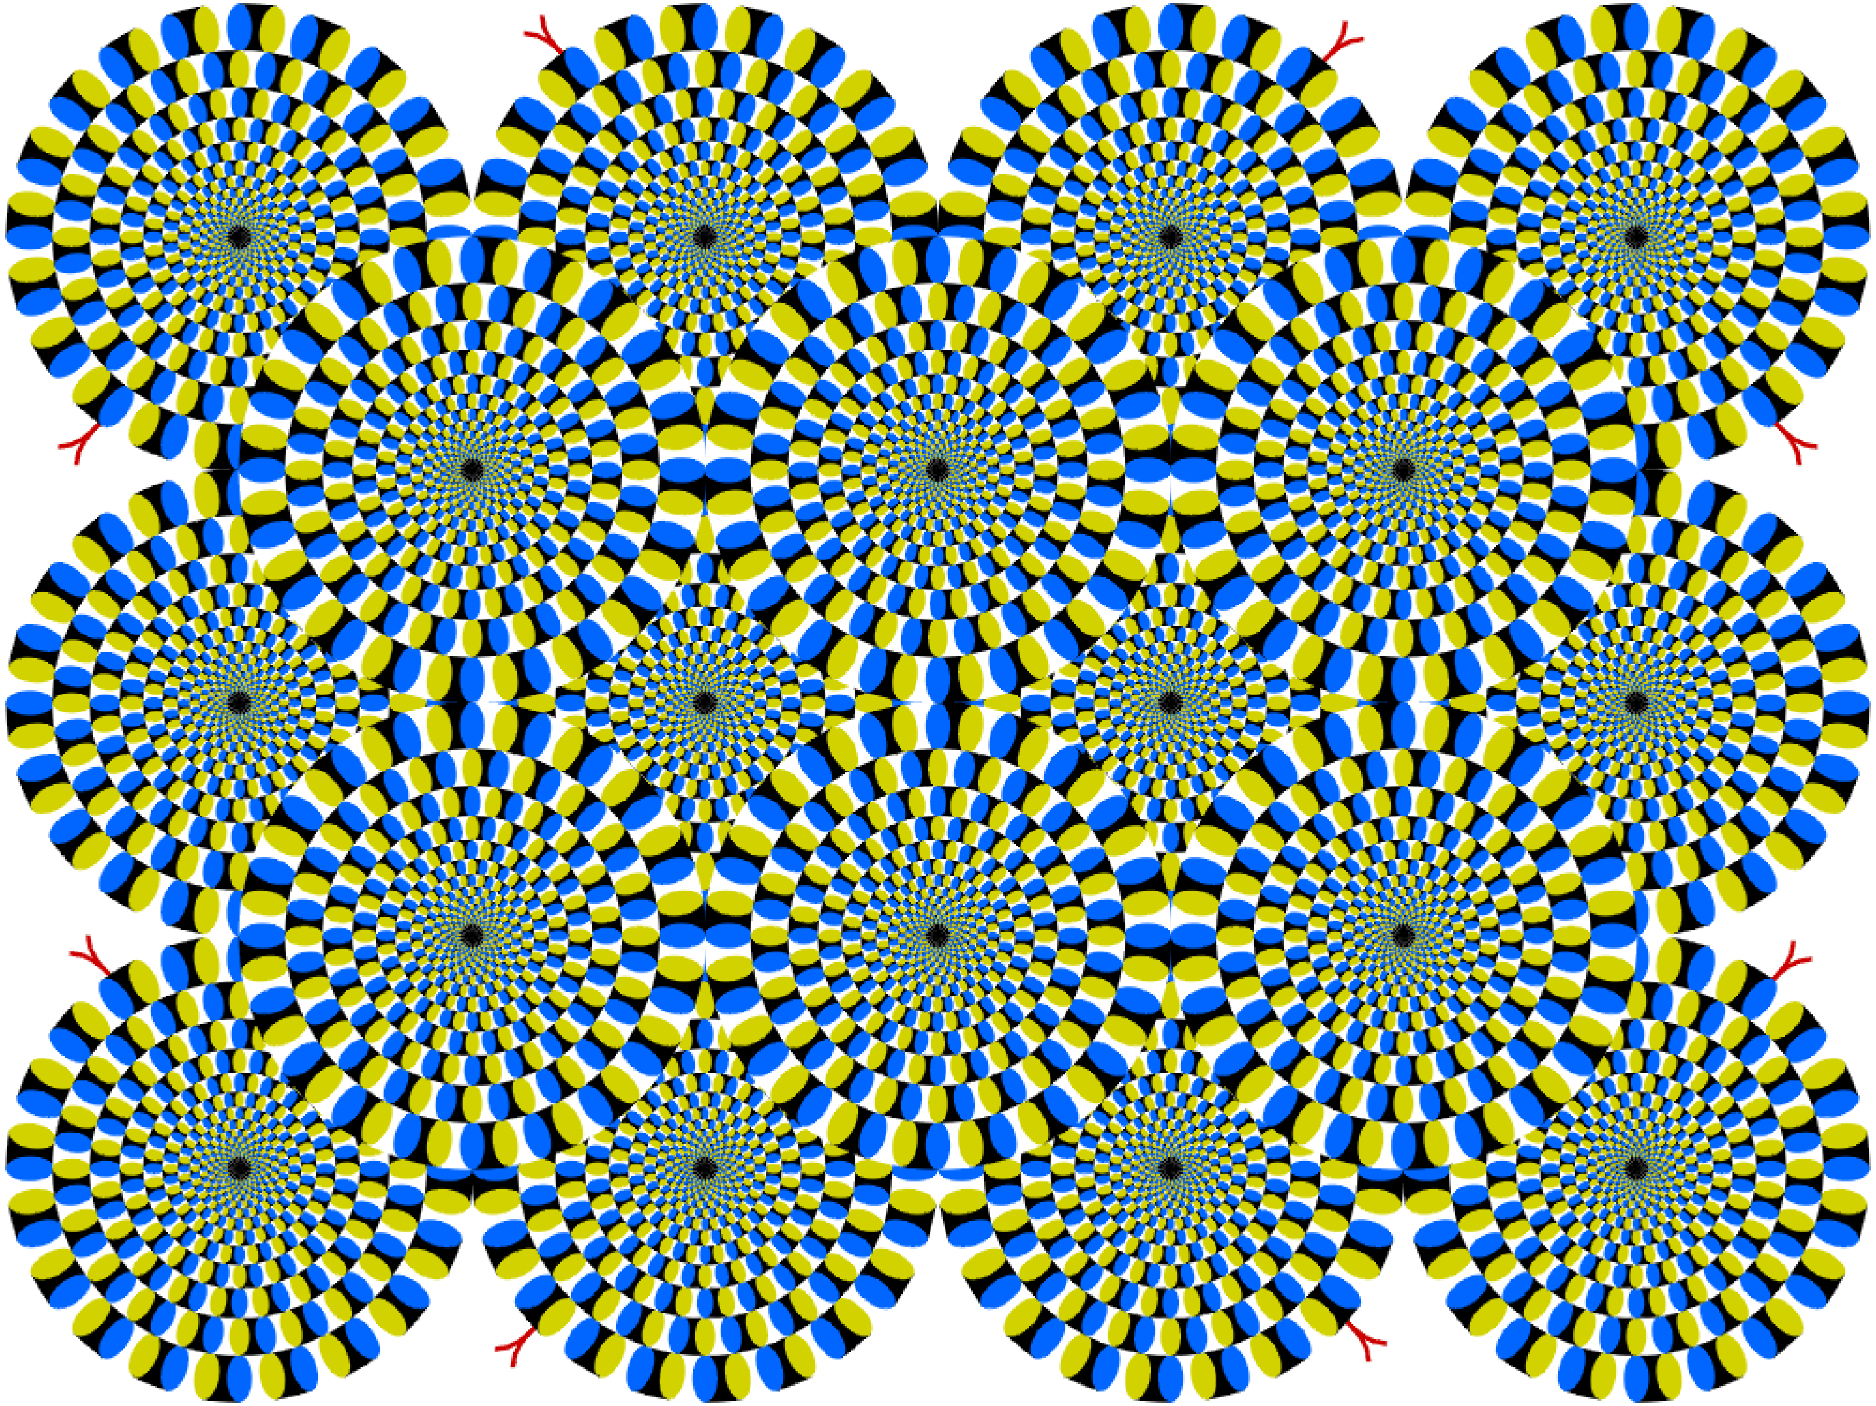

Recap: projection

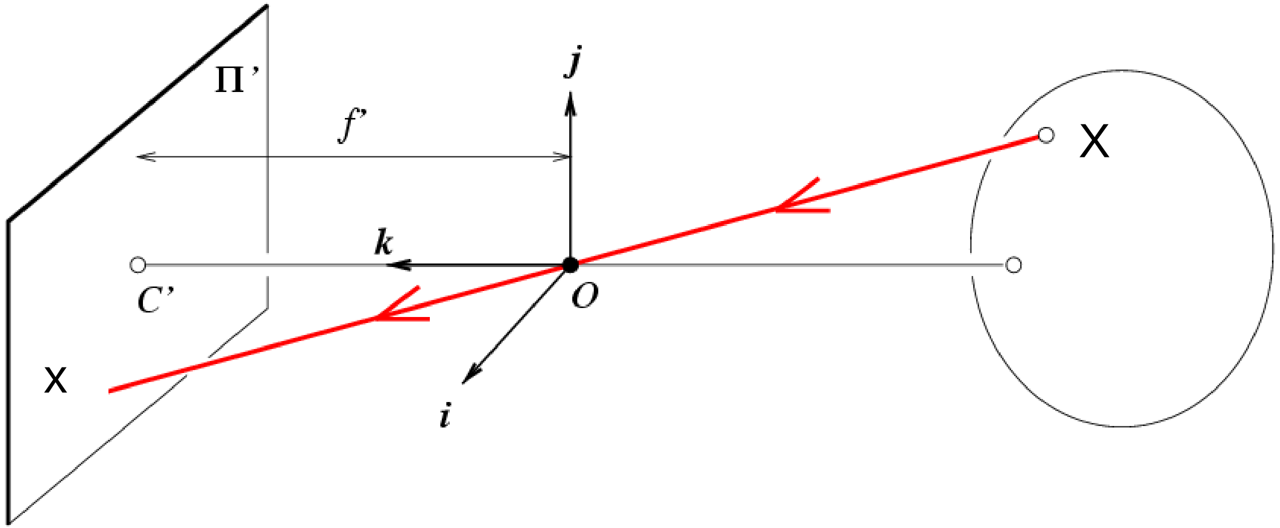

$$\mathbf{x} = \mathbf{K}\begin{bmatrix}\mathbf{R} & \mathbf{t}\end{bmatrix}\mathbf{X}$$

$$w\begin{bmatrix}u\\v\\1\end{bmatrix} = \begin{bmatrix}\alpha & s & u_0\\0 & \beta & v_0\\0 & 0 & 1\end{bmatrix}\begin{bmatrix}r_{11} & r_{12} & r_{13} & t_x\\r_{21} & r_{22} & r_{23} & t_y\\r_{31} & r_{32} & r_{33} & t_z\end{bmatrix}\begin{bmatrix}x\\y\\z\\1\end{bmatrix}$$

Relating multiple views

Figure Credit: Bundler: Structure from Motion (SfM) for Unordered Image Collections

Recap of Filtering

- Linear filtering is dot product at each position
  - Not a matrix multiplication
  - Can smooth, sharpen, translate (among many other uses)

$$\frac{1}{9}\begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

- We can use the Fourier transform to represent images in the frequency domain.
  - Filtering in the spatial domain is multiplication in the frequency domain.

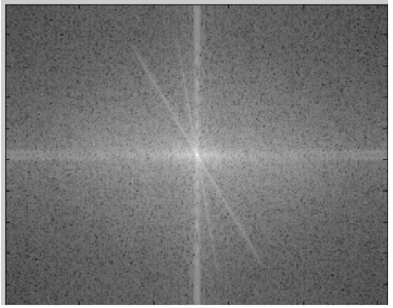

Fourier Bases

Teases away fast vs. slow changes in the image.

