# ECE 60141 / BME 66100
# Foundations of Computational Imaging

## Lecture 1: What is Computational Imaging?
## Chapter 1

Charles A. Bouman

Purdue University

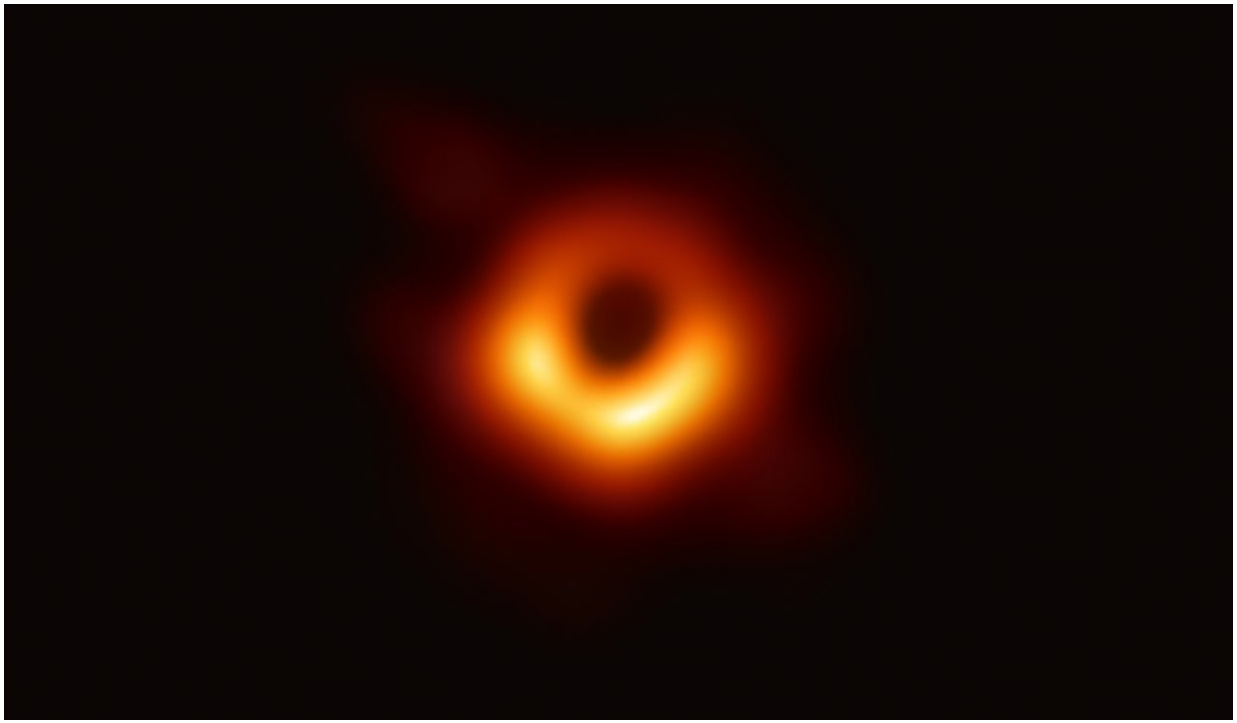
Monday, August 24, 2026

# Today's Big Ideas

- Sensors rarely give a usable image, so computation forms it.
- Recovering the image from measurements is an inverse problem.
- Model-based: a sensor (forward) model plus an image prior.
- Priors are needed: real images are a thin slice of all pixel arrays.
- Real images: 2-D, noncausal, non-Gaussian. We start 1-D Gaussian.

Reading: Chapter 1, *Foundations of Computational Imaging*.

# April 10, 2019: the first image of a black hole



Credit: Event Horizon Telescope Collaboration

- *Computed*, not photographed. So what did the telescope measure?
- The **Event Horizon Telescope** is an Earth-sized array of radio telescopes, aimed at the supermassive black hole in Messier 87. It measured radio-frequency data — and an algorithm turned that data into the picture above.

# What is computational imaging?

- Traditional: optics form the image; the sensor records a picture.
- Computational: the sensor records *data*, not a picture.
- So where does the image come from?

# What is computational imaging?

- Traditional: optics form the image; the sensor records a picture.
- Computational: the sensor records *data*, not a picture.
- So where does the image come from?

