
EECS 16A Imaging 2

TA, ASE, ASE, ASE


Agenda

- Quick overview + review
- Images as matrices and vectors
- Pixel-by-pixel scanning
- Reconstructing scans as images

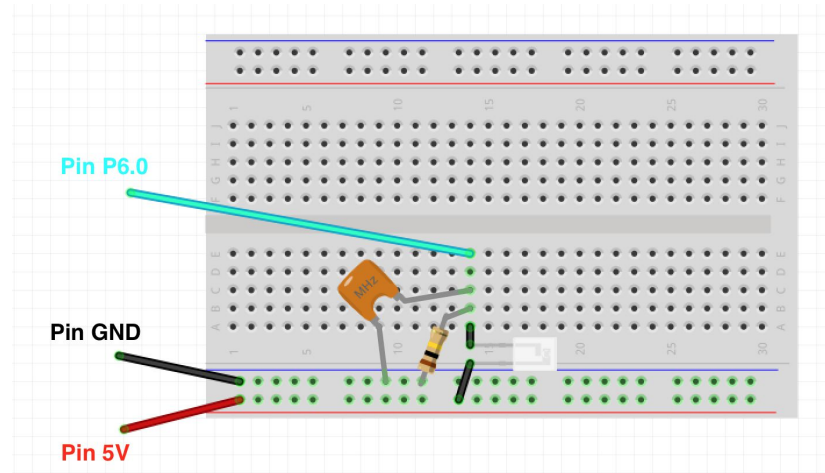
Last Week: Imaging 1

- Built our very first circuit!
 - What did this circuit do?

Circuit Diagram



Breadboard Diagram



Today's Lab: Single Pixel Scanning

- Circuit from last week measures **light** intensity
- Simulated projector illuminates image in a controlled way
- Python programming to reconstruct image

Why?

- Imaging 1:
 - Finding a link between physical quantities and voltage is powerful
 - If you can digitize it, you can do anything (IOT devices, internet, code, processing)
- Imaging 2:
 - What measurements are good measurements?
 - Remember Kody and Nara from Dis1A

Kody and Nara

2. Finding The Bright Cave

Nara the one-handed druid and Kody the one-handed ranger find themselves in dire straits. Before them is a cliff with four cave entrances arranged in a square: two upper caves and two lower caves. Each entrance emits a certain amount of light, and the two wish to find exactly the amount of light coming from each cave. Here's the catch: after contracting a particularly potent strain of ghoulish fever, our intrepid heroes are only able to see the total intensity of light before them (so their eyes operate like a single-pixel camera). Kody and Nara are capable adventurers, but they don't know any linear algebra – and they need your help.

Kody proposes an imaging strategy where he uses his hand to completely block the light from two caves at a time. He is able to take measurements using the following four masks (black means the light is blocked from that cave):




Figure 1: Four image masks.