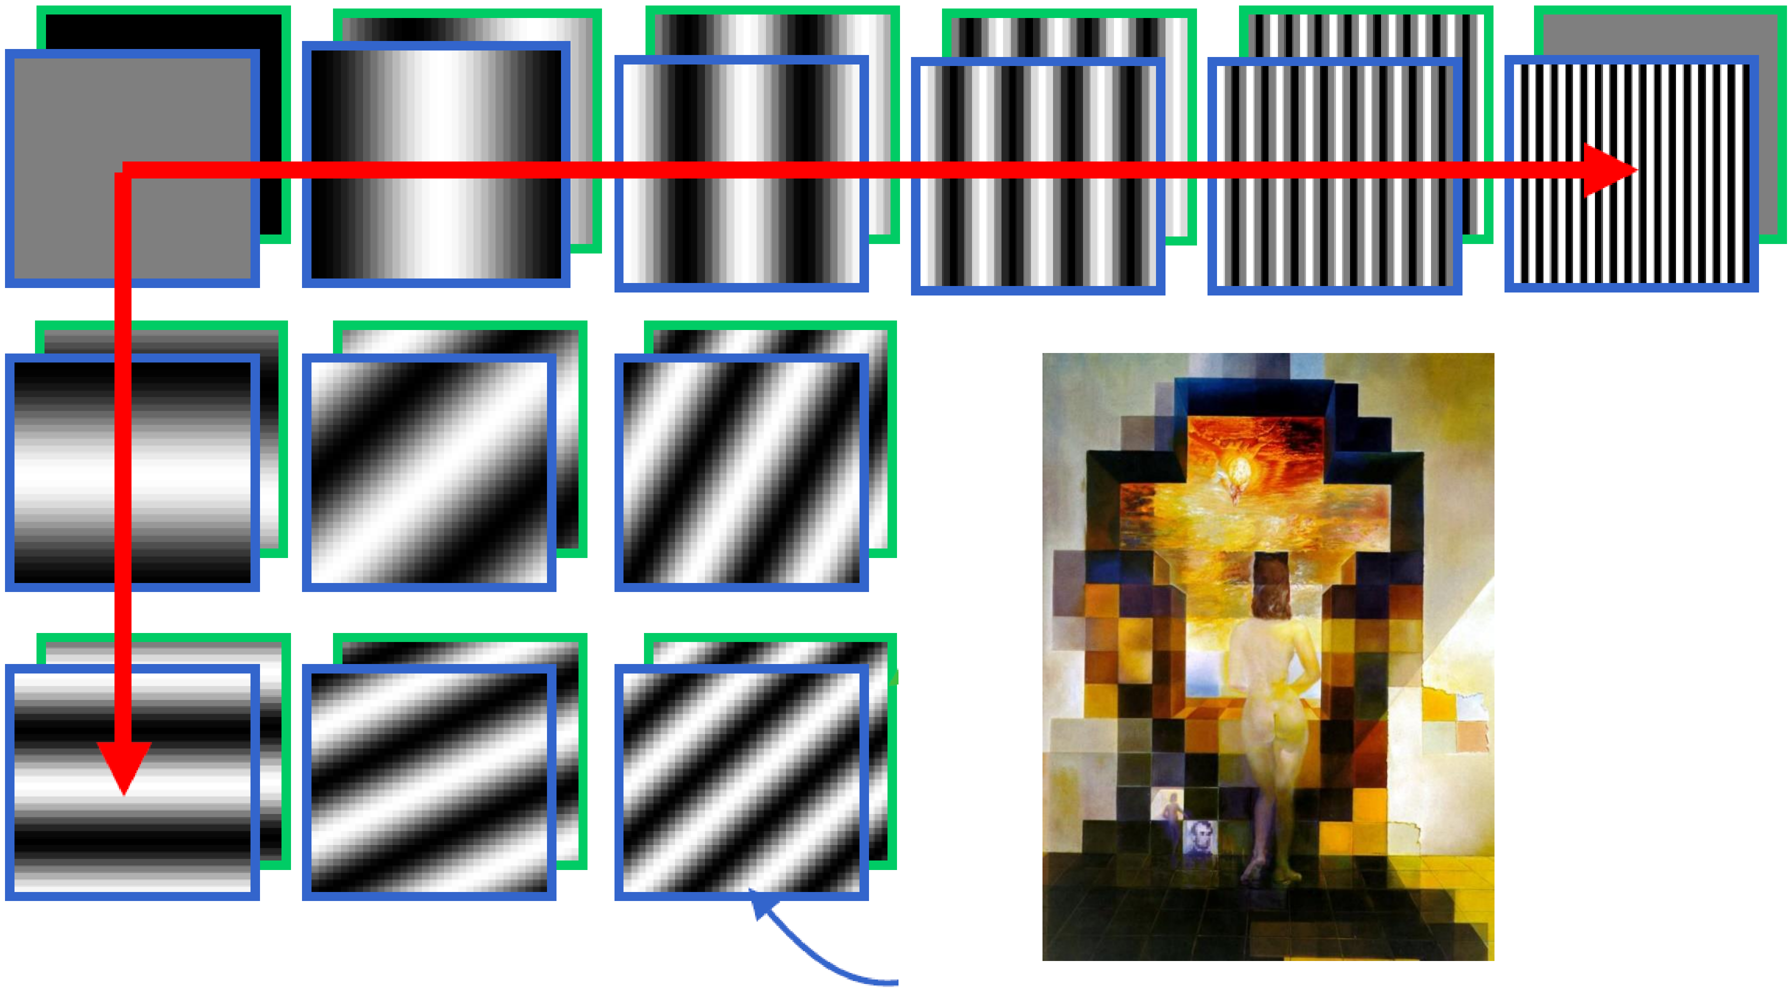

This change of basis is the Fourier Transform

Fourier Bases

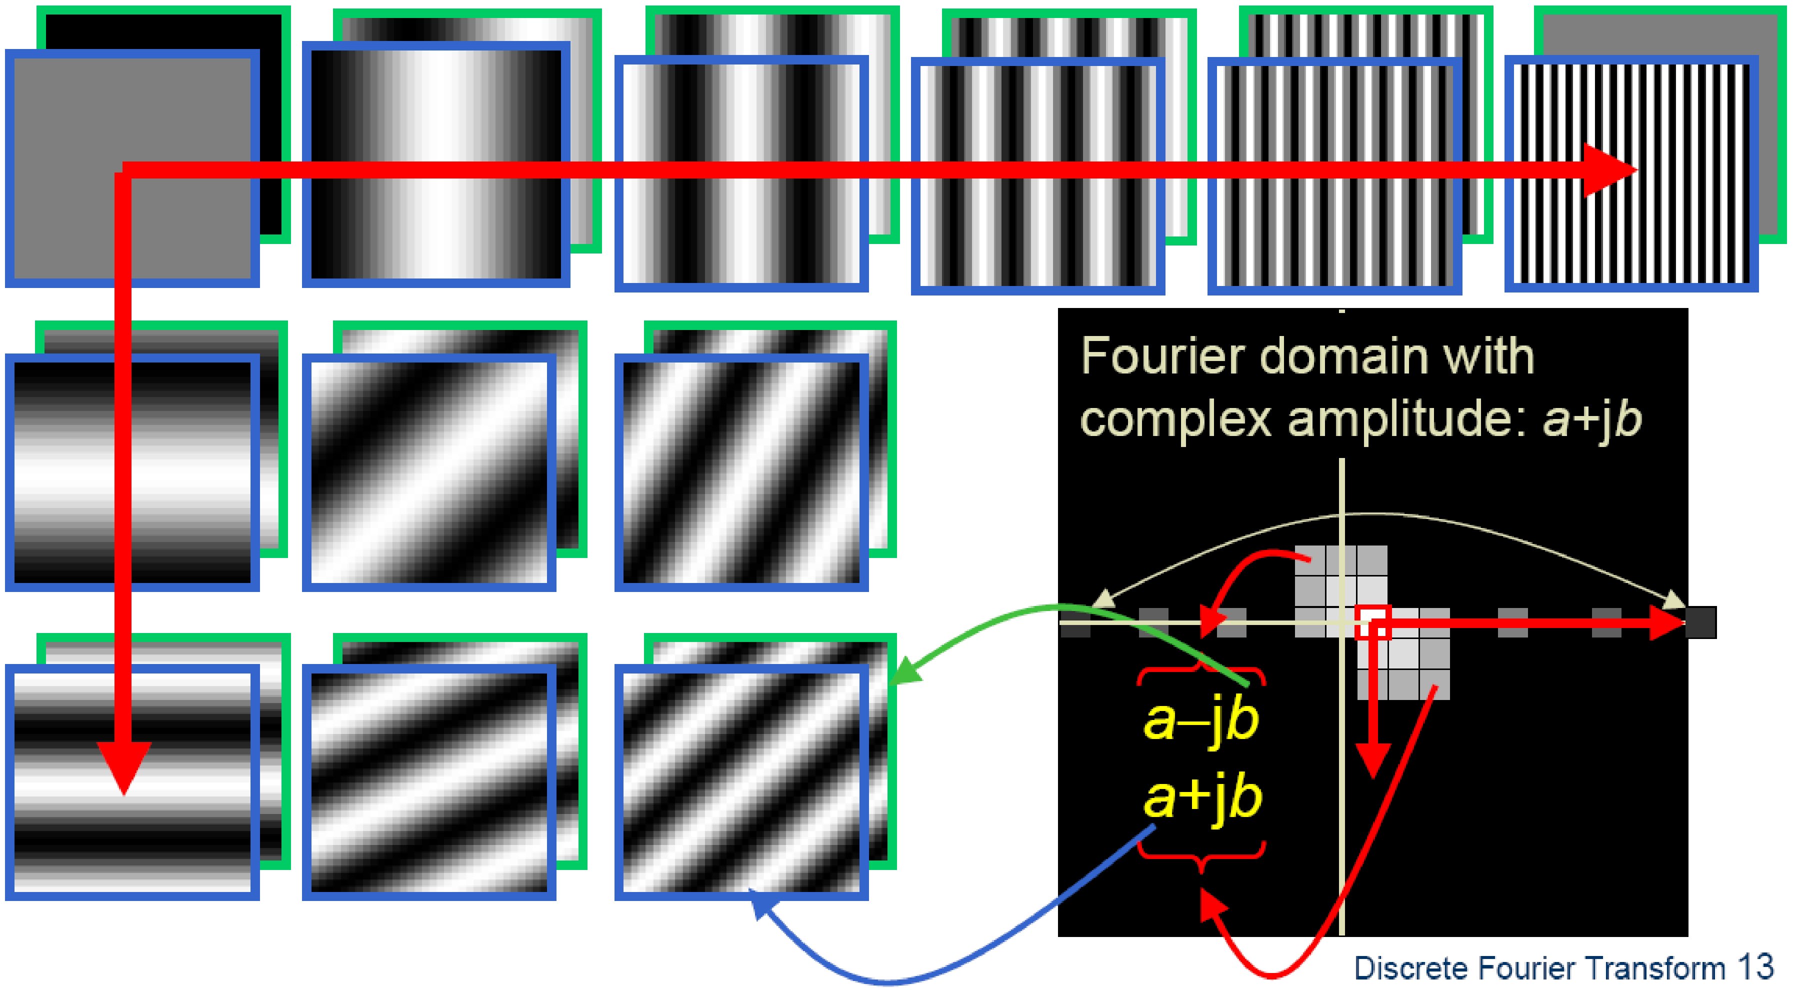

Low and High Pass filtering

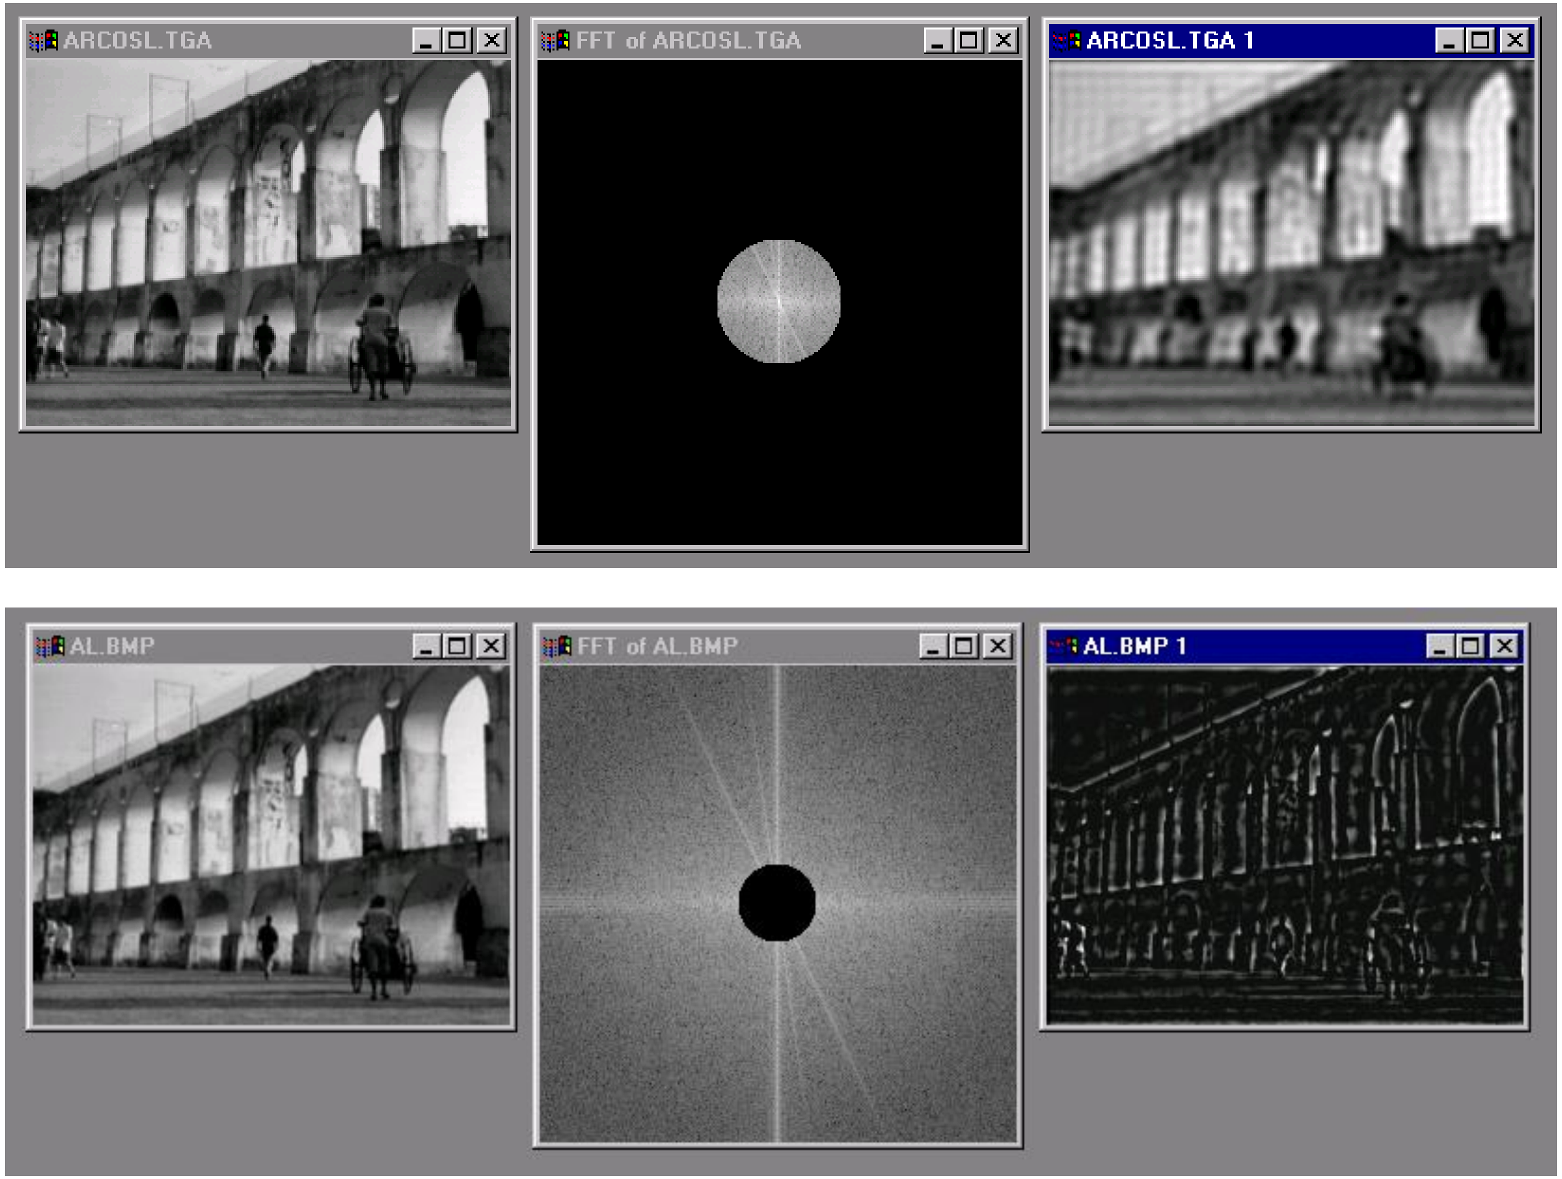

The Convolution Theorem

- The Fourier transform of the convolution of two functions is the product of their Fourier transforms

$$\mathrm{F}[g * h] = \mathrm{F}[g]\,\mathrm{F}[h]$$

- **Convolution** in spatial domain is equivalent to **multiplication** in frequency domain!

$$g * h = \mathrm{F}^{-1}[\mathrm{F}[g]\,\mathrm{F}[h]]$$

Filtering in spatial domain

1	0	-1
2	0	-2
1	0	-1

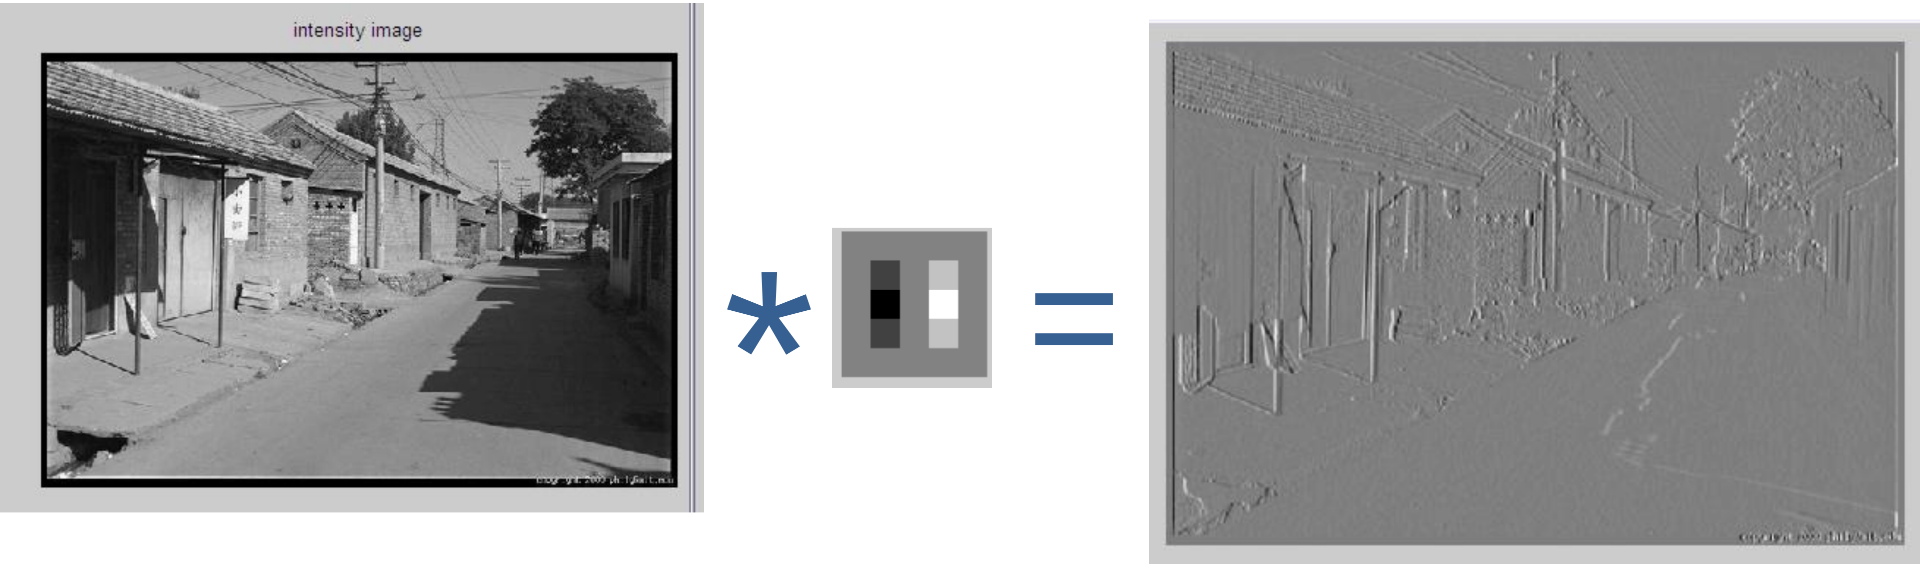

Filtering in frequency domain

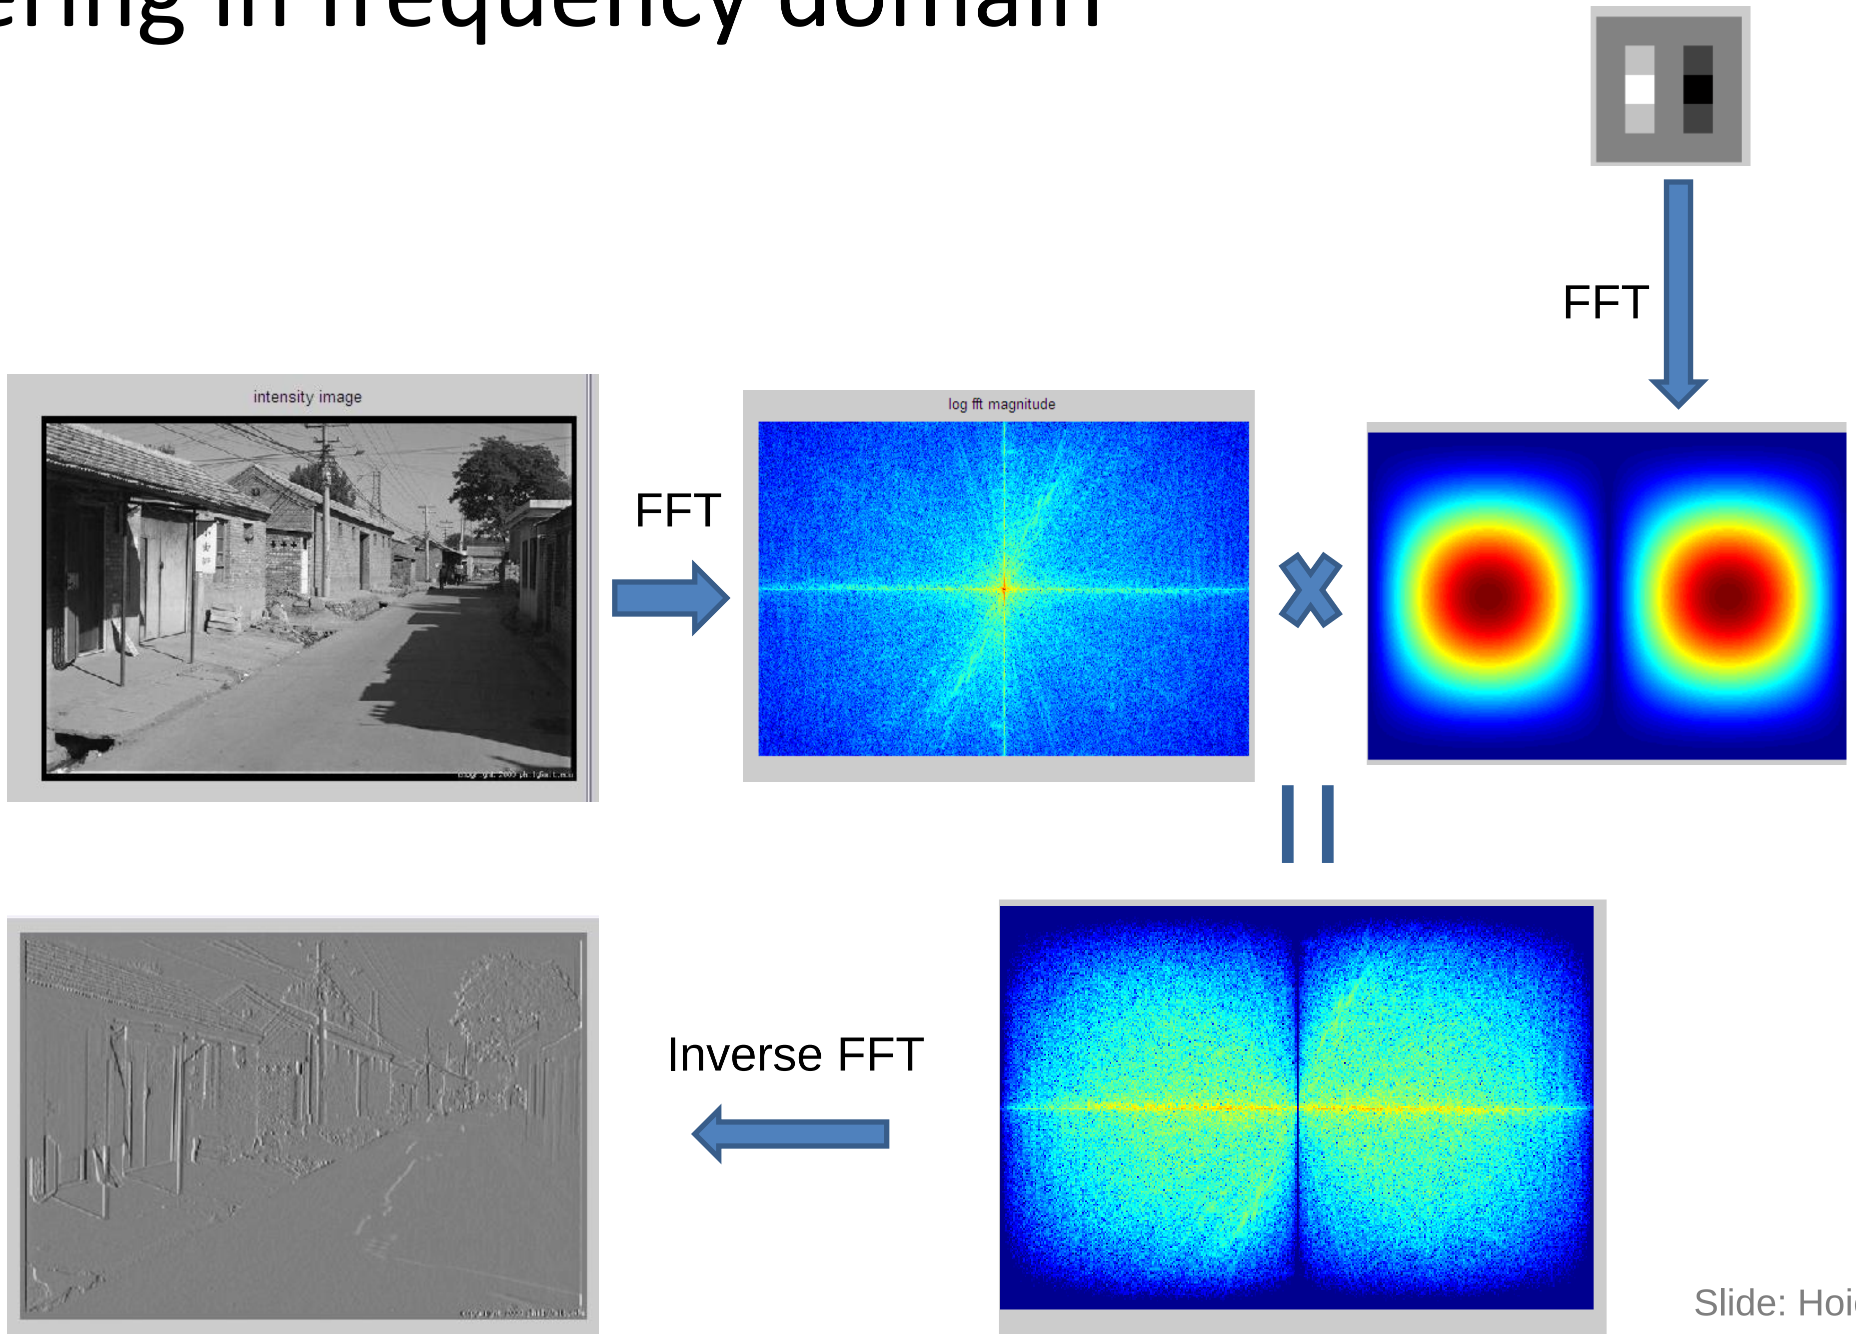

Slide: Hoiem

Filtering

Why does the Gaussian give a nice smooth image, but the square filter give edgy artifacts?

Gaussian 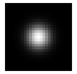

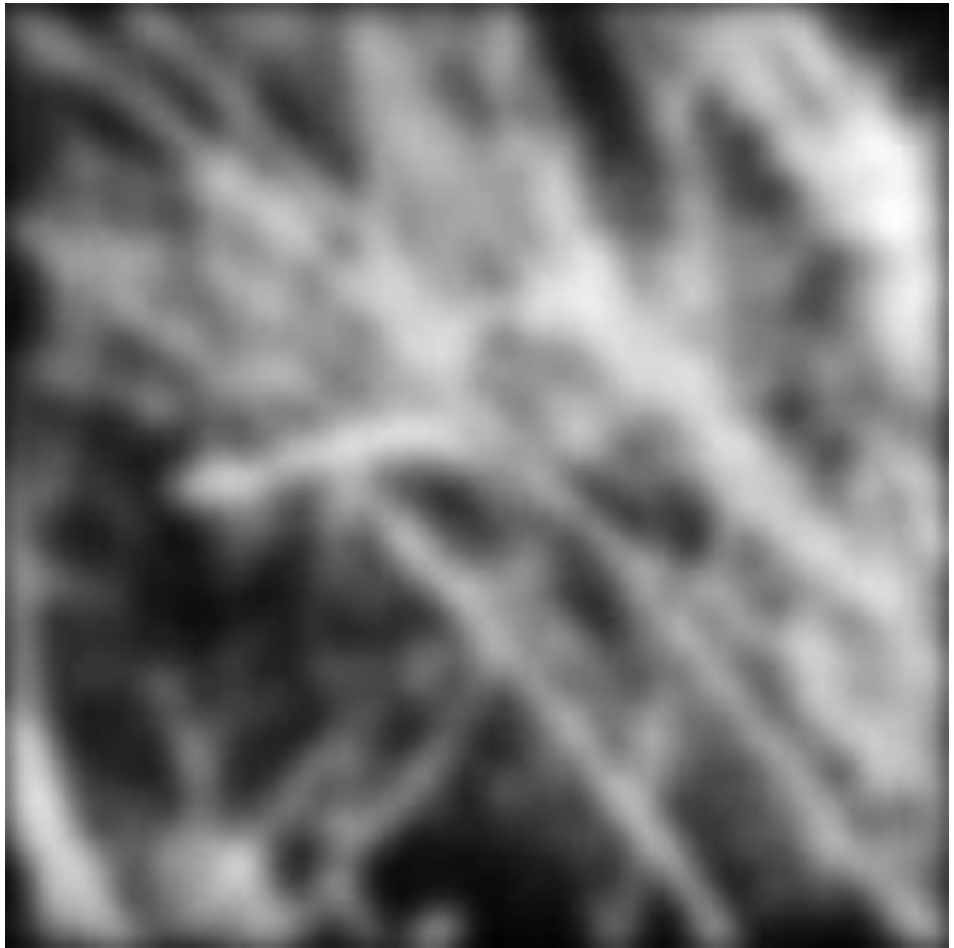

Box filter 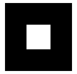

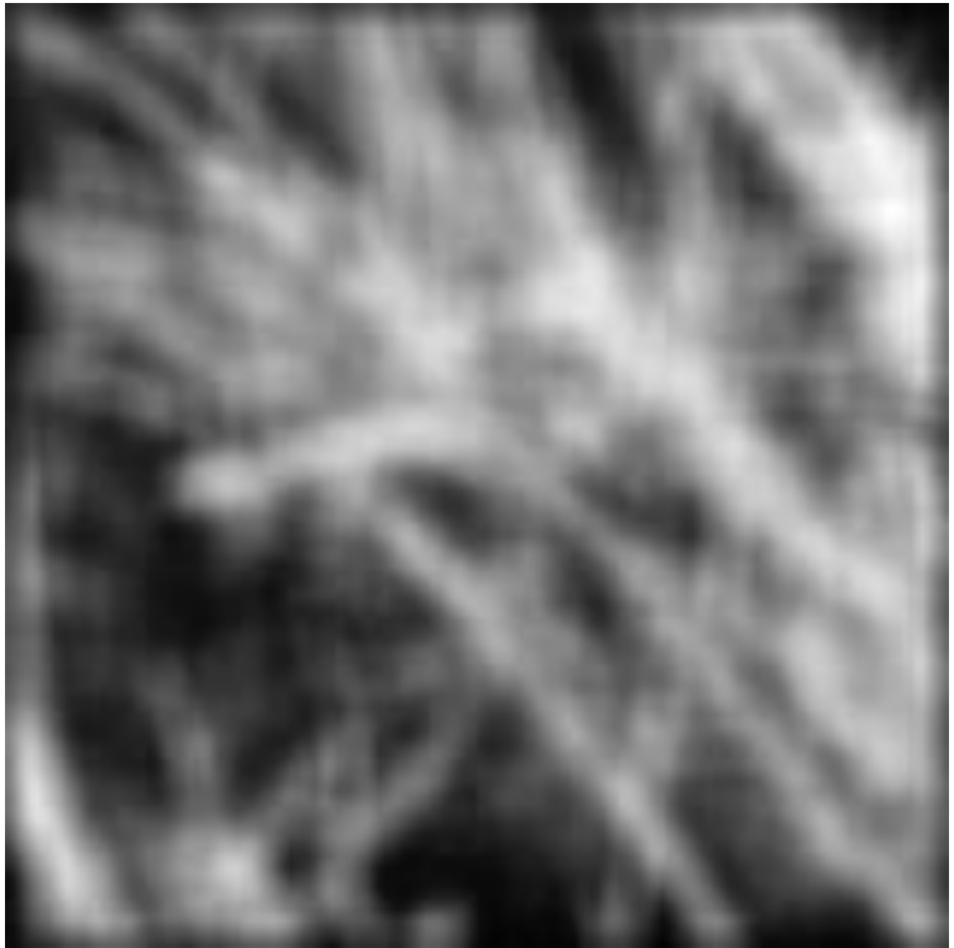

Gaussian

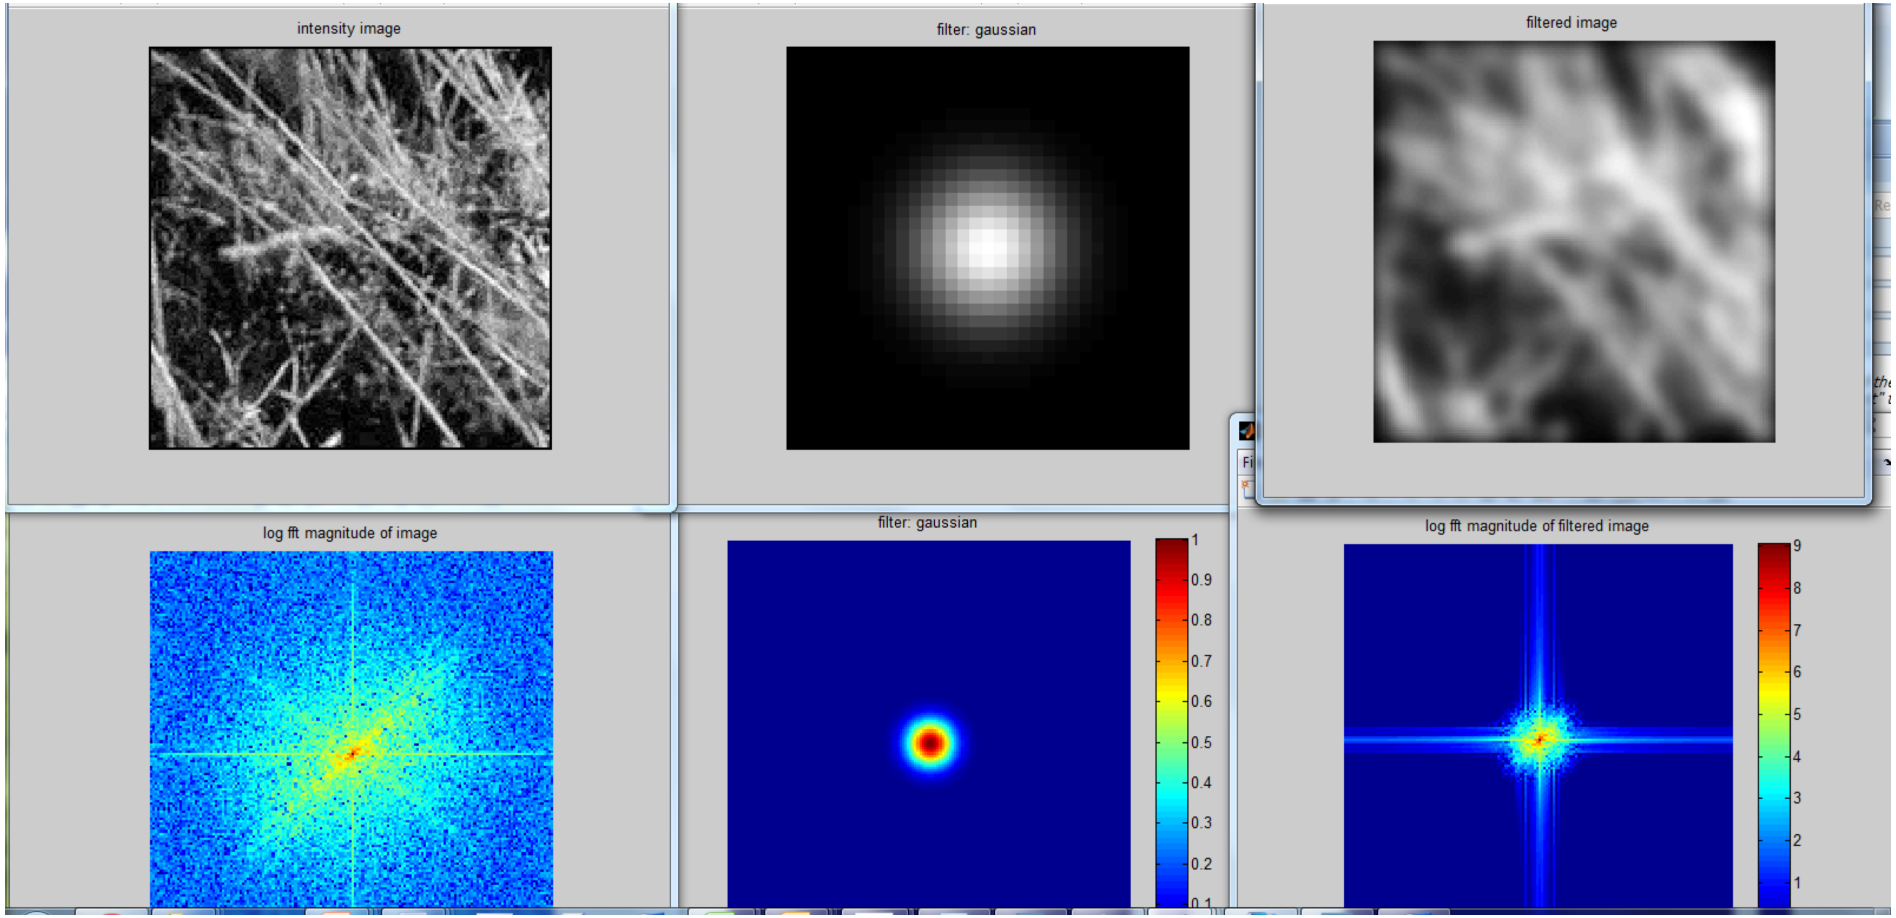

intensity image
filter: gaussian
filtered image
log fft magnitude of image
filter: gaussian
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
log fft magnitude of filtered image
9
8
7
6
5
4
3
2
1

