

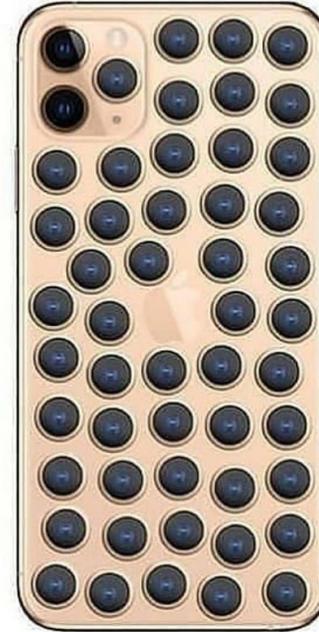
2019

iPhone 11 Pro



2029

iPhone 21 Pro



EECS 16A

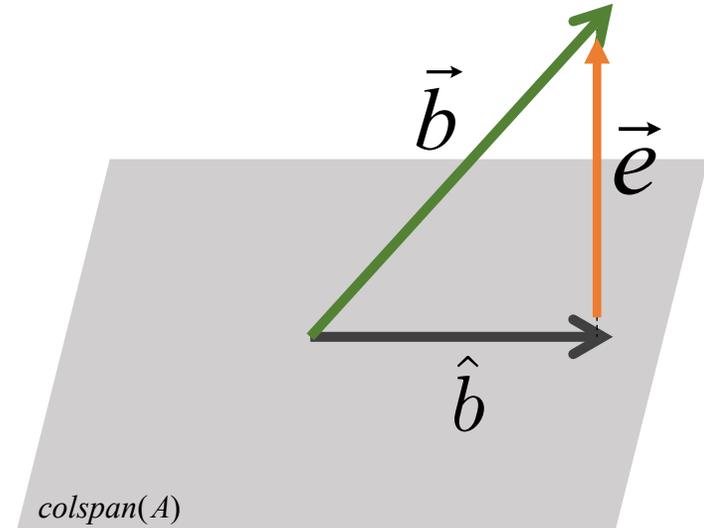
fun stuff: computational imaging

EECS16A: course evaluations

- course-evaluations.berkeley.edu

Overdetermined system: use least squares

$$\begin{matrix} \text{gray rounded rectangle} \\ A \end{matrix} \begin{matrix} \text{blue rounded rectangle} \\ x \end{matrix} = \begin{matrix} \text{green rounded rectangle} \\ b \end{matrix}$$



- the least-squares solution “minimally perturbs” b

$$\hat{x} = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \vec{b}$$

Underdetermined system: ????

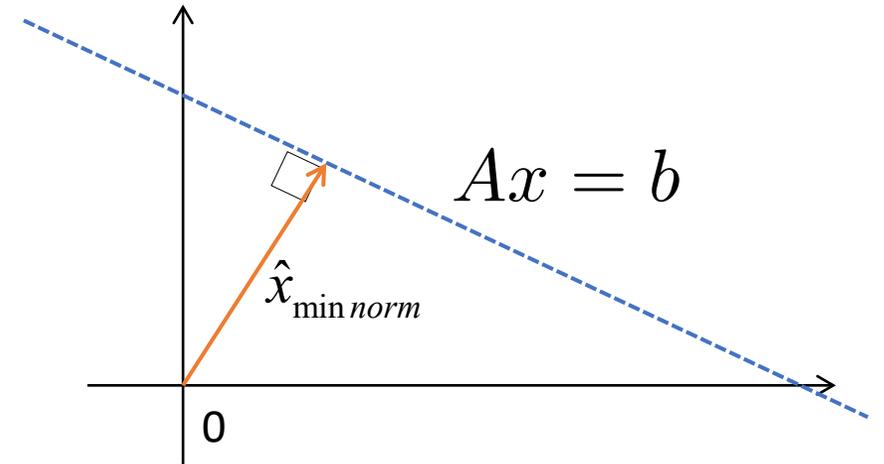
$$A x = b$$

IF TV SCIENCE WAS MORE LIKE REAL SCIENCE



Underdetermined system: ????

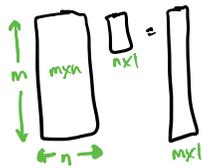
$$A x = b$$



- Can be infinite valid solutions!
- Ideas: pick the 'smallest' one? The 'sparsest'?
 - e.g. min norm:

$$\hat{x}_{\min norm} = \mathbf{A}^T (\mathbf{A}\mathbf{A}^T)^{-1} \vec{b}$$

So far, we've only talked about 'overdetermined systems':

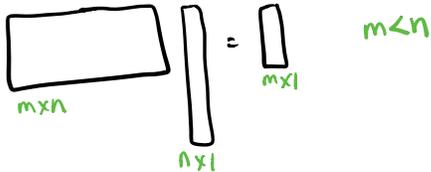


$m > n$ "overdetermined"

$A\vec{x} = \vec{b}$
 Recall: for least squares sol'n we minimized error:
 $\min_{\vec{x}} \|\vec{e}\|^2 = \min_{\vec{x}} \|A\vec{x} - \vec{b}\|^2$
 $\hookrightarrow \hat{\vec{x}} = (A^T A)^{-1} A^T \vec{b}$

What if problem is "underdetermined"?

Infinite sol'n's (more unknowns than eq's)
 \hookrightarrow (though 0 sol'n's if inconsistent)



But, can we find a good estimate?

\hookrightarrow Yes, in special cases!

e.g. sparsity (most entries are zero)

Example: We could look for 'smallest' \rightarrow min norm solution:

$$\begin{bmatrix} 1 & 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} 1 \end{bmatrix} \rightarrow x_1 + x_2 = 1$$

smallest sol'n:
 $x_1 = 0.5, x_2 = 0.5$

$\min_{\vec{x}} \|\vec{x}\|^2$ such that $A\vec{x} = \vec{b}$ (constrained optimization) $\rightarrow \hat{\vec{x}} = A^T (AA^T)^{-1} \vec{b}$

to understand this, take EE 127

$\min_{\vec{x}, \vec{\lambda}} \|\vec{x}\|^2 + \vec{\lambda}^T (\vec{b} - A\vec{x})$
scalars (Lagrange multiplier)

take $\frac{\partial}{\partial \vec{x}} (\vec{x}^T \vec{x} + \vec{\lambda}^T (\vec{b} - A\vec{x})) = 0$

$A^T 2\vec{x} - A^T \vec{\lambda} = 0$

$\vec{\lambda} = (AA^T)^{-1} 2A\vec{x}$

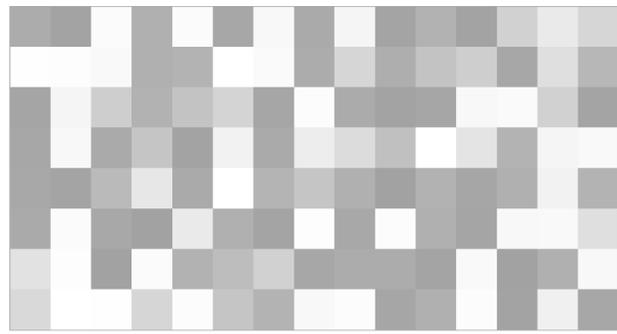
take deriv. $A\vec{x} = \vec{b}$

$\vec{\lambda} = (AA^T)^{-1} 2\vec{b}$

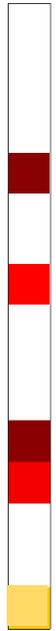
$\hat{\vec{x}} = A^T (AA^T)^{-1} \vec{b}$

Minimum norm sol'n!

'Sparsity' tells us how 'dense' the solution is

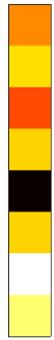


A



x

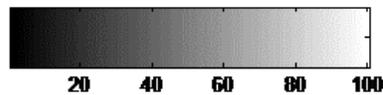
=



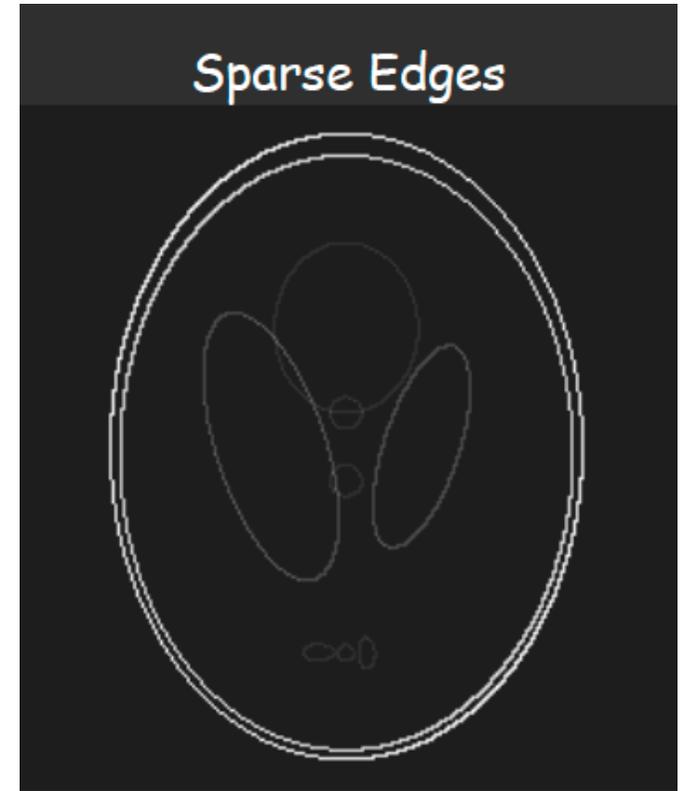
b

The fraction of non-zero elements in a matrix is called the *sparsity* (e.g. here $k=5$ for x vector)

Sometimes not-sparse things are sparse in a different basis



Take $|\text{derivatives}|$

A blue arrow pointing right, indicating the operation "Take $|\text{derivatives}|$ ". The text "Take $|\text{derivatives}|$ " is written in white inside the arrow.

How can we use sparsity for good?

Example: image compression

Reduce memory by smartly choosing which information to throw away



No compression



23:1 compression



144:1 compression

sparsity

Image compression:

$$\begin{matrix} \text{not} \\ \text{sparse} \end{matrix} \begin{bmatrix} | \\ | \\ \tilde{x} \\ | \\ | \end{bmatrix}_{n \times 1} = \begin{bmatrix} \Psi \end{bmatrix}_{n \times n} \begin{matrix} \text{sparse} \\ | \\ | \\ s \\ | \\ | \end{matrix}_{n \times 1}$$

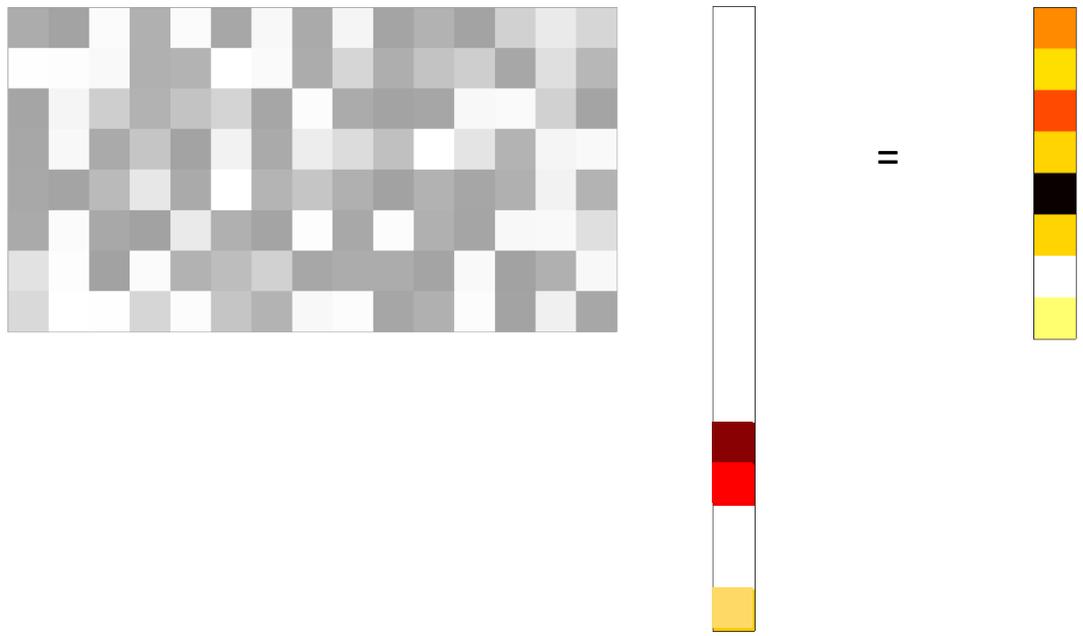
transform matrix
change of basis
↳ DCT, FFT, etc.

but mostly zeros!

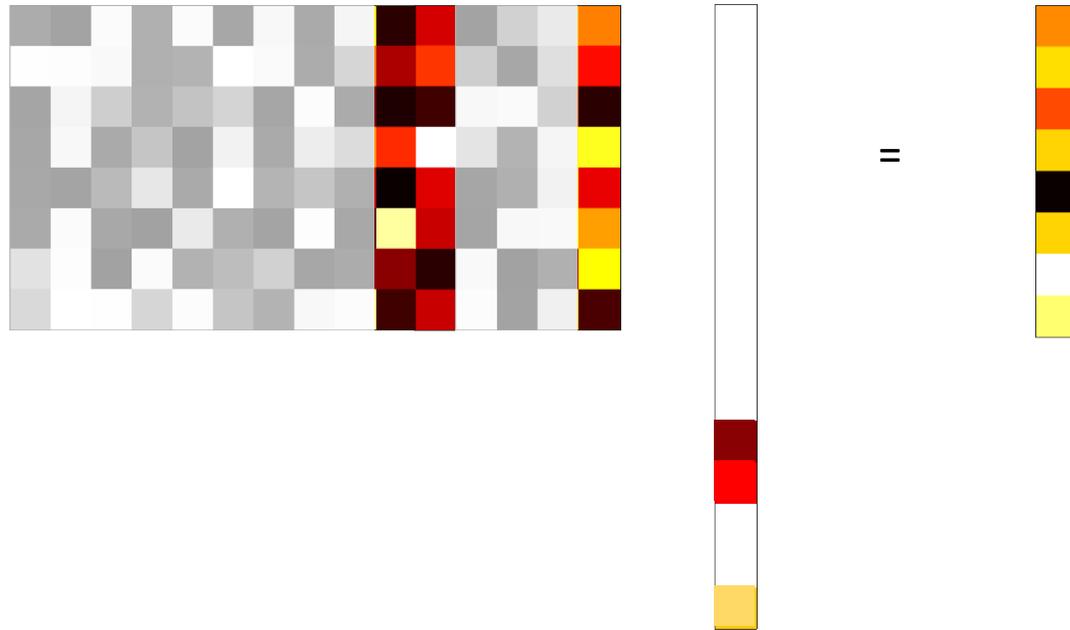
ML: dictionary learning is all about finding a good basis transform to make sparse.