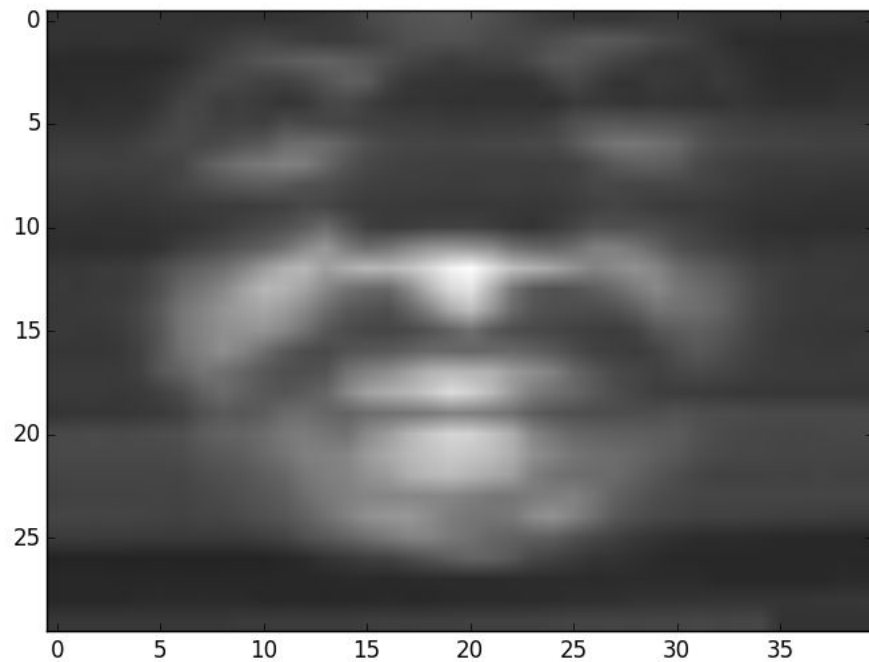
- Let \vec{x} be the four-element vector that represents the magnitude of light emanating from the four cave entrances. Write a matrix \mathbf{K} that performs the masking process in Figure 1 on the vector \vec{x} , such that $\mathbf{K}\vec{x}$ is the result of the four measurements.
- Does Kody's set of masks give us a unique solution for all four caves' light intensities? Why or why not?
- Nara, in her infinite wisdom, places her one hand diagonally across the entrances, covering two of the cave entrances. However, her hand is not wide enough, letting in 50% of the light from the caves covered and 100% of the light from the caves not covered. The following diagram shows the percentage of light let through from each cave:

Illuminating the Big Picture

- Linear dependence
 - When can you recover your image?
 - Does it matter what mask matrix you pick?
 - Does it matter how you cover the pixels?
- Invertibility
 - When can you solve $Ax = b$?
 - How does this relate to our system?
 - How does this affect the way we pick our masking matrix?

Sample Images

Nick:



Images, Matrices, Vectors



- What are the unknowns in our system?
 - Want to do an experiment to get information about these unknowns
- We can do a lot of interesting processing on vectors, but we need to convert the image into one first
 - In lecture and discussion, you have seen how to turn an image into a vector. How?

Images, Matrices, Vectors



[0]	[1]
[2]	[3]
[4]	[5]

Images, Matrices, Vectors



[0]	[1]
[2]	[3]
[4]	[5]



[0]
[1]

Images, Matrices, Vectors



[0]	[1]
[2]	[3]
[4]	[5]




[0]
[1]
[2]
[3]

Images, Matrices, Vectors



[0]	[1]
[2]	[3]
[4]	[5]



[0]
[1]
[2]
[3]
[4]
[5]

Images, Matrices, Vectors



[0]	[1]
[2]	[3]
[4]	[5]

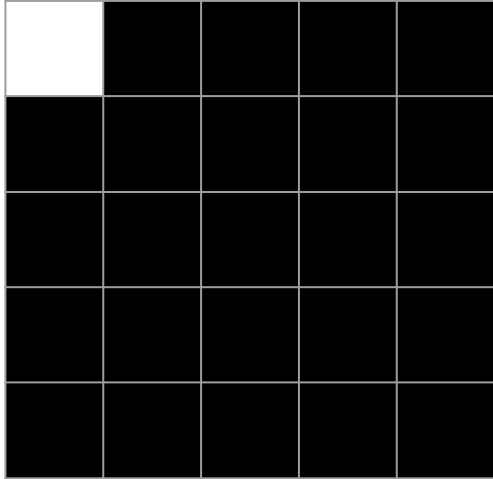


[0]
[1]
[2]
[3]
[4]
[5]

Pixel-by-Pixel Scan of an Image



Pixel-by-Pixel Scan of an Image



Pixel-by-Pixel Scan of an Image




Masked image



Image

Pixel-by-Pixel Scan of an Image




Masked image



Image

Pixel-by-Pixel Scan of an Image



Masked image



Image

Poll Time!

What would you expect the dimensions of a vector representing a 2x3 image to be?

- A. 2x3
- B. 3x2
- C. 6x1
- D. 5x1

To read all the pixels of a 4x4 image, how many pixel-by-pixel scans do we need to do?