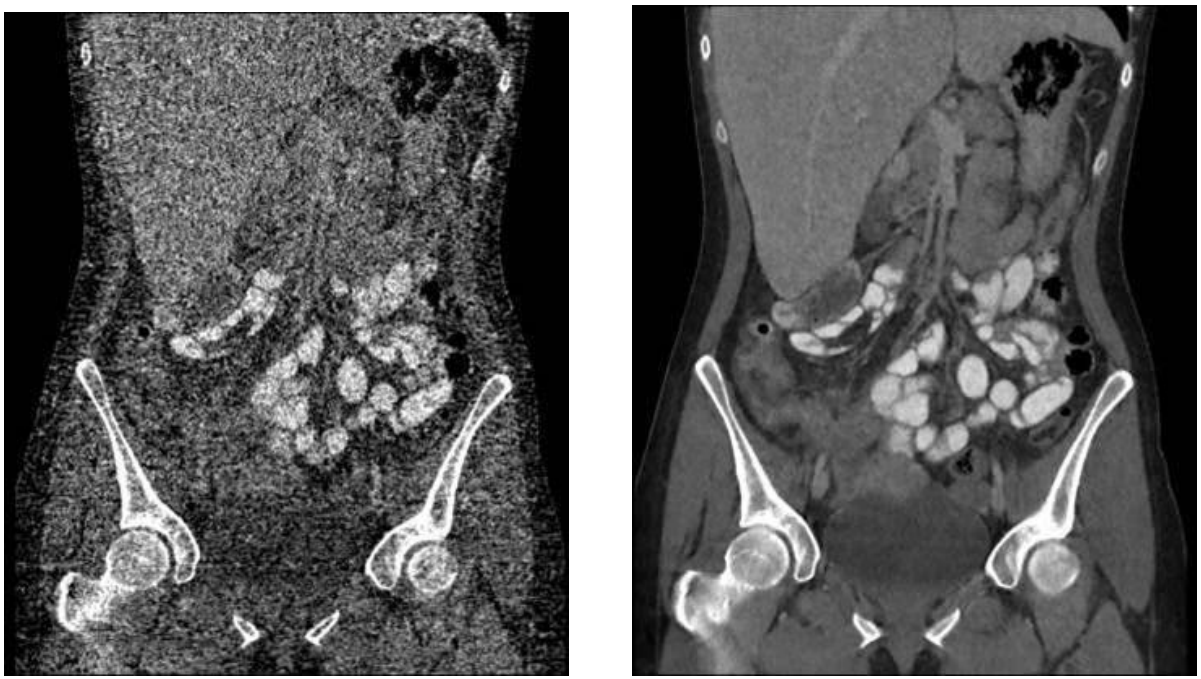
The image you see is the solution to an *inverse problem*.

CT, MRI, PET, electron microscopy, synthetic aperture radar, computational photography — and your cellphone camera — all work this way.

**The core idea of this course**

Model the sensor (the forward model), model the image (the prior), and compute the reconstruction by optimization or sampling.

