EECS 16A Imaging 3

Insert your names here

Announcements

- Buffer labs will be **3/2-3/5**
  - You can make up **one** missed lab from the Imaging Module, if needed (unless you have received approval to make-up multiple labs)
  - Fill out the sign-up form (linked at the end of the lab notebook) if you plan to attend a buffer section

Announcements

- Optional Imaging Lab sections on **3/3, 3/4**
  - We'll be hosting an optional lab which will give you the opportunity to try out your light sensor circuit from Imaging I with real images!
  - Fill out the sign-up form (linked at the end of the lab notebook) if you plan to attend an optional lab section
- See upcoming Piazza post for more details on Imaging Buffer and the Optional Imaging Lab

Last time: 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	...
				...				

i_1
i_2
i_3
i_n

=

s_1
s_2
s_3
s_n

Masking Matrix H

Unknown, vectorized image, $\vec{i}$

Recorded Sensor readings, $\vec{s}$

Last time: Single-pixel scanning

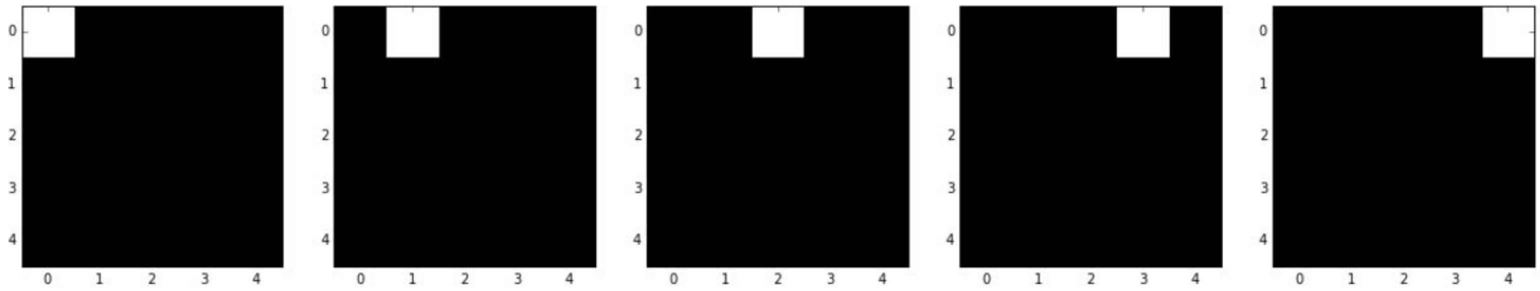

- Setup a masking matrix where each row is a mask
  - Measured each pixel individually once

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

- How did we reconstruct our image, once we had s?

Poll Time! (this is review)

What are the requirements of our masking matrix H? (multiple choice)

A. H is invertible.
B. H has linearly independent columns
C. H has a trivial nullspace.
D. Determinant of H is 0.

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

Our system

Poll Time! (this is review)

What are the requirements of our masking matrix H? (multiple choice)

A. H is invertible.
B. H has linearly independent columns
C. H has a trivial nullspace.
D. Determinant of H is 0.

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

Our system

Questions from Imaging 2

Goal: Understand which measurements are good measurements

- ✓ Can we always reconstruct our image → **need invertible H**
- ? Are all invertible matrices equally good as scanning matrices?
- ? What happens if we mess up a single scan?
- ? What if we use multiple pixel instead of single pixel scan?