Box Filter

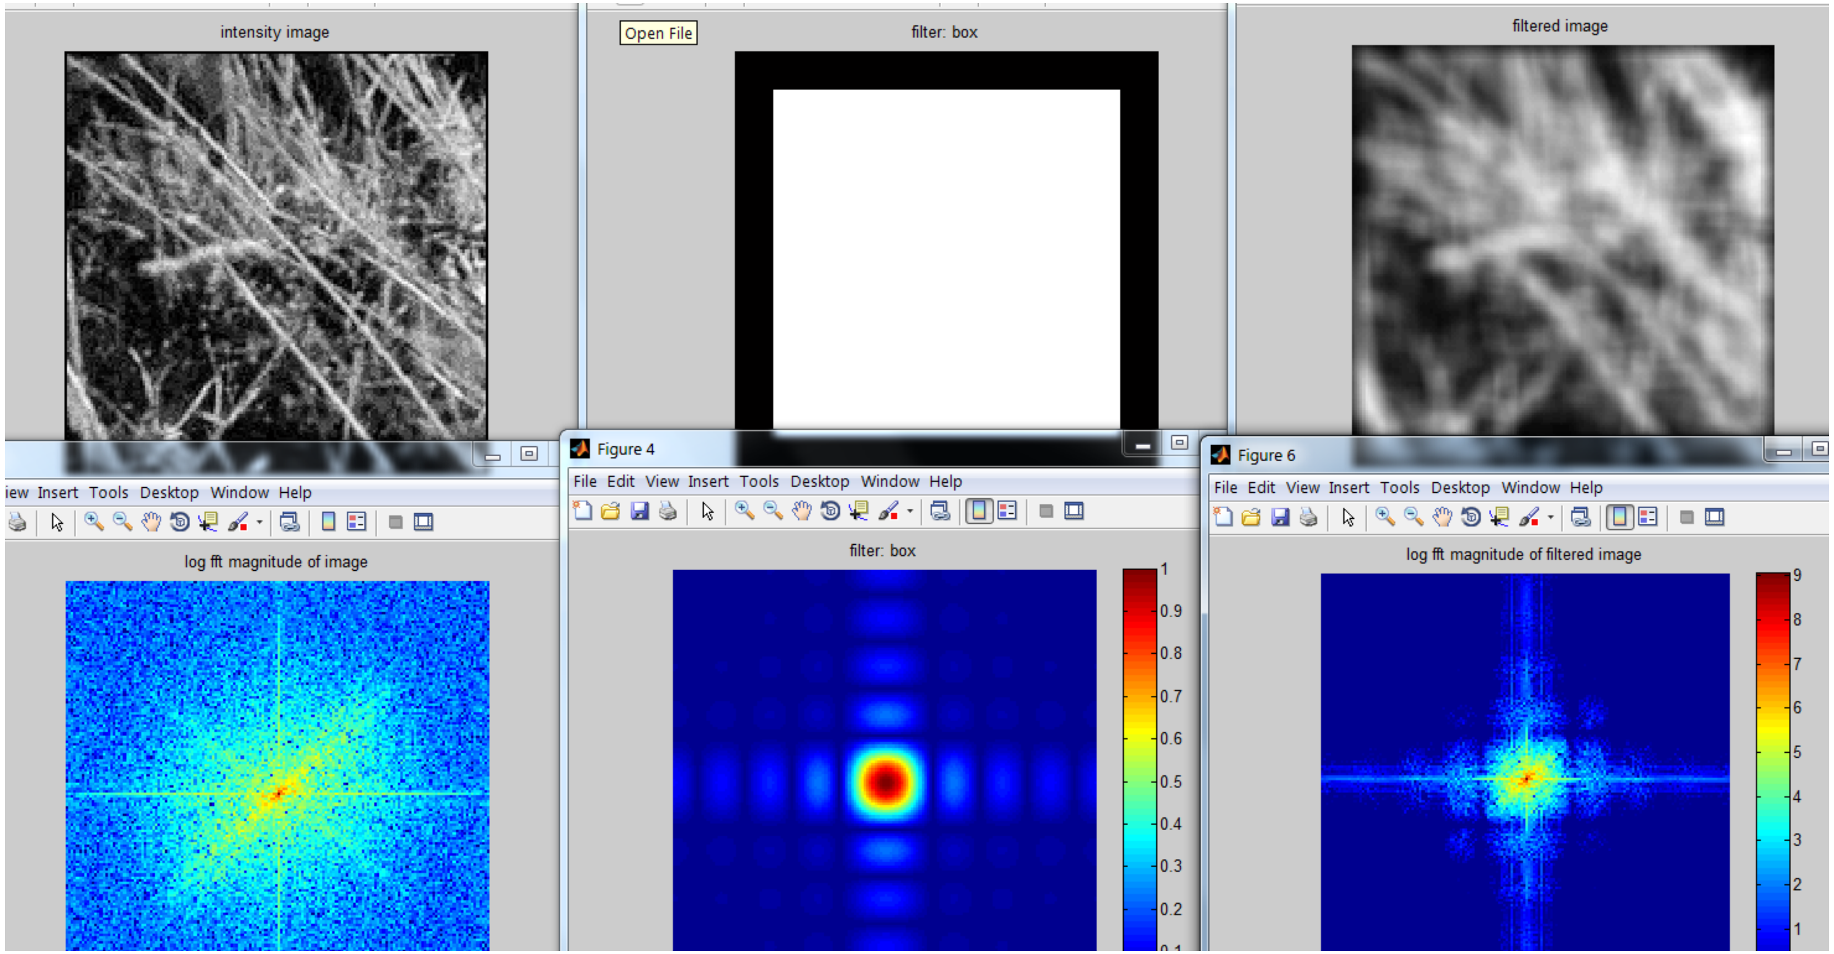

intensity image
Open File
filter: box
filtered image
iew Insert Tools Desktop Window Help
log fft magnitude of image
Figure 4
File Edit View Insert Tools Desktop Window Help
filter: box
1
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
Figure 6
File Edit View Insert Tools Desktop Window Help
log fft magnitude of filtered image
9
8
7
6
5
4
3
2
1

Is convolution invertible?

- If convolution is just multiplication in the Fourier domain, isn’t deconvolution just division?
- Sometimes, it clearly is invertible (e.g. a convolution with an identity filter)
- In one case, it clearly isn’t invertible (e.g. convolution with an all zero filter)
- What about for common filters like a Gaussian?

But you can't invert multiplication by 0

- But it's not quite zero, is it…

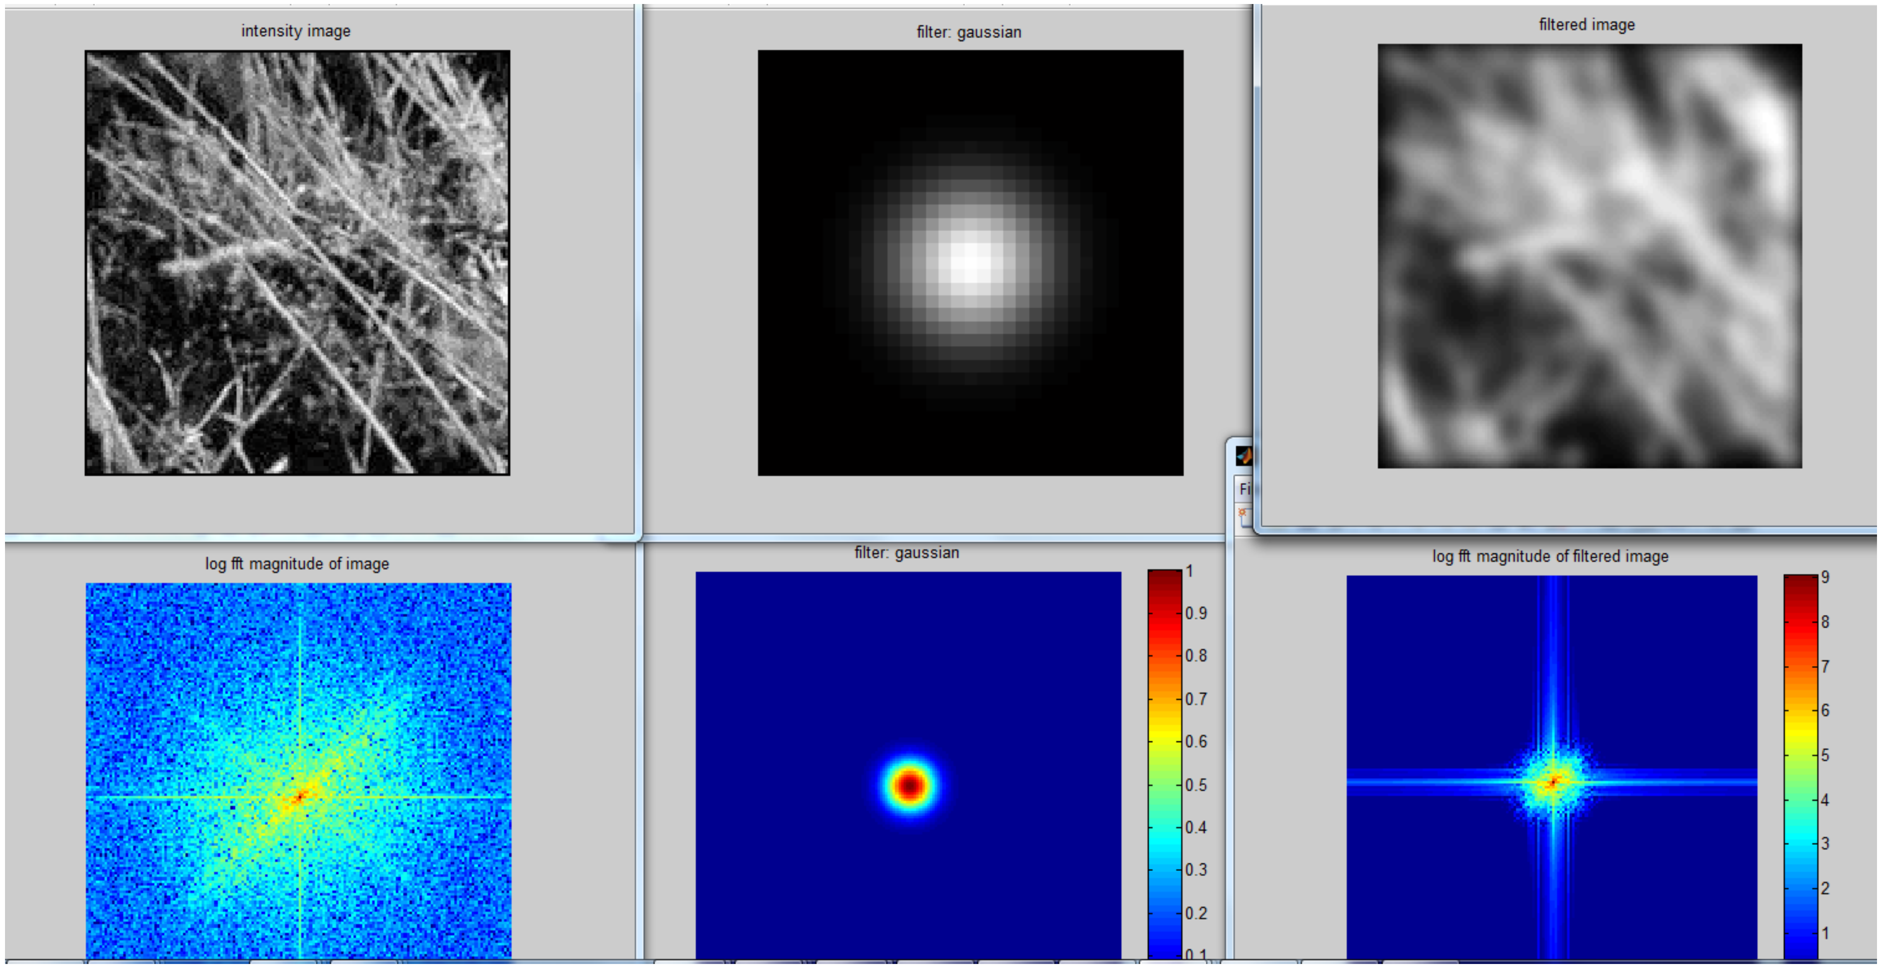

Let’s experiment on Novak

Convolution

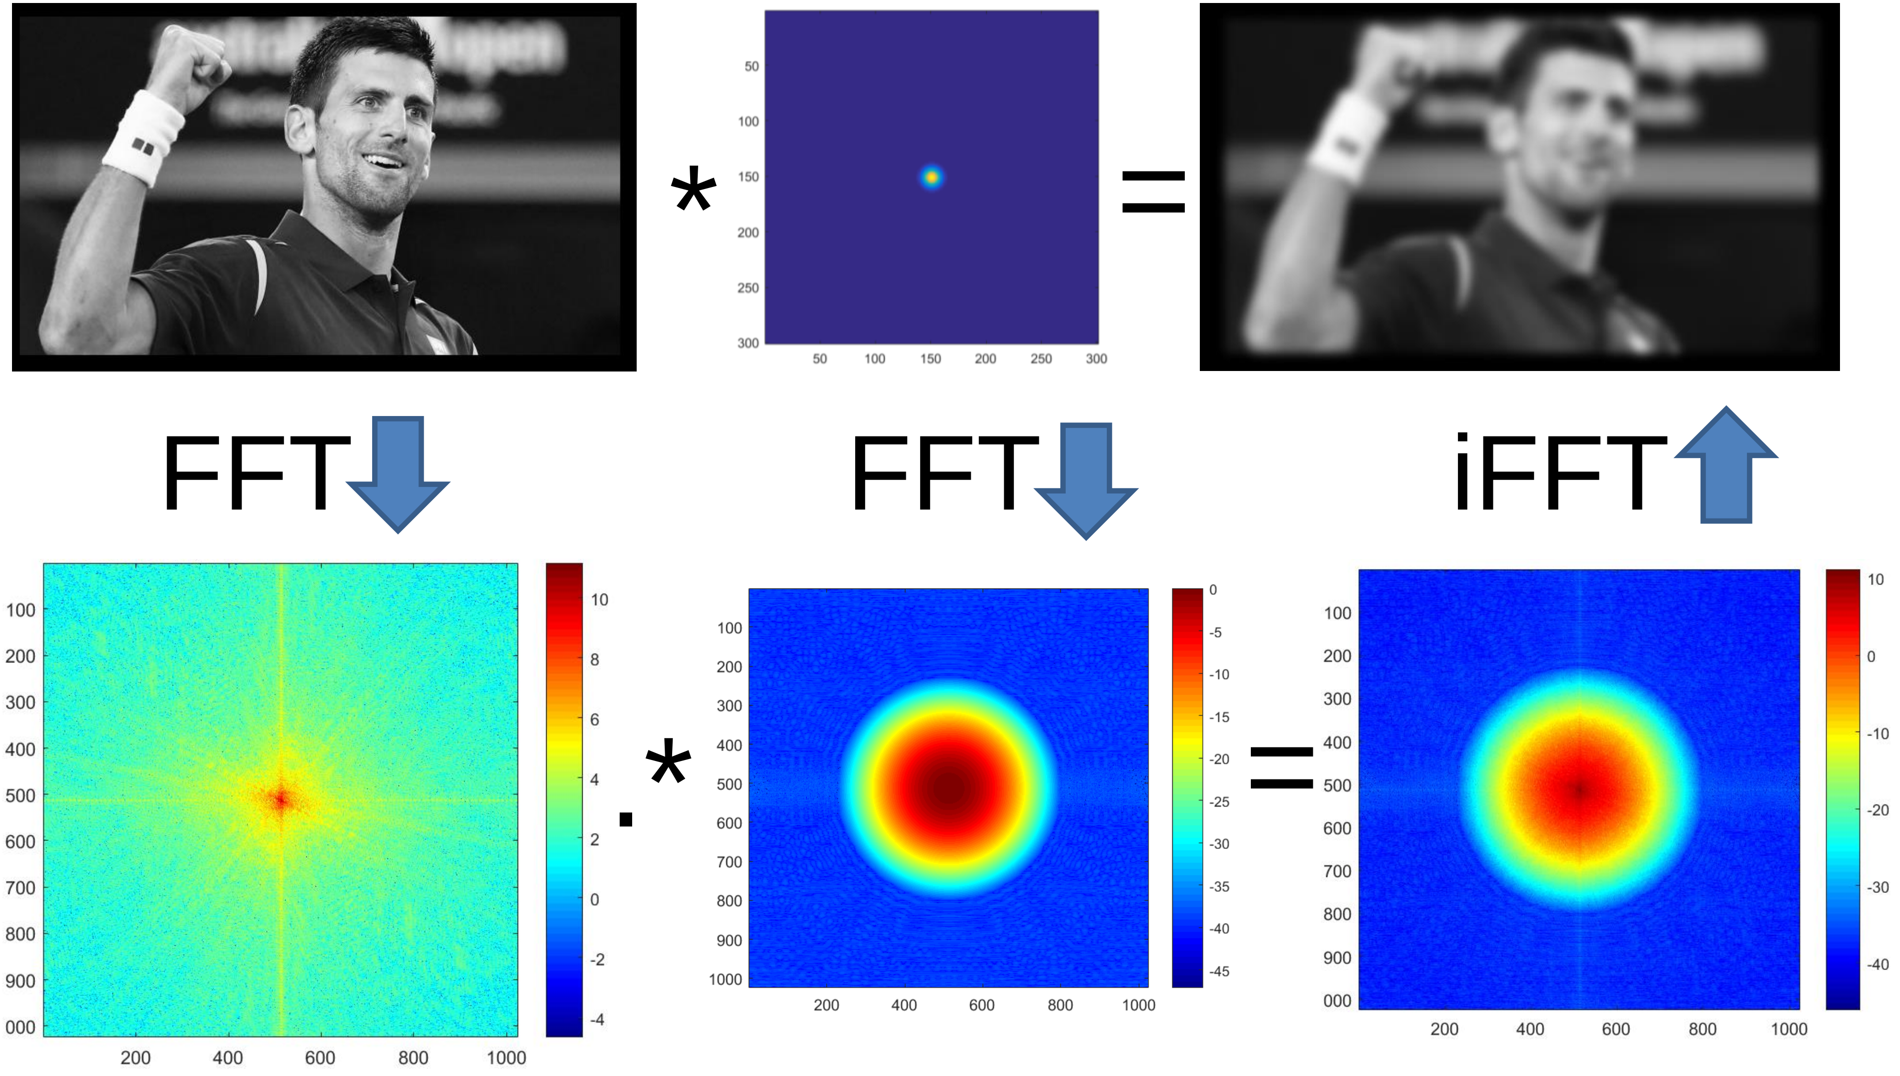

Deconvolution?