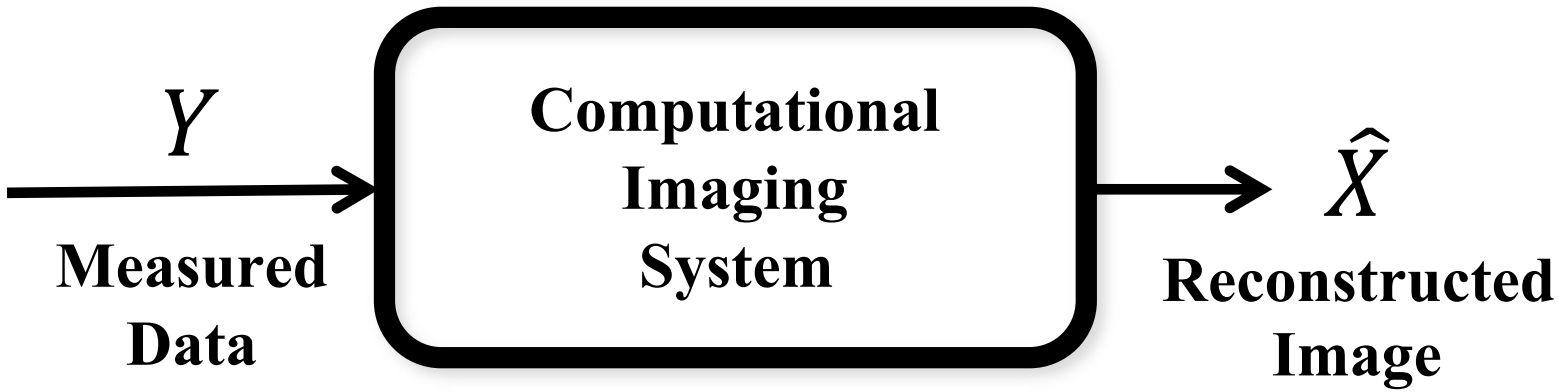
# A promise for the semester



- Same scan, same CT scanner: direct (left) vs. model-based iterative (right).
- Which image would you rather have your radiologist read?
- Same data, better *mathematics*: finer detail and less noise.

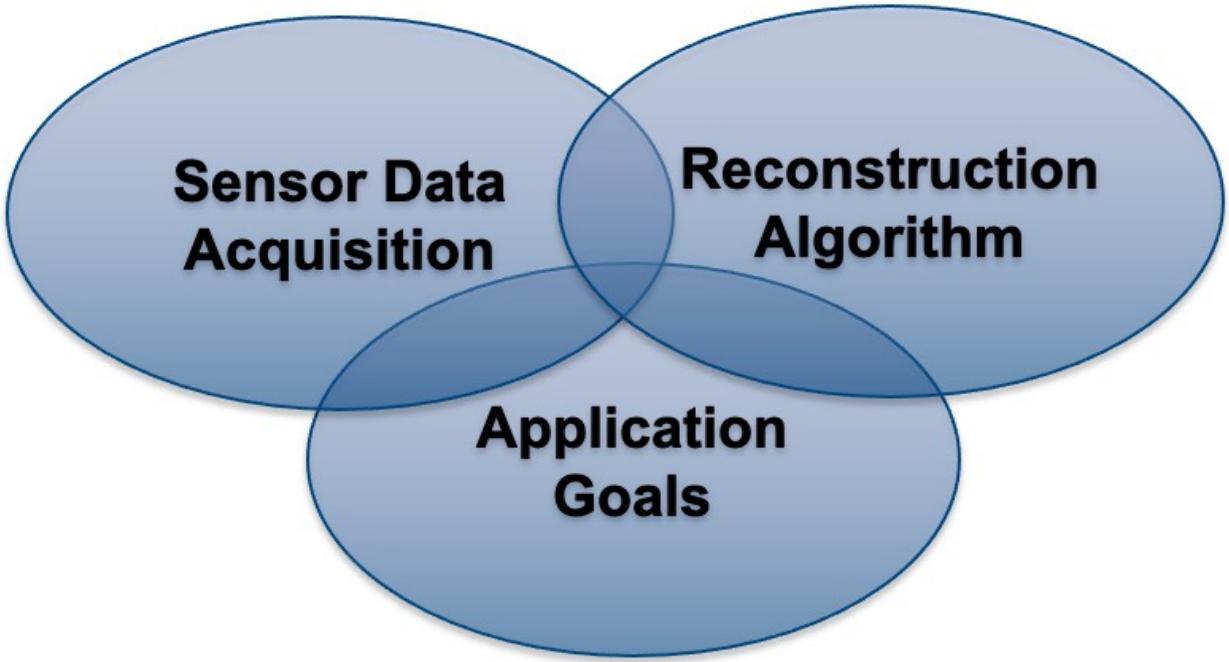
By the end of this course, you will understand exactly why the image on the right is possible.

# Sensors do not make images — algorithms do




- Practical sensors almost never directly generate usable images.
- A cellphone CMOS sensor needs color dematrixing, deblurring, denoising, and registration before you see a picture.
- In CT or SAR, the raw data is not even recognizable as an image.