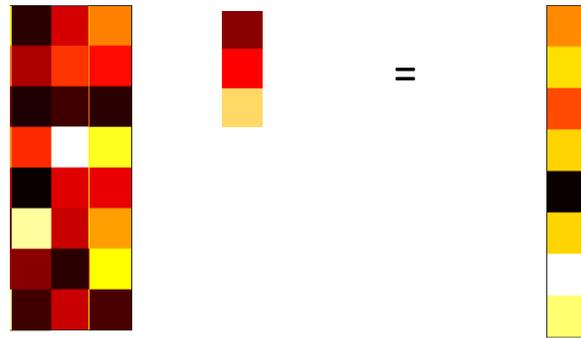
Sparse x means only a few columns of A 'matter'



Sparse x means only a few columns of A 'matter'

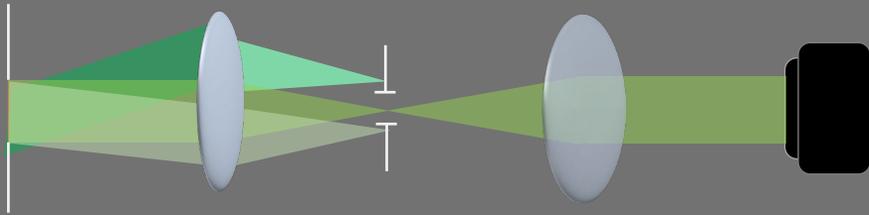


If we knew which elements were non-zero, we could solve a small least squares problem:



Computational Imaging

Hardware Design

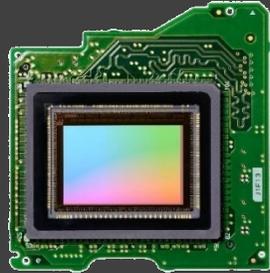
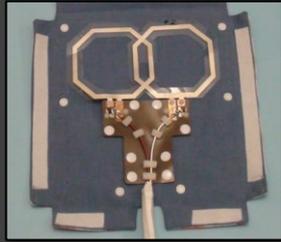


Computation Design

find \mathbf{x}
such that $\mathbf{y} = |\mathbf{Ax}|^2$

The hard part is *integration*

Hardware Toolbox



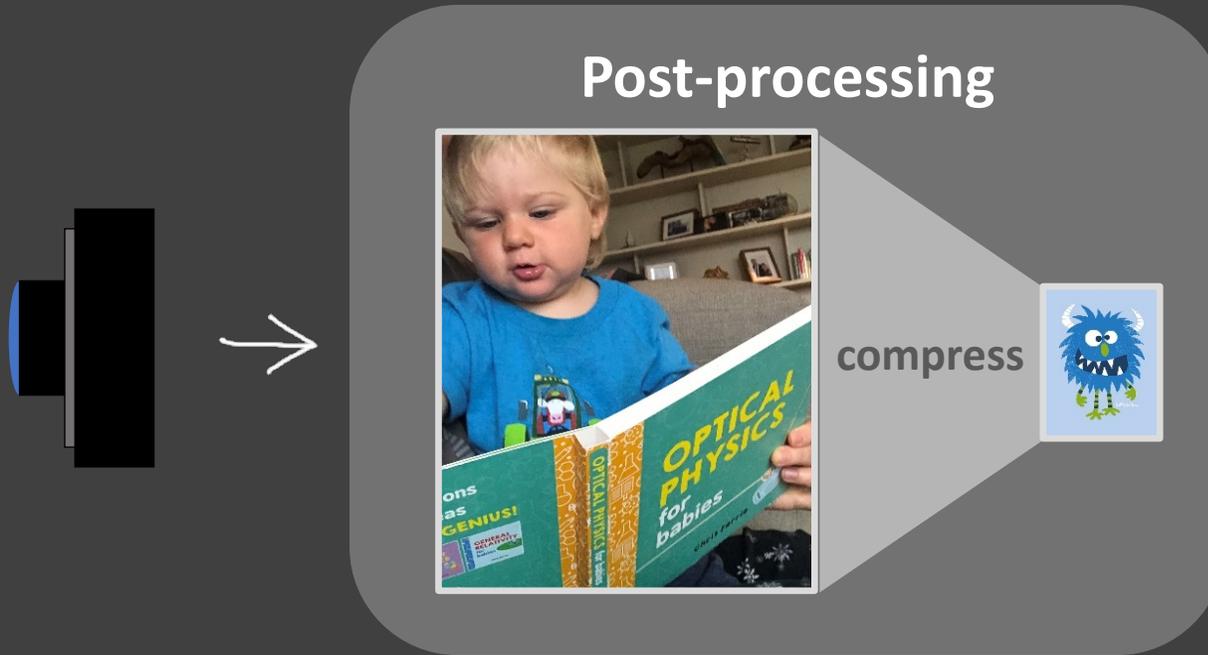
Computational Toolbox



A new way to approach imaging system design

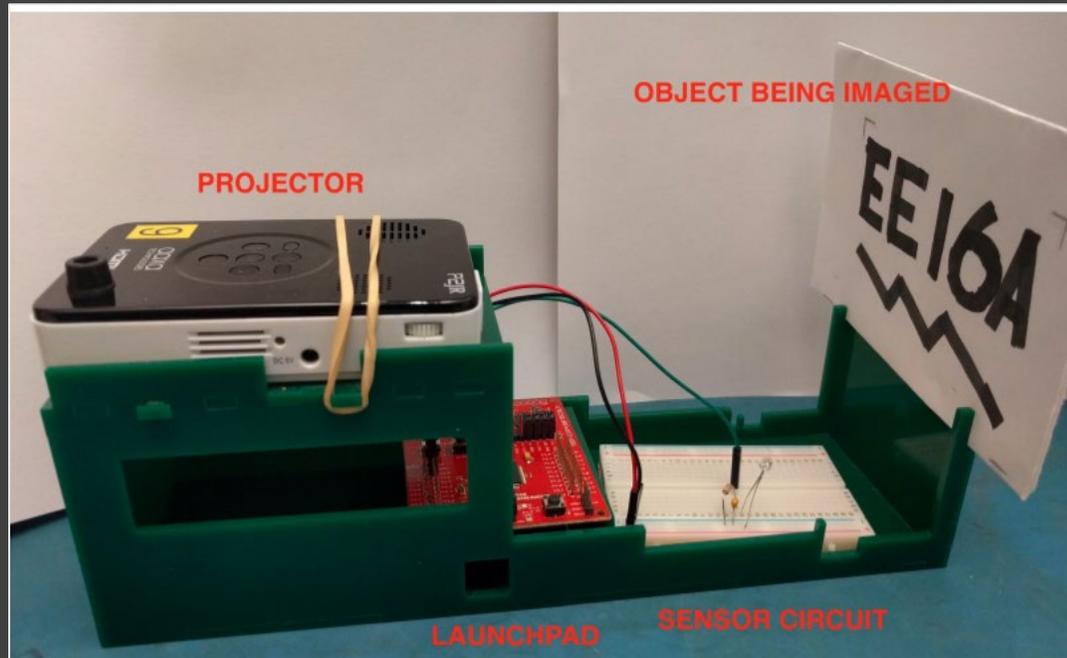
- leverage commodity hardware
- pave the way for design of new devices

Compressible images can be represented with less data

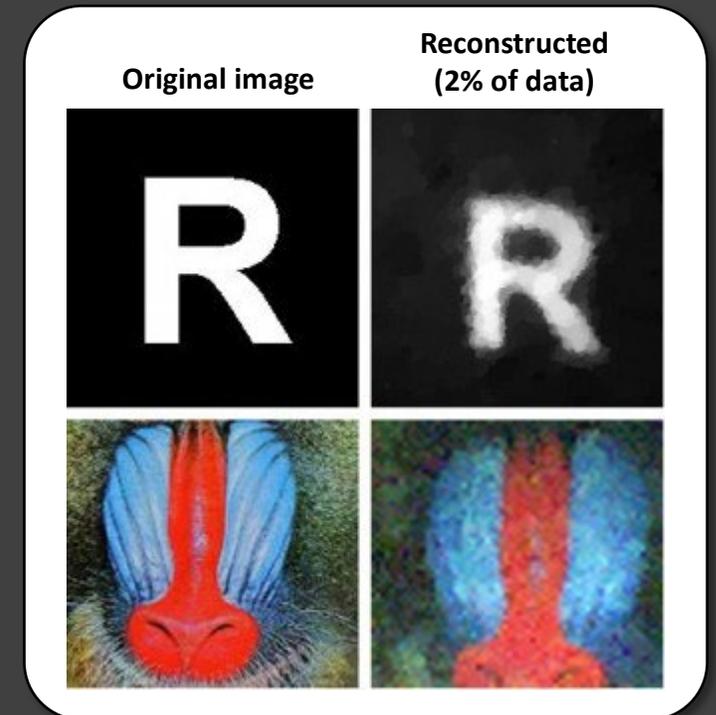


Can we compress data at the capture stage?

Yes! With compressed sensing! Example: single-pixel camera

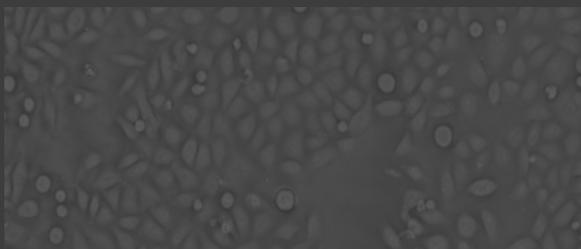


If you design the patterns on your imaging lab well, and images are compressible, you could solve with very little data!

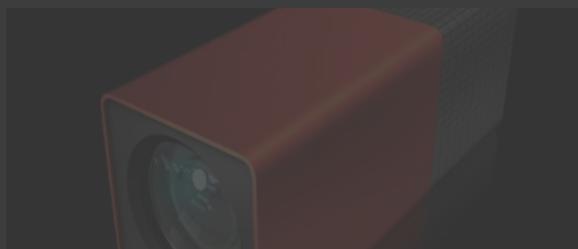


Berkeley Center for Computational Imaging

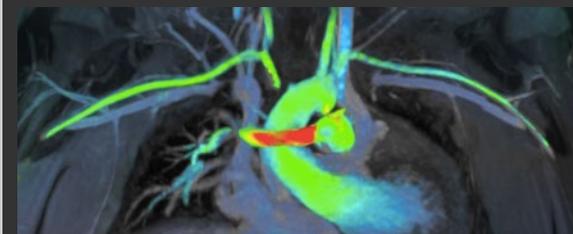
Computational Microscopy
Laura Waller



Computational Photography
Ren Ng



Computational MRI
Michael Lustig

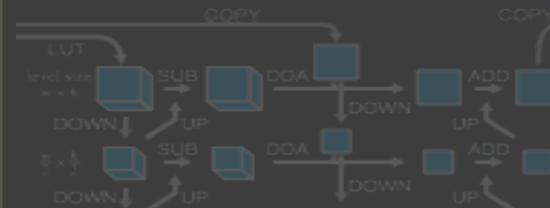


Algorithms & Optimization
Ben Recht



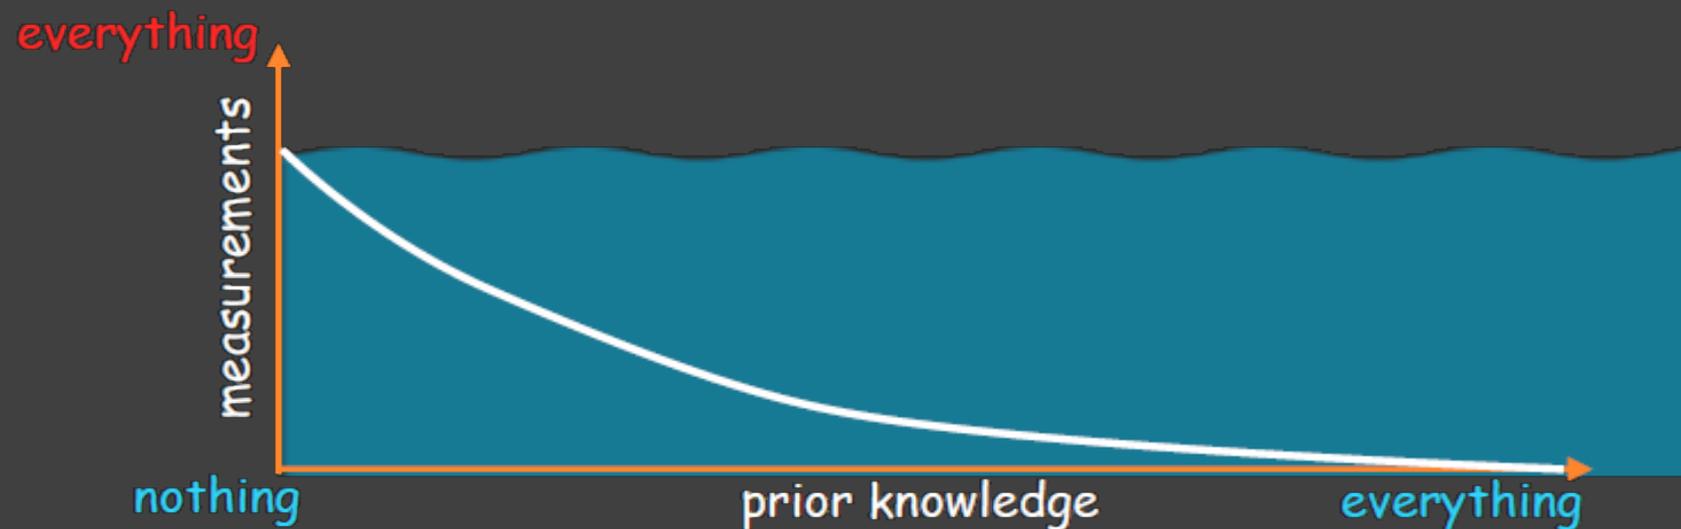
$$x(s) = \text{PSF} \circledast \left\{ \sum_{j=1}^k c_j \delta(s - s_j) \right\}$$