Today: Multi-pixel scanning

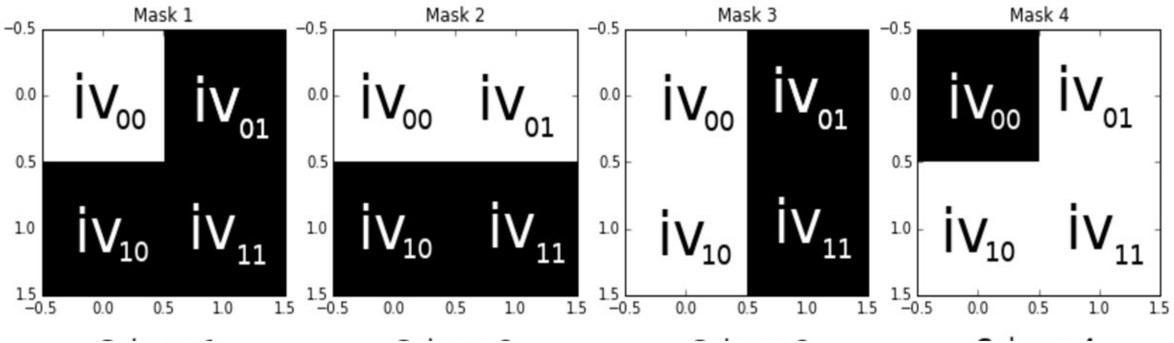

- **Can we measure multiple pixels at a time?**
  - Measurements are now linear combinations of pixels

- **How can we reconstruct our scanned image?**
  - Is multi-pixel mask still possible to be linearly independent, aka invertible?

Why do we care?

- Improve image quality by averaging
    - Good measurements → good average
- Redundancy is useful
    - Averaging measurements is better than using bad measurement values
    - Does not “solve” bad measurements, but makes us tolerant of some errors

How do we do it?

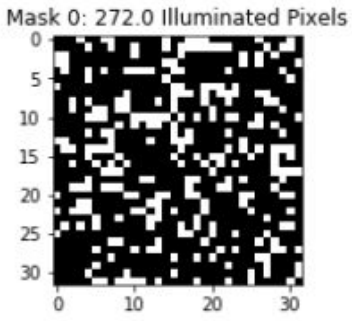

- Change masks to illuminate multiple pixels per scan
  - Multiple 1's in each row of masking matrix H
  - Measure linear combinations of pixels instead of single pixels
- BUT multiple pixels → more noise
  - Noise = random variation in our measurement that we don't want (ex: room light getting into box)
  - Signal = data that we do want (light from pixel illumination)
- Too much noise → hard to distinguish signal from noise
  - Want high signal, low noise
  - High signal-to-noise ratio (SNR)

What is noise?

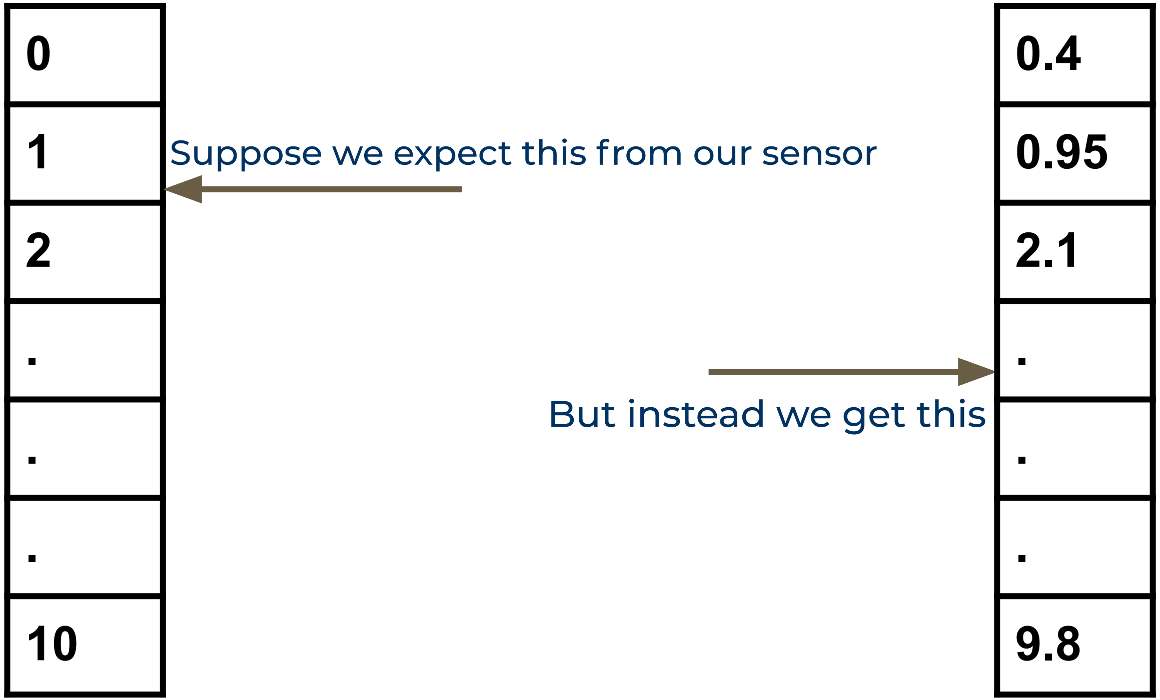

What is noise?

