Name : ______________________ © Jeffrey A. Anderson **Class Number**: _____

Engr 11: Introduction to MATLAB
Student Exam 2 Questions

How long is this exam?

- This exam is scheduled for a 110 minute period.
- Make sure you have 4 sheets of paper (8 pages front and back) including this cover page.
- There are a total of 6 separate questions on this exam

How will your written work be graded on these questions?

- Your work should show evidence of original thought and deep understanding. Work that too closely resembles the ideas presented in Jeff's lesson notes will likely NOT earn top scores. Work that does not demonstrate individualized, nuanced understanding will likely NOT earn top scores.
- Read the directions carefully. Your work will be graded based on what you are being asked to do.
- In order to earn a top score, please show all your work. In most cases, a correct answer with no supporting work will NOT earn top scores. What you write down and how you write it are the most important means of getting a good score on this exam.
- Neatness and organization are IMPORTANT! Do your best to make your work easy to read.
- You will be graded on accurate use of the notation we studied in this class.

What can you use on this exam?

- You may use less-than or equal to SIX double-sided note sheets (or 12 singled-sided sheets).
- Each note sheet is to be no larger then 11-inches by 8.5-inches (standard U.S. letter-sized paper).
- You may write on both sides of your note sheets.
- Your note sheet must be handwritten.
- PLEASE SUBMIT ALL OF YOUR NOTE SHEETS WITH YOUR EXAM.
- You are allowed to use calculators for this exam. Examples of acceptable calculators include TI 83, TI 84, and TI 86 calculators. You are not allowed to use any calculator with a Computer Algebra System including TI 89 and TI NSpire. If you have a question, please ask your instructor about this.

What other rules govern your participation during this exam?

- PLEASE DO NOT TURN THIS PAGE UNTIL TOLD TO DO SO!
- It is a violation of the Foothill Academic Integrity Code to, in any way, assist another person in the completion of this exam. Please keep your own work covered up as much as possible during the exam so that others will not be tempted or distracted. Thank you for your cooperation.
- No notes (other than your note sheets), books, or classmates can be used as resources for this exam.
- Please turn off your cell phones during this exam. No cell phones will be allowed on your desk.

Table 1: Various Powers of Two

n	2^n	=	Decimal
0	2^0	=	1
1	2^1	=	2
2	2^2	=	4
3	2^3	=	8
4	2^4	=	16
5	2^5	=	32
6	2^6	=	64
7	2^7	=	128
8	2^8	=	256
9	2^9	=	512
10	2^{10}	=	$1,024$
11	2^{11}	=	$2,048$
12	2^{12}	=	$4,096$
13	2^{13}	=	$8,192$
14	2^{14}	=	$16,384$
15	2^{15}	=	$32,768$
16	2^{16}	=	$65,536$
22	2^{22}	=	$4,194,304$
23	2^{23}	=	$8,388,608$
31	2^{31}	=	$2,147,483,648$
32	2^{32}	=	$4,294,967,296$
52	2^{52}	=	$4,503,599,627,370,496$
53	2^{53}	=	$9,007,199,254,740,992$
63	2^{63}	=	$9,223,372,036,854,775,808$
64	2^{64}	=	$18,446,744,073,709,551,616$

Table 2: Hexadecimal Nibble Chart

Decimal	4-bit binary niblbe	Lowercase Hexadecimal
0	0000	0
1	0001	1
2	0010	2
3	0011	3
4	0100	4
5	0101	5
6	0110	6
7	0111	7
8	1000	8
9	1001	9
10	1010	a
11	1011	b
12	1100	c
13	1101	d
14	1110	e
15	1111	f

Lesson 5: Introduction to Fixed-Point Numbers

1. (Long T.) Suppose we have the following rational number:

$$x = \frac{235}{32}$$

   A. Explain why x has a finite representation in both radix 10 and radix 2 number systems.
   B. Translate this number into the $\texttt{ufixed8}_{2(3,5)}$ data class by computing the binary division problem.
   C. Translate this number into the $\texttt{ufixed8}_{2(3,5)}$ data class via an expansion in terms of powers of 2.
   D. What is the total range of $x \in \mathbb{Q}_I$ that can be stored in the $\texttt{ufixed8}_{2(3,5)}$ data class?

2. (Adam D.) What type of real numbers can be stored exactly in a finite binary expansion? What types of real numbers are impossible to store exactly using a finite binary expansion?