High Performance Visual Computing
Jonathan Ragan-Kelley

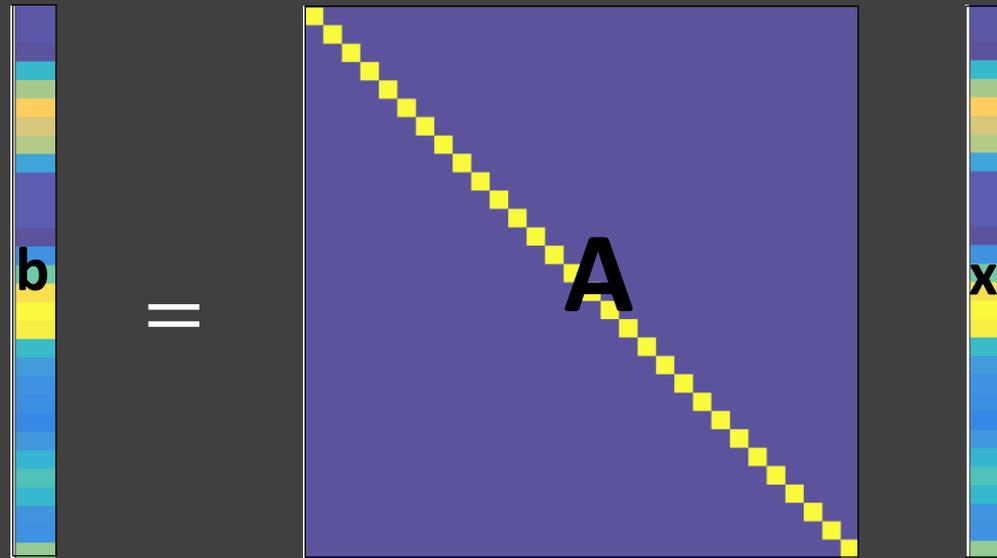


Compressed sensing is all about using prior knowledge

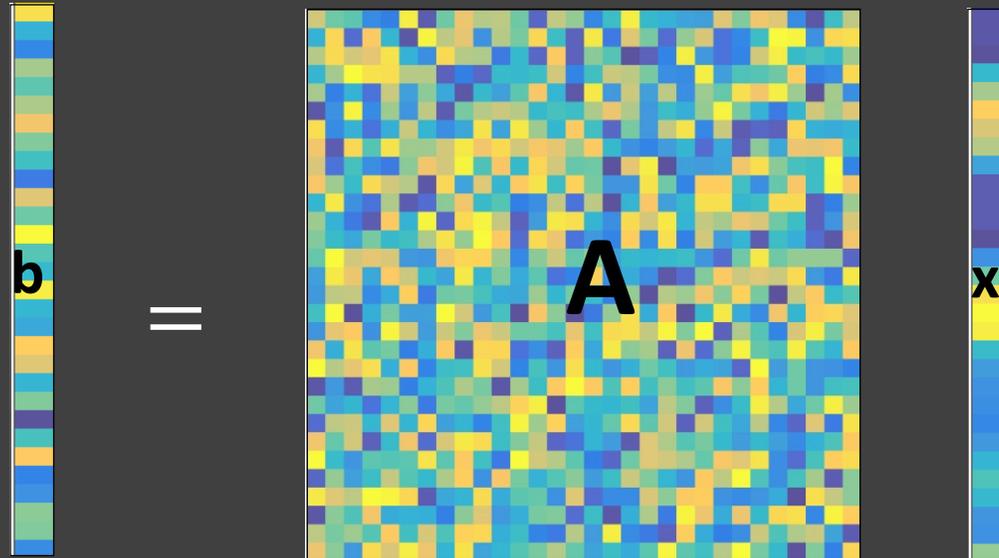
- Redundancy reduces sampling requirements
(The more you know, the less you need)



Traditional sensing: direct measurements



Multiplexed measurements

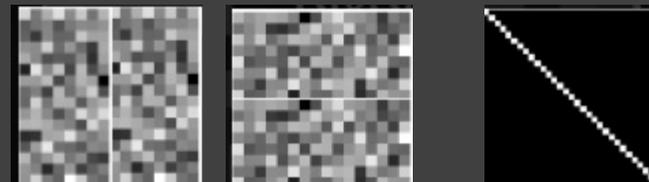


The diagram illustrates the equation $b = Ax$. On the left is a vertical column vector labeled b , containing 16 colored segments. In the center is a square matrix labeled A , which is a 16x16 grid of colored pixels. On the right is another vertical column vector labeled x , also containing 16 colored segments. An equals sign is placed between b and A , and between A and x .

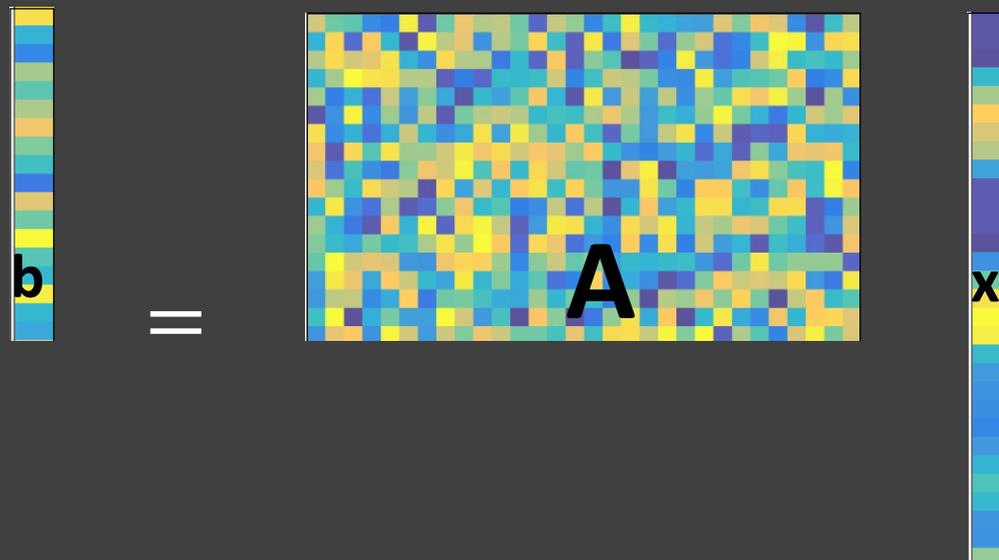
What makes a good sensing matrix?

A is orthogonal

$$A^T \quad A \quad = \quad I$$



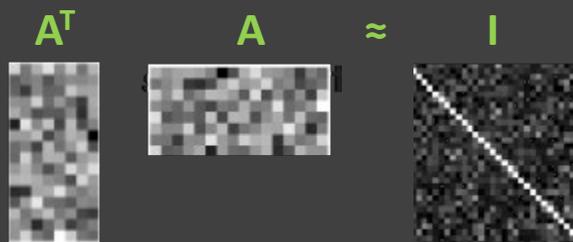
Compressed sensing solves underdetermined problems using sparsity



The diagram illustrates the equation $b = Ax$. On the left is a vertical vector b with 15 colored segments. In the center is a square matrix A with a complex, noisy pattern of colors. On the right is another vertical vector x with 15 colored segments, where only one segment is bright yellow, representing a sparse vector.

What makes a good sensing matrix?

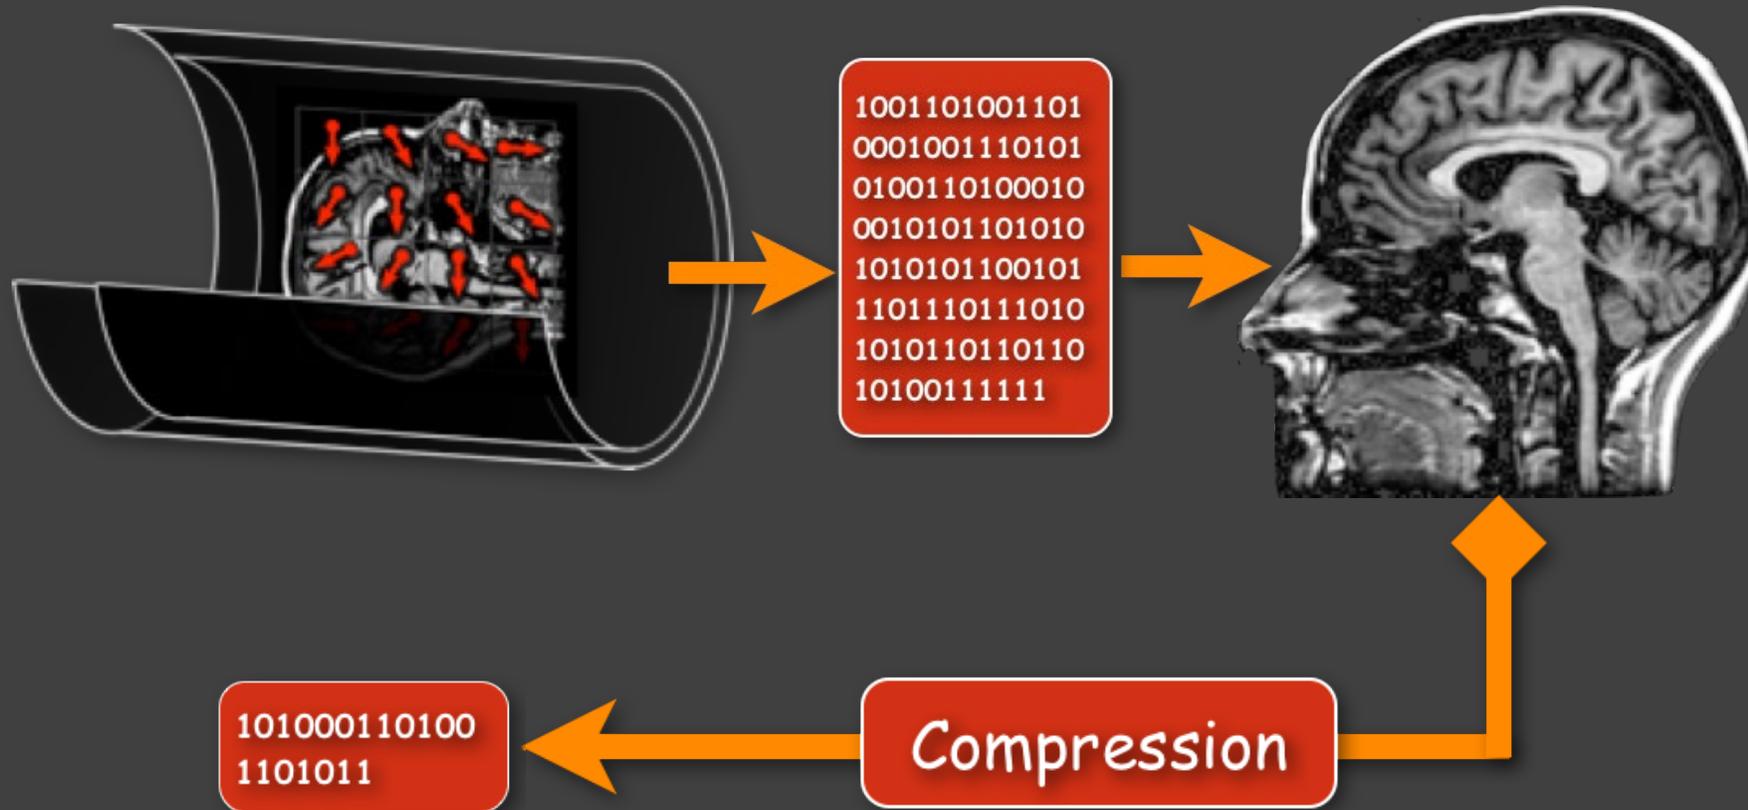
A is (almost) orthogonal



The diagram shows three square matrices: A^T , A , and I . A^T and A are grayscale images with a noisy, random appearance. I is a grayscale image with a clear diagonal line from the top-left to the bottom-right. The matrices are arranged as A^T , A , and I with an approximation symbol \approx between A and I .

Compressed Sensing MRI

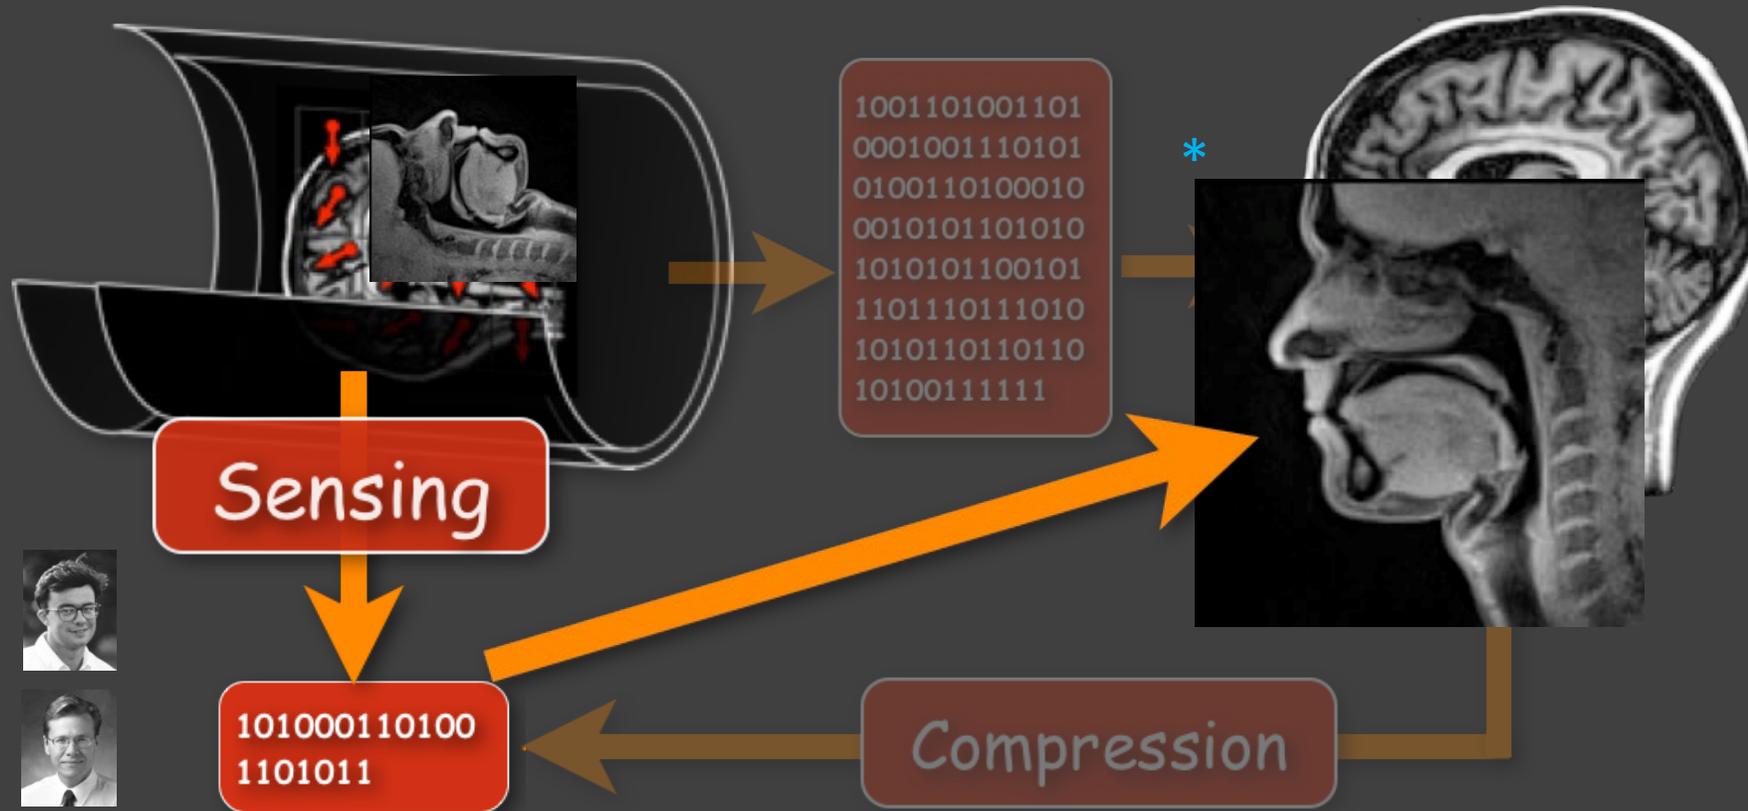
Medical images are compressible
Standard approach: First collect, then compress