- A. 4
- B. 8
- C. 16
- D. 32

Poll Time!


What would you expect the dimensions of a vector representing a 2x3 image to be?

- A. 2x3
- B. 3x2
- C. 6x1
- D. 5x1

To read all the pixels of a 4x4 image, how many pixel-by-pixel scans do we need to do?


- A. 4
- B. 8
- C. 16
- D. 32

Representing our Masks in Python



mask0 =
np.array([[[1, 0, 0, 0, 0]
[0, 0, 0, 0, 0]
[0, 0, 0, 0, 0]
[0, 0, 0, 0, 0]
[0, 0, 0, 0, 0]]])

Representing our Masks in Python




mask1 =
np.array([[[0, 1, 0, 0, 0]
[0, 0, 0, 0, 0]
[0, 0, 0, 0, 0]
[0, 0, 0, 0, 0]
[0, 0, 0, 0, 0]]])

Turning the Masks Into Vectors

5x5 mask to 25x1 vector


mask0 =

1	0	0	0	0
0	0	0	0	0
0	0	0	0	0
0	0	0	0	0
0	0	0	0	0



Row 0	1	T
Row 1	0	
Row 2	0	
Row 3	0	
Row 4	0	
	⋮	
Row 24	0	

Turning the Masks Into Vectors




Generating the Masking Matrix from the Masks


mask0

1,	0,	0,	0,	0
0,	0,	0,	0,	0
0,	0,	0,	0,	0
0,	0,	0,	0,	0
0,	0,	0,	0,	0

Generating the Masking Matrix from the Masks



Generating the Masking Matrix from the Masks




Generating the Masking Matrix from the Masks

mask1


0,	1,	0,	0,	0
0,	0,	0,	0,	0
0,	0,	0,	0,	0
0,	0,	0,	0,	0
0,	0,	0,	0,	0

1	0	0	0	0	0	0	0	...
---	---	---	---	---	---	---	---	-----

Generating the Masking Matrix from the Masks



Generating the Masking Matrix from the Masks



Generating the Masking Matrix from the Masks

1	0	0	0	0	0	0	0	...
0	1	0	0	0	0	0	0	...
0	0	1	0	0	0	0	0	...

Generating the Masking Matrix from the Masks

1	0	0	0	0	0	0	0	...
0	1	0	0	0	0	0	0	...
0	0	1	0	0	0	0	0	...
0	0	0	1	0	0	0	0	...

Generating the Masking Matrix from the Masks

$$H =$$

1	0	0	0	0	0	0	0	...
0	1	0	0	0	0	0	0	...
0	0	1	0	0	0	0	0	...
0	0	0	1	0	0	0	0	...
0	0	0	0	1	0	0	0	...
0	0	0	0	0	1	0	0	...
0	0	0	0	0	0	1	0	...
0	0	0	0	0	0	0	1	...
...								

Measuring a Pixel is Matrix-Vector Multiplication

1	0	0	0	0	0	0	0	...
0	1	0	0	0	0	0	0	...
0	0	1	0	0	0	0	0	...
0	0	0	1	0	0	0	0	...
0	0	0	0	1	0	0	0	...
0	0	0	0	0	1	0	0	...
0	0	0	0	0	0	1	0	...
...								

Masking Matrix H

i_1
i_2
i_3
i_n

Unknown,
vectorized
image, \vec{i}

=

s_1
s_2
s_3
s_n

Recorded
Sensor
readings, \vec{s}

Measuring a Pixel is Matrix-Vector Multiplication

$$\vec{s} = H\vec{i}$$

- We know H and we have the sensor readings, how do we get the image?
- How do we solve this?
- When can we solve this?
 - Conditions on H

Poll Time!

$$\vec{s} = H\vec{i}$$

Select all of the following that must be true for the image vector i to be recoverable from the sensor vector s .

1. H must be invertible
2. H must have linearly independent rows
3. H must be a square matrix
4. H must be the identity matrix

Poll Time!