3. (Seungheon L.) Consider the following fractional representations of the rational number

$$x = \frac{647}{64}$$

   Using pen-and-paper analysis, convert x into a finite binary expansion. Please solve this problem using an expansion of x into terms of powers of 2. Then, assuming we set

$$\hat{y} = \texttt{ufixed8}_{2(5,3)}(x)$$

   find the binary representation for $\hat{y}$.

4. (Kenny Q.) Suppose $\ell, f \in \mathbb{N}_0$ such that $\ell + f = 16$. For a specific choice of ℓ and f that satisfy this condition, please determine:

   A. The largest number that can be stored in the $\texttt{ufixed16}_{2(\ell,f)}$ data class.
   B. The smallest number that can be stored in the $\texttt{ufixed16}_{2(\ell,f)}$ data class.
   C. The distance between two consecutive numbers stored in the $\texttt{ufixed16}_{2(\ell,f)}$ data class.

5. (Jonathan F., Jingwen P.) Using the "chop" rule for rounding (literally chop off the extra bits), set

$$x_1 = \texttt{ufixed8}_{2(5,3)}\left(\frac{349}{32}\right) \qquad \text{and} \qquad x_2 = \texttt{ufixed8}_{2(4,4)}\left(\frac{349}{32}\right)$$

   Then, calculate the absolute error $|x_1 - x_2|$ and explain why your answer makes sense?

6. (James S.) Consider all the possible data classes that might arise from the `ufixed8` function we created in lab 4:

$$\texttt{ufixed8}_{2(8,0)}$$
$$\texttt{ufixed8}_{2(7,1)}$$
$$\texttt{ufixed8}_{2(6,2)}$$
$$\texttt{ufixed8}_{2(5,3)}$$
$$\texttt{ufixed8}_{2(4,4)}$$
$$\texttt{ufixed8}_{2(3,5)}$$
$$\texttt{ufixed8}_{2(2,6)}$$
$$\texttt{ufixed8}_{2(1,7)}$$
$$\texttt{ufixed8}_{2(0,8)}$$

A. What is the minimum and maximum positive values associated with each of these representations?
B. How do these values relate to your understanding of floating-point numbers?
C. Why should floating-point representations be the next natural step to the fixed-point data classes.

Lesson 6: Introduction to Floating-Point Numbers

7. (Michelle T.) Consider the following two variables:

$$x = 1.0011 \times 2^2 \qquad \text{and} \qquad y = 1.1101 \times 2^1.$$

Let $a = x + y$ be a mathematical object that store the exact value of the sum. Now, typecast a as an 8-bit floating point number:

$$z = \text{binary8(a)}.$$

What is the 8-bit, raw binary code stored in memory for the value of z?

8. (Bintang G., Jingwen P.) Consider the following raw, uninterpreted 8-bit binary word

$$B = 1101\ 0111$$

What decimal value does the B have if:

A. we interpret this number as an unsigned binary integer?
B. we interpret this number as a signed integer in signed-magnitude representation?
C. we interpret this number as a signed integer in twos complemment representation?
D. we interpret this number as a $\texttt{ufixed8}_{2(2,6)}$ representation?
E. we interpret this number as a `binary8` representation?

9. (Ethan L.) Create a MATLAB function

```
function b = uint8_to_binary16(u)

end %function (uint8_to_binary16)
```

that takes as input a single unsigned integer typecast in MATLAB's `uint8` data class with

$$u = \text{ an unsigned integer where } 0 \leq u \leq 255.$$

This function should convert the input integer u into a raw bit string that encodes the corresponding to the value of u but stored in the `binary16` data class we discussed this quarter. The output of this function will be a 1×16 row-vector $\mathbf{b} \in \mathbb{R}^{1\times 16}$ whose individual entries are either 0 or 1. This row vector should represent the raw, uninterpreted $16-$bit binary word that stores the proper bit string to represent the value of `binary16`(u).

---

10. (Seungheon L.) Consider the following number

$$y = (15)_{10}$$

A. Convert y into an unsigned binary integer.

B. Represent y as a normalized floating point number.

C. Find the raw bit string used to store y in the `binary8` data class.

---

11. (James S.) Consider the uninterpreted bit string associated with the hexadecimal string `c0de`. What decimal value would be encoded by this bit string represent if we interpreted this number as:

A. an unsigned binary integer?

B. a signed integer in signed-magnitude representation?

C. a signed integer in twos complemment representation?

D. a fixed point number using data class $\texttt{ufixed16}_{2(8,8)}$?

E. a half-precision `binary16` floating point number with 1 bit for the sign, 5 bits for the biased exponent (with bias $K = 2^{5-1} - 1$), and 10 bits for the significand?