Compressed Sensing MRI

Medical images are compressible

New approach: Acquire “compressed” data directly!



powered by



 Lucile Packard
Children's Hospital
Stanford



Prof. Shreyas Vasanawala

12 month male (114 bpm)

Scan time: 11:05 min; Total acceleration

Resolution: 0.9 x 0.9 x 1.4 mm³

VENC: (150, 150, 150) cm/s

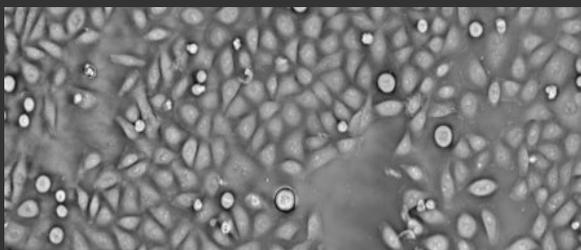
Contrast: ferumoxytol

Generated using Arterys (San Francisco, CA)



Berkeley Center for Computational Imaging

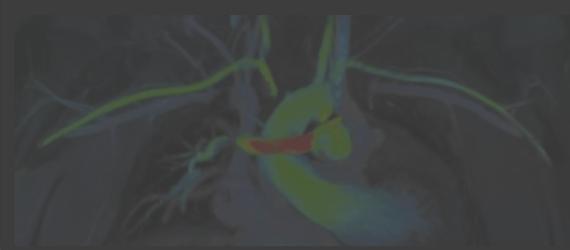
Computational Microscopy
Laura Waller



Computational Photography
Ren Ng



Computational MRI
Michael Lustig

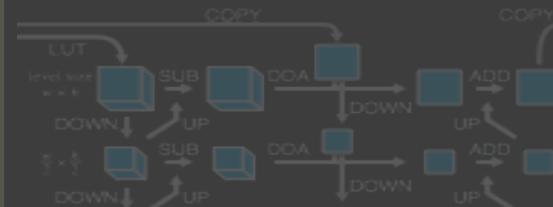


Algorithms & Optimization
Ben Recht



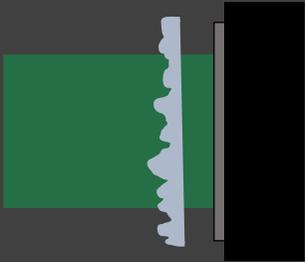
$$x(s) = \text{PSF} \circledast \left\{ \sum_{j=1}^k c_j \delta(s - s_j) \right\}$$

High Performance Visual Computing
Jonathan Ragan-Kelley

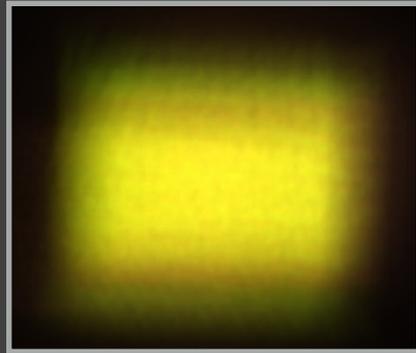


Computational imaging pipeline

Hardware design



Take picture



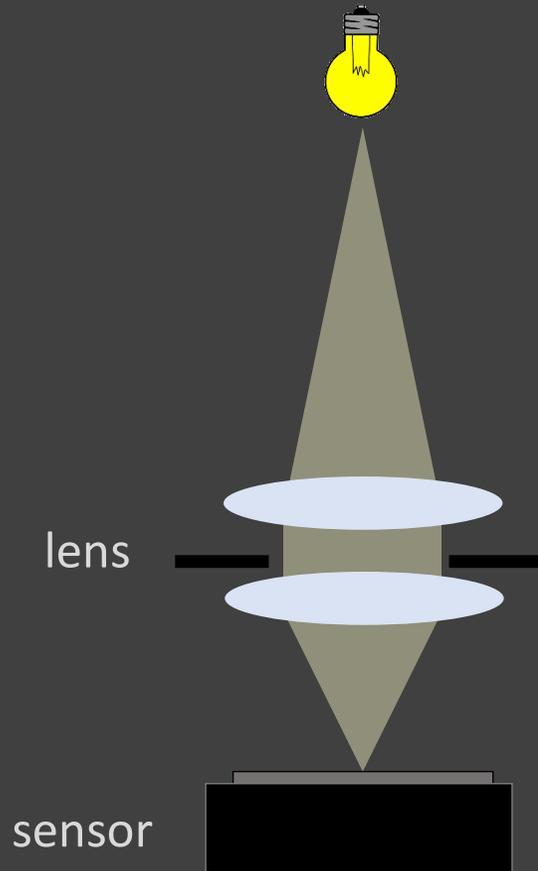
Crunch Data



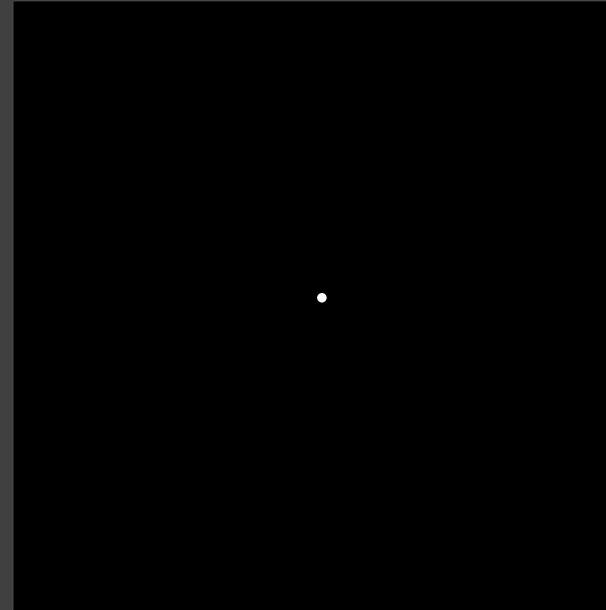
Final result



Lenses map points to points



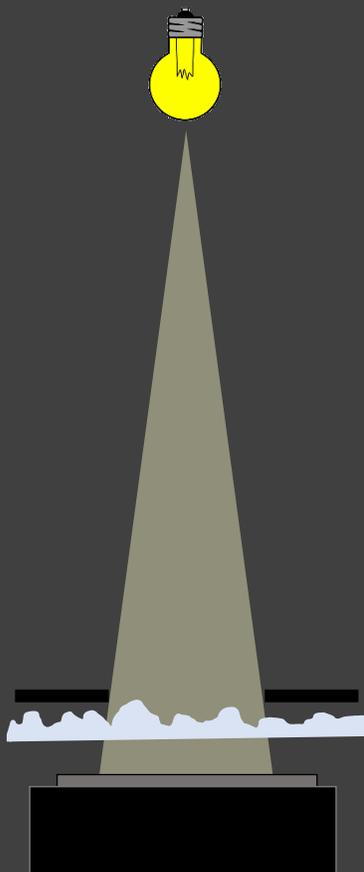
System response to point source



DiffuserCam: stick a scatterer on a sensor

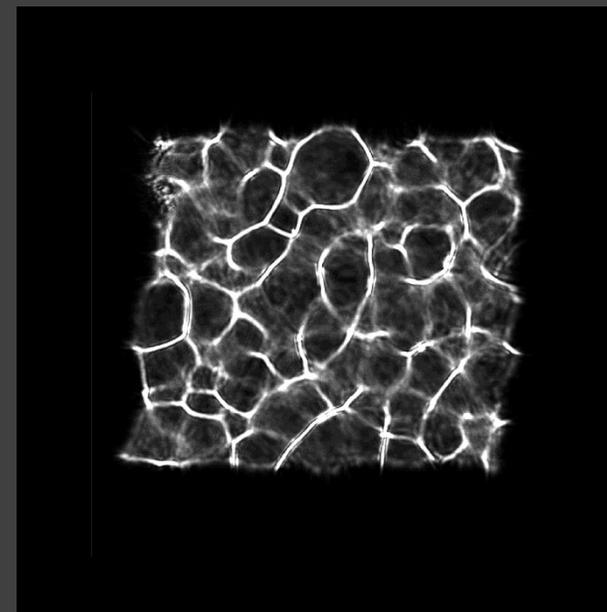


diffuser



sensor

System response to point source

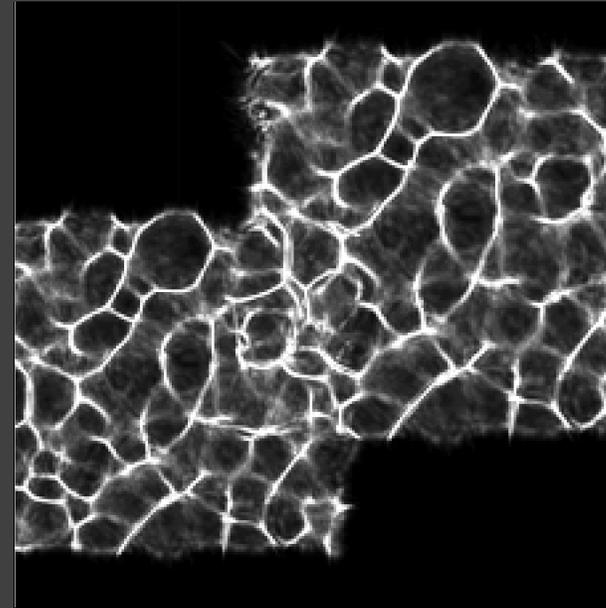
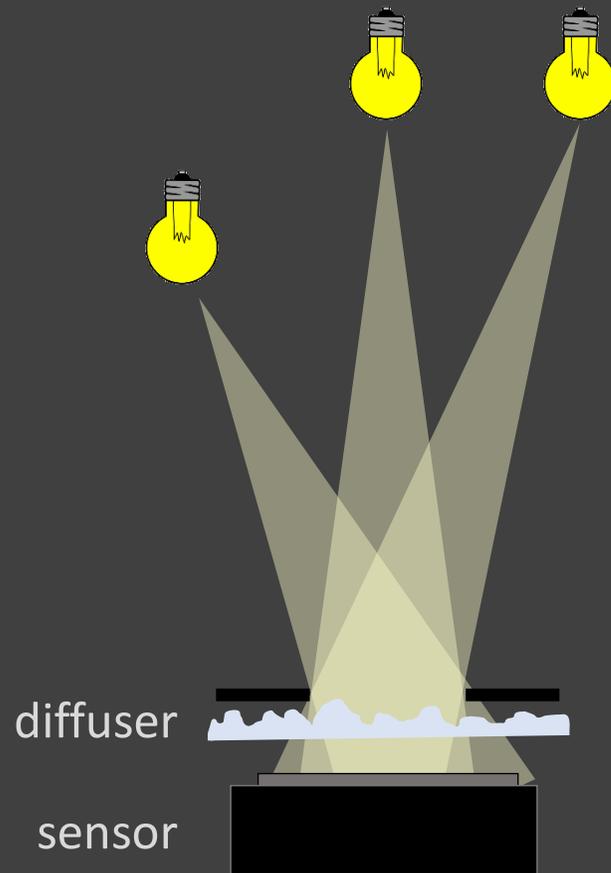