Measured values =
ideal vector + noise vector (ω)

$$\begin{bmatrix} 0.4 \\ 0.95 \\ 2.1 \\ . \\ . \\ . \\ 9.8 \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \\ 2 \\ . \\ . \\ . \\ 10 \end{bmatrix} + \begin{bmatrix} 0.4 \\ -0.05 \\ 0.1 \\ . \\ . \\ . \\ -0.2 \end{bmatrix}$$

$$\vec{s}_{real} \qquad \vec{s}_{ideal} \qquad \vec{w}$$

How does noise affect our system?

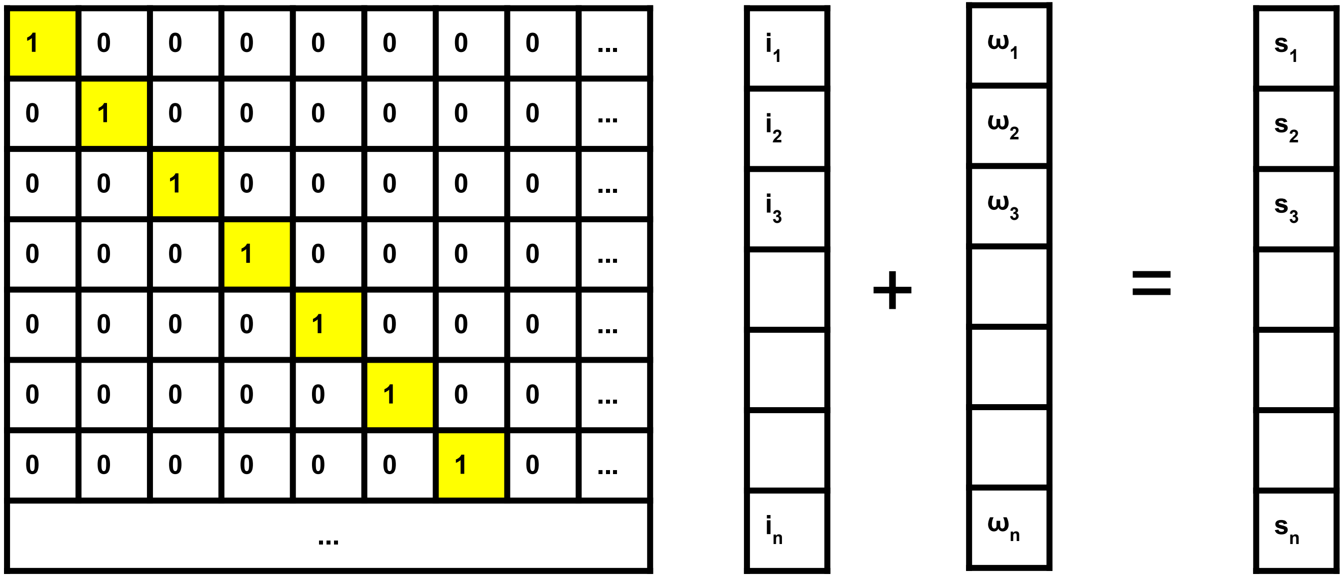

Masking Matrix H

Unknown, vectorized image, $\vec{i}$

Random noise vector, $\vec{w}$

Recorded Sensor readings, $\vec{s}$

A more realistic system

- Sensor readings = image vectors applied to H + noise vector

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

- We can't reconstruct **i**, but we can estimate it

$$\vec{i}_{est} = H^{-1}\vec{s} = \vec{i} + \boxed{H^{-1}\vec{w}}$$

Be careful about the noise term or else it could blow up !!