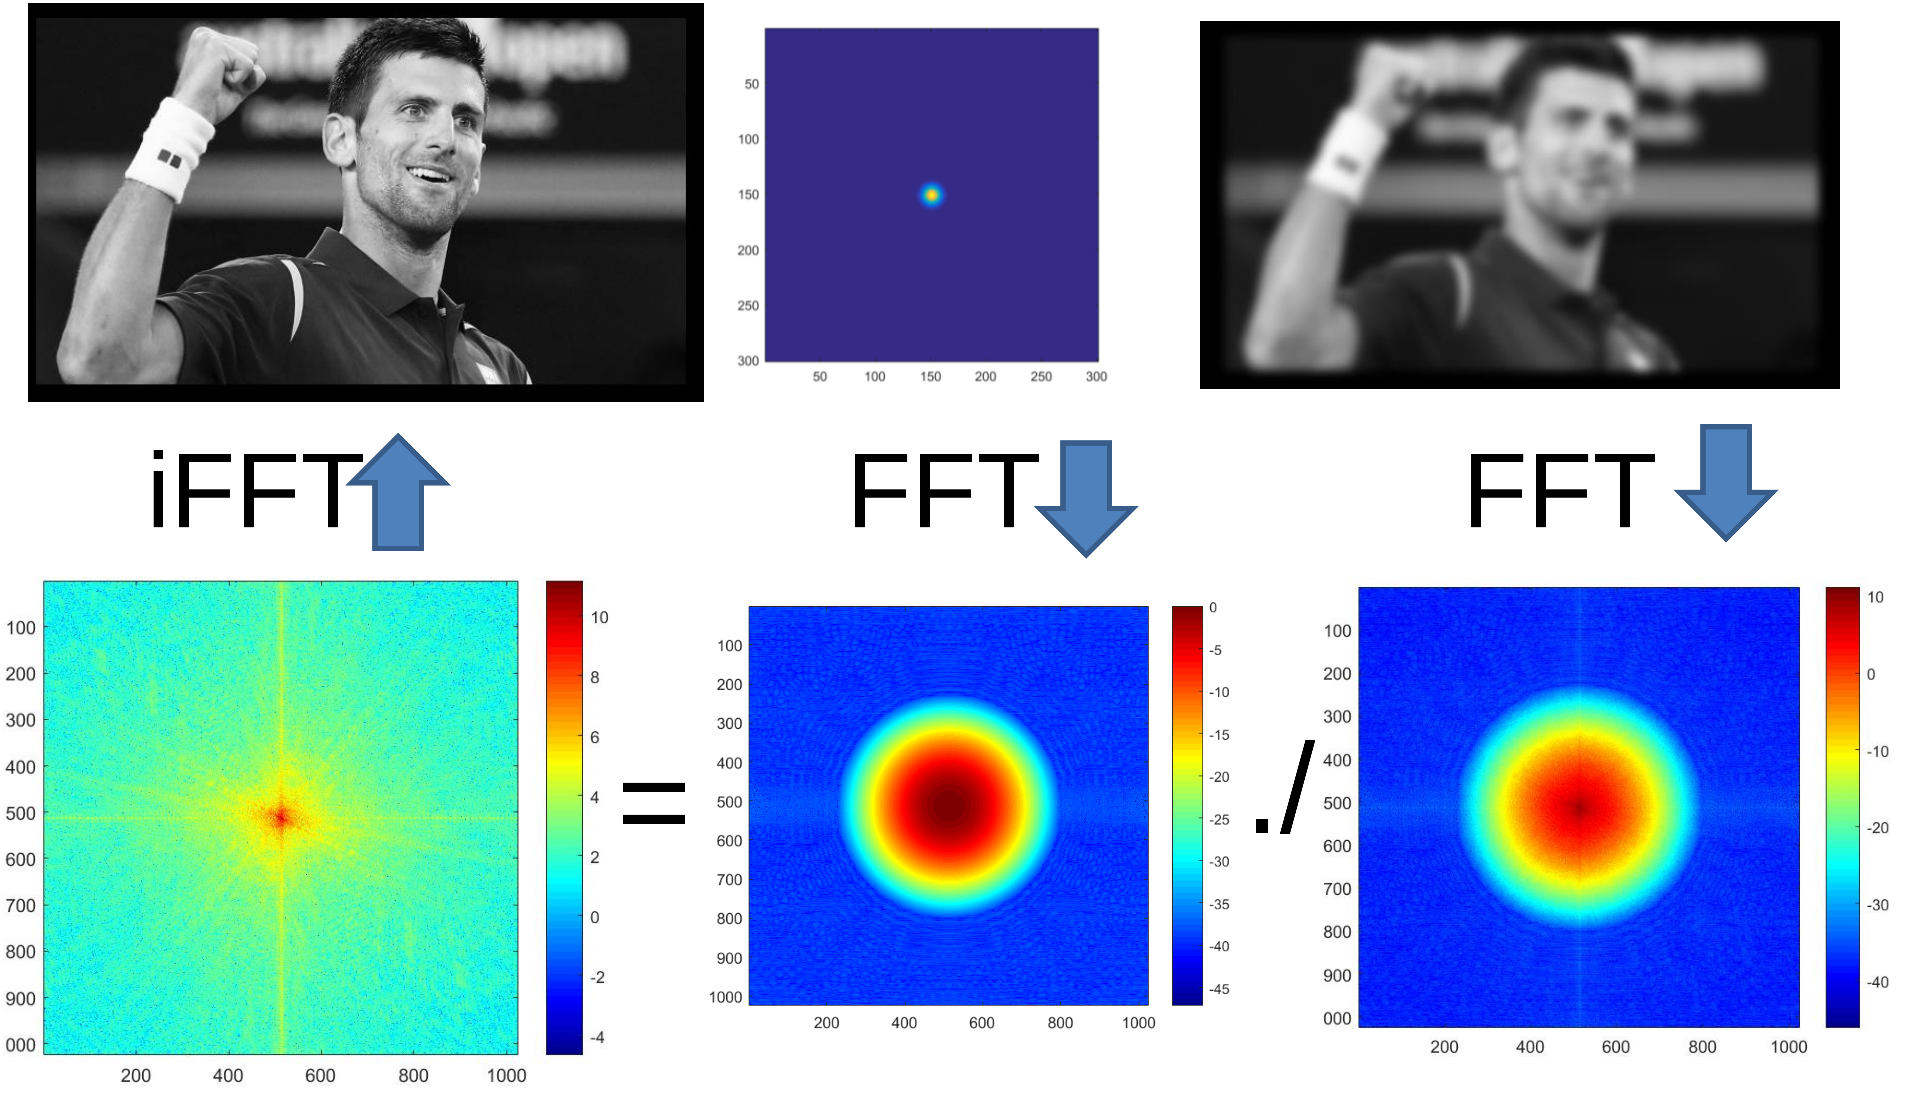

But under more realistic conditions

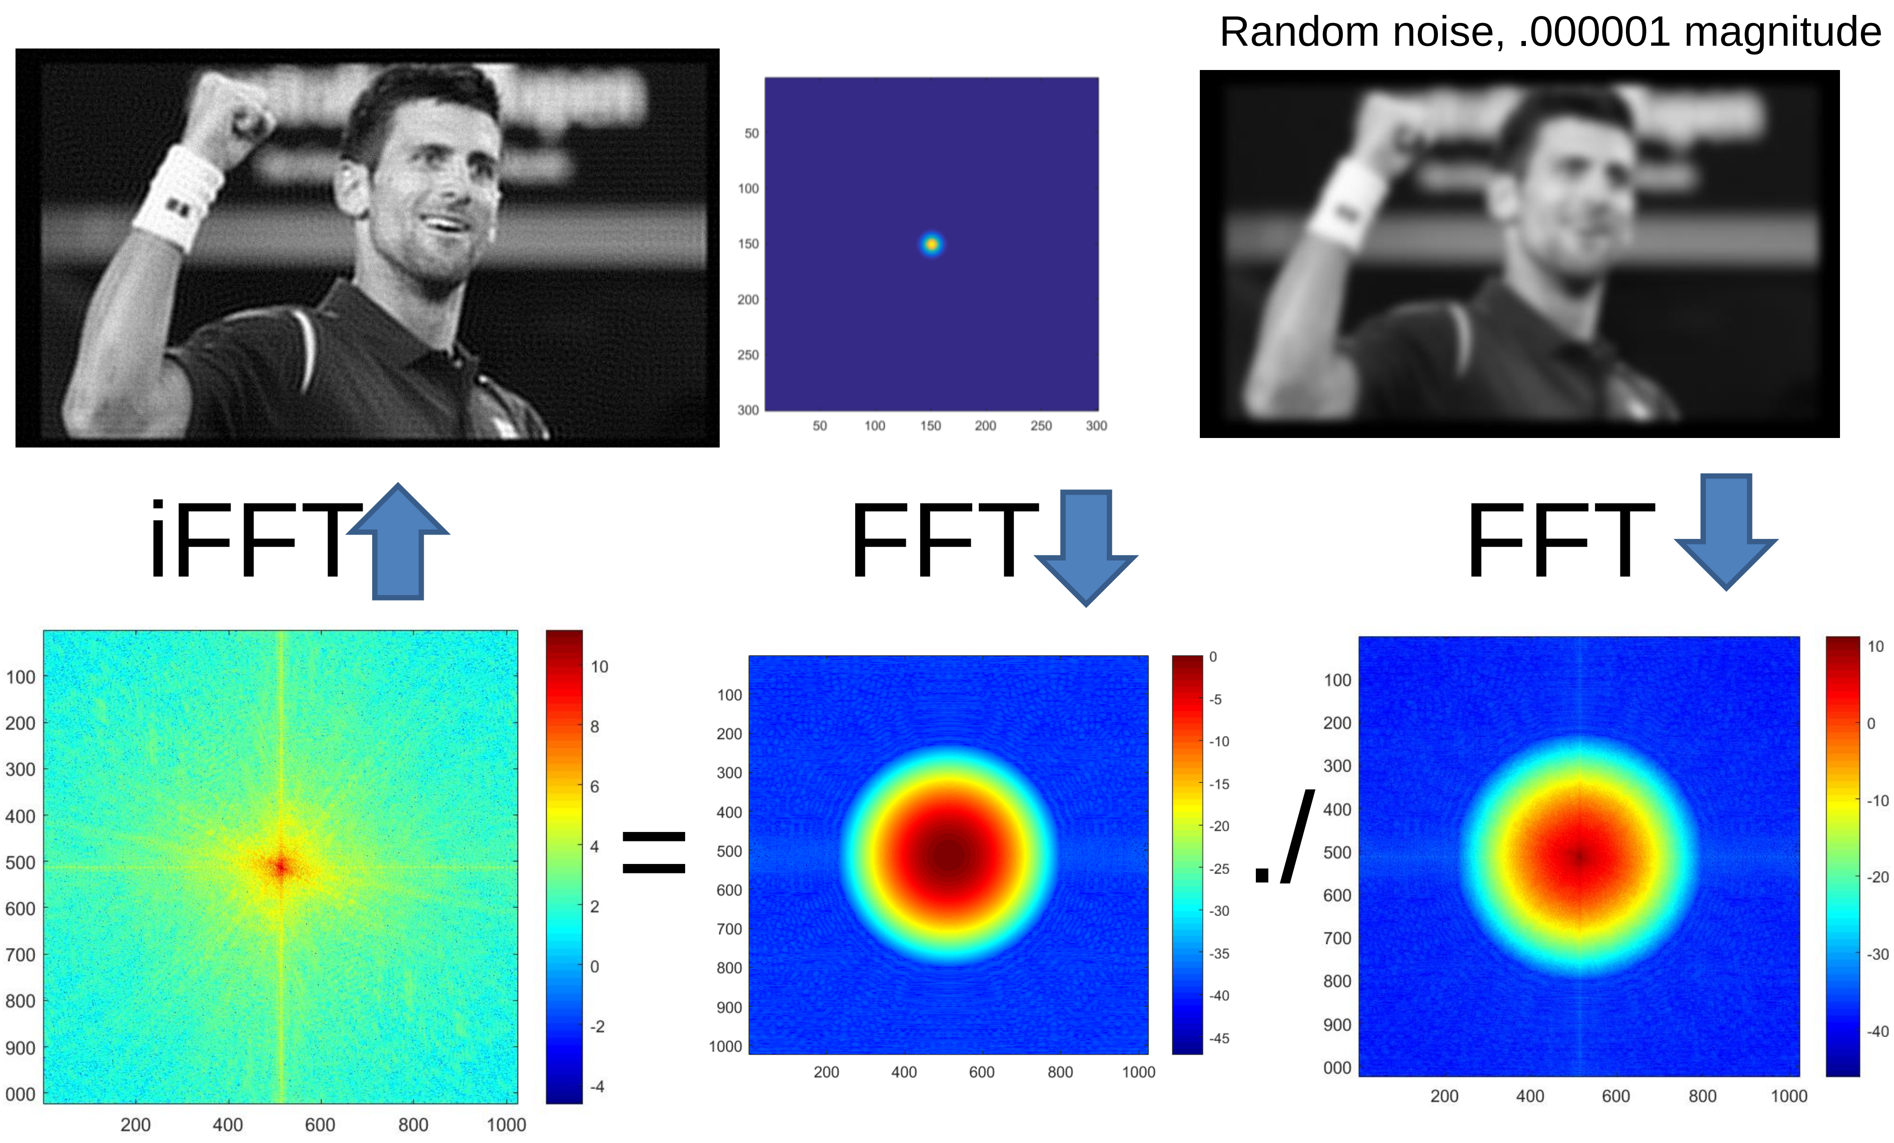

But under more realistic conditions

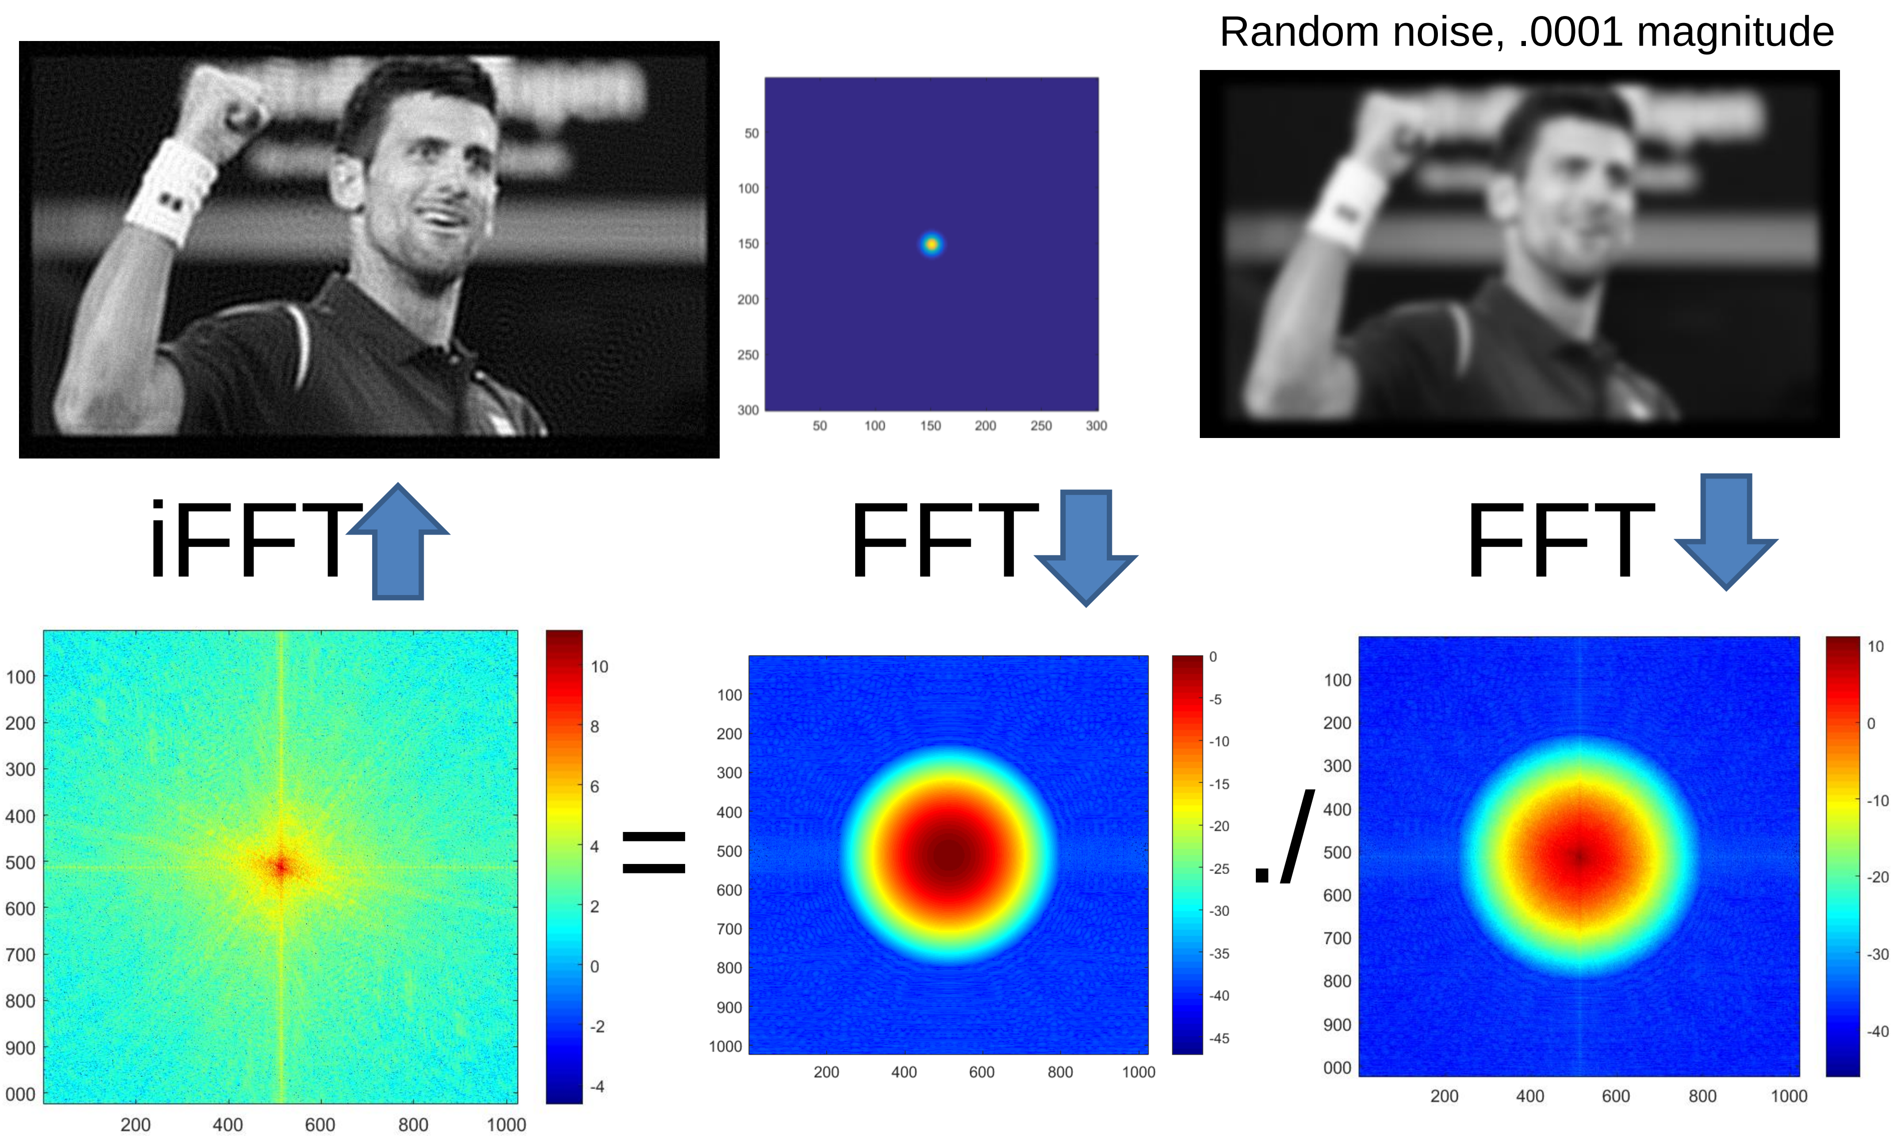

But under more realistic conditions

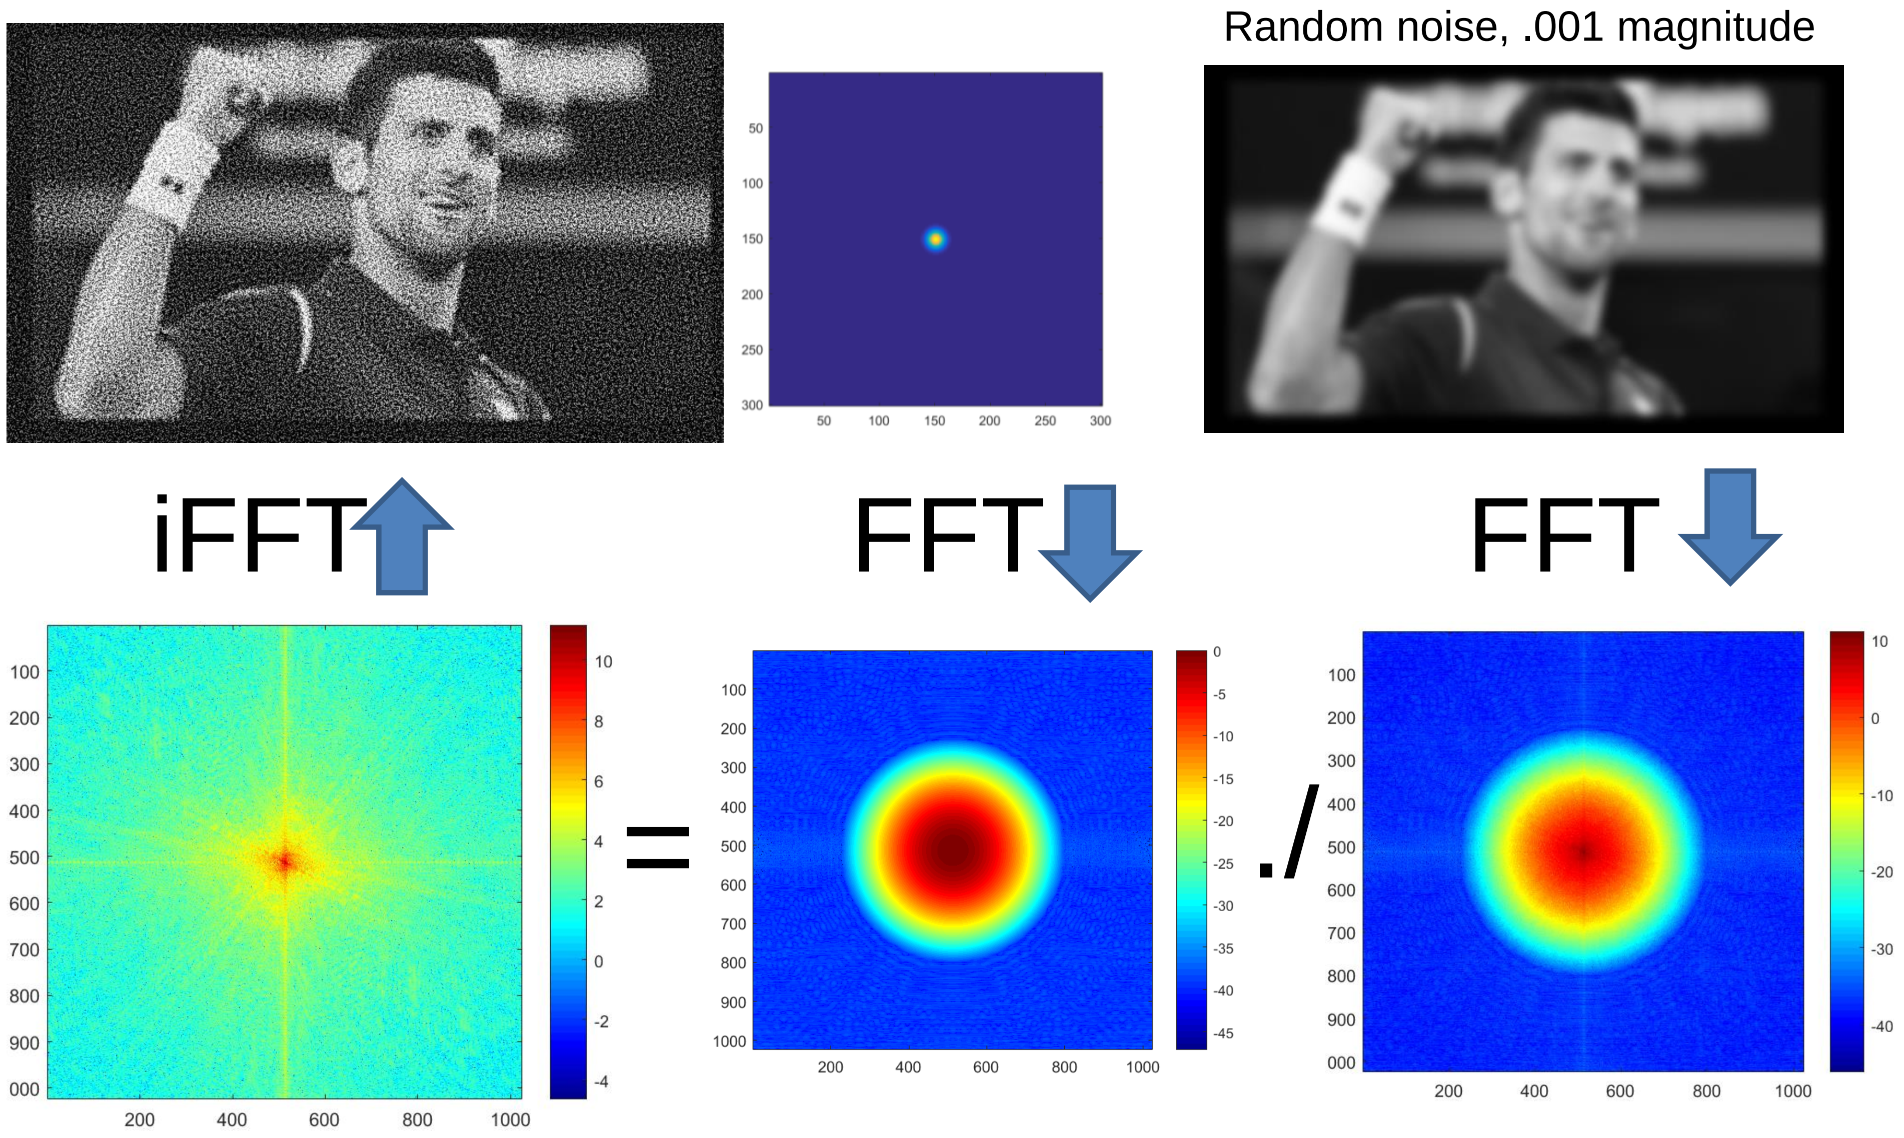

With a random filter…

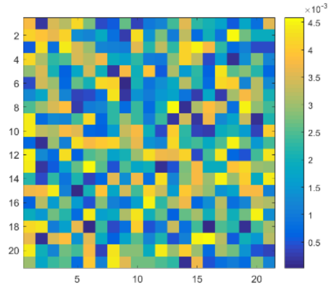

Random noise, .001 magnitude

iFFT

FFT

FFT

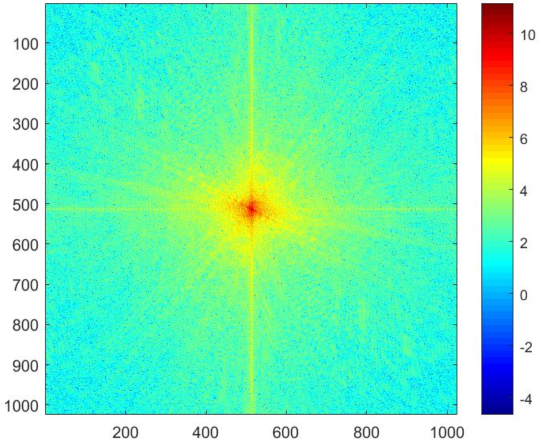

=

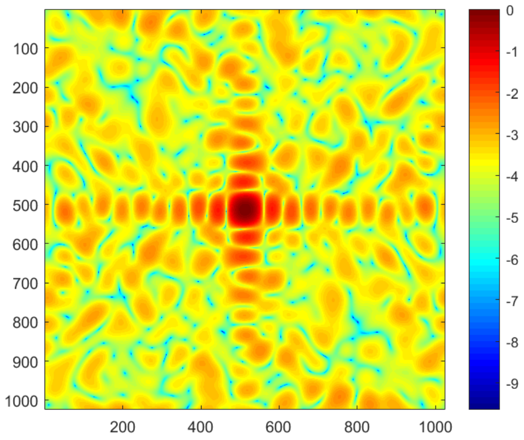

./

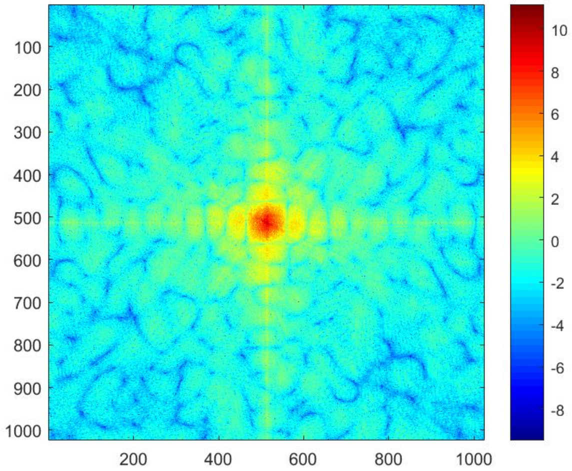

Deconvolution is hard

- Active research area.
- Even if you know the filter (non-blind deconvolution), it is still very hard and requires strong *regularization.*
- If you don’t know the filter (blind deconvolution) it is harder still.

