Please provide complete and well-written solutions to the following exercises.

Due September 1, 10AM PST, to be uploaded as a single PDF document to blackboard (under the Assignments tab).

Homework 2

Exercise 1. Let $\mathcal{F}$ be the set of all positive double precision floating point numbers (except for `NaN`s and `Inf`s), that have the exponent `7fe` (in their hexadecimal representation in Matlab). (For example, after entering the command `format hex` in Matlab, we can see that the number `realmax` is in $\mathcal{F}$, since its hexadecimal representation in Matlab is `7fefffffffffffff`)

- How many elements are in $\mathcal{F}$? That is, what is the cardinality $|\mathcal{F}|$ of $\mathcal{F}$.
- What fraction of elements of $\mathcal{F}$ are in the interval $[2^{1023}, 2^{1024})$?
- What fraction of elements of $\mathcal{F}$ are in the interval $[2^{1023}, \frac{3}{2}2^{1023})$?
- Using e.g. Matlab's `rand` function, write a program that estimates the fraction of $x \in \mathcal{F}$ that satisfy the Matlab expression `x * (1/x)==1`. (It would take a pretty long time to check how many elements of $\mathcal{F}$ satisfy this equation, so you should not do that.)

Warning: Matlab's `rand` function tries to find a uniformly random chosen number in the interval $(0, 1)$ and then round it to the nearest floating point number. This operation is different than choosing a floating point number uniformly over all (positive) floating point numbers with a fixed exponent. (This is the point of the second and third items of this exercise, and the point of the `floatgui` program.) For this reason, your answer to the last part of the question should be much different from the output of the program: `x=rand(1,1000); sum(x.*(1./x)==1)/1000` .

Exercise 2. Do the following plot in Matlab.

```
x = 0.988:.0001:1.012;
y = x.^7-7*x.^6+21*x.^5-35*x.^4+35*x.^3-21*x.^2+7*x-1;
plot(x,y)
```

This is the function $y(x) = (x-1)^7$ for $x \in [.988, 1.012]$. Does the plot look like a polynomial? Explain why or why not.

Exercise 3. Suppose we want to solve the linear system of equations

$$17x_1 + 5x_2 = 22,$$
$$1.7x_1 + .5x_2 = 2.2.$$

Note that $(x_1, x_2) = (1, 1)$ is a solution to this system of equations.

Matlab can numerically solve this system with the following program

```
A = [17 5; 1.7 0.5];
b = [22; 2.2];
x = A\b
```

- What is the solution x that is output from the program?
- Is the output of the program an actual solution of the original system of equations?
- What is the determinant of A? What does Matlab output from the command `det(A)`?

Warning: for a 2×2 matrix A and a scalar $t > 0$, we have $\det(tA) = t^2 \det(A)$. So, the value of a determinant does not necessarily say anything about how well we can solve a linear system of equations of the form $Ax = b$.

Exercise 4. The sin function, like other special functions such as cos, exp, log, etc., cannot be computed exactly on a computer. A common way to compute these special functions is via power series. Recall that sin has the following power series that is absolutely convergent for all $x \in \mathbf{R}$:

$$\sin(x) = \sum_{k=0}^{\infty} (-1)^k \frac{x^{2k+1}}{(2k+1)!}.$$

With this power series in mind, run the following program when $x = \pi/2, 11\pi/2, 21\pi/2$ and $31\pi/2$. (Before you run the program, set x to a specific value.)

```
s = 0;
t = x;
n = 1;
while s+t ~= s;
    s = s + t;
    t =-x.^2/((n+1)*(n+2)).*t;
    n = n + 2;
end
```

When the program terminates, the value of s is the computed value of $\sin(x)$. For each value of x stated above, answer the following:

- What is the absolute error of the computation of $\sin(x)$?
- How many terms of the power series were used in the computation of $\sin(x)$?
- What is the largest term in the power series expansion of $\sin(x)$? (Hint: consider using the `max` command)

Exercise 5. Suppose we want to compute the quantity

$$x - \sin(x)$$

for any real $x \in \mathbf{R}$. For x near zero, there will be a loss of significance error, so we should perhaps try to find a better way to compute this quantity.

- Find the loss of significance (i.e. the number of zero bits at the end of the binary mantissa) when $x - \sin(x)$ is computed directly in double precision floating point arithmetic in Matlab, when $x = 2^{-25}$.
- Find the loss of significance (i.e. the number of zero bits at the end of the mantissa) when $x - \sin(x)$ is computed as
$$\frac{x^3}{3!} - \frac{x^5}{5!},$$
when $x = 2^{-25}$. (Your answer can be off by one or two from the true value.)
- Estimate the relative error when $x - \sin(x)$ is computed as
$$\frac{x^3}{3!} - \frac{x^5}{5!},$$
when $x = 2^{-25}$. (Your answer does not have to exactly correct. It is okay to be approximately correct.)

Exercise 6. This exercise examines an unstable recurrence computation.

Consider the following recursion with $x_1 := 1$ and $x_2 := 1/3$.
$$x_{n+1} = \frac{13}{3}x_n - \frac{4}{3}x_{n-1}, \qquad \forall\, n \geq 2.$$

- Verify that the recurrence is solved by $x_n := (1/3)^{n-1}$ for all $n \geq 1$.
- Using Matlab, solve for x_{40}. For example, use

```
x(1)=1; x(2)=1/3;
for i=3:40
    x(i)=(13/3)*x(i-1) - (4/3)*x(i-2);
end
x(40)
```

Is the answer what you expected to get? (Hint: examine a logarithmically scaled plot in the y-axis, using `semilogy(abs(x))`.)
- With a different initial condition, the above recurrence can have other solutions. To find them, rewrite the recurrence as
$$\begin{pmatrix} 13/3 & -4/3 \\ 1 & 0 \end{pmatrix} \begin{pmatrix} x_n \\ x_{n-1} \end{pmatrix} = \begin{pmatrix} x_{n+1} \\ x_n \end{pmatrix}, \qquad \forall\, n \geq 2.$$
Then note that the eigenvalues of the matrix $A := \begin{pmatrix} 13/3 & -4/3 \\ 1 & 0 \end{pmatrix}$ are $1/3$ and 4, so iterating the recurrence shows that
$$\begin{pmatrix} 13/3 & -4/3 \\ 1 & 0 \end{pmatrix}^{n-1} \begin{pmatrix} x_2 \\ x_1 \end{pmatrix} = \begin{pmatrix} x_{n+1} \\ x_n \end{pmatrix}, \qquad \forall\, n \geq 1.$$
Since A has two distinct eigenvalues, it is diagonalizable, so if $\begin{pmatrix} x_2 \\ x_1 \end{pmatrix}$ is written as a linear combination $a_1v_1 + a_2v_2$ of the corresponding eigenvectors $v_1, v_2 \in \mathbf{R}^2$ of A, the recurrence becomes
$$a_1(1/3)^{n-1}v_1 + a_2 4^{n-1} v_2 = \begin{pmatrix} x_{n+1} \\ x_n \end{pmatrix}, \qquad \forall\, n \geq 1.$$

- Show that in the case $x_1 = 1$ and $x_2 = 1/3$, we have $a_2 = 0$. However, small numerical errors that occur in the computation of the recurrence correspond to a_2 being computed to be nonzero. Explain how this relates to the logarithmic plot `semilogy(abs(x))` you examined above.