Lesson 7: Introduction to IEEE 754 Format

12. (Bintang G., Natalie T.) Consider the following snippet of code:

```
x = single(-122.8);
y = -122.8;
format hex
x, y
```

Show the output that MATLAB will produce in the Command Window after we execute this code. Note, you can assume that MATLAB rounds up.

13. (Natalie T., Long T., Jonathan F.) Convert the number $x \in \mathbb{Q}_I$ given by

$$x = \frac{1365}{32} \qquad \text{where} \qquad 1365 = 2^{10} + 2^8 + 2^6 + 2^4 + 2^2 + 2^0$$

into `single` and `double` format using pen and paper analysis. Note: for a take-home style exam, please allow students to use MATLAB to check their work and confirm. They can attach a screen shot of the command window as evidence of checking.

14. (Ethan L.) Suppose we use IEEE 754 format to create a variable using the following command:

$$x = \texttt{binary32}\left(\frac{47}{5}\right).$$

Write out the 32 bits used to encode x in memory and explain what each bit represents. Make explicit connections between the bit string you provide in your answer and the binary-to-decimal map defined by the `binary32` data class.

15. (Nick K.) Consider the following lines of code:

```
format hex
for i = 0:15
    x(i+1, 1) = 2^53 +i;
end %for (i)
x
```

What pattern(s) do you notice? How do your observations of patterns relate to the way we store double precision numbers? Compare the results generated in the command window with the exact values of these variables in infinite precision. Why don't these values match up? Remember that MATLAB stores these variables in double precision by default.

Lesson 10: Introduction to MATLAB Graphics

16. (Kenny Q.) Create a figure with four subplots. Details of the plots are labeled below.

    A. For the first subplot, define the following variables:

    ```
    x1 = [-5:5];    y1 = sin(x1);
    ```

    Then, plot `x1` versus `y1` using a green dotted line with × markers. Make sure the legend is present. Also label the horizontal axis using the string "x-axis" and label the vertical axis using the string "y-axis." Do not label the entire plot with a title for the first subplot.

    B. For the second subplot, define the following variables:

    ```
    x2 = [1:2:49];  y2 = x2 + x2^0.5;
    ```

    Then, plot `x2` versus `y2` using a red dash-dot line with triangle markers. This time, make sure that NO legend is present. Label the entire plot using the title "Is this a line?" Also, label the horizontal axis a "x" and the vertical axis as "y."

    C. For the second subplot, define the following variables:

    ```
    x3 = [1:0.0001:2];  y3 = exp(x3);
    ```

    Then, plot `x3` versus `y3` using a magenta dashed line with pentagram markers. Make sure that NO legend is present and label the entire plot using the title "The study of sharp increase in the exponential function over an interval." Also, label the horizontal axis with the string "the independent variable" and the vertical axis as "the dependent variable."

    D. For the final subplot, define the following variables:

    ```
    x4 = [0:0.01:2*pi]; y4 = cos(x4);
    ```

    Then, plot `x4` versus `y4` using a blue solid line with no markers. Create a legend and label the entire plot using the title "A study of the shape of the Cosine curve." Label the horizontal axis with the string "the horizontal axis" and the vertical axis as "the vertical axis."

Challenge Problem

17. (Josh H.) Consider the following rational number

$$x = \frac{737}{21}.$$

A. Convert this number into a half-precision `binary16` floating point string with 16 bits total. In this case, we use 1 bit for the sign, 5 bits for the biased exponent, and 10 bits for the significand. Specifically show the unadjusted binary string prior to conversion to floating point. Make sure to explicitly state the value of the exponent bias K. If necessary, use round to the nearest number.

B. What is the actual decimal value of the approximation stored in your `binary16` floating point string from part A above.

C. If we were to store x in the fixed-point data class $\texttt{ufixed16}_{2(6,10)}$, what binary string would we use to approximate this value?

D. Which approximation, the `binary16` or $\texttt{ufixed16}_{2(6,10)}$ is closer to the exact value of x?

E. Compare and contrast the advantages and disadvantages of choosing a fixed-point versus a floating-point representation for a given number of bits. Explain your reasoning using this example.