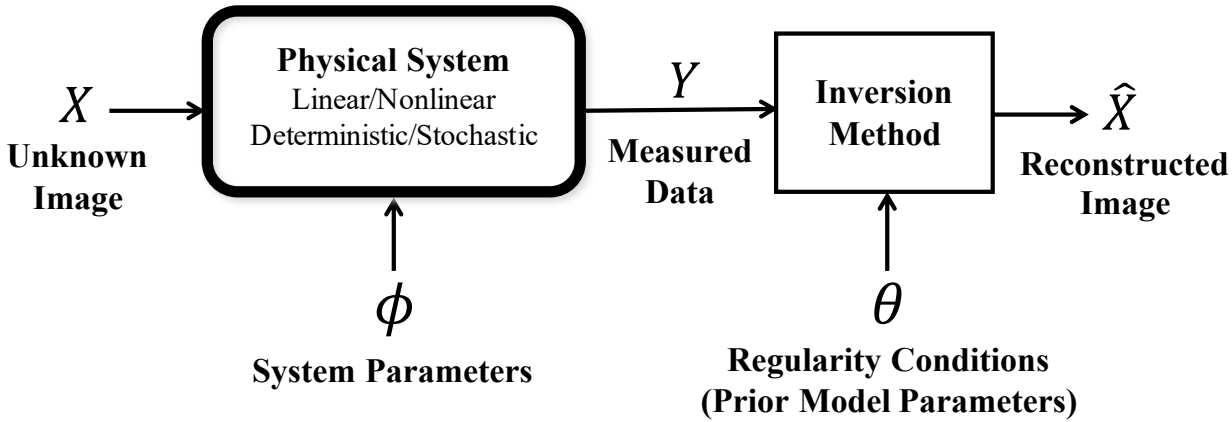
# Co-design of sensor and algorithm




- Suppose perfect measurements cost \$10 million, and imperfect measurements plus a good algorithm cost \$100,000. Which system should we build?
- Design the sensor and the reconstruction algorithm *together* to meet the application goals at minimum cost. The goal is the most *informative* measurements, not *perfect* measurements.
- **Mick Jagger's theorem** states that "you can't always get the sensor data you want, but if you try sometimes, you get what you need."

# The inverse problem

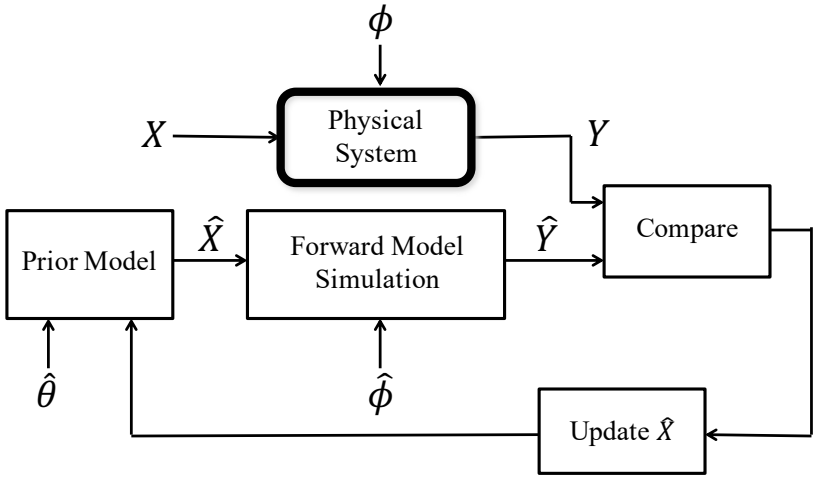
- There is something we want to know: an image $X$.
- Nature hands us measurements instead: $Y$.
- The **inverse problem**: estimate $\hat{X}$ of $X$ from $Y$.




- Often also estimate $\phi$ (calibration) and $\theta$ (regularization).
- Microscopes, phone cameras, radio telescopes, PET: same structure.

# The model-based approach

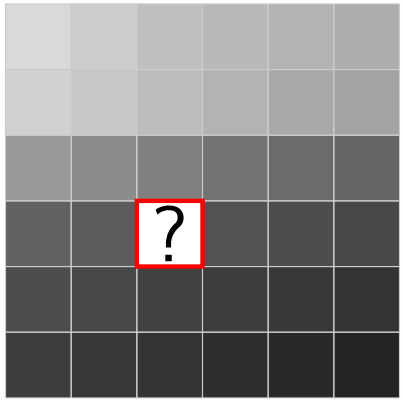
If we guessed an image, how could we tell whether the guess was right?




- The **forward model** $p(y|x)$ predicts the data a guess would produce.
- Compare with the real measurements, update the guess, repeat.
- The **prior model** $p(x)$ keeps every guess looking like a real image.

