Since the first value is negative, we cannot use the normalized fitness values.

    (b) Rank order, normalized: SOLUTION: The order is already from smallest to largest, so the probabilities (in the given ordering) is:

$$\frac{1}{10}, \frac{2}{10}, \frac{3}{10}, \frac{4}{10}$$

    (c) Softmax with $\tau = 1$. SOLUTION: Use a calculator to get the approximate values:

$$[0.012, 0.089, 0.242, 0.656]$$

19. Suppose we have a genetic algorithm using three dimensional real coordinates. Suppose the pair of coordinates for mating are given as $(1, 2, 1)$ and $(3, -1, 1)$. Calculate the children if we choose the middle coordinate and $\beta = 1/2$.

    SOLUTION: The y coordinate for child 1 will be:

$$(1 - \beta)2 + \beta(-1) = 1 - \frac{1}{2} = \frac{1}{2}$$

so child 1 will be the triple: $(1, 1/2, 1)$

The y coordinate for child 2 will be:

$$(1-\beta)(-1) + \beta(2) = -\frac{1}{2} + 1 = \frac{1}{2}$$

so child 2 will be the triple: $(3, 1/2, 1)$

20. Let

$$f(x_1, x_2) = 3x_1x_2 + x_1^2$$

(a) Linearize f about the point $(1, 1)$.

SOLUTION: We need $f(1, 1)$ and $\nabla f(1, 1)$. We see that $f(1, 1) = 3 + 1 = 4$. For the gradient,

$$\nabla f = [3x_2 + 2x_1, 3x_1]|_{(1,1)} = [5, 3]$$

Therefore, the linearization of f about the point $(1, 1)$ is the tangent plane:

$$4 + 5(x-1) + 3(y-1)$$

(b) Compute the Hessian of f

SOLUTION: The Hessian will be:

$$\begin{bmatrix} 2 & 3 \\ 3 & 0 \end{bmatrix}$$

(c) Show that Newton's Method, started at $(1, 1)$, converges in one step.

SOLUTION: TYPO- I meant to specify that Newton's Method was to optimize the function (or to find the zeros of the gradient). With that in mind, we can compute the necessary values:

$$\vec{x}_1 = \begin{bmatrix} 1 \\ 1 \end{bmatrix} - Hf^{-1}(1,1)\nabla f(1,1) = \begin{bmatrix} 1 \\ 1 \end{bmatrix} + \frac{1}{9}\begin{bmatrix} 0 & -3 \\ -3 & 2 \end{bmatrix}\begin{bmatrix} 5 \\ 3 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

And that is our sole critical point (setting the gradient to zero).

21. Let

$$f(x, y) = 3x^2 + xy - y^2 + 3x - 5y$$

(a) At the point $(1, 1)$, in which direction is f increasing the fastest? What is its rate of change in the fastest direction (Hint: Directional derivative)

SOLUTION: In general, the function will increase at the fastest rate if we move in the direction of the gradient. In that case, the rate of change will be the directional derivative in the direction of the gradient. At the point $(1, 1)$, we have:

$$D_u f = \nabla f(1,1) \cdot \frac{\nabla f(1,1)}{\|\nabla f(1,1)\|} = \|\nabla f(1,1)\|$$

In this case:

$$\nabla f = [6x + y + 3, x - 2y - 5]|_{(1,1)} = [10, -6]$$

so the magnitude is: $\sqrt{10^2 + (-6)^2} = \sqrt{136} \approx 11.66$

(b) Compute the Hessian of f.

SOLUTION: The Hessian will be $\begin{bmatrix} 6 & 1 \\ 1 & -2 \end{bmatrix}$

(c) Use the "second derivatives" test to determine if we have a local extremum at the critical point.

SOLUTION: The determinant of the Hessian is negative, so we will have a saddle point at the critical point (which we could find by setting the gradient to zero- But in this case is not needed).