Blind Deconvolution Example

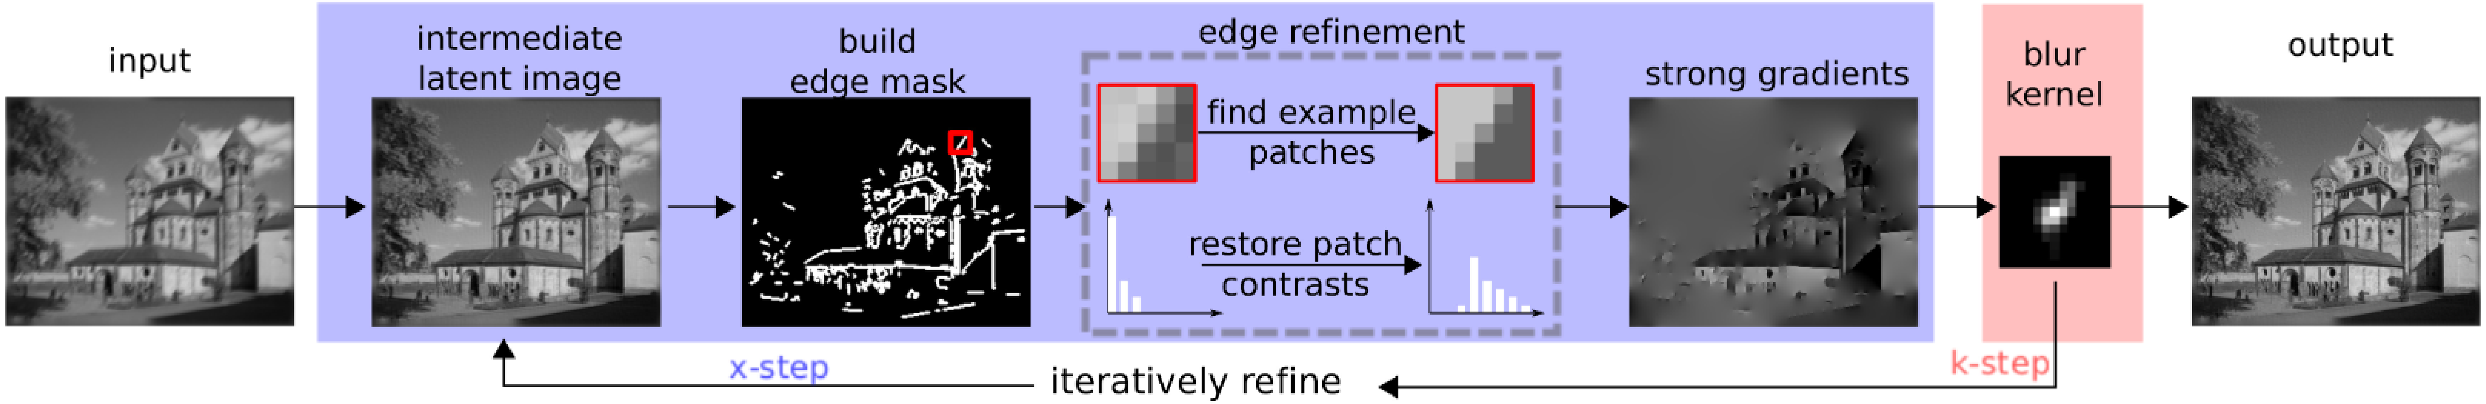

Figure 1. Algorithm pipeline. Our algorithm iterates between x-step and k-step with the help of a patch prior for edge refinement process. In particular, we coerce edges to become sharp and increase local contrast for edge patches. The blur kernel is then updated using the strong gradients from the restored latent image. After kernel estimation, the method of [20] is used for final non-blind deconvolution.

Edge-based Blur Kernel Estimation Using Patch Priors.
Libin Sun, Sunghyun Cho, Jue Wang, and James Hays.
IEEE International Conference on Computational Photography 2013.

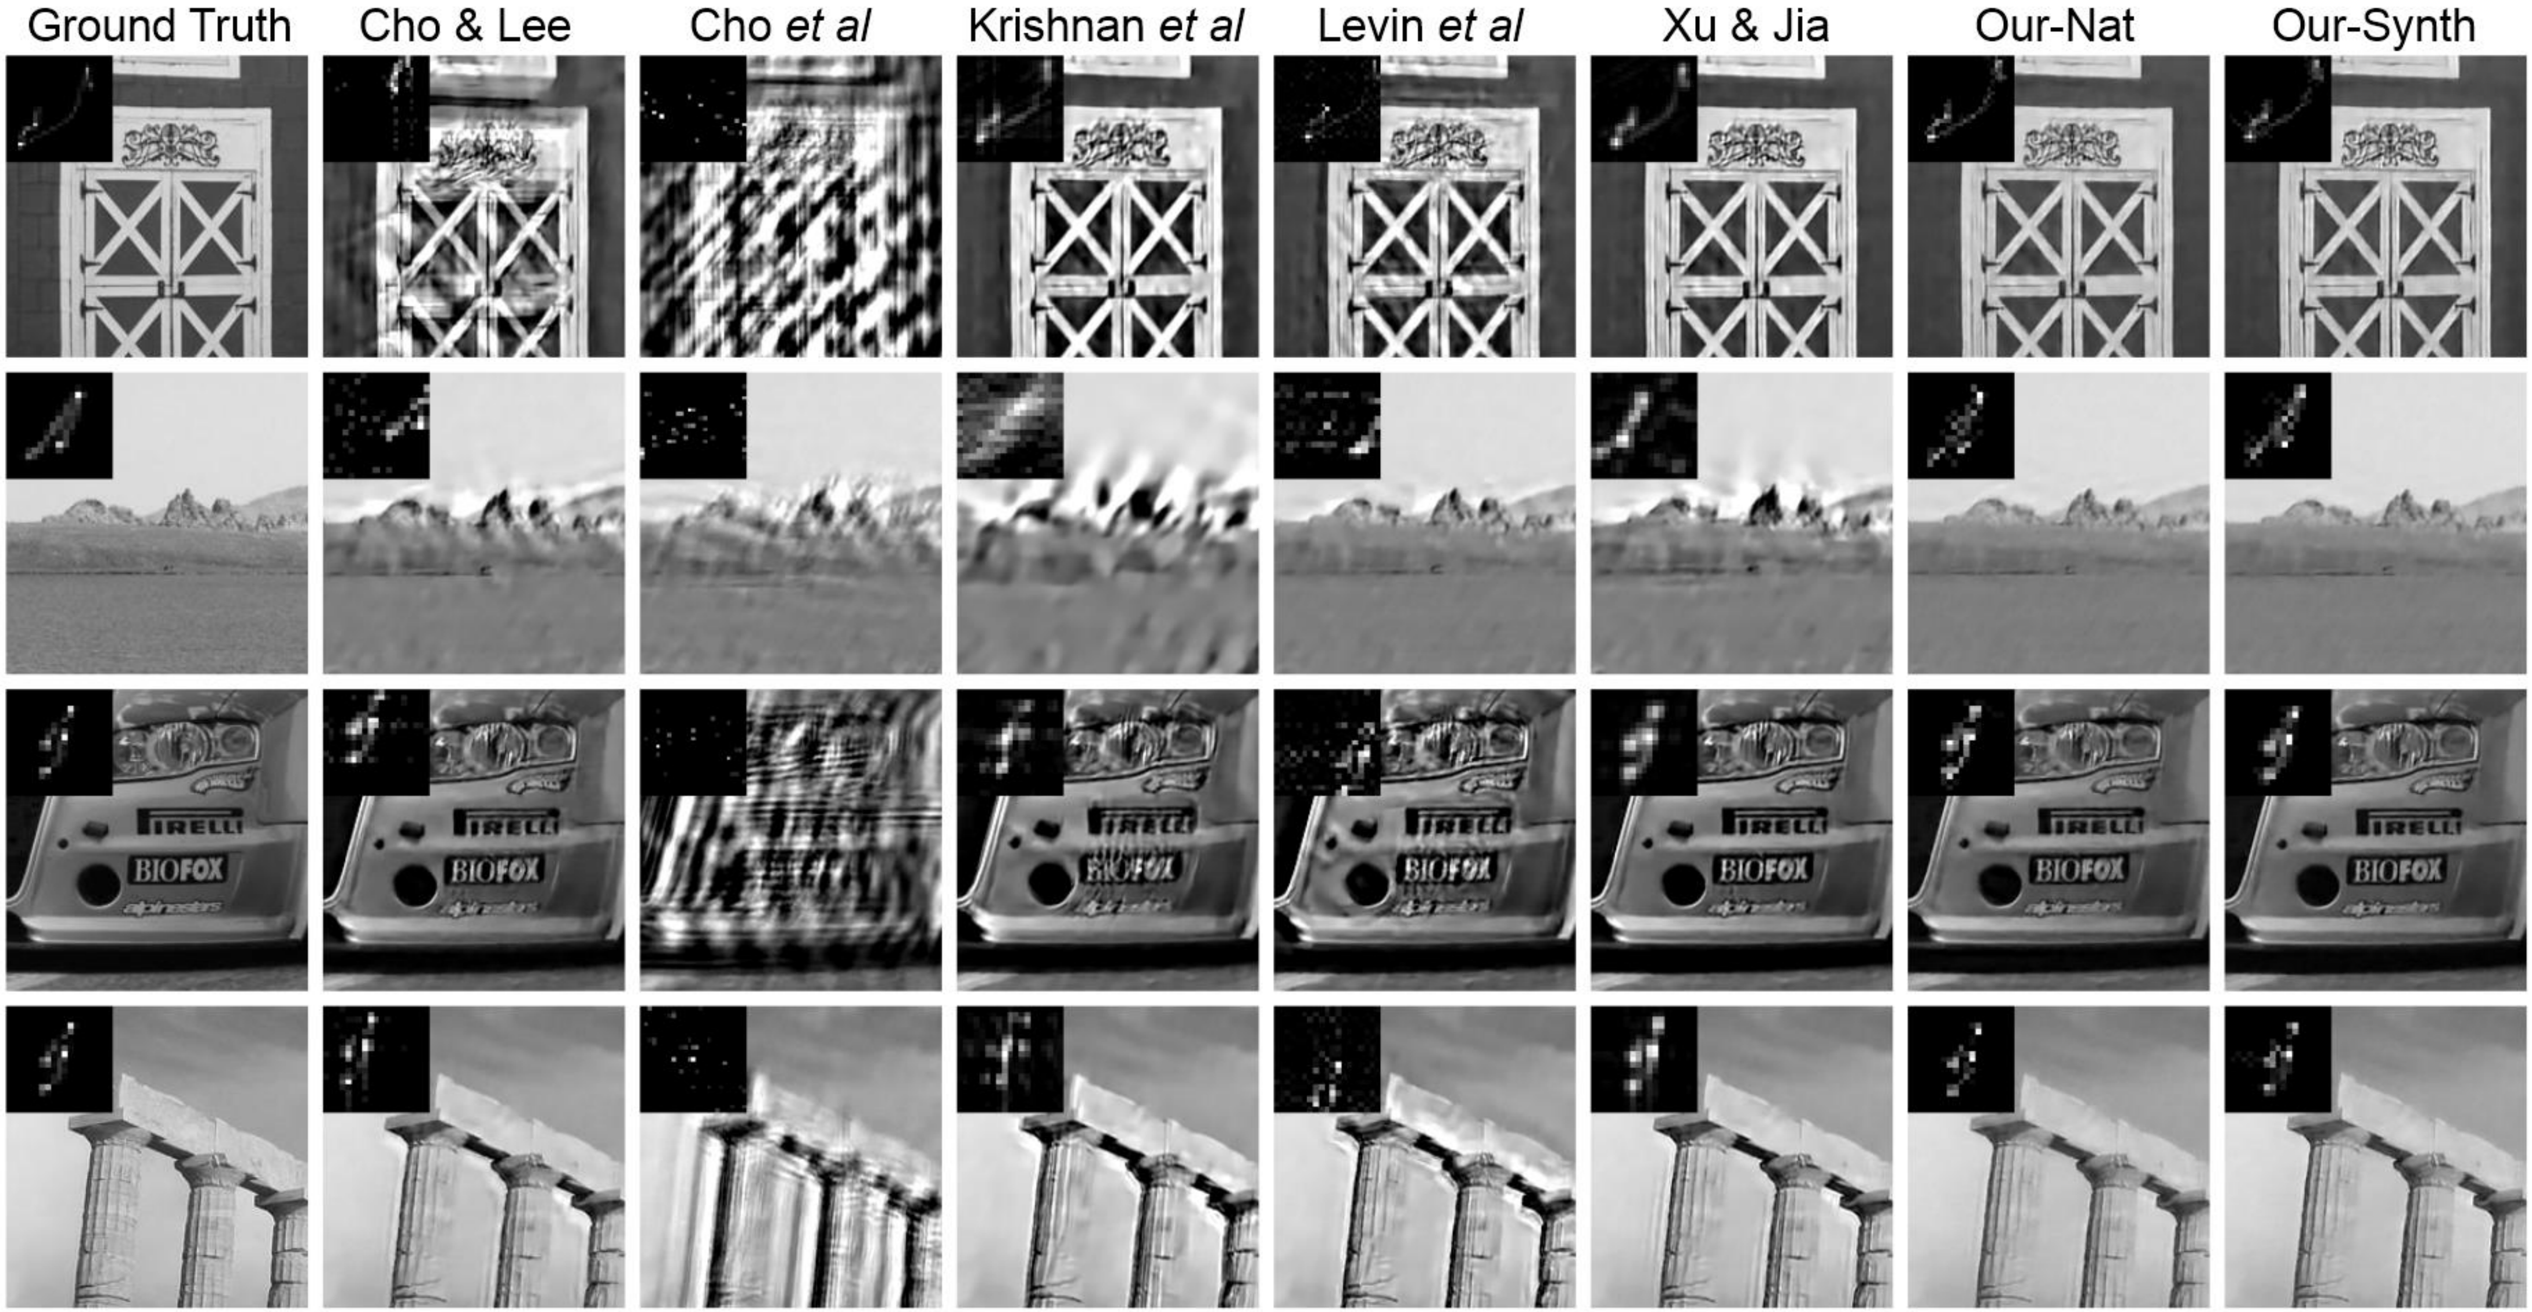

Edge-based Blur Kernel Estimation Using Patch Priors.
Libin Sun, Sunghyun Cho, Jue Wang, and James Hays.
IEEE International Conference on Computational Photography 2013.

Canvas Quiz

Fill in the blanks:

Filtering Operator

1) _ = D * B
2) A = _ * C
3) F = D * _
4) _ = D * D

A

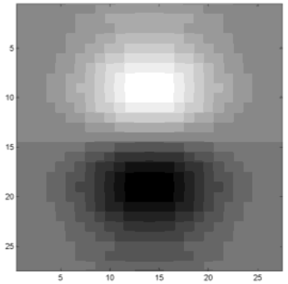

B

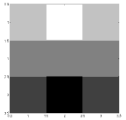

C

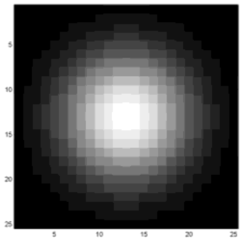

E

F

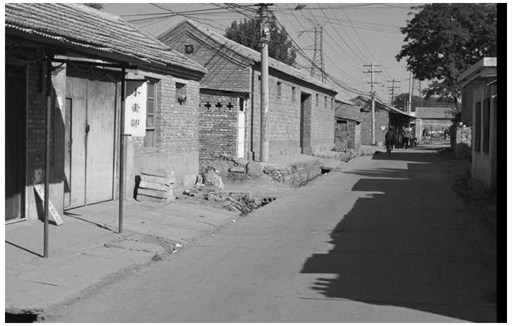

G

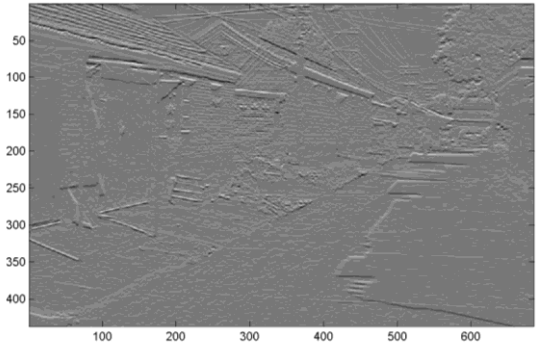

H

I

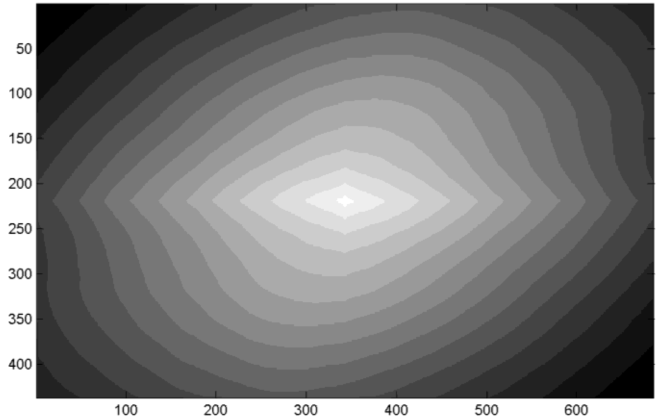

D

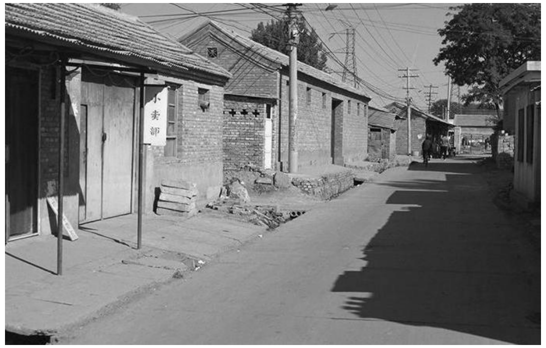

Slide: Hoiem

Project 1: Convolution and Hybrid Images

CS 6476

Spring 2025

Logistics

- Due: Check Canvas for up to date information.
- Project materials including report template: Project 1
- Hand-in: Gradescope
- Required files: `<your_gt_username>.zip`, `<your_gt_username>_project1.pdf`

Figure 1: Look at the image from very close, then from far away.