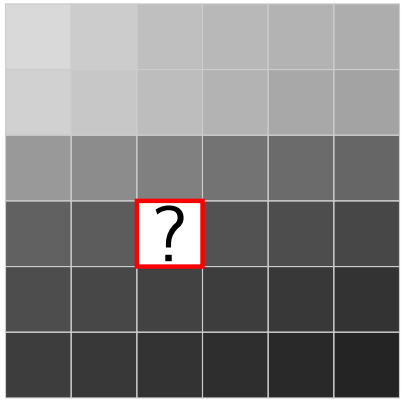
We compute the inverse by modeling the *forward* process — so it works even when no analytical inverse exists.

# A thought experiment



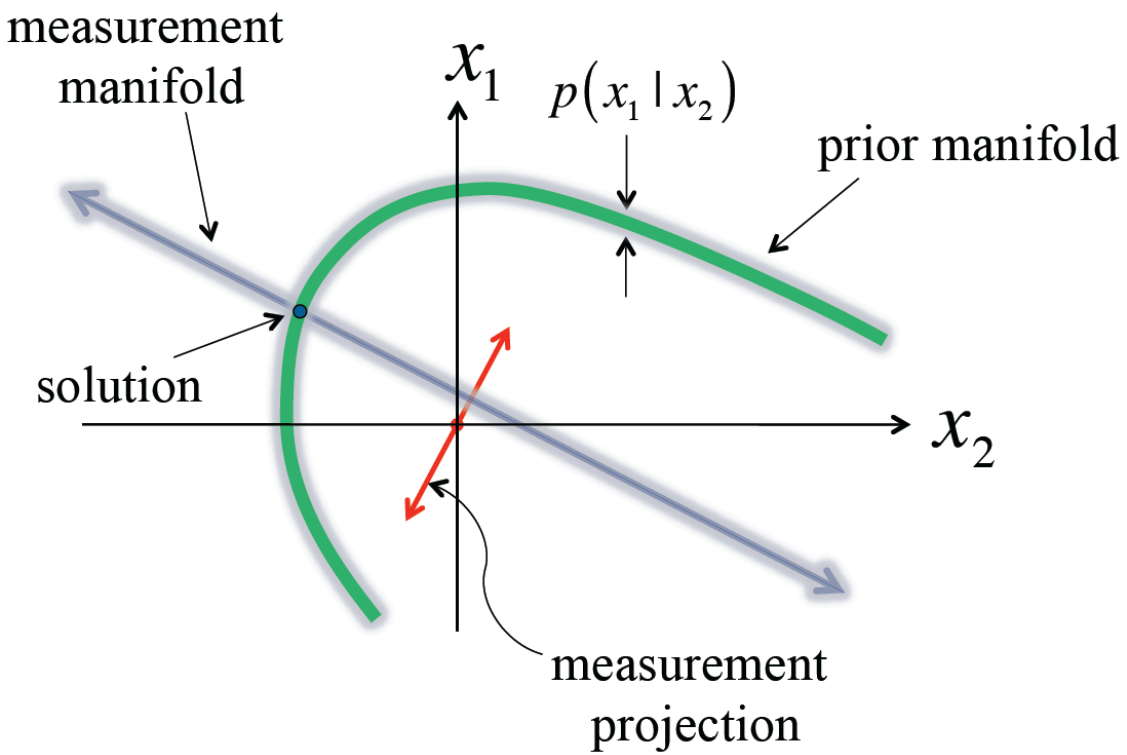

- One pixel is covered. Can you guess roughly what belongs there?

# A thought experiment




- One pixel is covered. Can you guess roughly what belongs there?
- Of course you can. Why?
- Most pixel arrays are not pictures; neighbors predict each pixel.

# Why priors work: images lie on a thin manifold




- A **thin manifold**: the curved set of $[0,1]^N$ real images occupy.
- Measurements $Y = AX + W$ define another thin manifold.
- Even if $M < N$, the good solution is at the *intersection*.