Select all of the following that must be true for the image vector i to be recoverable from the sensor vector s .

1. H must be invertible
2. H must have linearly independent rows
3. H must be a square matrix¹
4. H must be the identity matrix

¹A tall matrix with redundant equations could be made to work, but noise in the system might cause the equations to become inconsistent (More in Imaging 3)

How Scanning Works: iPython

H =


1	0	0	0	0	0	0	0	...
0	1	0	0	0	0	0	0	...
0	0	1	0	0	0	0	0	...
0	0	0	1	0	0	0	0	...
0	0	0	0	1	0	0	0	...
0	0	0	0	0	1	0	0	...
0	0	0	0	0	0	1	0	...
...								

How Scanning Works: iPython

$H =$

1	0	0	0	0	0	0	0	...
0	1	0	0	0	0	0	0	...
0	0	1	0	0	0	0	0	...
0	0	0	1	0	0	0	0	...
0	0	0	0	1	0	0	0	...
0	0	0	0	0	1	0	0	...
0	0	0	0	0	0	1	0	...
...								

How Scanning Works: iPython



How Scanning Works: iPython

H =

1	0	0	0	0	0	0	0	...
0	1	0	0	0	0	0	0	...
0	0	1	0	0	0	0	0	...
0	0	0	1	0	0	0	0	...
0	0	0	0	1	0	0	0	...
0	0	0	0	0	1	0	0	...
0	0	0	0	0	0	1	0	...
...								



0
1
0
0
0
⋮
0



0,	1,	0,	0,	0
0,	0,	0,	0,	0
0,	0,	0,	0,	0
0,	0,	0,	0,	0
0,	0,	0,	0,	0

How Scanning Works: iPython

H =


1	0	0	0	0	0	0	0	...
0	1	0	0	0	0	0	0	...
0	0	1	0	0	0	0	0	...
0	0	0	1	0	0	0	0	...
0	0	0	0	1	0	0	0	...
0	0	0	0	0	1	0	0	...
0	0	0	0	0	0	1	0	...
...								



0
1
0
0
0
...
0



0,	1,	0,	0,	0
0,	0,	0,	0,	0
0,	0,	0,	0,	0
0,	0,	0,	0,	0
0,	0,	0,	0,	0






What Makes a Mask Good?


- Linearly independent columns → Invertible
 - Can't get a solution without this
 - There is a unique solution
- What would be a bad mask?
- Food for thought: Are all invertible matrices equally as good?
 - Find out in Imaging 3 next week

Cumulative Scanning


Identity
(H)




Alternating
(H_Alt)




Setup



Power strip to power your projector



Sample Setup Images



Setup (cont.)

1. Draw a “simple” image
2. Project masks (rows of H) onto it in a dark environment
3. Measure ambient light sensor reading s
4. Multiply by H inverse to find i ($= H^{-1}.s$)

Tips for a Good Image

- READ CLOSELY. There are many small directions that help you get a good setup
- Focus projector using dial on the side
- Close the box firmly & scan under dark conditions

Poll Time

Select all of the following that describe the relationship between H (the masking matrix), s (the sensor vector), and i (the image vector)?

1. $Hs = i$
2. $Hi = s$
3. $H^{-1}i = s$
4. $H^{-1}s = i$
5. $i * s = H$

Poll Time

Select all of the following that describe the relationship between H (the masking matrix), s (the sensor vector), and i (the image vector)?

1. $Hs = i$
2. $Hi = s$
3. $H^{-1}i = s$
4. $H^{-1}s = i$
5. $i * s = H$

Debugging



1. Make sure wires/resistors/light sensor are not loose
2. Light sensor orientation: short leg goes into +ve
3. Check COM Port
4. Reupload code to launchpad after making any change in circuit
5. Check Baud Rate in Serial Monitor (115200)
6. Projector might randomly restart in the middle of the lab. Make sure brightness 0 contrast 100.
7. Cover box with jacket for dark scanning conditions
8. If you see a very bright corner in the scan, move the light sensor away from the projector