Eigenvalues for inverse matrices

- H Is an NxN matrix that we know is linearly independent (invertible).
    - No eigenvalue = 0
- Assume H has N linearly independent eigenvectors
- $\mathrm{H}v_i = \lambda_i v_i \text{ for } i = 1 \ldots \mathrm{N}$
- N lin. ind. vectors can span $\mathbb{R}^N$
    - They span the noise vector
- The inverse of H has eigenvalues $\frac{1}{\lambda_1} \cdots \frac{1}{\lambda_N}$ (as proven in homework)

$$\mathrm{H}^{-1} v_i = \frac{1}{\lambda_i} v_i \text{ for } i = 1 \ldots \mathrm{N}$$

How do eigenvalues affect noise?

The noise vector can be written as:

$$\vec{\omega} = \alpha_1 \overrightarrow{v_1} + \alpha_2 \overrightarrow{v_2} + \cdots \alpha_n \overrightarrow{v_n}$$

Including effect of H^{-1}

$$H^{-1}\vec{\omega} = H^{-1}(\alpha_1 \overrightarrow{v_1} + \alpha_2 \overrightarrow{v_2} + \cdots \alpha_n \overrightarrow{v_n})$$

Rewritten with eigenvalues:

$$H^{-1}\vec{\omega} = \frac{1}{\lambda_1}\alpha_1 \overrightarrow{v_1} + \frac{1}{\lambda_2}\alpha_2 \overrightarrow{v_2} + \cdots \frac{1}{\lambda_n}\alpha_n \overrightarrow{v_n}$$

Linking it all together

$$\vec{i}_{est} = H^{-1}\vec{s} + \boxed{H^{-1}\vec{\omega}}$$

$$\boxed{H^{-1}\vec{\omega}} = \frac{1}{\lambda_1}\alpha_1\vec{v_1} + \frac{1}{\lambda_2}\alpha_2\vec{v_2} + \cdots \frac{1}{\lambda_n}\alpha_n\vec{v_n}$$

- Remember: want small noise term for high signal-to-noise ratio
- The noise is directly related to the eigenvalues.

Poll Time!

- **Do we want small or large eigenvalues for the H matrix in order to get a good image?**
  A. Large
  B. The magnitude doesn't matter
  C. Small

- **Which of the following equations correctly model our imaging system? (multiple choice)**
  A. $s_{ideal} = H.i$
  B. $s_{real} = s_{ideal} + w = H.i + w$
  C. $s_{real} = s_{ideal} + w = H.i + H.w$
  D. $i_{est} = H^{-1}.s_{real} = H^{-1}.s_{ideal} + H^{-1}.w$
  E. $i_{est} = H^{-1}.s_{real} = H^{-1}.s_{ideal} + w$

Poll Time!

- Do we want small or large eigenvalues for the H matrix in order to get a good image?
  A. Large
  B. The magnitude doesn't matter
  C. Small

- Which of the following equations correctly model our imaging system? (multiple choice)
  A. $s_{ideal} = H.i$
  B. $s_{real} = s_{ideal} + w = H.i + w$
  C. $s_{real} = s_{ideal} + w = H.i + H.w$
  D. $i_{est} = H^{-1}.s_{real} = H^{-1}.s_{ideal} + H^{-1}.w$
  E. $i_{est} = H^{-1}.s_{real} = H^{-1}.s_{ideal} + w$

