

Please answer the following questions in complete sentences in submit the solution on Blackboard Fri. Sept. 30th by Midnight.

Homework 3

- 2016-09-27: Small clarifications based on Piazza questions.

Problem 1: Warm up problems (15 points)

These questions are most similar to what I'd ask on an exam or on future quizzes.

1. Suppose \mathbf{A} is an $r \times s$ matrix. How many multiplication and addition operations are required for a general matrix vector product \mathbf{Ax} .
2. Suppose \mathbf{A} is an $m \times n$ matrix and \mathbf{D} is an $m \times m$ diagonal matrix (that means that it has zeros everywhere except on the diagonal elements). Write down the simplest expression for the ij th entry of \mathbf{DA} you can. (This will end up being used in the future for some problems!)
3. This is just like the previous problem, but now do the same thing where \mathbf{A} is an $m \times n$ matrix and \mathbf{D} is an $n \times n$ diagonal matrix. Write down the simplest expression for the ij th entry of \mathbf{AD} .

Problem 2: Implement MatMul

(If you wish, you may perform this task in a language besides Julia that provides a built-in matrix-matrix multiplication routine – such as Matlab or Python, but you must have the same comparison.)

1. Implement the following function:

```
"""
`matmul`
=====

Compute Matrix-matrix multiplication or MatMul.

Functions
-----
- `C = matmul(A,B)` returns the matrix-matrix product of C=A*B
"""
function matmul(A,B)
end
```

using only scalar operations. (i.e. without any of Julia's built in matrix operations.)

2. Report the output of your function for the following commands

```
@show matmul([1 2], [2; 1])
@show matmul([1 2; -2 4; 0 3], [-1; 1])
```

```

@show matmul([-1 1], [1 -2 0; 2 4 3])
@show matmul(3, 5)
t = pi/2
@show matmul([cos(t) -sin(t); sin(t) cos(t)],
             [cos(t) sin(t); -sin(t) cos(t)]);

```

3. Compare the time and accuracy of your code to Julia's built in matrix-matrix operator `*` on matrices of random normals with sizes ranging between 10 and 1000. To evaluate accuracy, use:

```

C = matmul(A,B)
D = A*B
diff = vecnorm(C-D,);

```

To evaluate time, use:

```
dt = @elapsed C = matmul(A,B);
```

Use 50 different random samples for each size and report the mean time for Julia vs. your own code as one plot. (So on the x-axis we have problem size, and on the y axis, we have one line for the mean time over 50 trials of Julia's built-in operation, and a second line we have the mean time over 50 trials of our function.)

Also prepare the same type of plot on the *maximum difference* over 50 trials between your code and Julia's code to show the accuracy.

4. Write a new function:

```

"""
`matmul2`
=====

Compute Matrix-matrix multiplication or MatMul faster

Functions
-----
- `C = matmul(A,B)` returns the matrix-matrix product of C=A*B
"""
function matmul(A,B)
end

```

That uses some of the ideas we talked about in class to make a faster matrix-matrix product. You may use Julia's built-in matrix-matrix multiplication for up to 16-by-16 matrices. Or matrix-vector products up to $256 - by - 1$. Compare the time and accuracy of your new `matmul2` code to your original.

Problem 3: Make Yoda Spin!

1. G&C Chapter 2 Problem 11. The Julia code for this example is

```

# create matrix whose columns contain the coordinates of
# each vertex.
U = [1.0 0 -1 0 1.0; 0 1 0 -1 0]

theta = pi/4.0

# Create a red unit square
# Note U(1,:) denotes the first row of U
plot(U[1,:]',U[2,:]',fill=(0,:red))

```

```

# Perform rotation.
R = [cos(theta) -sin(theta); sin(theta) cos(theta)];
V = R*U;
# Plot the blue square
plot!(V[1,:]', V[2,:]', fill=(0,:blue))

```

2. G&C Chapter 2 Problem 12 (not assigned)

```

**I really wanted to assign this one, but I still haven't
figured out how to translate the Matlab code into Julia
code. This one might get assigned in the future.**

```

Problem 4: Make an image small

Consider the following problem. We have a 32×32 pixel image. Each pixel is a real-valued number between 0 and 1. However, I want to show this on an iPhone and I only have a 16×16 pixel area to show the image. In order to reduce the size of the image, I want to average groups of 4 pixels. In this problem, we'll create a Julia program to do this