Next Topics

- Image Formation
- Biological Vision
- Light and Color

From the 3D to 2D

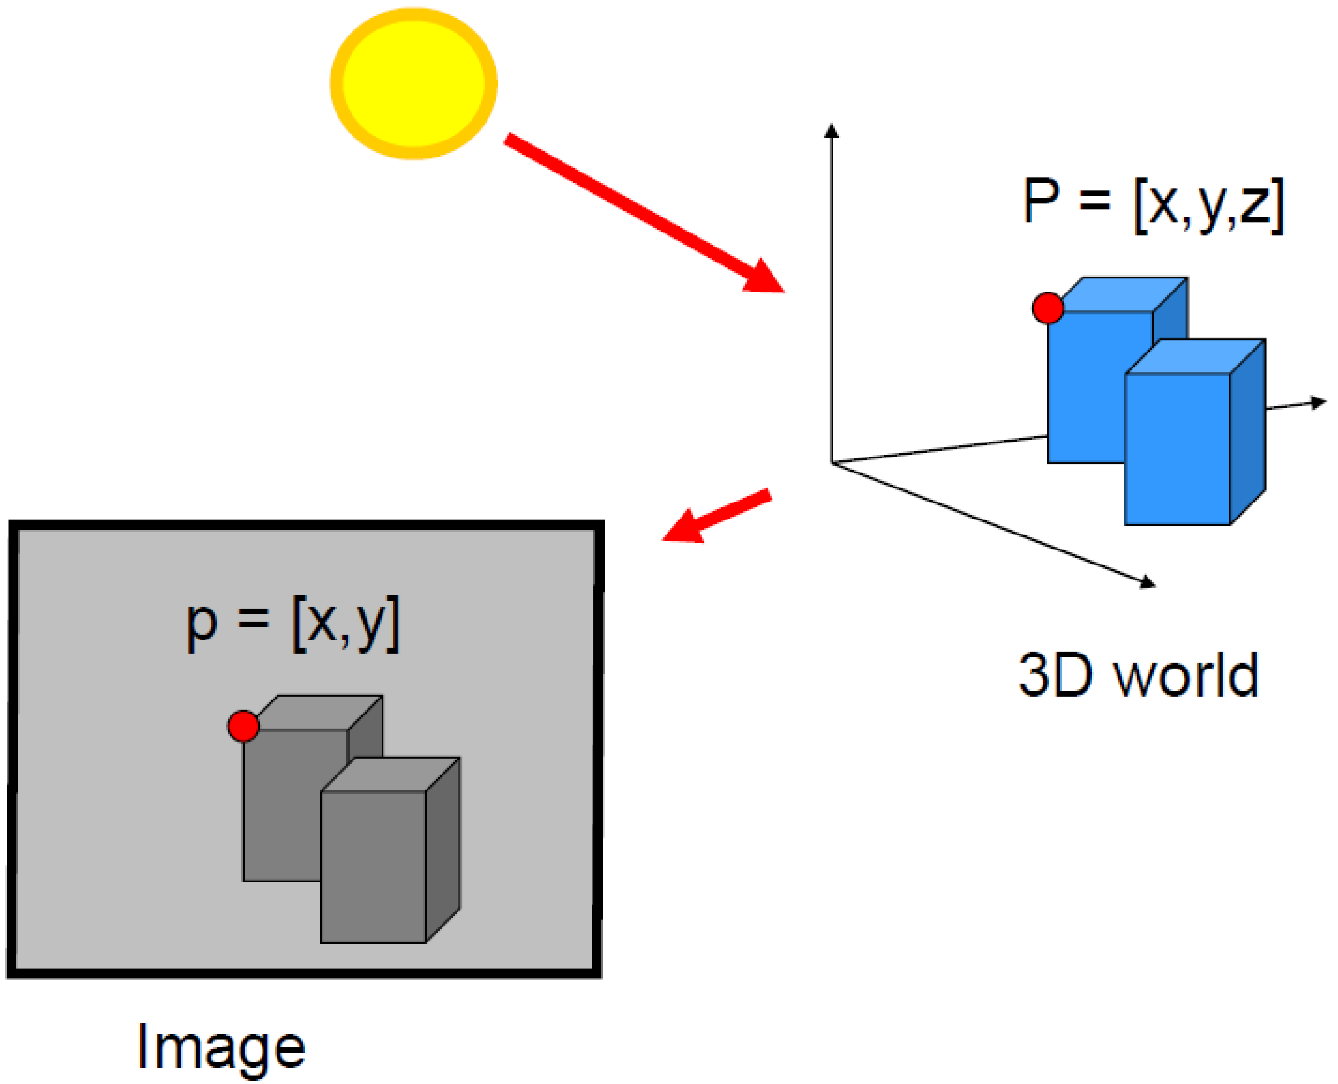

Slide credit Fei Fei Li

Image Formation

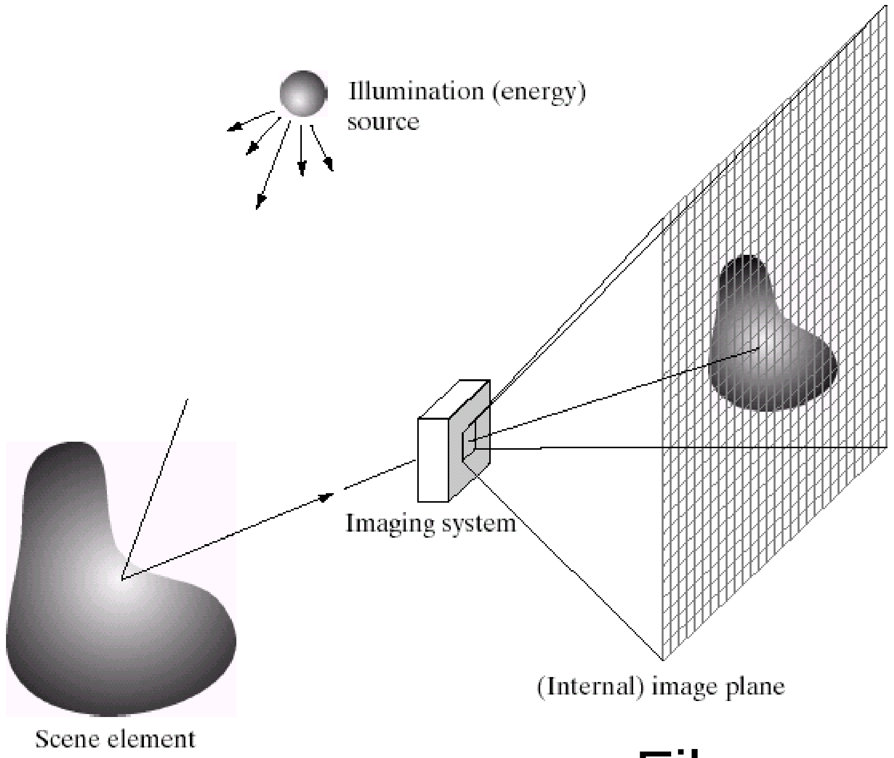

Film

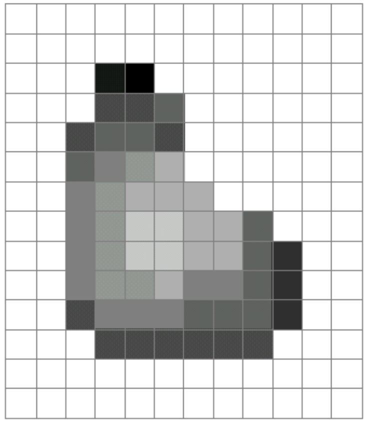

Digital Camera

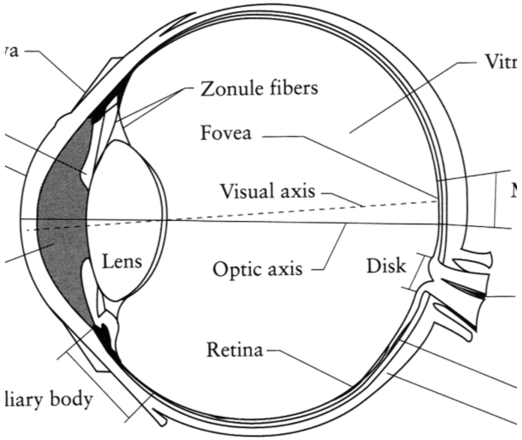

The Eye

Vision is so biologically ubiquitious that it is interesting when it is not present, e.g. the blind cava tetra

Blind cave form [edit]

Blind cave fish form

A. mexicanus is famous for its blind cave form, which is known by such names as **blind cave tetra**, **blind tetra** (leading to easy confusion with the Brazilian *Stygichthys typhlops*), **blind cave characin** and **blind cavefish**. Depending on the exact population, cave forms can have degenerated sight or have total loss of sight and even their eyes, due to down-regulation of the protein αA-crystallin and consequent lens cell death.[15] Despite losing their eyes, cavefish cells respond to light responsive and show an endogenous circadian rhythm.[16] During the start of development, larvae still exhibit a shadow response which is controlled by the pineal eye.[17] The fish in the Pachón caves have lost their eyes completely whilst the fish from the Micos cave only have limited sight.[18] Cave fish and surface fish are able to produce fertile offspring.[18]

These fish can still, however, find their way around by means of their lateral lines, which are highly sensitive to fluctuating water pressure.[19] Blindness in A. *mexicanus* induces a disruption of early neuromast patterning, which further causes asymmetries in cranial bone structure. One such asymmetry is a bend in the dorsal region of their skull, which is propounded to increase water flow to the opposite side of the face, functionally enhancing sensory input and spatial mapping in the dark waters of caves.[20] Scientists suggest that gene cystathionine beta synthase-a mutation restricts blood flow to cavefish eyes during a critical stage of growth so the eyes are covered by skin.[21]

A photon's life choices

- Absorption
- Diffusion
- Reflection
- Transparency
- Refraction
- Fluorescence
- Subsurface scattering
- Phosphorescence

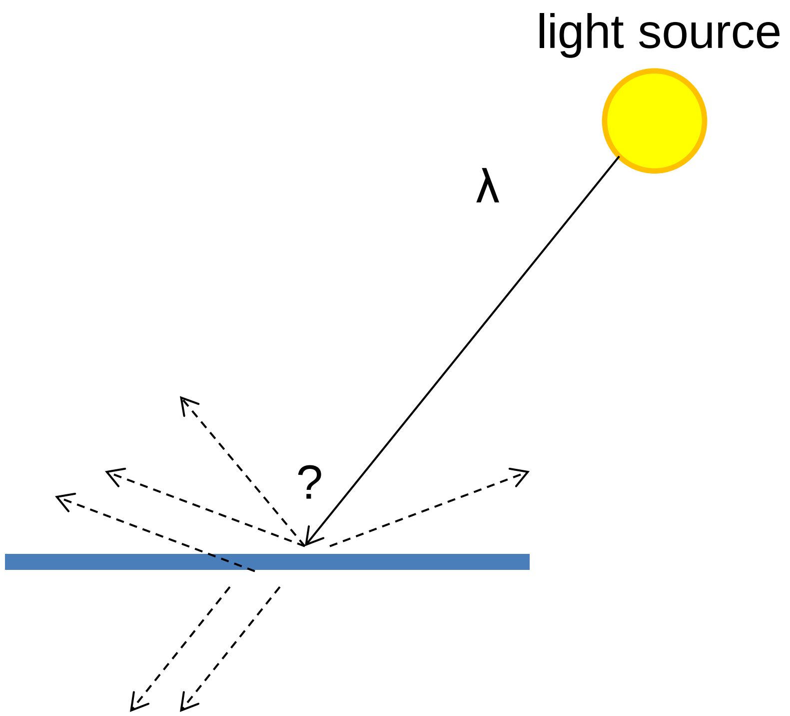

A photon's life choices

- **Absorption**
- Diffusion
- Reflection
- Transparency
- Refraction
- Fluorescence
- Subsurface scattering
- Phosphorescence

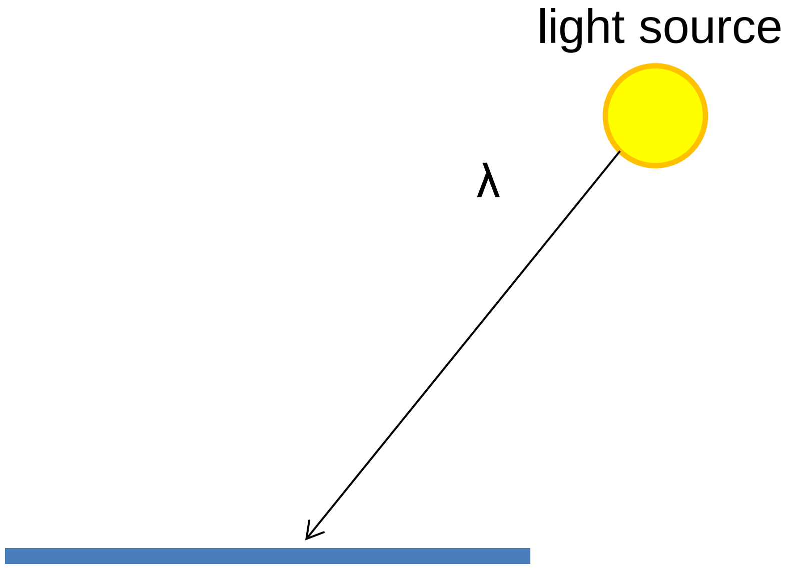

A photon’s life choices

- Absorption
- **Diffuse Reflection (e.g. matte surface)**
- Reflection
- Transparency
- Refraction
- Fluorescence
- Subsurface scattering
- Phosphorescence

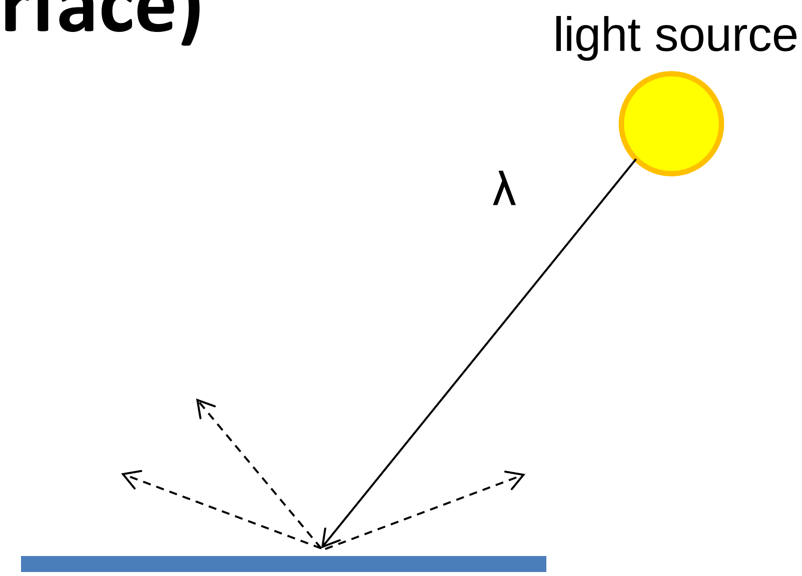

A photon’s life choices

- Absorption
- Diffusion
- **Reflection (e.g. specular surface)**
- Transparency
- Refraction
- Fluorescence
- Subsurface scattering
- Phosphorescence

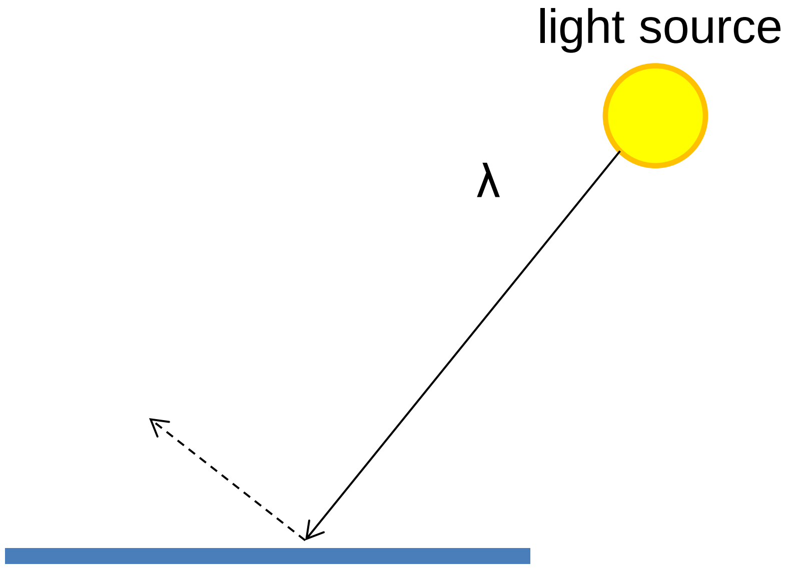

A photon’s life choices

- Absorption
- Diffusion
- Reflection
- **Transparency**
- Refraction
- Fluorescence
- Subsurface scattering
- Phosphorescence

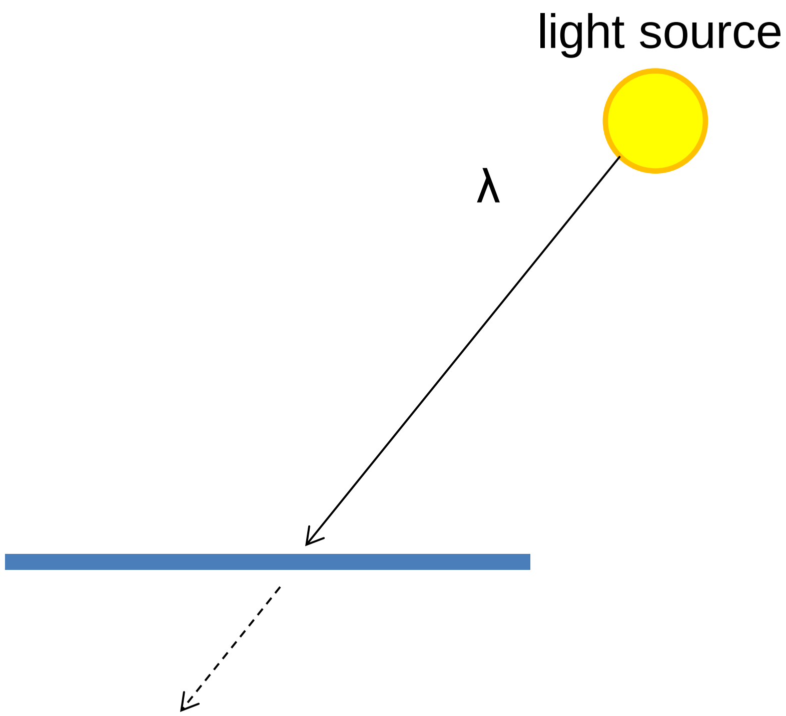

A photon’s life choices

- Absorption
- Diffusion
- Reflection
- Transparency
- **Refraction**
- Fluorescence
- Subsurface scattering
- Phosphorescence

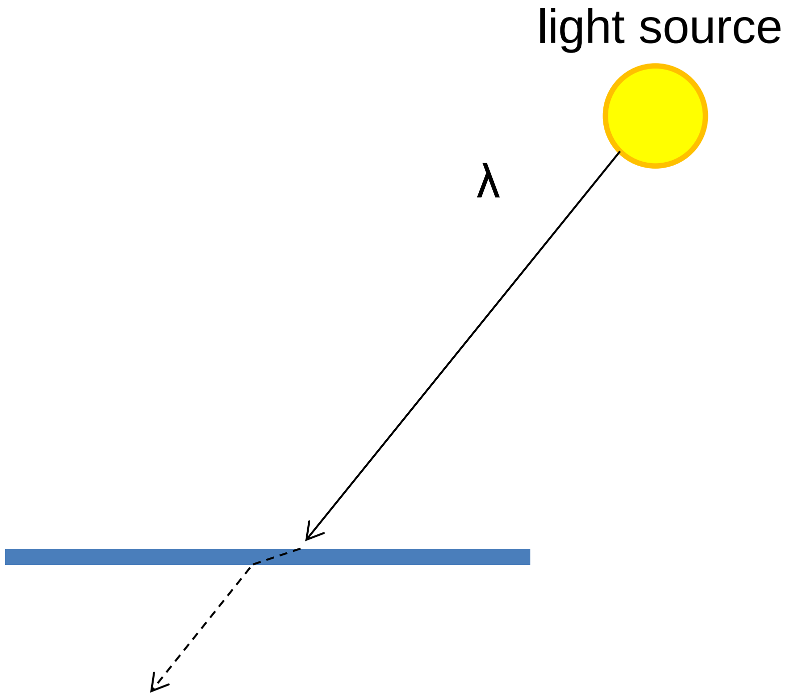

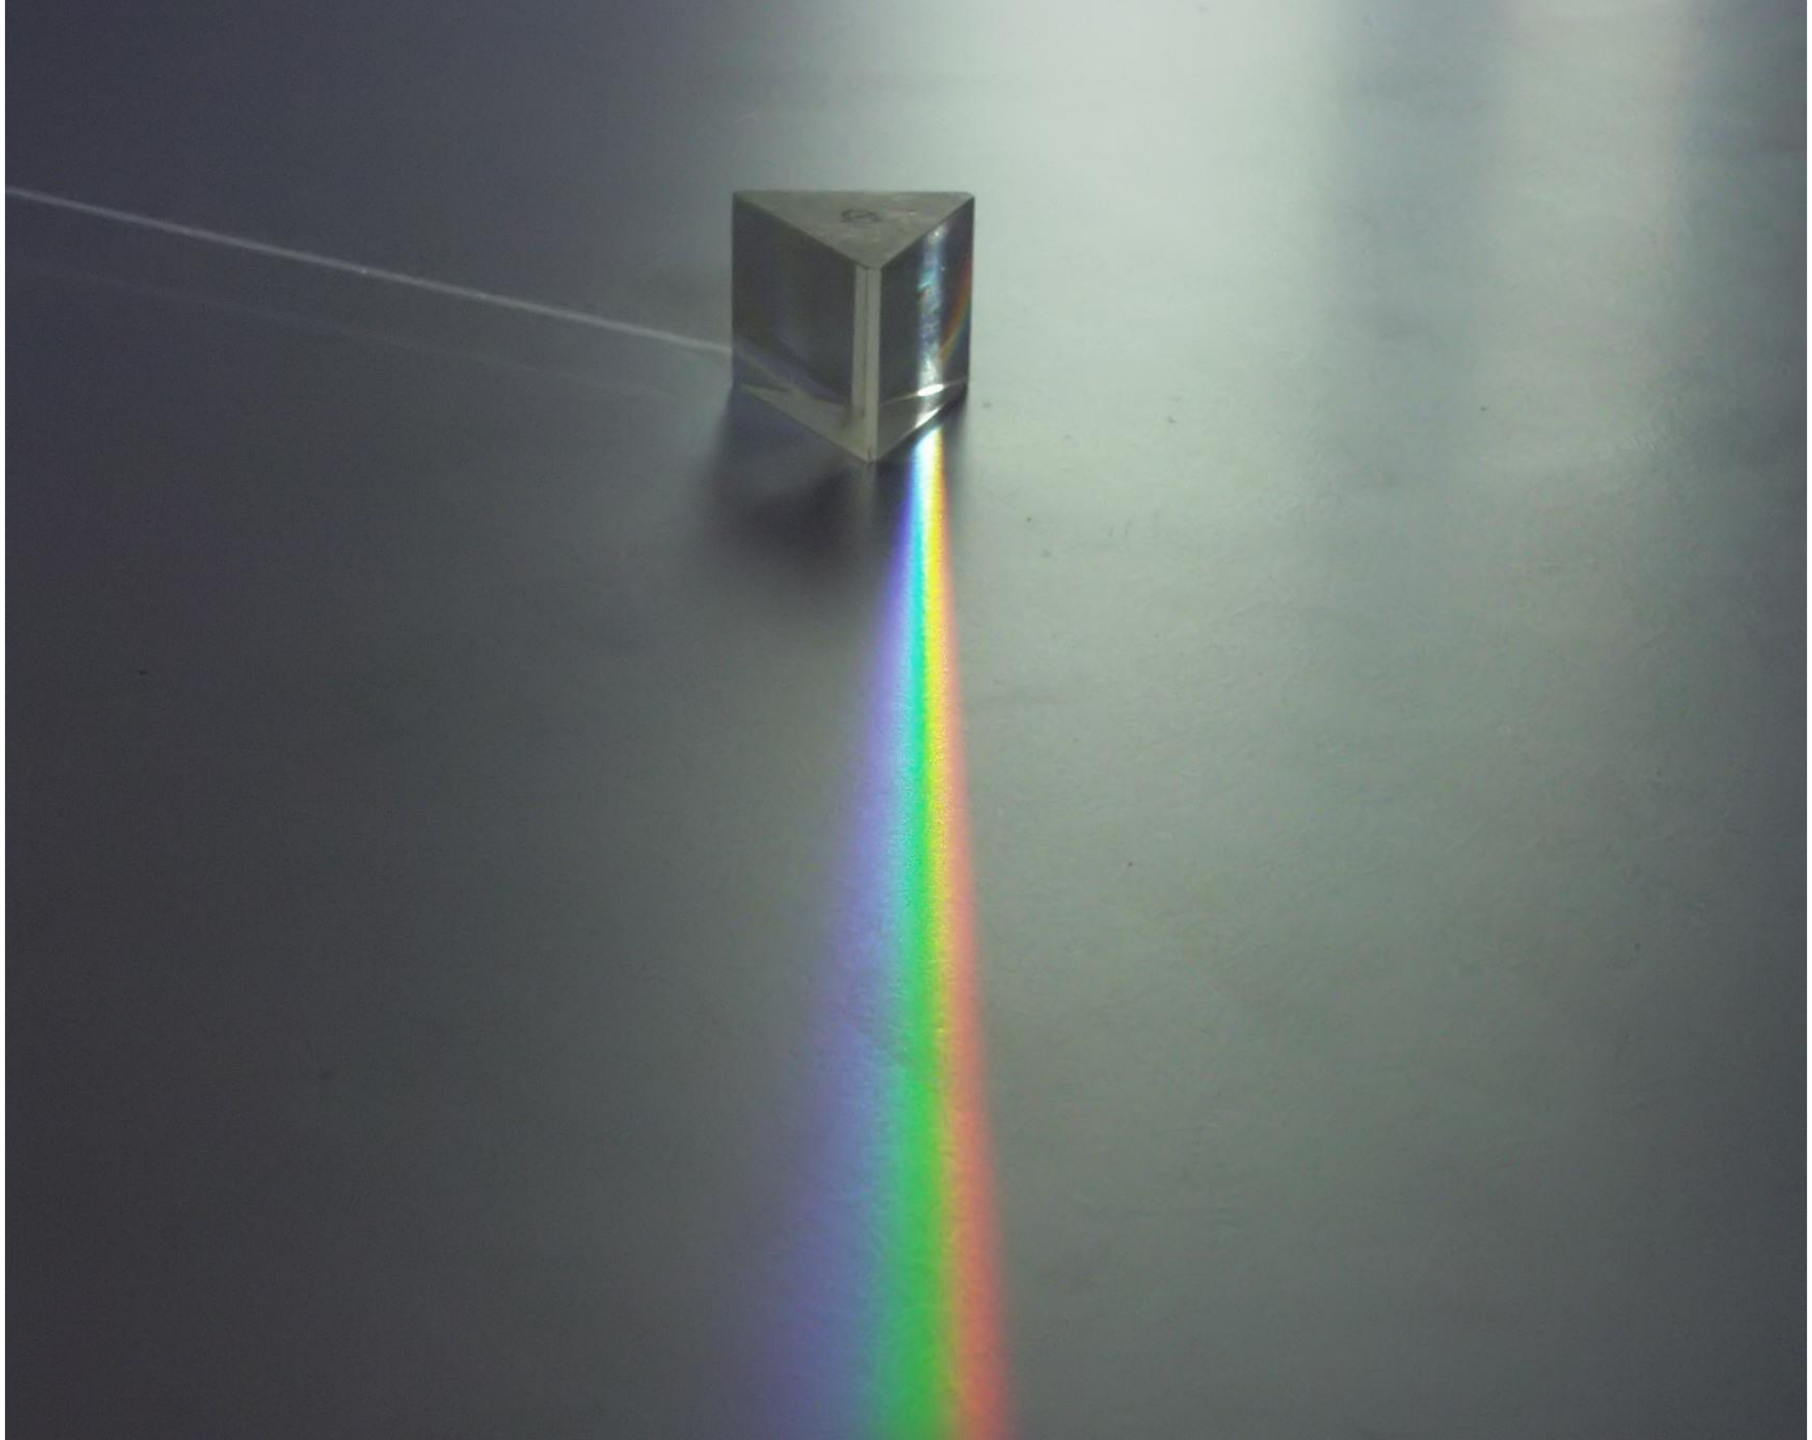

A photon’s life choices

- Absorption
- Diffusion
- Reflection
- Transparency
- Refraction
- **Fluorescence**
- Subsurface scattering
- Phosphorescence