<https://laurawaller.com/opensource>

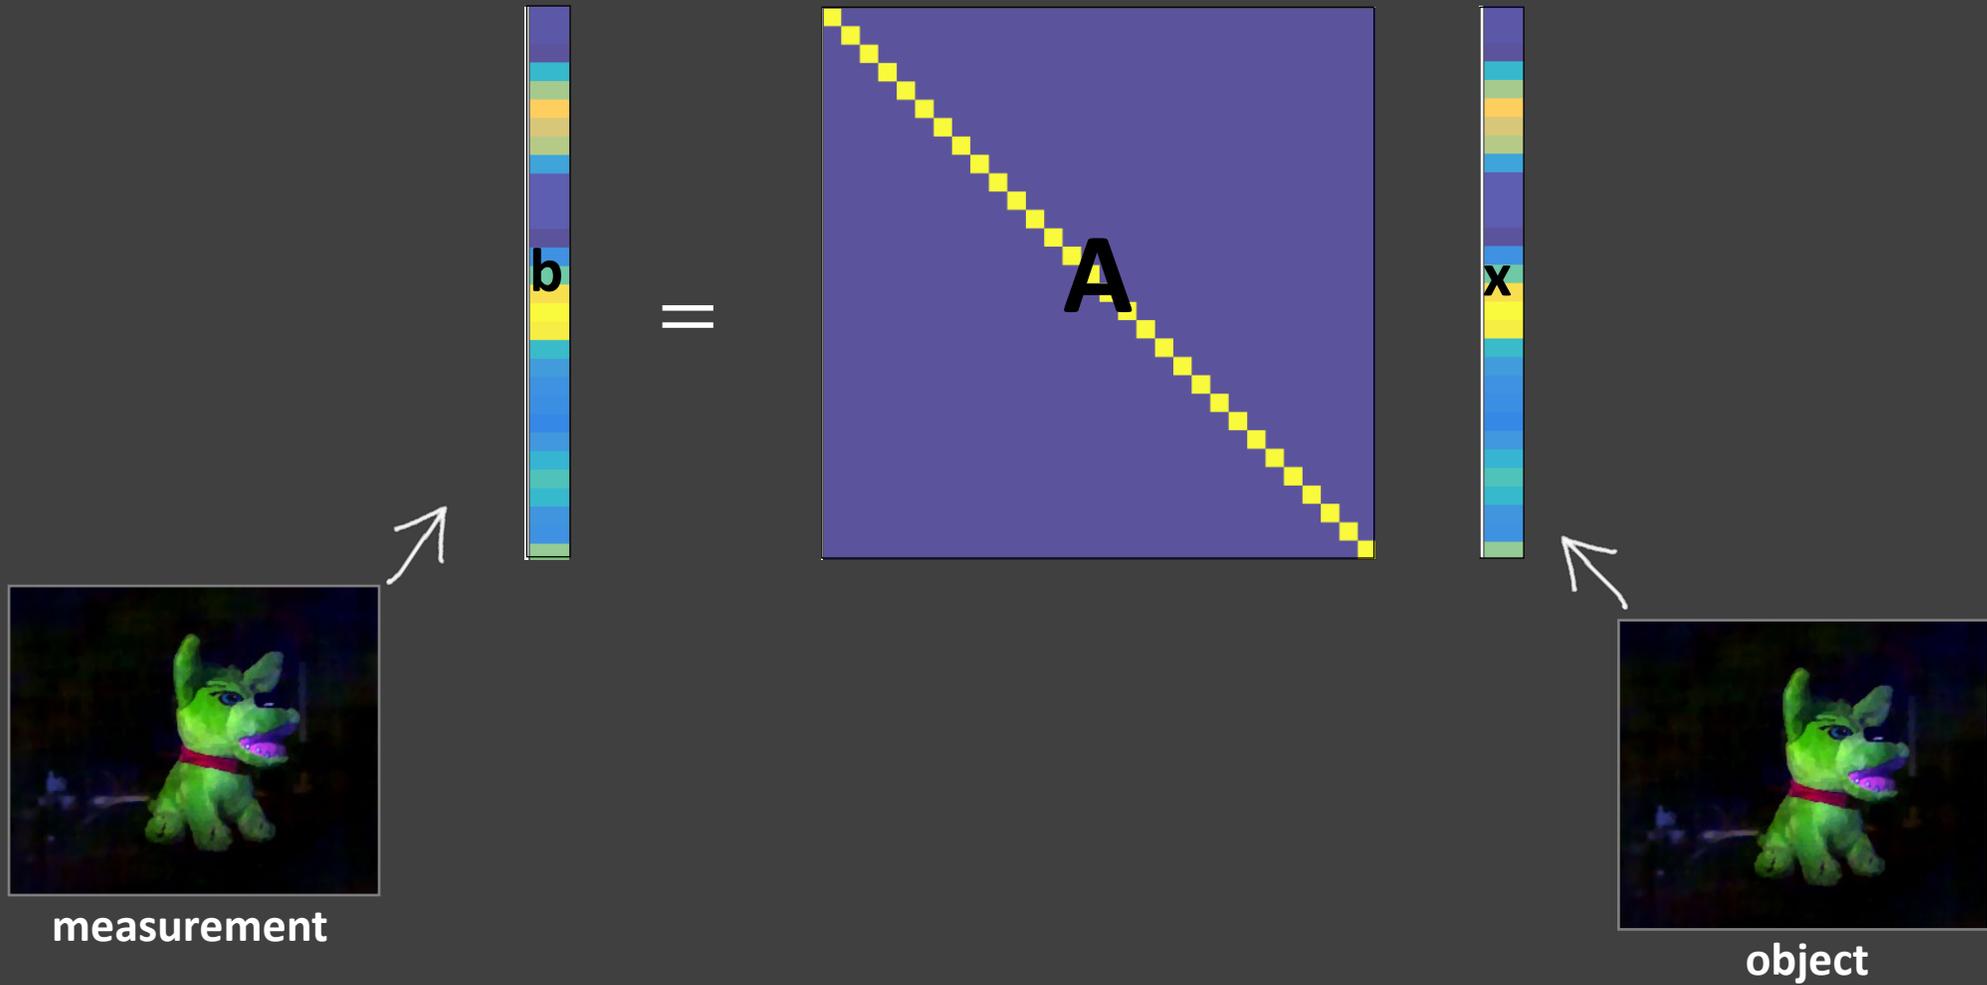
Camille Biscarrat
Shreyas Parthasarathy



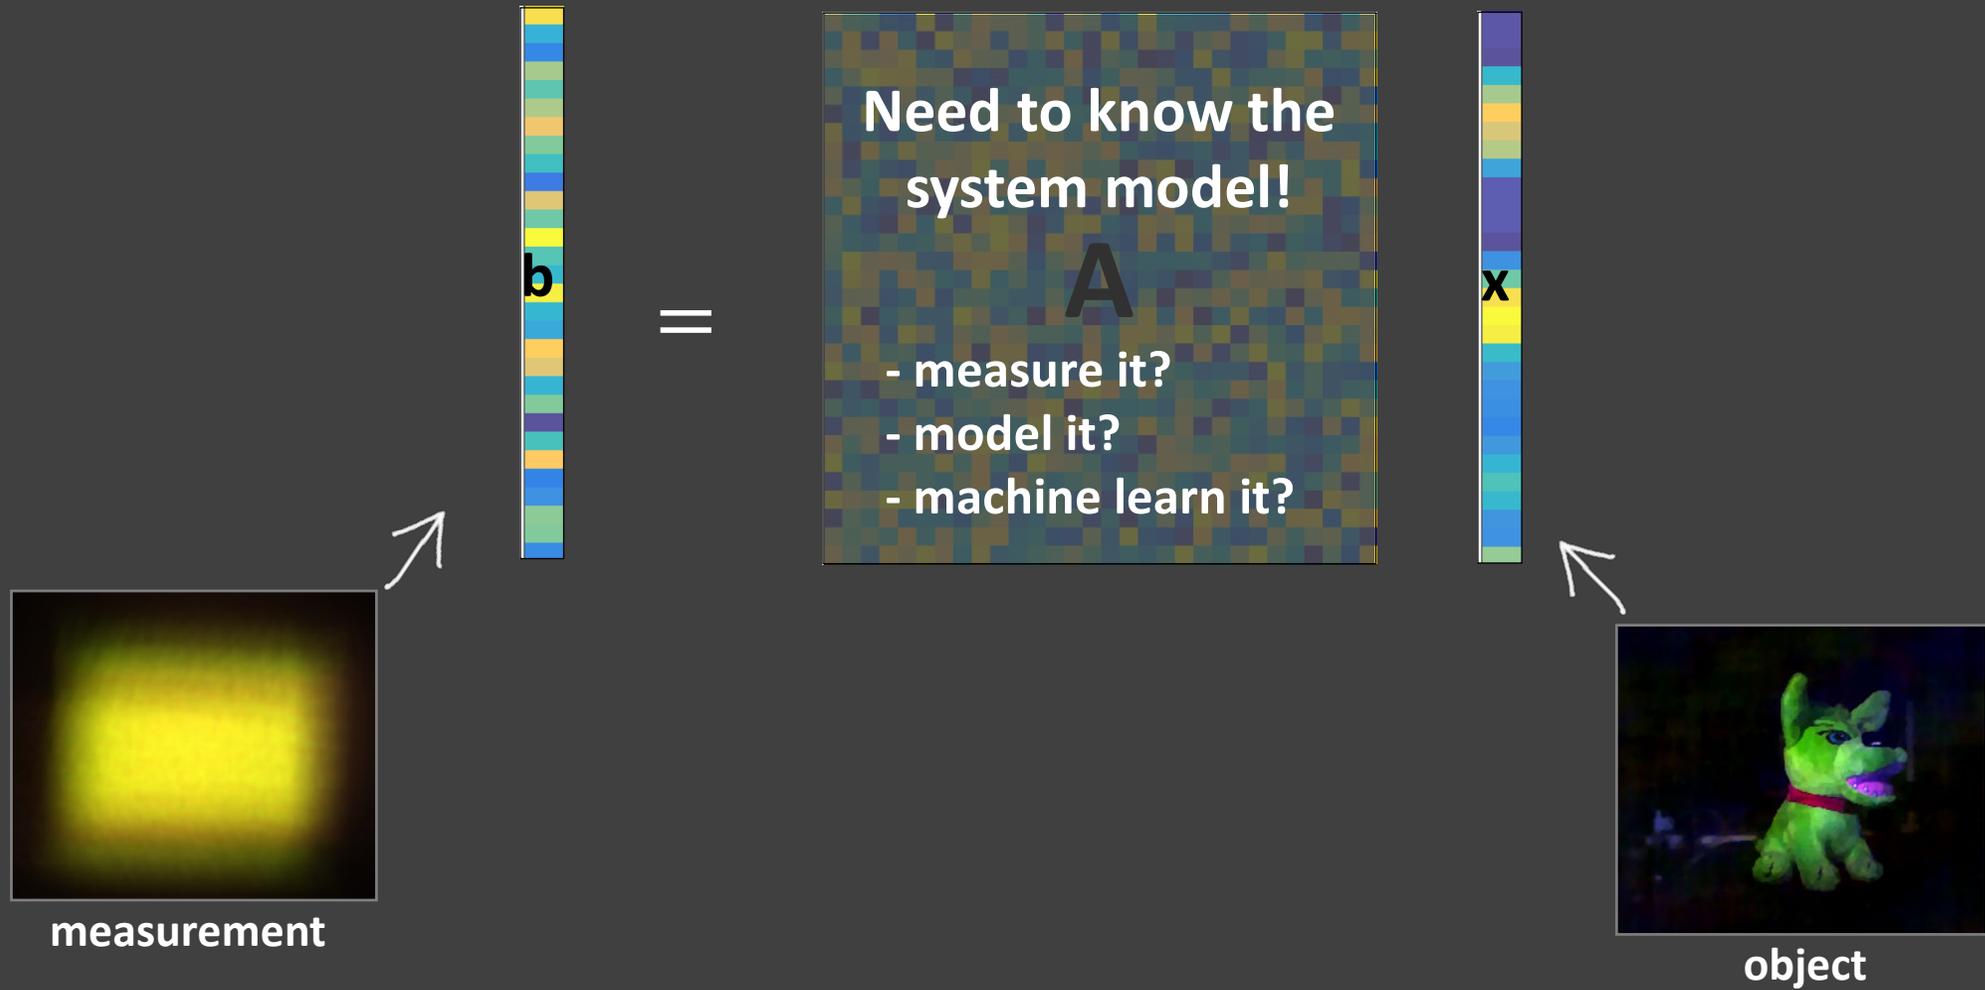
DiffuserCam: stick a scatterer on a sensor



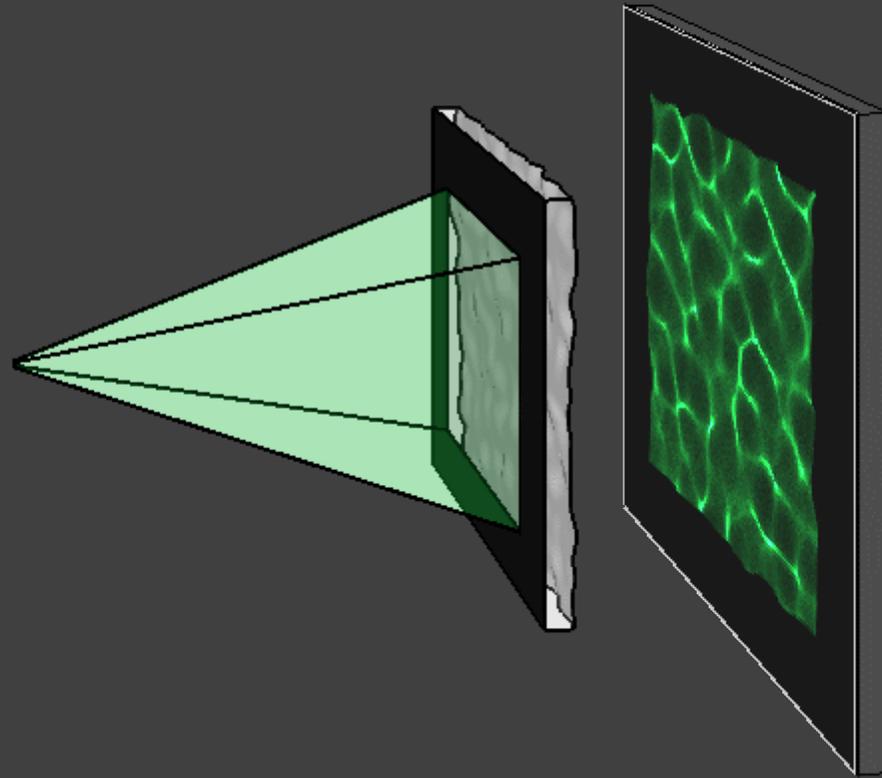
Traditional cameras take direct measurements



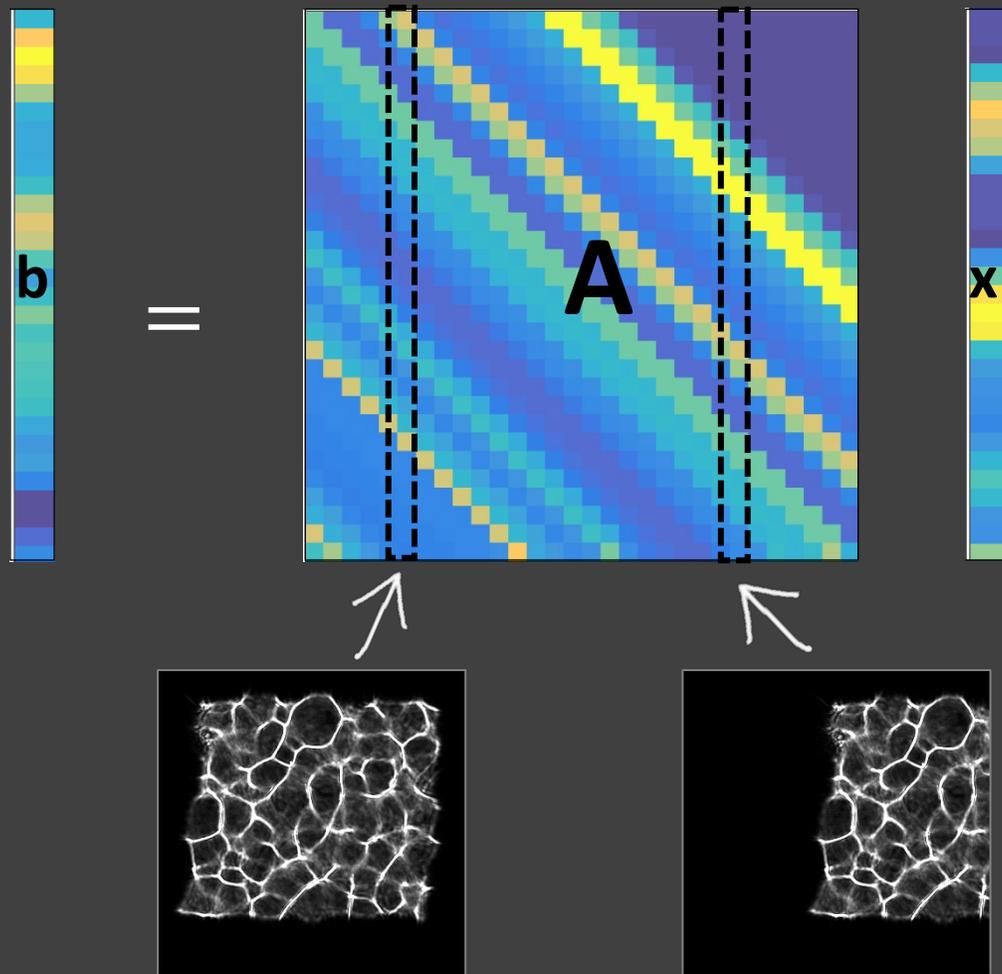
Computational cameras can multiplex



System response shifts with position

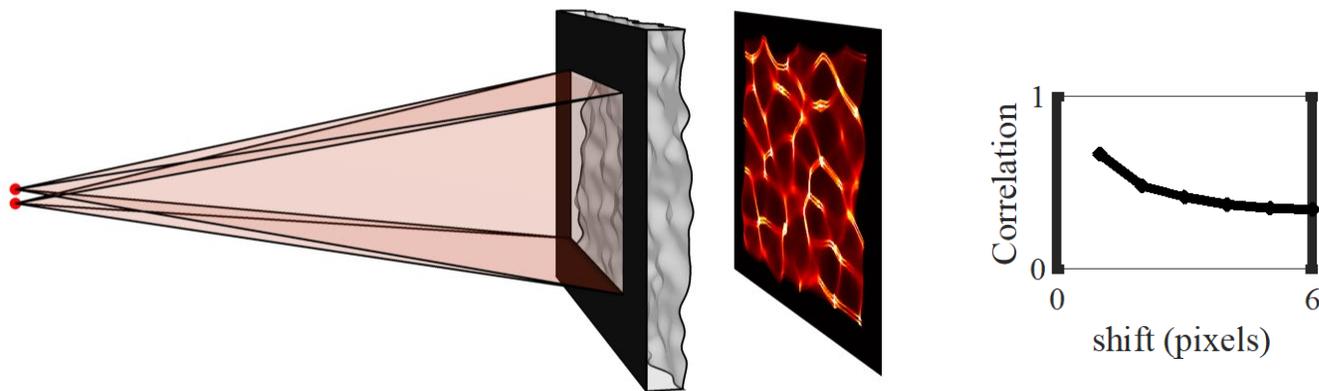


DiffuserCam system model is a shift-invariant

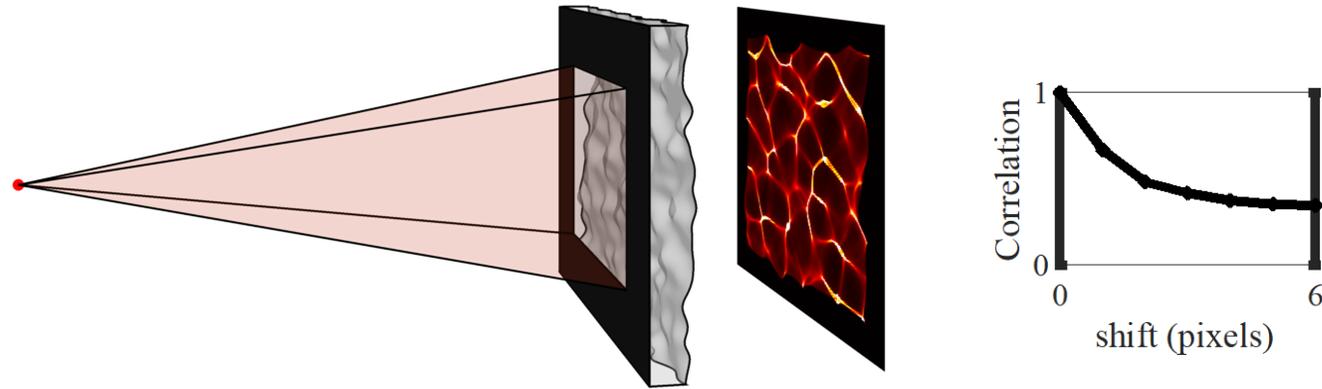


System response is same but shifted
for different image pixels

We could find location of a point by correlating image captured with shifts in system response!