Let's work through a smaller example first. For the 4×4 image, represented here by a matrix-like array:

$$\begin{bmatrix} x_1 & x_2 & x_3 & x_4 \\ x_5 & x_6 & x_7 & x_8 \\ x_9 & x_{10} & x_{11} & x_{12} \\ x_{13} & x_{14} & x_{15} & x_{16} \end{bmatrix}$$

I want to compute

$$\begin{bmatrix} (x_1 + x_2 + x_5 + x_6)/4 & (x_3 + x_4 + x_7 + x_8)/4 \\ (x_9 + x_{10} + x_{13} + x_{14})/4 & (x_{11} + x_{12} + x_{15} + x_{16})/4 \end{bmatrix}.$$

We will solve this problem using a matrix-vector product.

1. Suppose I call:

$$\begin{aligned} y_1 &= (x_1 + x_2 + x_5 + x_6)/4 & y_2 &= (x_3 + x_4 + x_7 + x_8)/4 \\ y_3 &= (x_9 + x_{10} + x_{13} + x_{14})/4 & y_4 &= (x_{11} + x_{12} + x_{15} + x_{16})/4. \end{aligned}$$

Further, suppose we consider the vectors $\mathbf{y} \in \mathbb{R}^4$ and $\mathbf{x} \in \mathbb{R}^{16}$ to be the new image and old image, respectively. Write down the matrix \mathbf{A} such that $\mathbf{y} = \mathbf{A}\mathbf{x}$.

In the remainder of the problem, we'll work through how to do this for a particular image in Julia.

2. Download <http://www.cs.purdue.edu/homes/dgleich/cs314-2016/homeworks/smallicon.csv> and type

```

download("http://www.cs.purdue.edu/homes/dgleich/cs314-2016/homeworks/smallicon.csv", "smallicon.csv")
X = readcsv("smallicon.csv")

```

in Julia. You should get a 32×32 matrix \mathbf{X} . What is the sum of diagonal elements of \mathbf{X} ?

3. Show the image

```

using Images
grayim(X)

```

Or write it to a file

```
using FileIO
using ImageMagick
save("icon1.png",X)
```

But that picture doesn't look right, does it? To get full points, make sure you adjust the figure so that it looks correct.!

Show your final code and the image you get.

4. In what follows, we'll talk about two different types of indices. The image index of a pixel is a pair (i, j) that identifies a row and column for the pixel in the image. The vector index of a pixel is the index of that pixel in a linear ordering of the image elements. For instance, in the small little example, pixel $(3,2)$ has linear index 10. Also, pixel $(1,4)$ has index 4. Julia can help us built a map between pixel indices and linear or vector indices:

```
N = reshape(1:(4*4), 4, 4)'
```

This creates the pixel index to linear index for the problem above because

```
N(1,4)
N(3,2)
```

return the appropriate entry.

In your own words, explain what the `reshape` operation does. What happens if you omit the final transpose above and try:

```
N = reshape(1:(4*4), 4, 4)
```

instead?

5. Now we need to construct the matrix \mathbf{A} in order to reduce the size of a 32×32 image to a 16×16 image as we did in part 1. Suppose we call the output vector \mathbf{y} and the output image \mathbf{Y} . I'm giving you the following template, that I hope you can fill in. Feel free to construct an \mathbf{A} that accomplishes our image reduction task any way you choose, but the following should provide some guidance.

```
NX = <fill in> # the map between pixel indices and linear indices for X
NY = <fill in> # the map between pixel indices and linear indices for Y
A = <fill in> # the matrix we are trying to build
for i=1:32
    for j=1:32
        xi = <fill in> # the index of the pixel i,j in the vector x
        yi = <fill in> # hint, div(i,k) is integer division!

        A[yi,xi] = 1/4 # place the entry of the matrix
    end
end
end
```

6. In order to use the matrix \mathbf{A} we created, we need to convert the matrix \mathbf{X} into a vector! The reshape operation helps here:

```
x = reshape(X', 32*32, 1)
```

We can now rescale the image \mathbf{X} by multiplying by \mathbf{A} and reorganizing back into a matrix \mathbf{Y} .

```
y = A*x
Y = reshape(y, 16, 16)'
```

Show the image of \mathbf{Y} . Does that look correct? (Hint, apply the same correction you did before!)