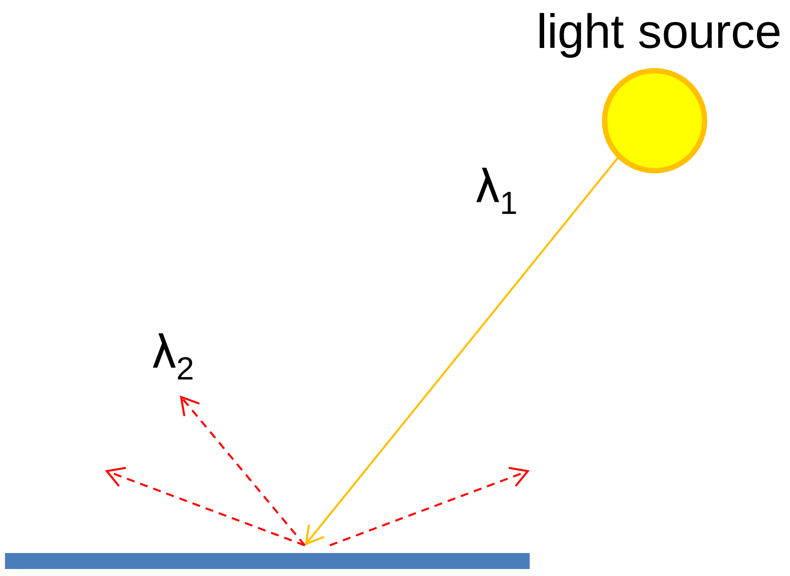

A photon's life choices

- Absorption
- Diffusion
- Reflection
- Transparency
- Refraction
- Fluorescence
- **Subsurface scattering**
- Phosphorescence

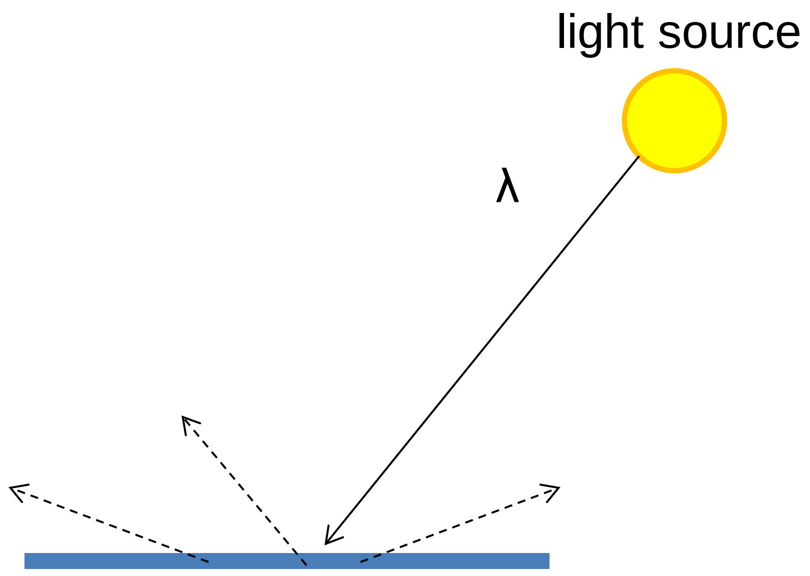

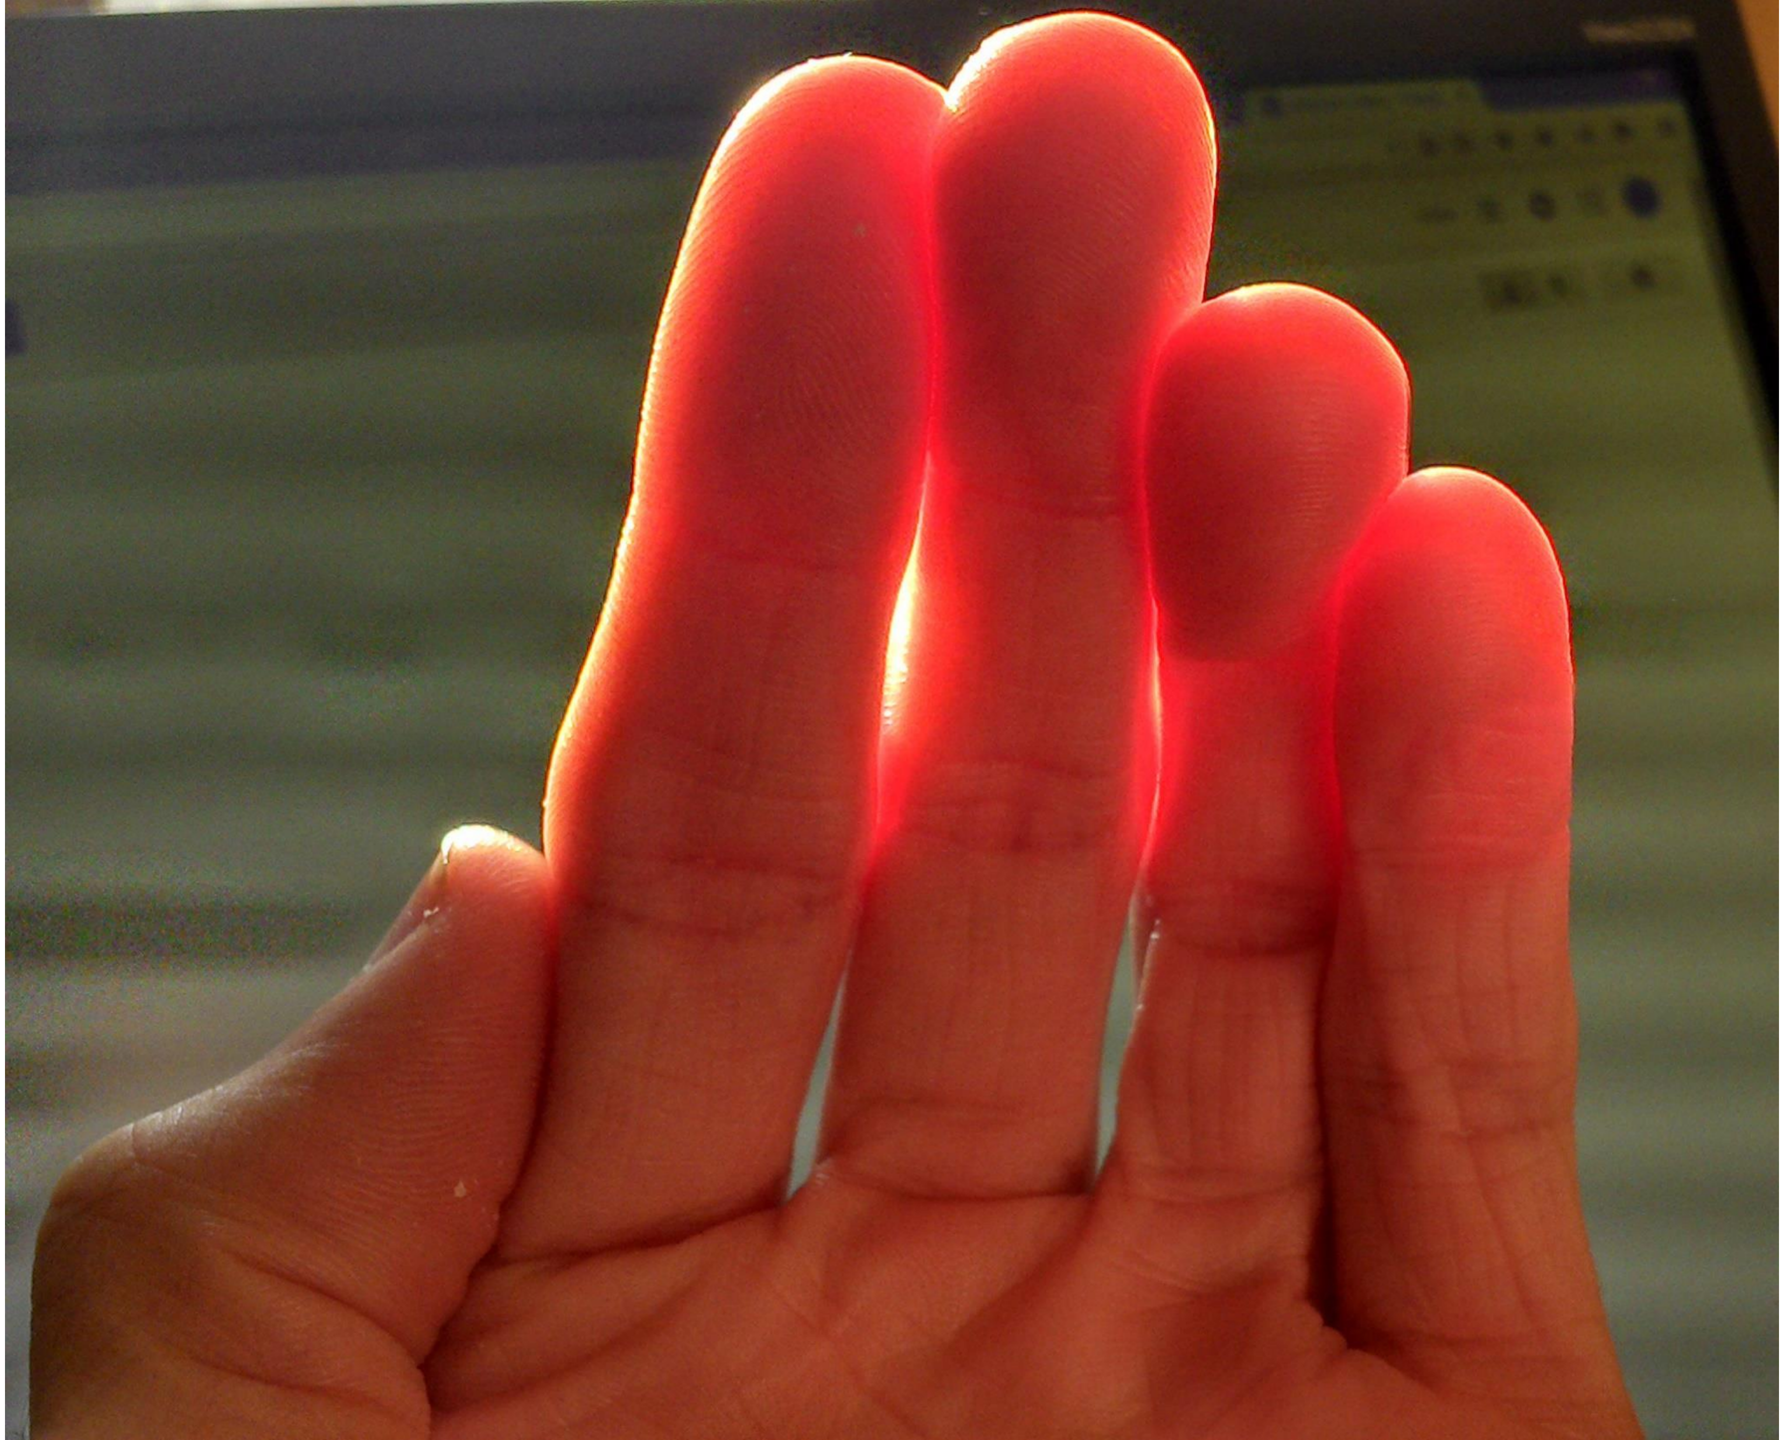

A photon’s life choices

- Absorption
- Diffusion
- Reflection
- Transparency
- Refraction
- Fluorescence
- Subsurface scattering
- **Phosphorescence**

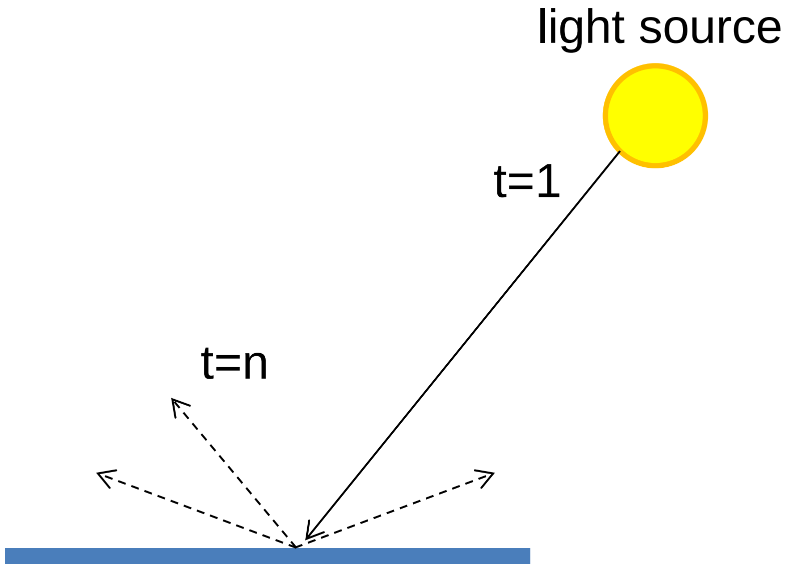


Sleepytime
STICK THESE STARS ON YOUR BEDROOM CEILING AND FALL ASLEEP UNDER THE STARS!

Lambertian Reflectance

- In computer vision, the complexity of light transport is mostly ignored.
- Surfaces are often assumed to be ideal diffuse reflectors with no dependence on viewing direction.

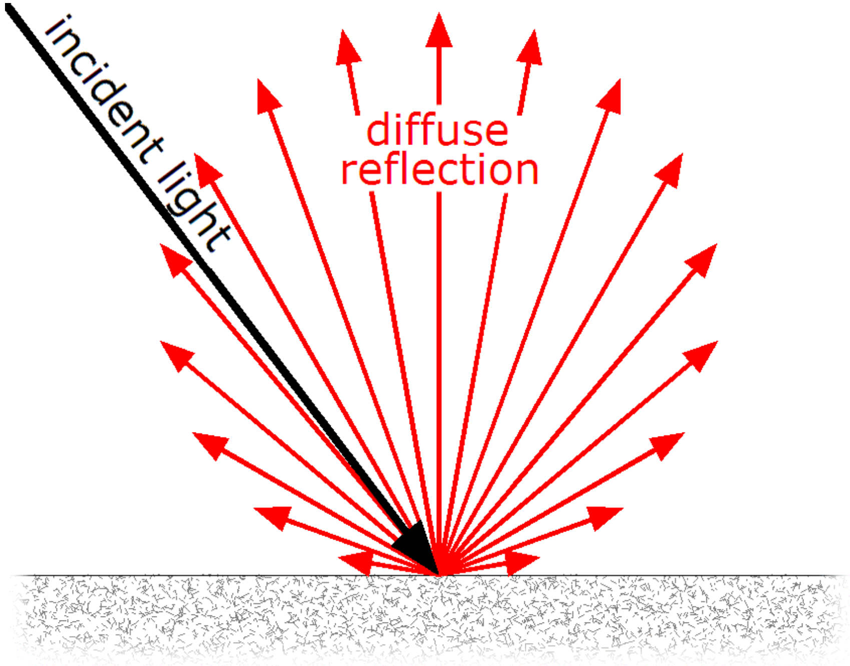

https://en.wikipedia.org/wiki/Lambertian_reflectance

Image Formation

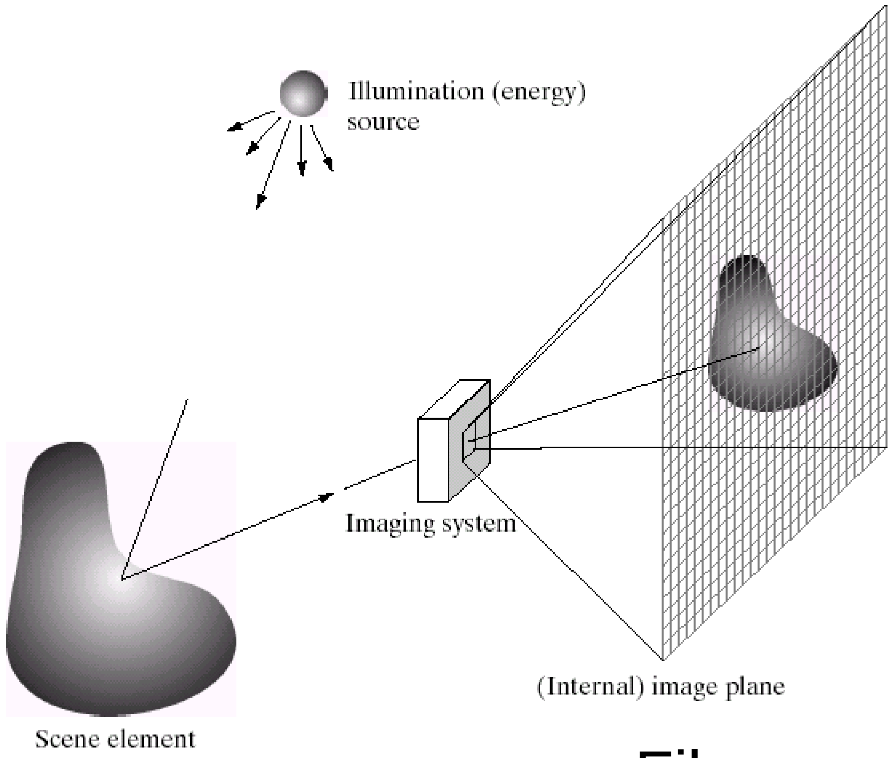

Film

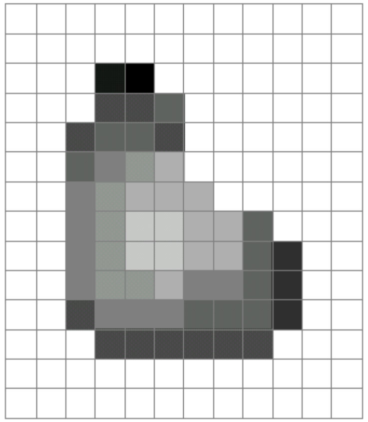

Digital Camera

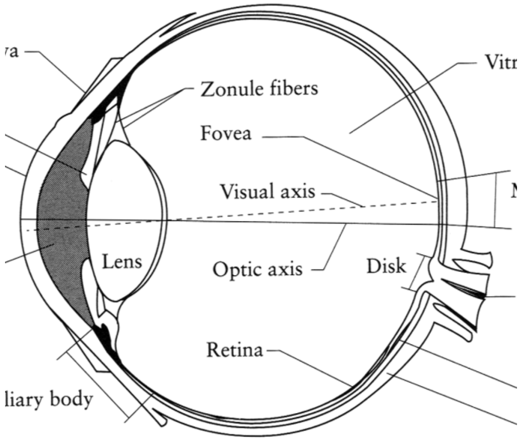

The Eye

Digital camera

- A digital camera replaces film with a sensor array
  - Each cell in the array is light-sensitive diode that converts photons to electrons
  - Two common types
    - Charge Coupled Device (CCD)
    - CMOS
  - http://electronics.howstuffworks.com/digital-camera.htm

Slide by Steve Seitz

Sensor Array

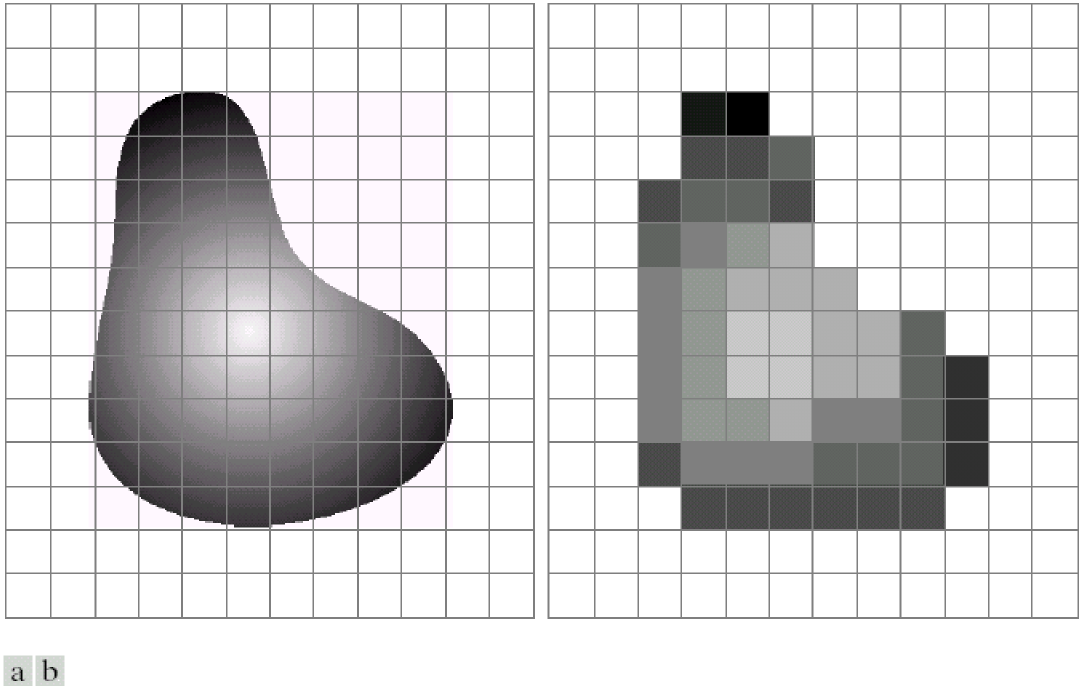

FIGURE 2.17 (a) Continuos image projected onto a sensor array. (b) Result of image sampling and quantization.