Possible scanning matrix: Random

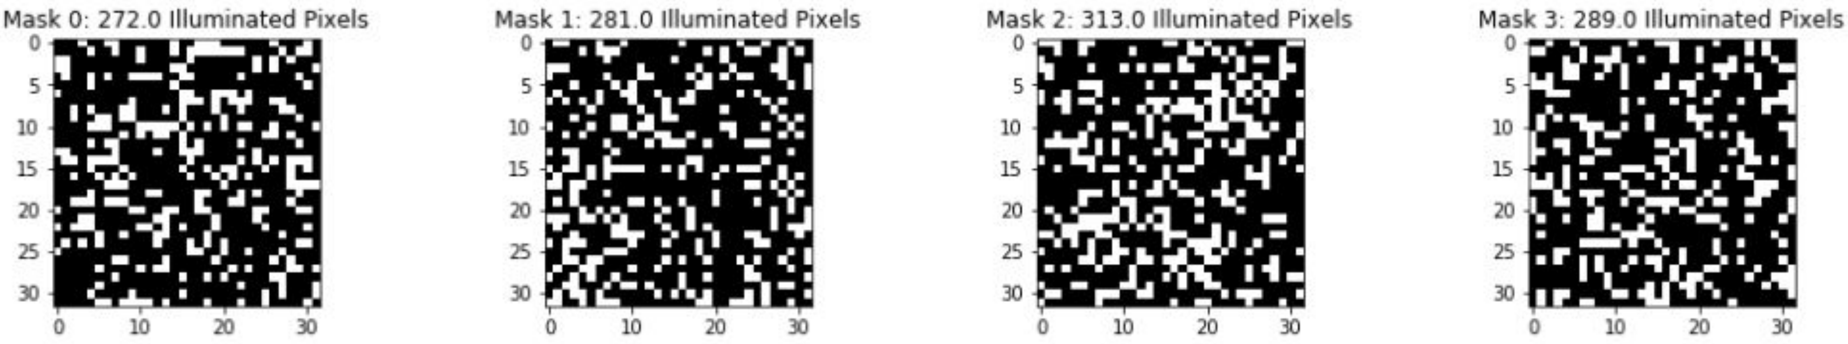

- Illuminate ~300 pixels per scan
  - *Usually* invertible
  - But what are its eigenvalues?

¯_(ツ)_/¯

A more systematic scanning matrix

- Hadamard matrix!
- Constructed to have large eigenvalues
  - Just what we need!

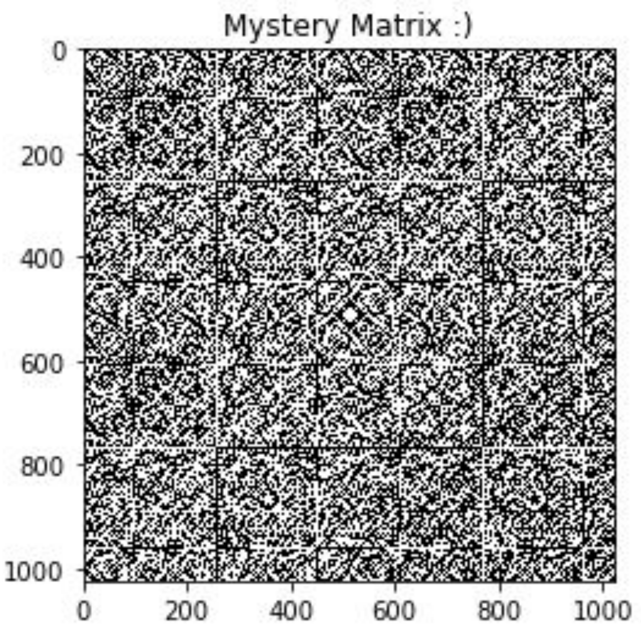

Software simulator setup

- Upload ("simple") image of an object (or use 16A logo)
- Convert image to 32x32 form for the simulator
- Project masks (rows of H) onto it and "measure" **s** using matrix multiplication
- Multiply with H inverse to find **i** ($=H^{-1}\mathbf{s}$)

The simulator handles noise addition in step 3 using a parameter sigma

Using the software simulator

- Start display view in another browser tab
- Enter the imagePath and run the simulator + shift to display tab
- Observe masks being projected onto the image + return to notebook tab
- Observe generated sensor reading
- Reconstruct image by multiplying with H inverse

Repeat steps 2-5 for each imaging experiment

Pointers

1. READ CAREFULLY - Long lab with lots of reading; heavily tests understanding of eigen-stuff (important for the exam)
2. You see the noisy sensor reading generated at the end instead of being generated entry by entry (i.e. just one masking simulation visual per experiment, no more cumulative simulation)
3. Choose an image that focuses on a single object and is not too detailed
4. Use a simple imagePath name
5. Before starting the imaging experiments, launch the display view in a separate tab using the link in the notebook
6. Enter imagePath correctly for each simulation block
7. Shift to the display tab as soon as you run a simulation block and return to the notebook once the visual has finished executing
8. In case the kernel crashes, simply save your notebook and restart it. You should navigate to the previous import block and run all blocks starting there.

P.S. The masking simulation visual can be super trippy ;)