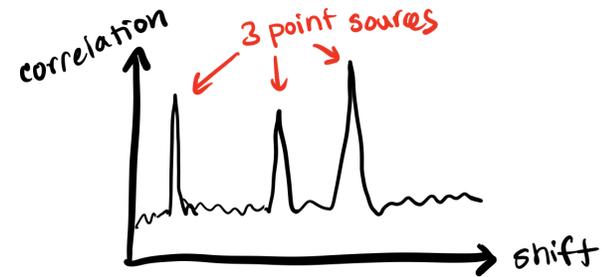
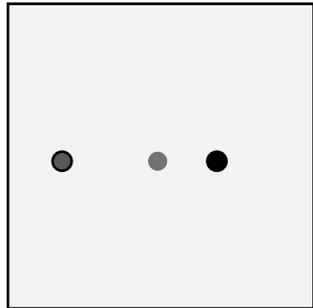
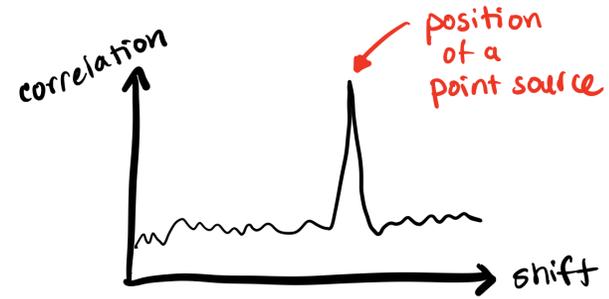
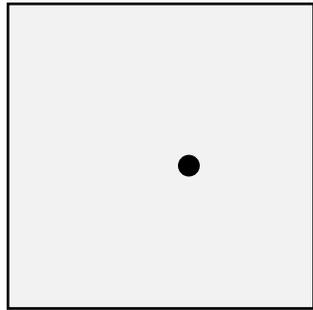


We could find location of a point by correlating image captured with shifts in system response!



Reconstructing image amounts to finding strength of each 'point source':

Looks a lot like our GPS problem! (especially if image is sparse)





raw sensor data



recovered scene

*solver is ADMM with TV reg in Halide

Grace Kuo
Nick Antipa





raw sensor data



recovered scene

*solver is ADMM with TV reg in Halide

Grace Kuo
Nick Antipa



Image reconstruction is nonlinear optimization

$$\arg \min_{x \geq 0} \left\| \begin{array}{c} \text{b} \\ \text{A} * \text{x} \end{array} \right\|_2^2 + \lambda \left\| \Phi \begin{array}{c} \text{x} \end{array} \right\|_1$$

↑
Sparsity basis

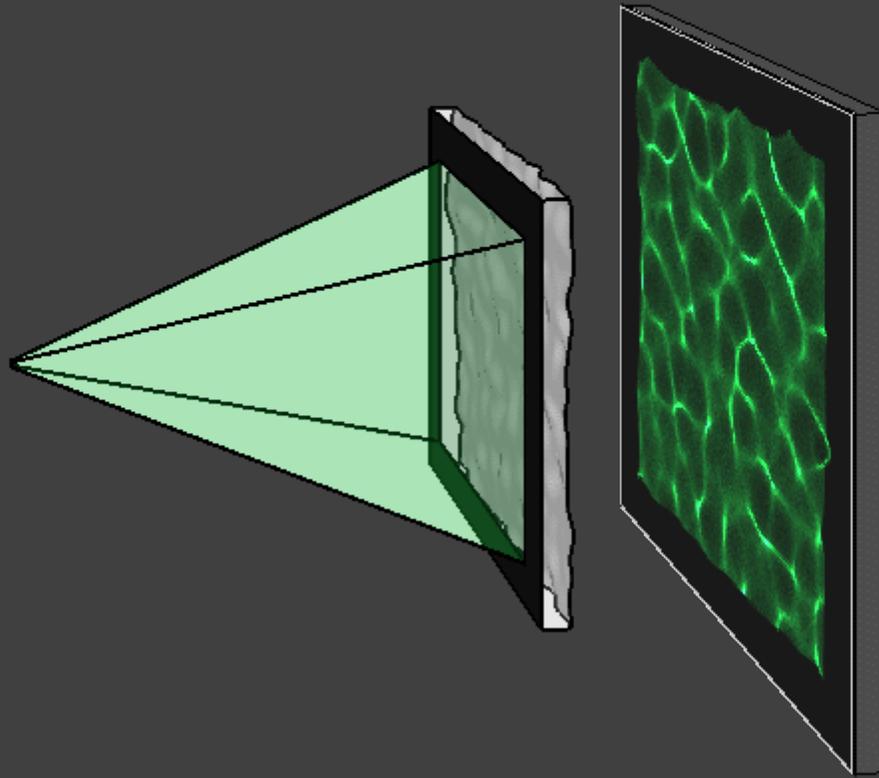
***solved with ADMM in Halide**

S. Boyd, et al. *Foundations and Trends in Machine Learning* (2011)

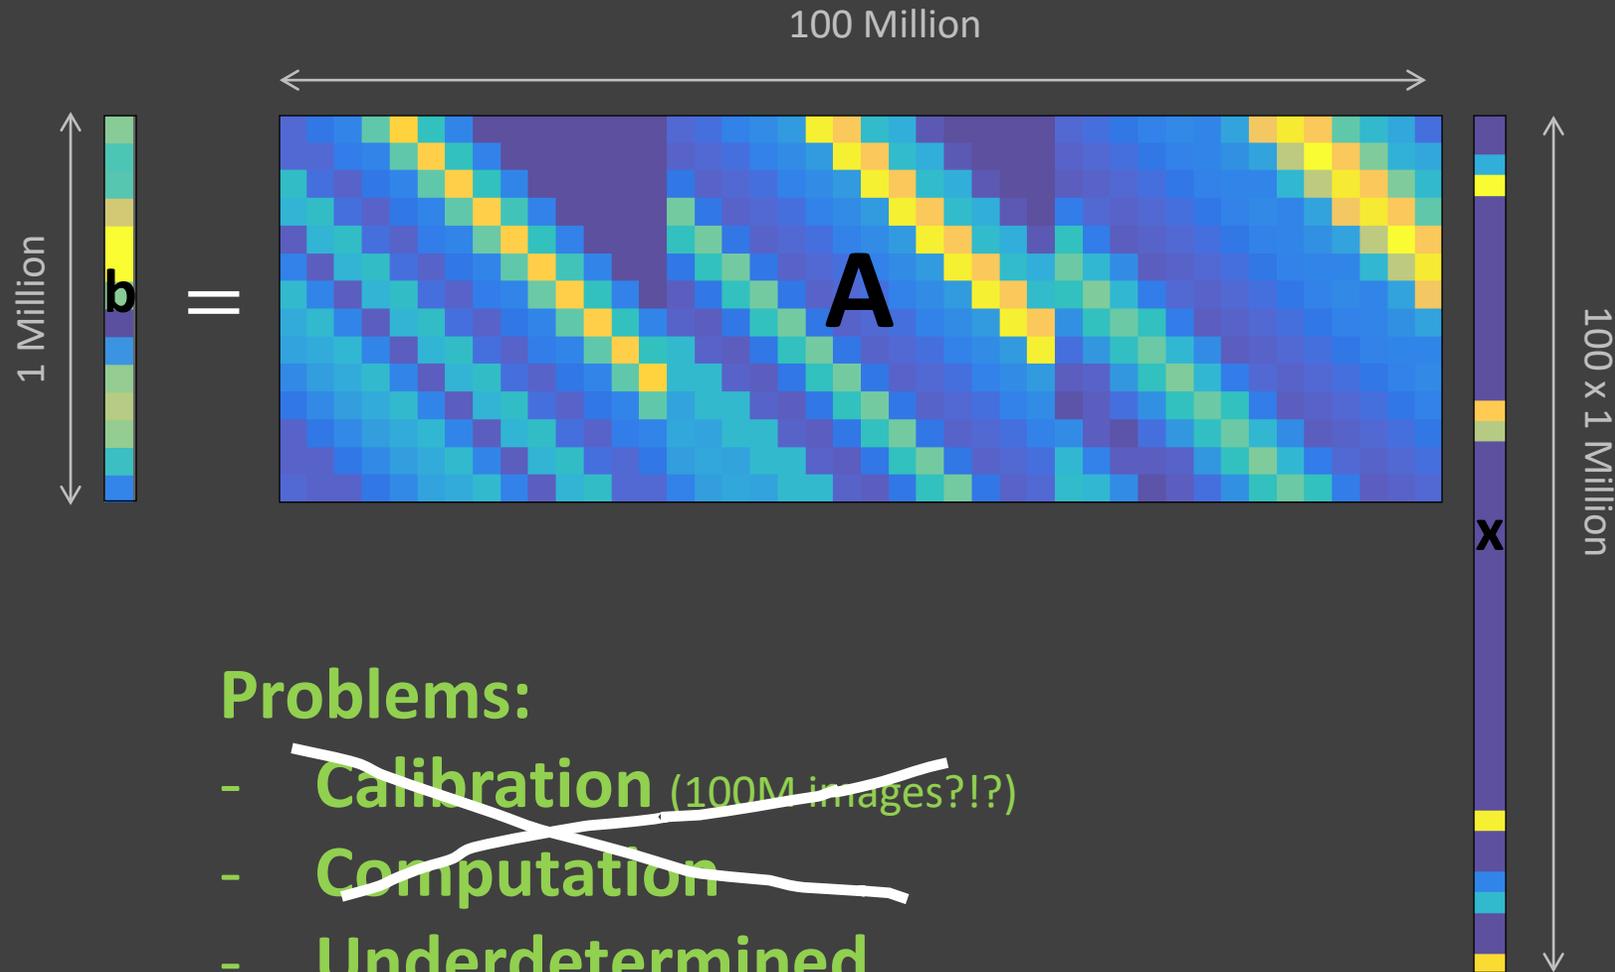
J. Ragan-Kelley, et al. *AMC SIGPLAN* (2013)

2D → 3D

Point spread function scales with depth

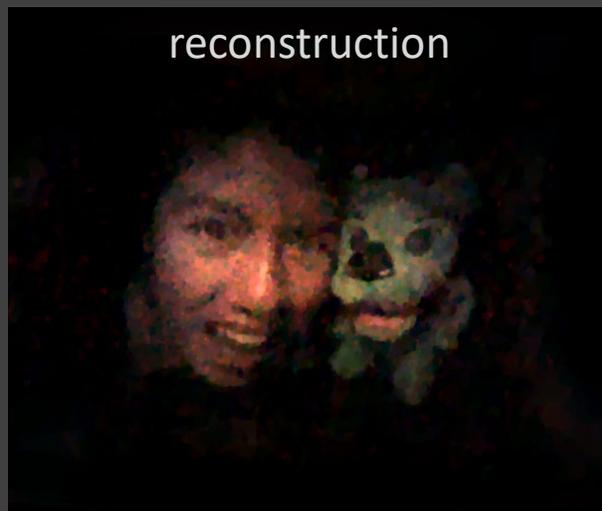
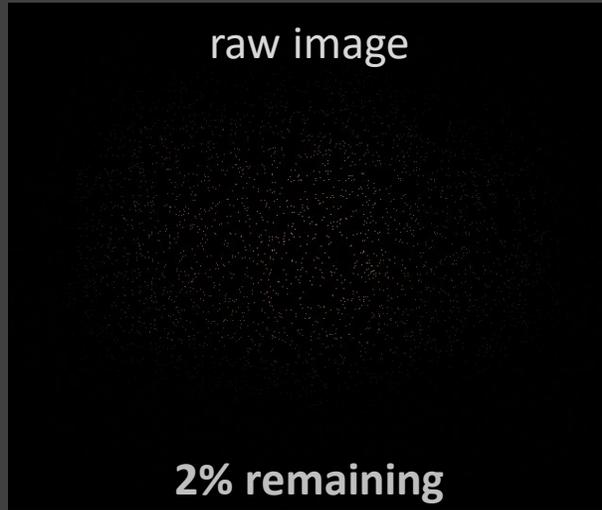


Single-shot 3D is underdetermined

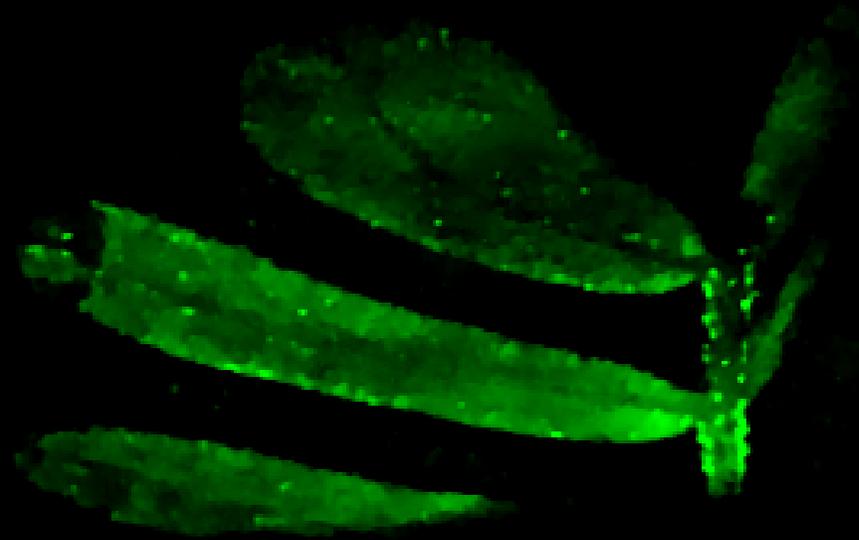
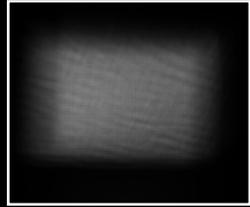


Compressed sensing to the rescue!