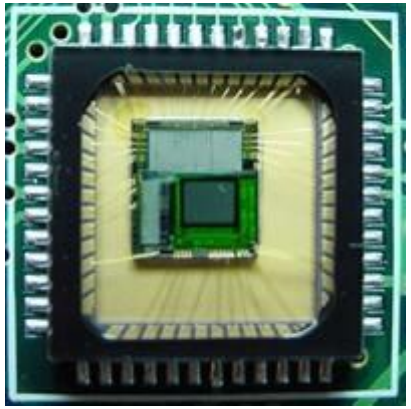

CMOS sensor

Sampling and Quantization

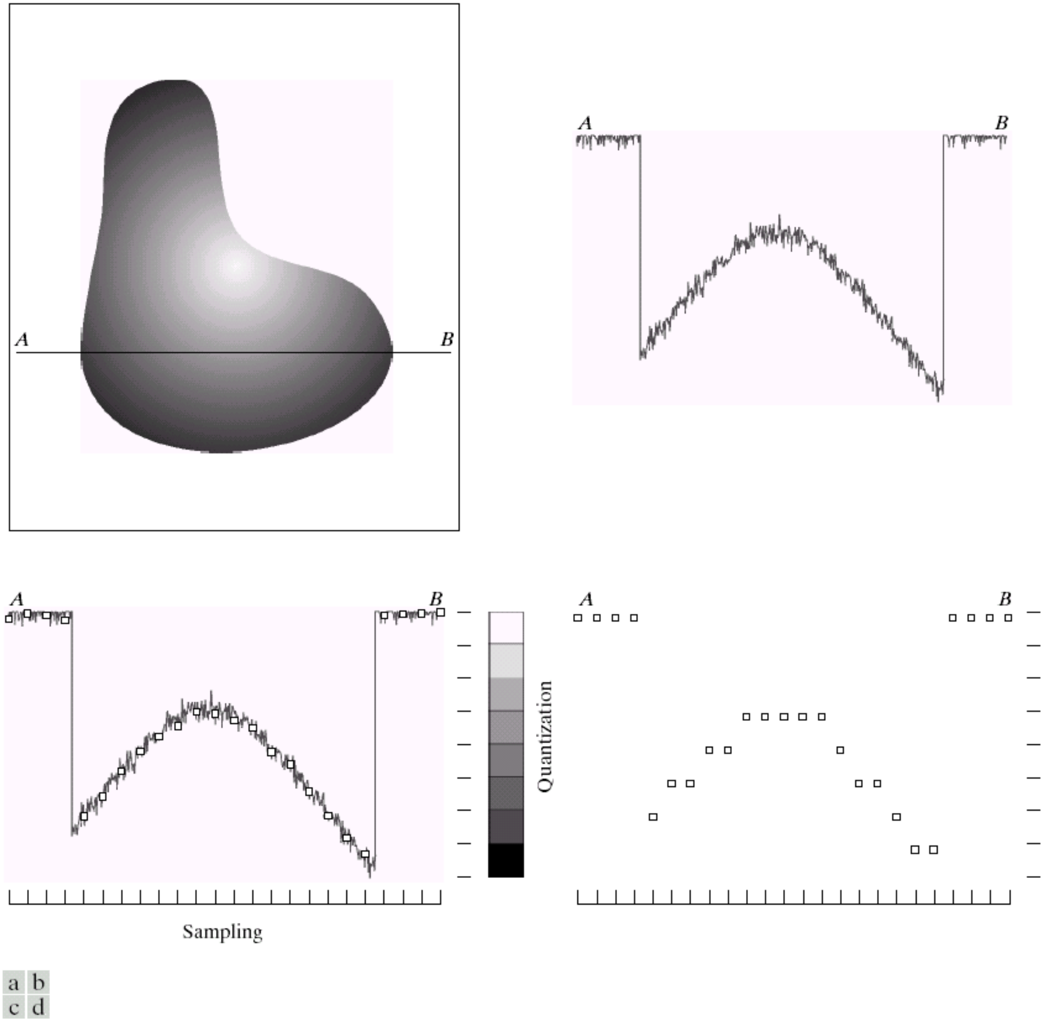

FIGURE 2.16 Generating a digital image. (a) Continuous image. (b) A scan line from A to B in the continuous image, used to illustrate the concepts of sampling and quantization. (c) Sampling and quantization. (d) Digital scan line.

Interlace vs. progressive scan

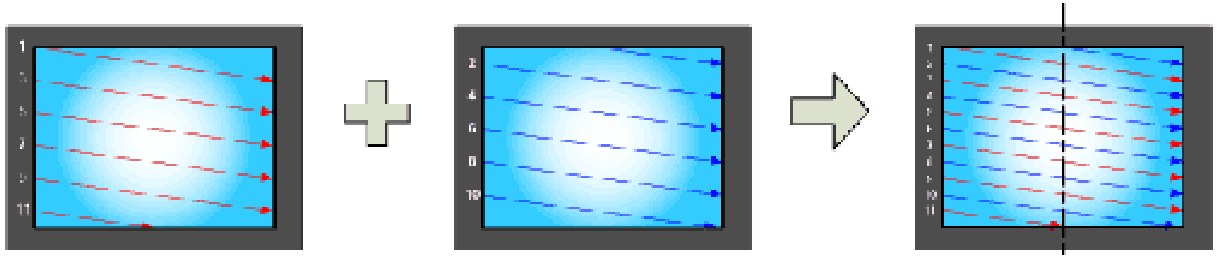

1st field: Odd field   2nd field: Even field   One complete frame using interlaced scanning

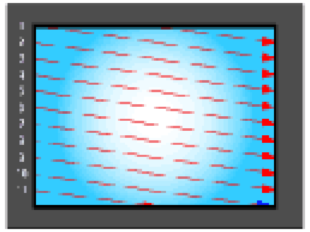

One complete frame using progressive scanning

http://www.axis.com/products/video/camera/progressive_scan.htm

Slide by Steve Seitz

Progressive scan or Global shutter

http://www.axis.com/products/video/camera/progressive_scan.htm

Slide by Steve Seitz

Interlaced

http://www.axis.com/products/video/camera/progressive_scan.htm

Slide by Steve Seitz

Slow mo guys – CRTs

- https://youtu.be/3BJU2drrtCM

Rolling Shutter

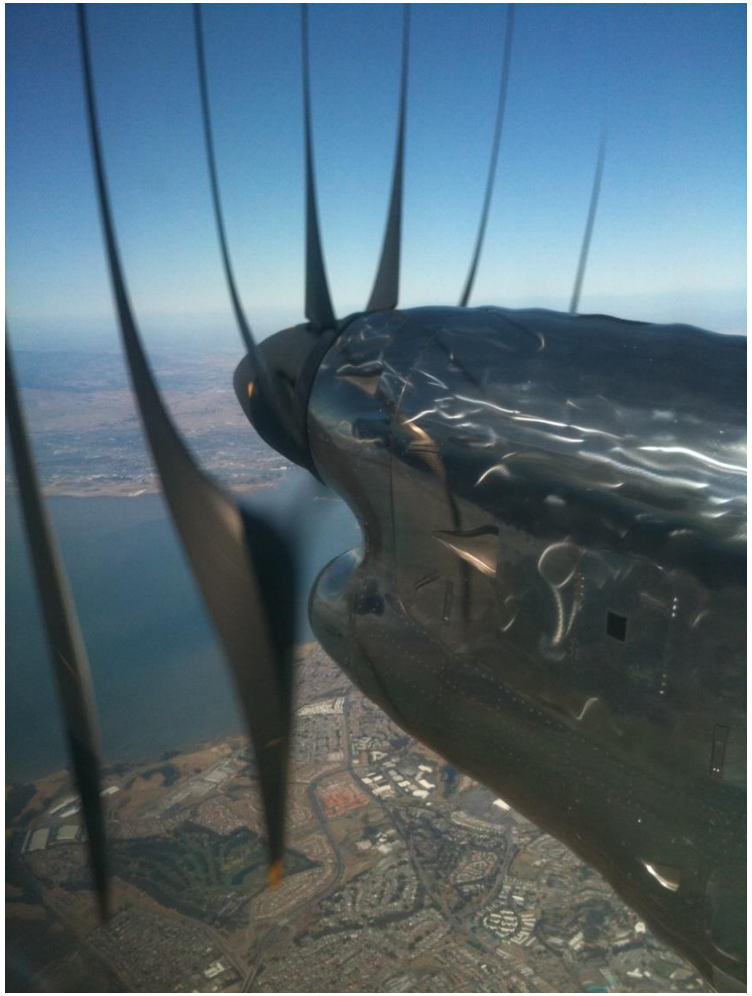

The Eye

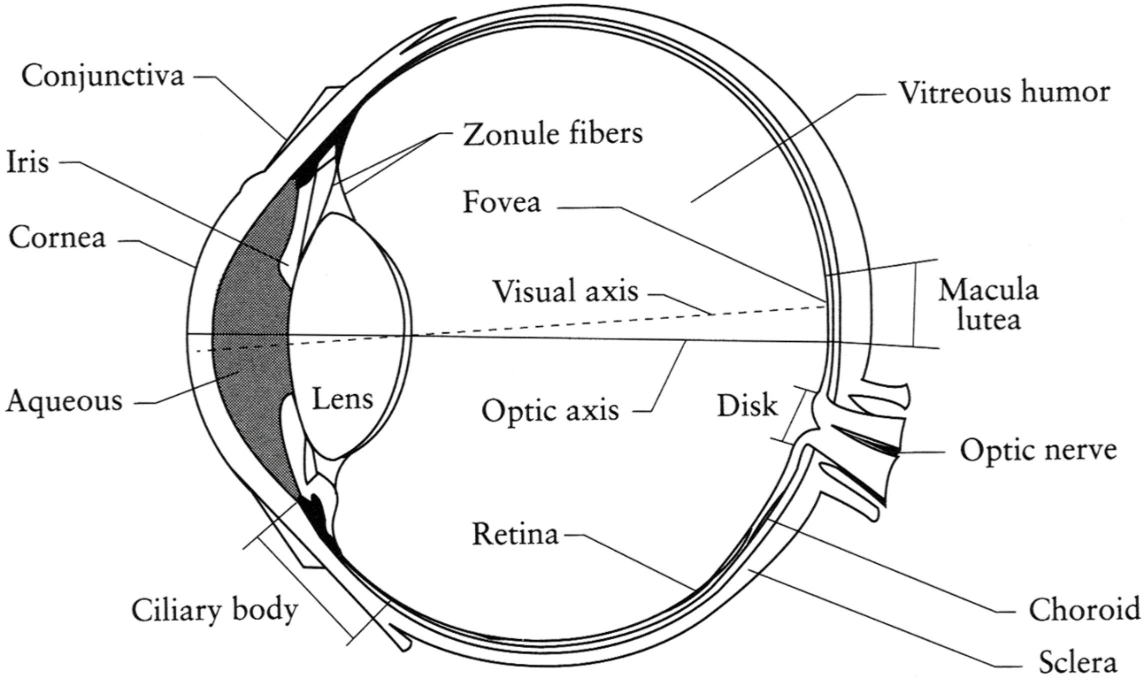

- The human eye is a camera!
  - **Iris** - colored annulus with radial muscles
  - **Pupil** - the hole (aperture) whose size is controlled by the iris
  - What's the "film"?
    - photoreceptor cells (rods and cones) in the **retina**

Slide by Steve Seitz

Aside: why do we care about human vision in this class?

- We don’t, necessarily.

Ornithopters

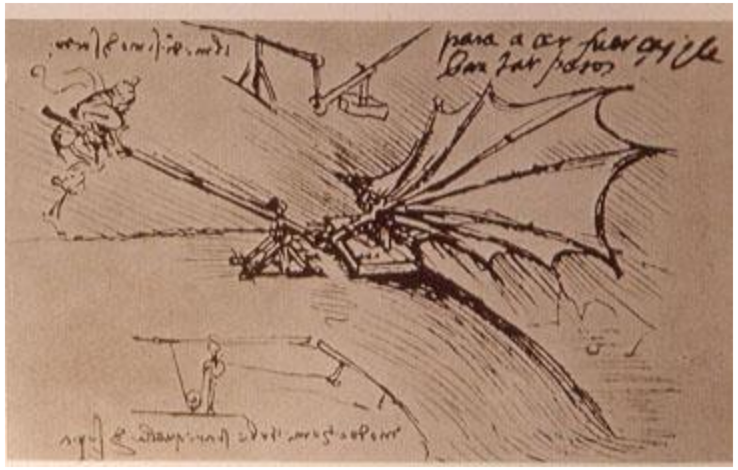

Why do we care about human vision?

- We don’t, necessarily.
- But cameras necessarily imitate the frequency response of the human eye, so we should know that much.
- Also, computer vision probably wouldn’t get as much scrutiny if biological vision (especially human vision) hadn’t proved that it was possible to make important judgements from 2D images.

Does computer vision “understand” images?

"Can machines fly?" The answer is yes, because airplanes fly.

"Can machines swim?" The answer is no, because submarines don't swim.

"Can machines think?" Is this question like the first, or like the second?

Source: Norvig

The Retina

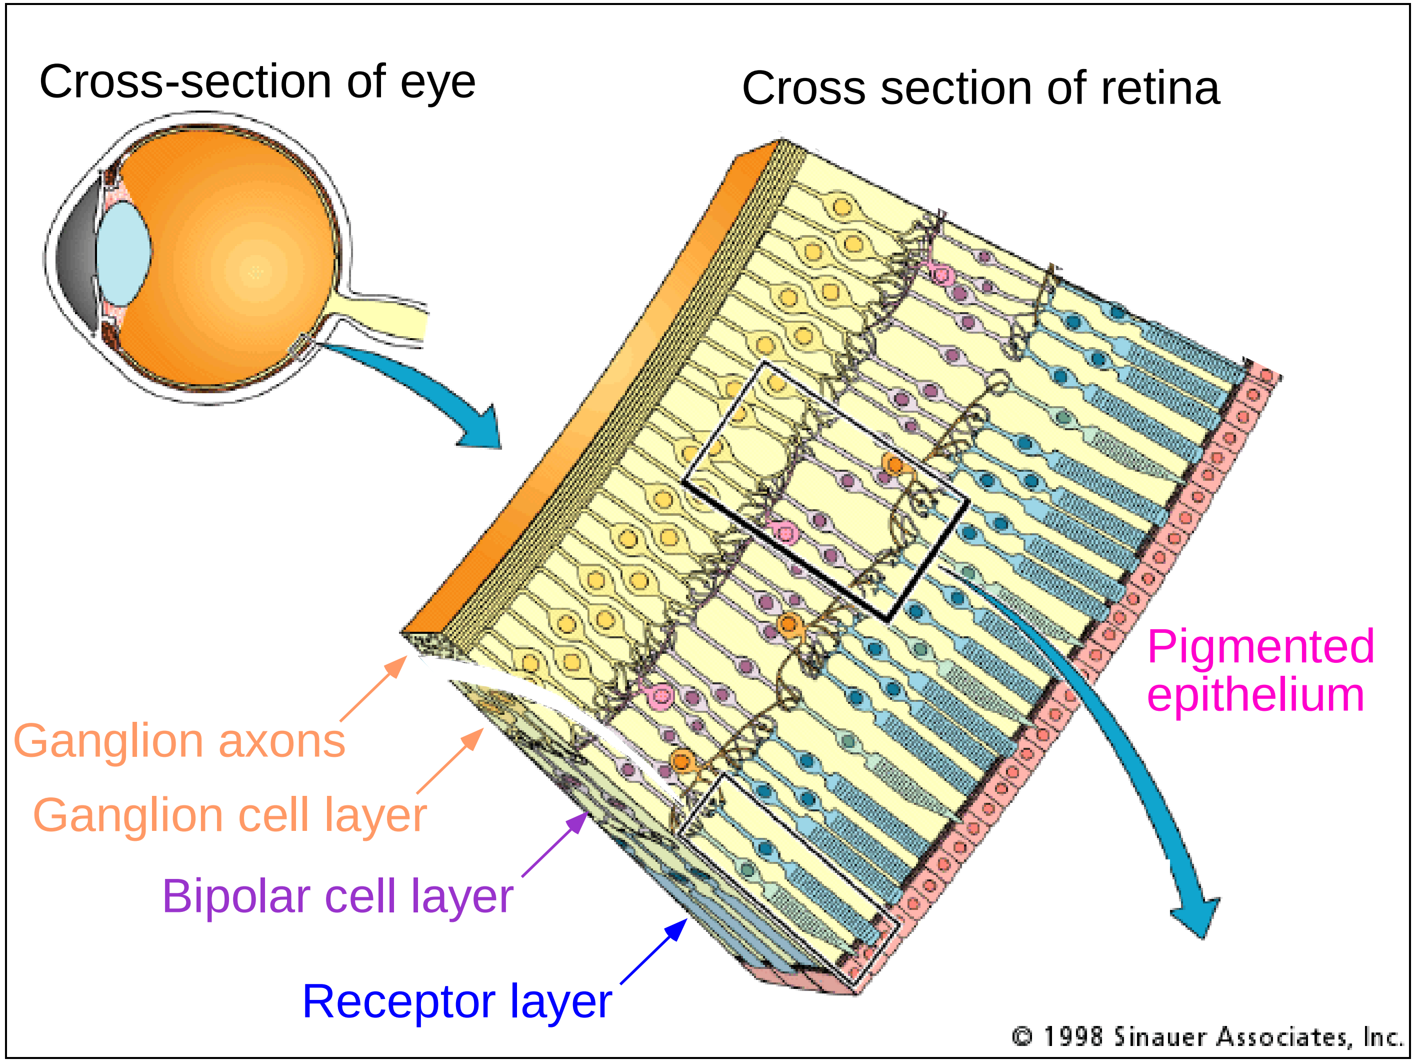
Cross-section of eye
Cross section of retina
Pigmented epithelium
Ganglion axons
Ganglion cell layer
Bipolar cell layer
Receptor layer
© 1998 Sinauer Associates, Inc.

Retina up-close

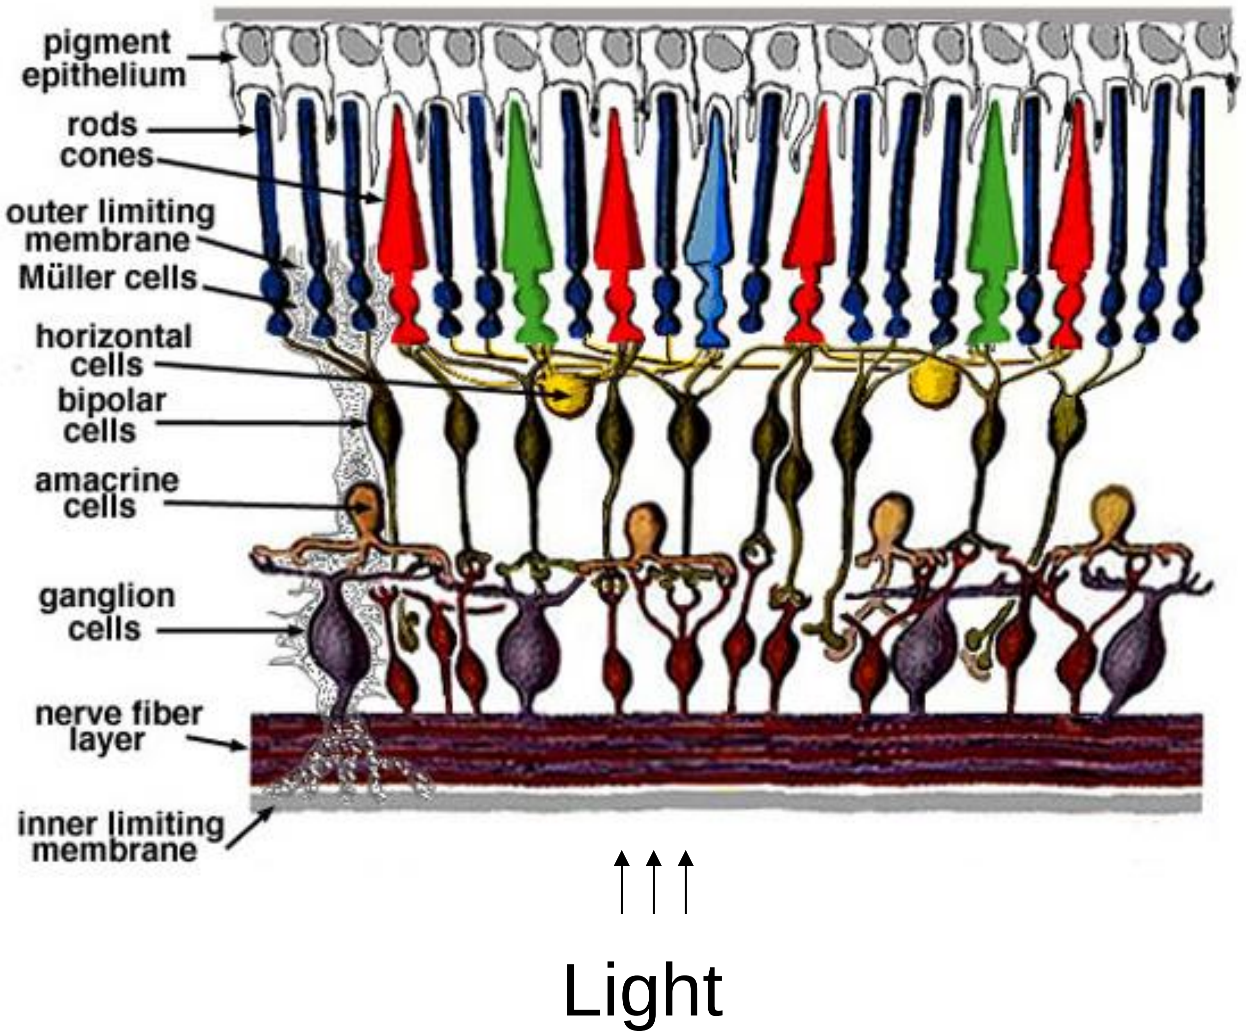

What humans don’t have: tapetum lucidum

Human eyes can reflect a tiny bit and blood in the retina makes this reflection red.

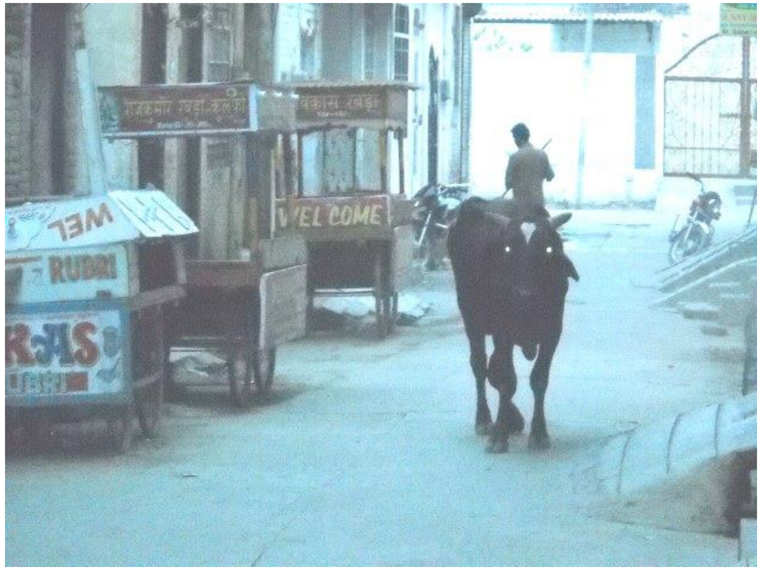

Wait, the blood vessels are in front of the photoreceptors??

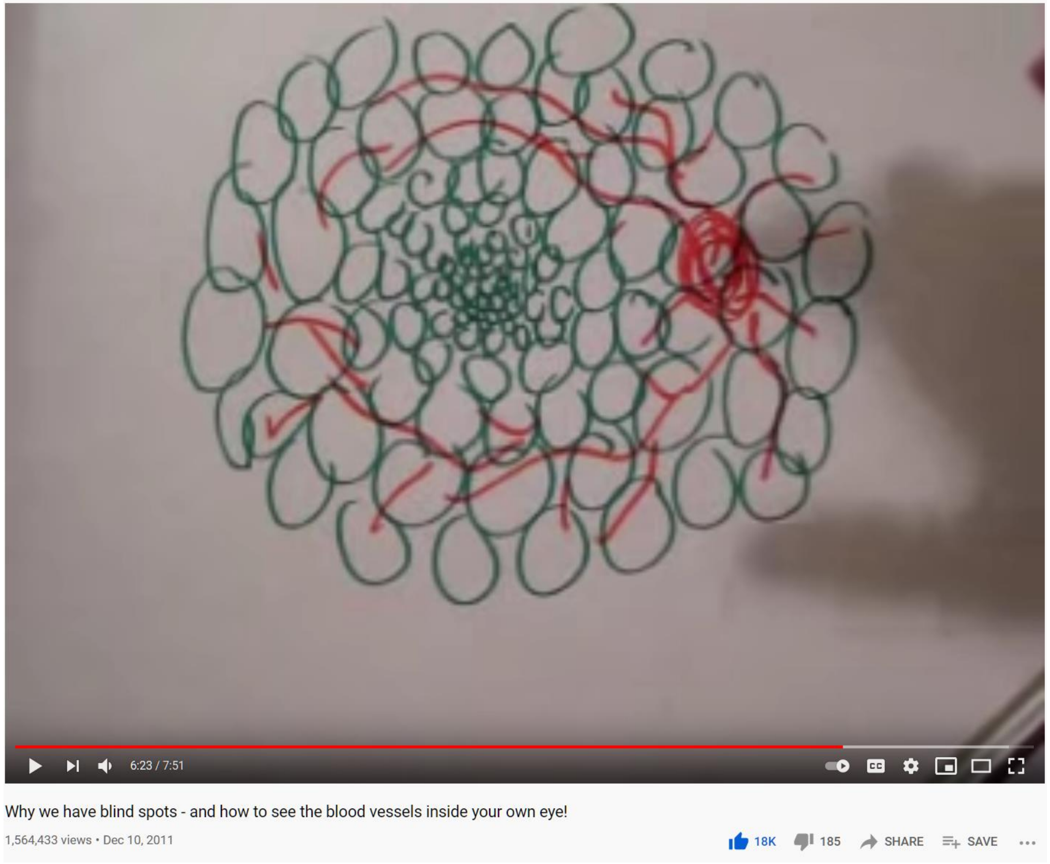

https://www.youtube.com/watch?v=L_W-IXqoxHA