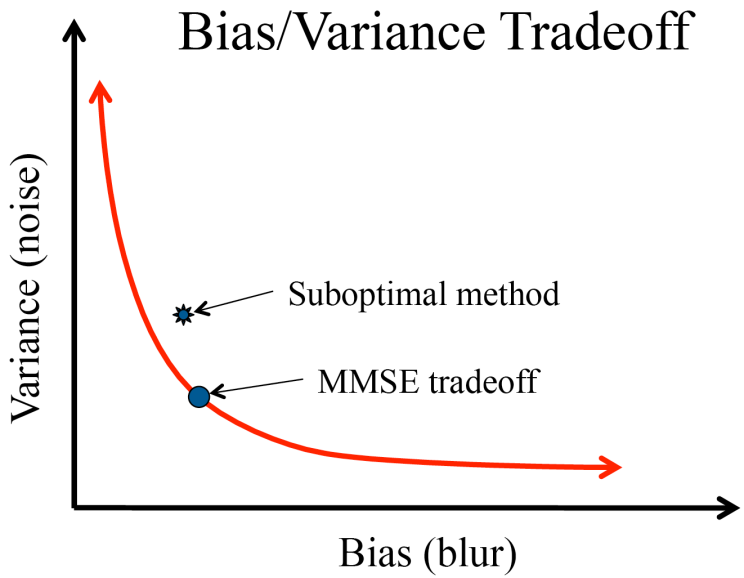
# “But aren’t priors cheating?”

*Shouldn’t we just reconstruct what the data say — no assumptions?*

# “But aren’t priors cheating?”

*Shouldn’t we just reconstruct what the data say — no assumptions?*

There is no such thing as “no assumptions.”




- Every inversion assumes: a resolution limit, a reconstruction filter.
- Zero bias often means very large (even infinite) variance. Variance looks like noise; bias looks like blur or artifacts. You saw both in the CT pair earlier.
- Goal: the most *accurate* prior, with the least harmful bias.

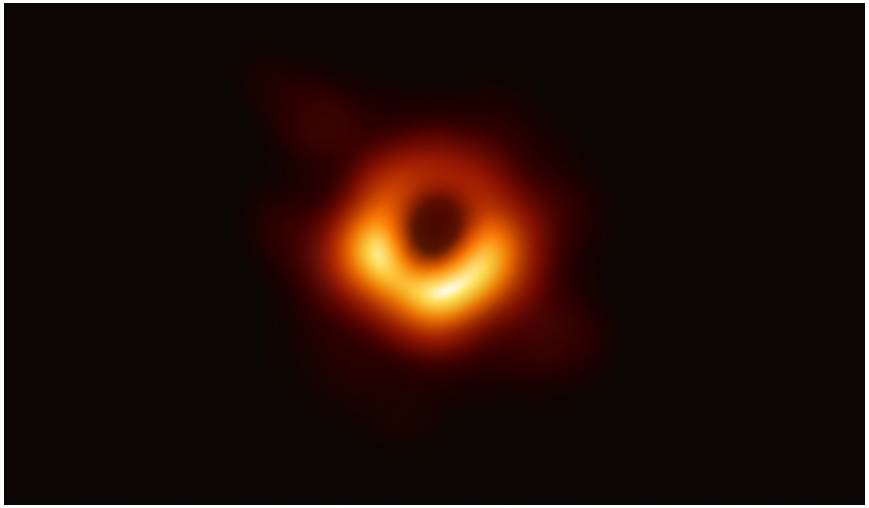
# Roadmap: from simple models to real images

- Real images are best modeled as **2-D, noncausal, and non-Gaussian** — adjacent-pixel differences are Laplacian, not Gaussian. That is the hardest box in this table:

| | **1-D** | **2-D** |
|---|---|---|
| Gaussian | causal AR | GMRF |
| Non-Gaussian | Markov process | MRF |
| Discrete | Markov chain | discrete MRF |

- We start in the simplest box (1-D, causal, Gaussian) and build up.
- Early examples will not look like CT scans. This table explains why.
- Noncausal models in 2-D have loops, which break the standard 1-D tools — that is why Markov random fields are a central theme.

# Where we are headed





- This semester: how measured data become images like these.
- Logistics and this week's to-dos: **Lecture 0**. Read it this week.

Welcome to the course!

# Summary

- Sensor data $Y$ are not images; computation gives $\hat{X}$.
- Estimate $X$ from $Y$, also estimating nuisance parameters $\phi$, $\theta$.
- Model-based methods combine $p(y|x)$ with $p(x)$ to return $\hat{X}$.
- Real images lie on a thin manifold in $[0,1]^N$, so measurements $Y = AX + W$ can determine a good solution even when $M < N$.
- Real images are 2-D, noncausal, and non-Gaussian, so the course starts with 1-D causal Gaussian models and works toward Markov random fields.

**Reading:** Chapter 1 **Next lecture:** Chapter 2.1, Random variables and expectation.