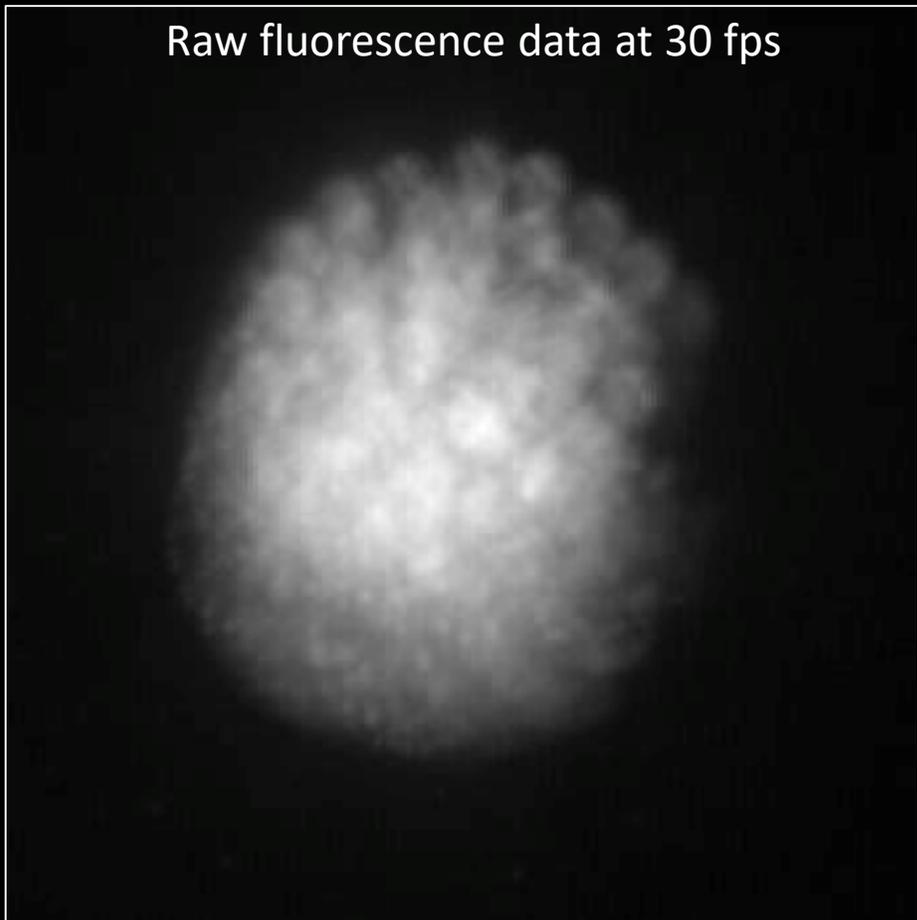
→ solves under-determined problems via a sparsity prior



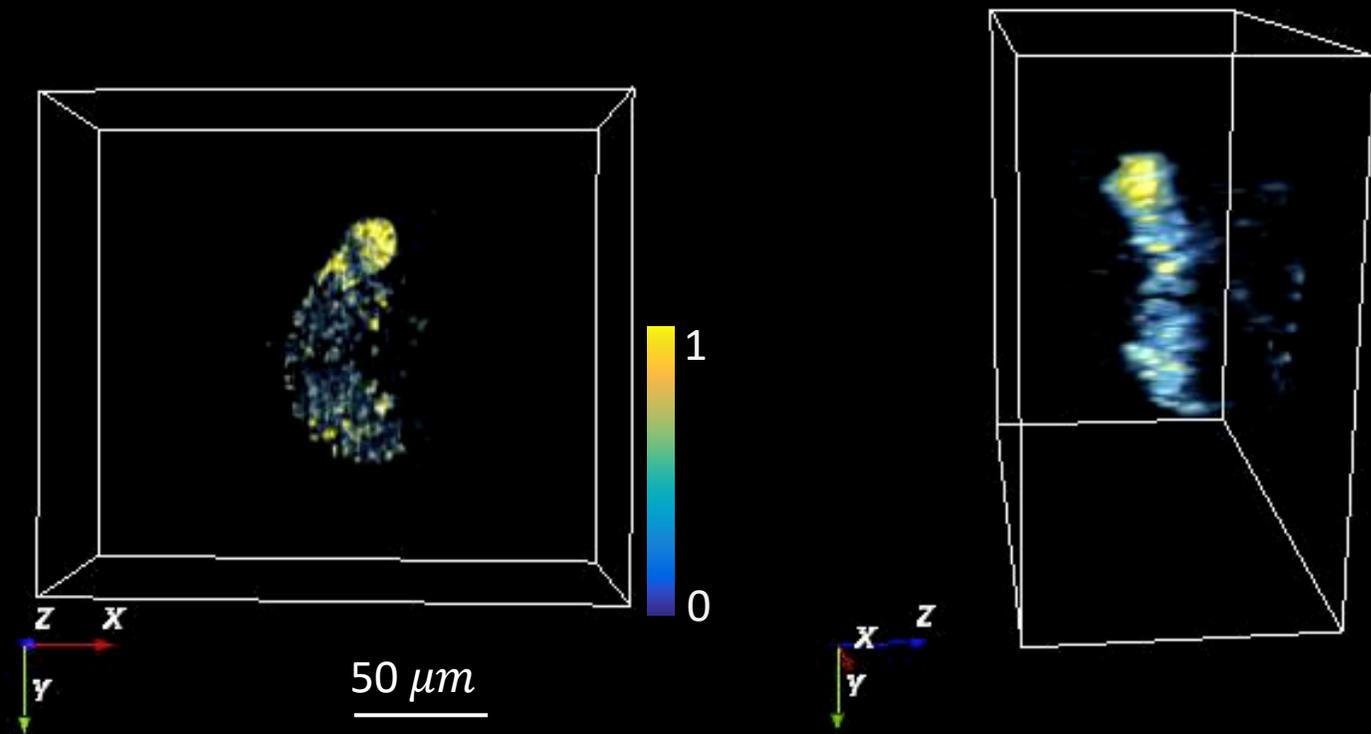
128x more voxels for **FREE!**



Raw fluorescence data at 30 fps



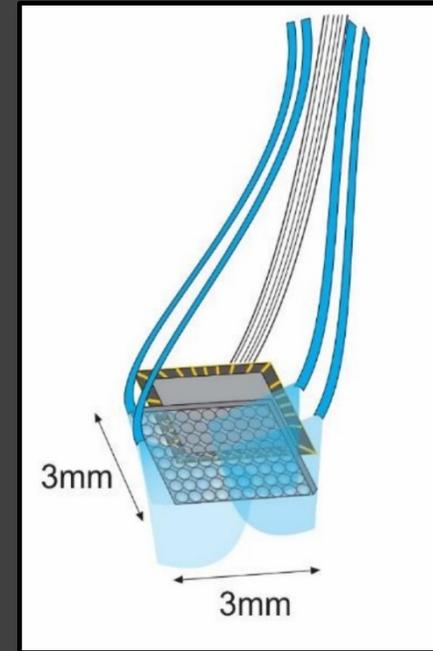
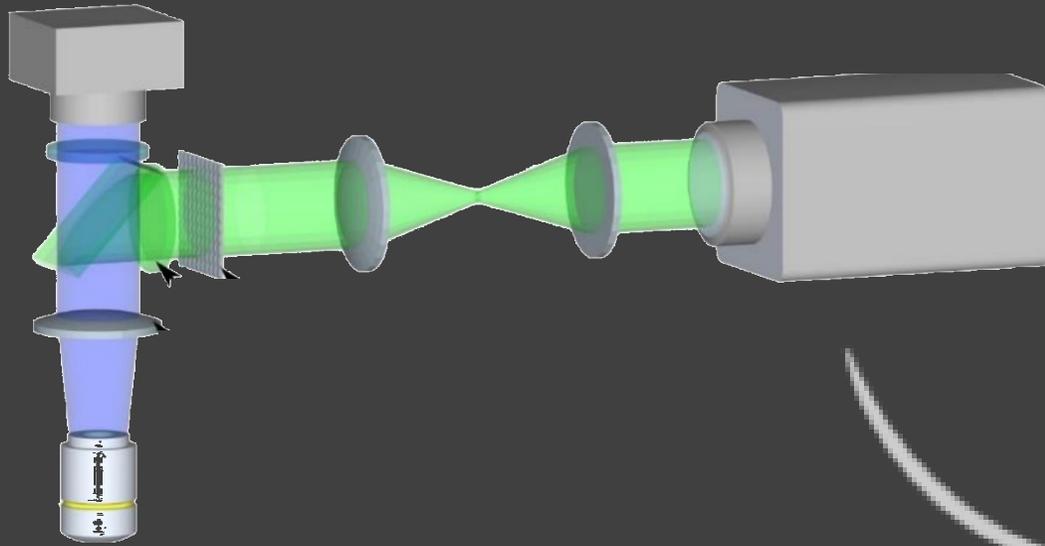
3D video reconstruction



Kyrollos Yanny
Nick Antipa



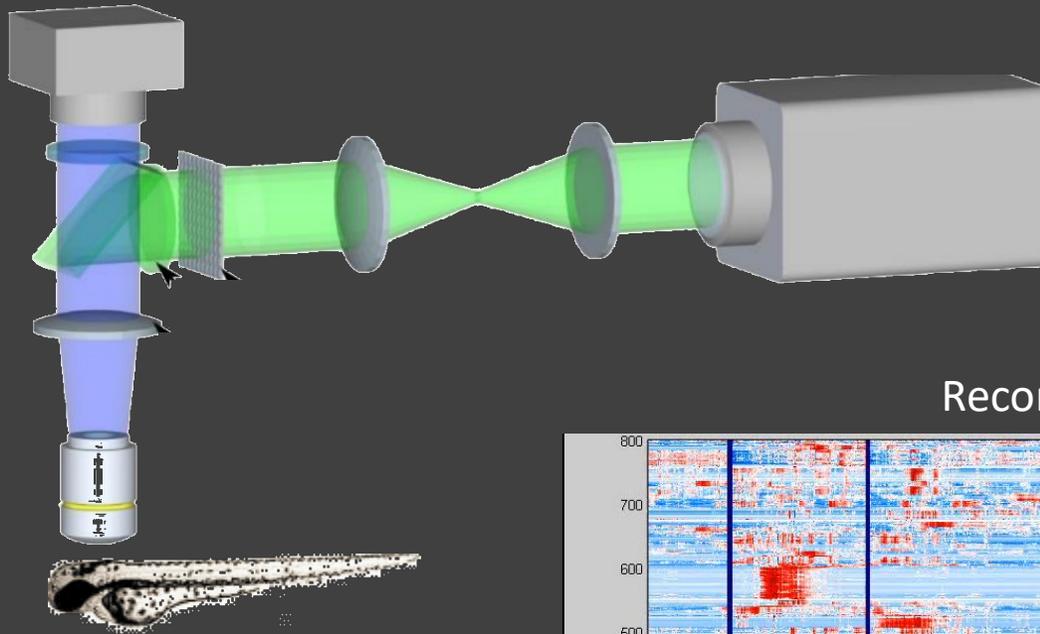
Towards lensless 3D microscopy



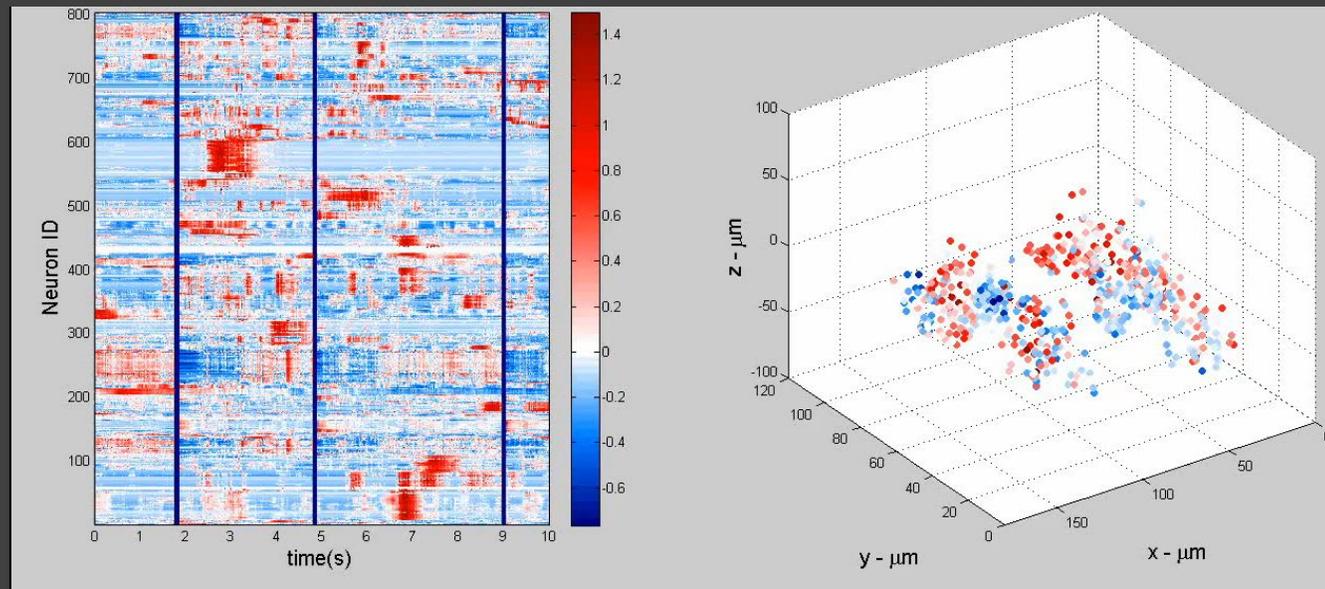
- Lensless imager:
- small
 - inexpensive
 - enables tiling



3D neural activity tracking



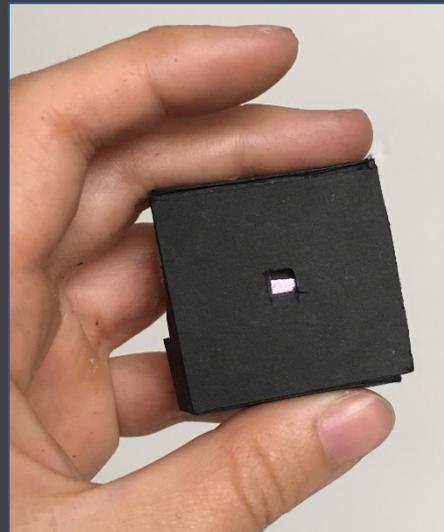
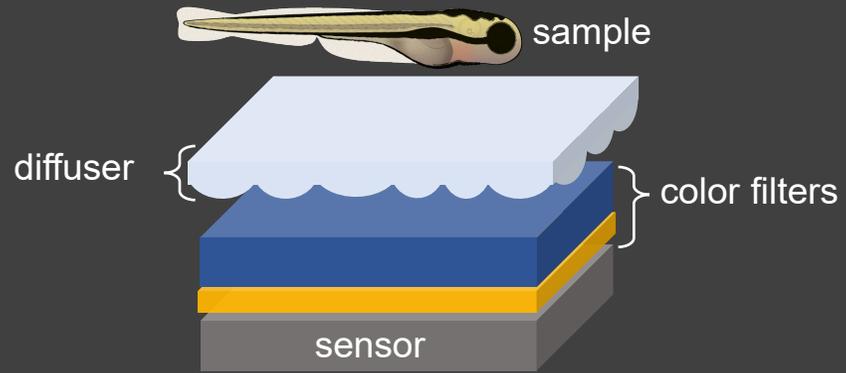
Reconstructed neural activity



Nico Pegard



Neural activity tracking with flat DiffuserScope



Grace Kuo

Outlook: Computers and optics should talk more

Hardware Toolbox



Computational Toolbox

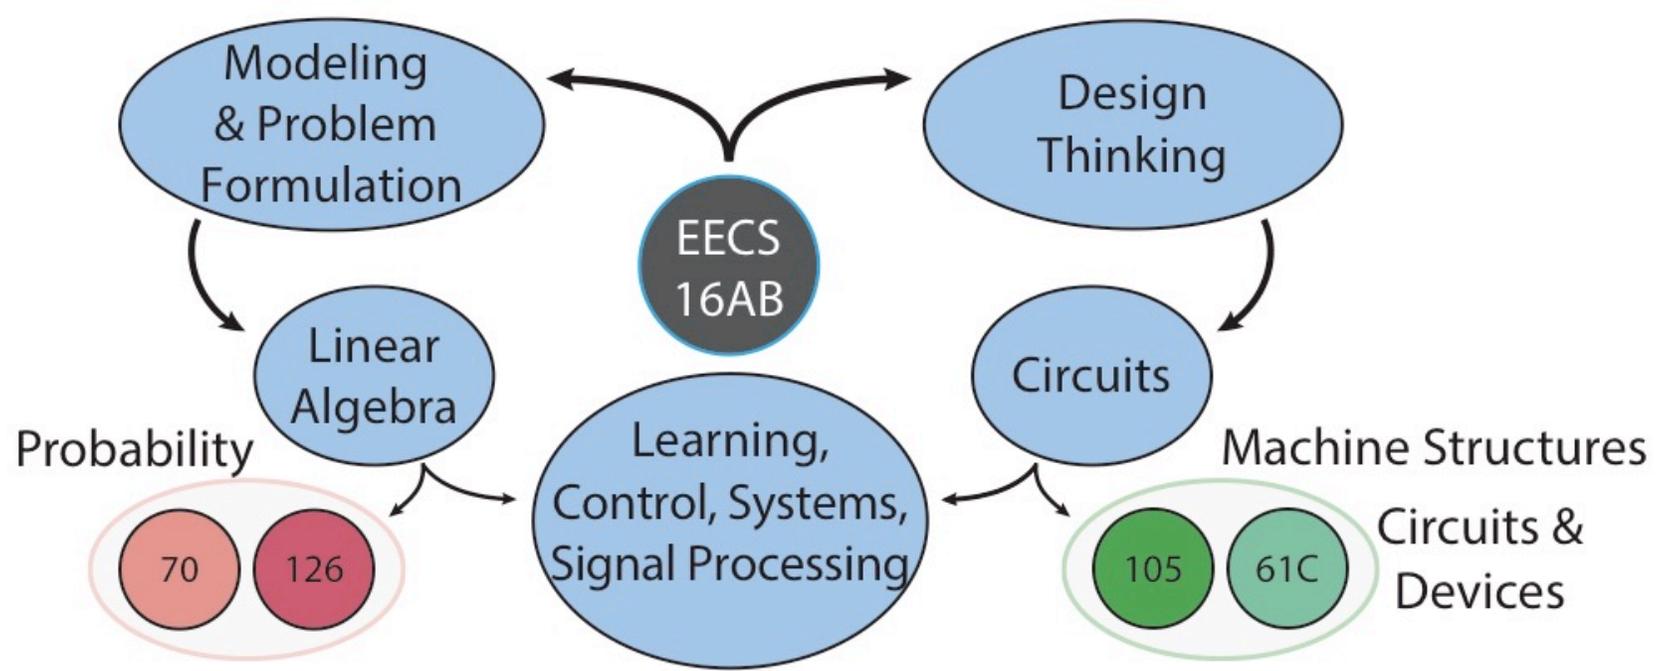


Pushing the limits of imaging can be done with existing physics and commodity hardware.

Enough about me...



- Congrats! This year has not been easy!
- What you have accomplished this semester:
 - Built a camera
 - Built two types of touchscreens
 - Built your own GPS system
- If you liked the class, please:
 - Thank your TAs!
 - apply to become one! (see piazza post, due May 2)
- This is just the beginning...

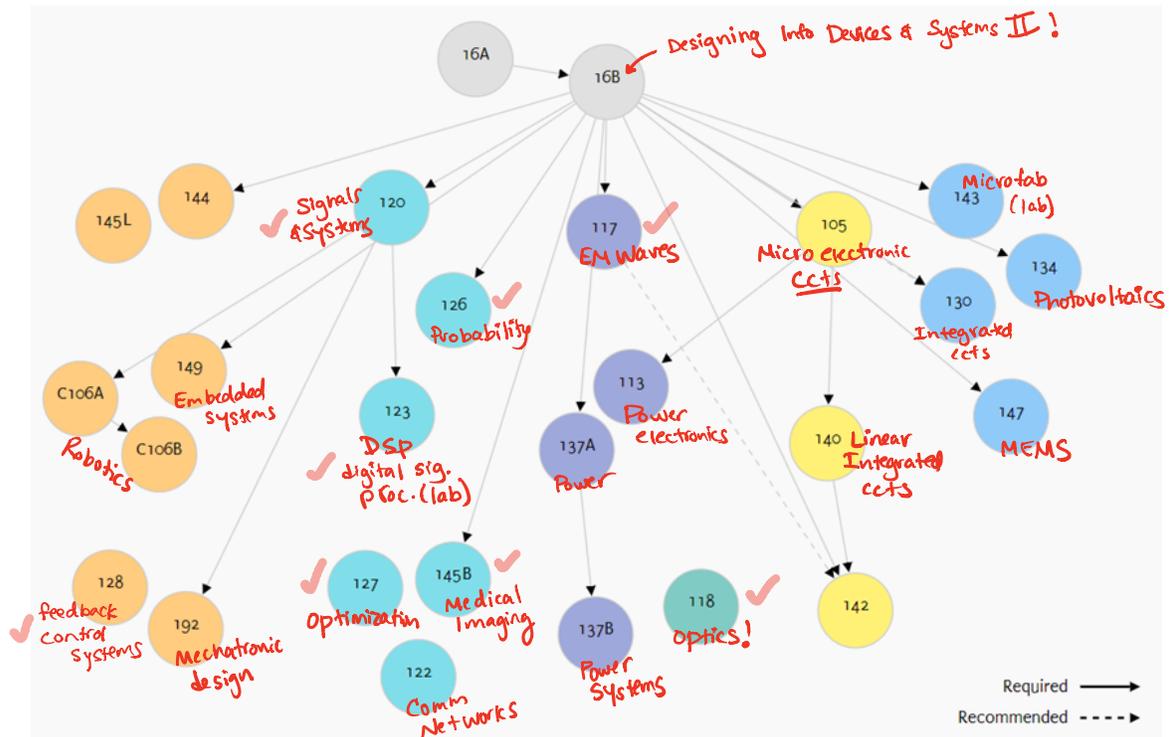


How to approach something unfamiliar
and systematically build understanding

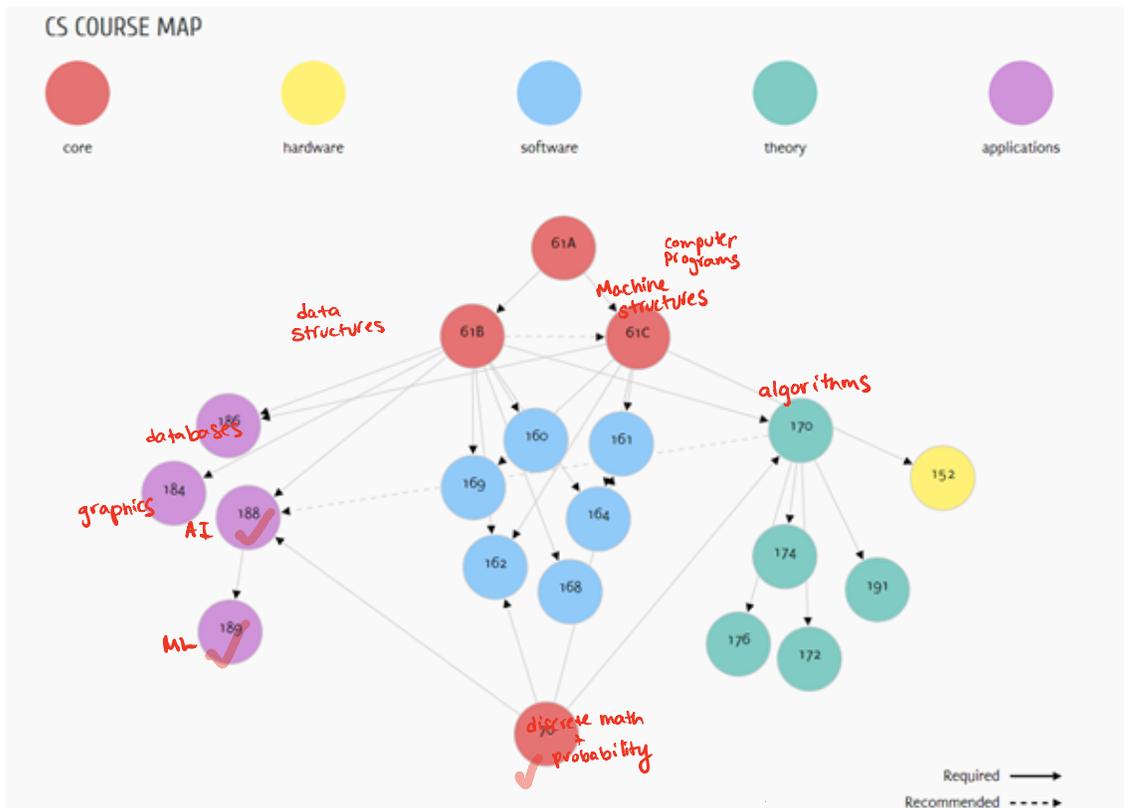
Linear Algebra: conceptual tools to model
Circuits: How to go from model to design, grounded in physical world

Intro to foundational concepts in Machine Learning

ECS course maps:



✓ classes useful to my research



If you liked Linear Algebra, you'll love:
(M1, M3)

- EECS 70: probability
- EE120: signals & systems
 - ↳ EE123 Digital Signal Processing (lab) → e.g. gold codes, correlations, communications
- EECS 126: graph transitions, wireless systems, using probability
 - ↳ e.g. Page Rank, pumps
 - ↳ like GPS
- EE127: optimization theory → how to make things optimal! (mathematical)
- CS 188: } AI/ML ← practical implementations for big data
- CS 189: } uses optimization,
- EECS C106AB: } controls + Robotics ← eigenvalues!
- C128: }

If you liked circuits, you'll LOVE:

EE105:

EE140:

CS 61C:

EE130:

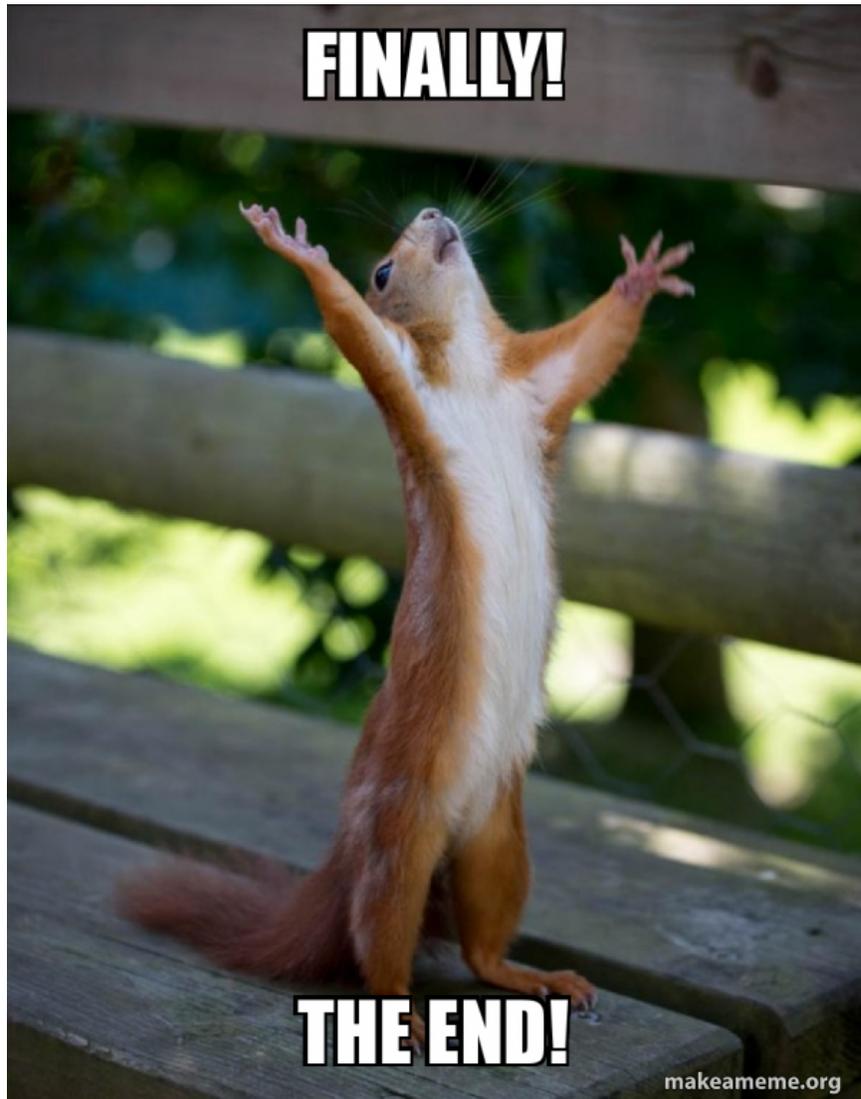
EE134:

EE137AB:

EE143:

Physical laws behind the devices { Phys 7B:
EE117:
EE118:
↳ also, imaging!

If you liked my jokes, you'll LOVE:
↳ sorry, not teaching next year



The End

Oh, except for the
final exam...

