Deconvolution & Wavefront Coding

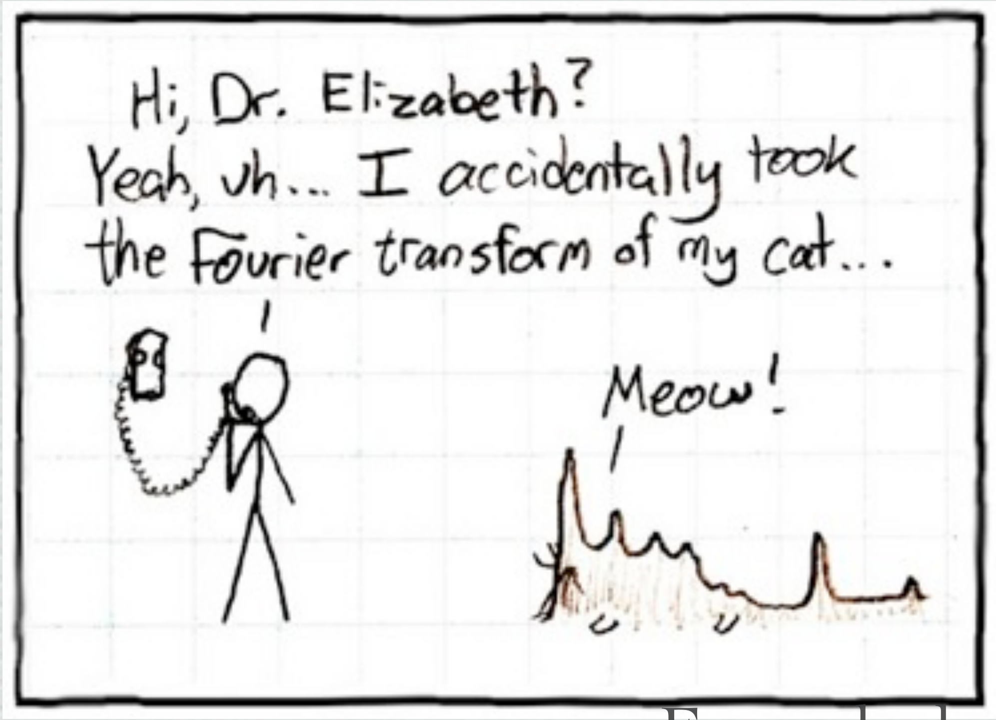

From xkcd

Frédo Durand
MIT EECS

Depth of field challenge

Hardware depth of field solution

- Stop down! (use smaller aperture)
- problem: noise

http://www.cdm-optics.com/site/publications.php

Why you need shallow depth of field

Photo with fake shallow depth of field

Why you need shallow depth of field

Original photo

Today's plan

- Remove blur computationally
  - Understand blur
  - A bit of Fourier analysis
  - Deconvolution
  - Noise, optimal deconvolution, frequency response
- Computational optics for DoF extension
  - single image capture + deconvolution
- Focal stack
  - Multiple-exposure solution

Blur: Linear shift-invariant filtering

- Replace each pixel by a linear combination of its neighbors.
  - only depends on relative position of neighbors
- The prescription for the linear combination is called the “convolution kernel”.

10	5	3
4	5	1
1	1	7

Local image data

0	0	0
0	0.5	0
0	1	0.5

kernel

	7	

Modified image data (shown at one pixel)

Convolution

original

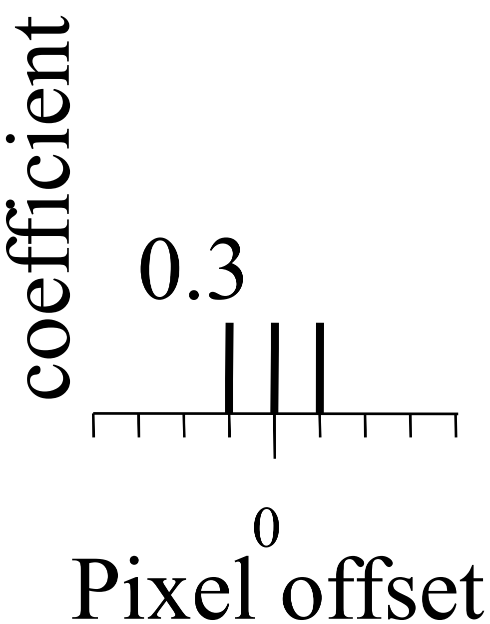

?

Blurring

original

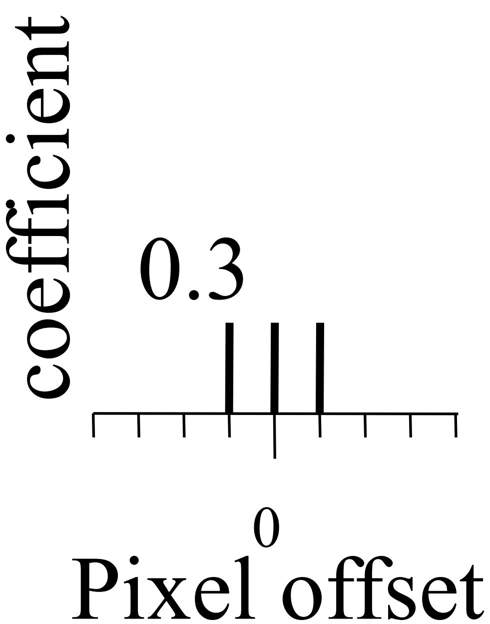

Blurred (filter applied in both dimensions).

Studying convolution

- Convolution is complicated
  - But at least it’s linear
    $(f+kg)- h = f-h +k\,(g-h)$
- We want to find a better expression
  - Let’s study functions whose behavior is simple under convolution

Blurring: convolution

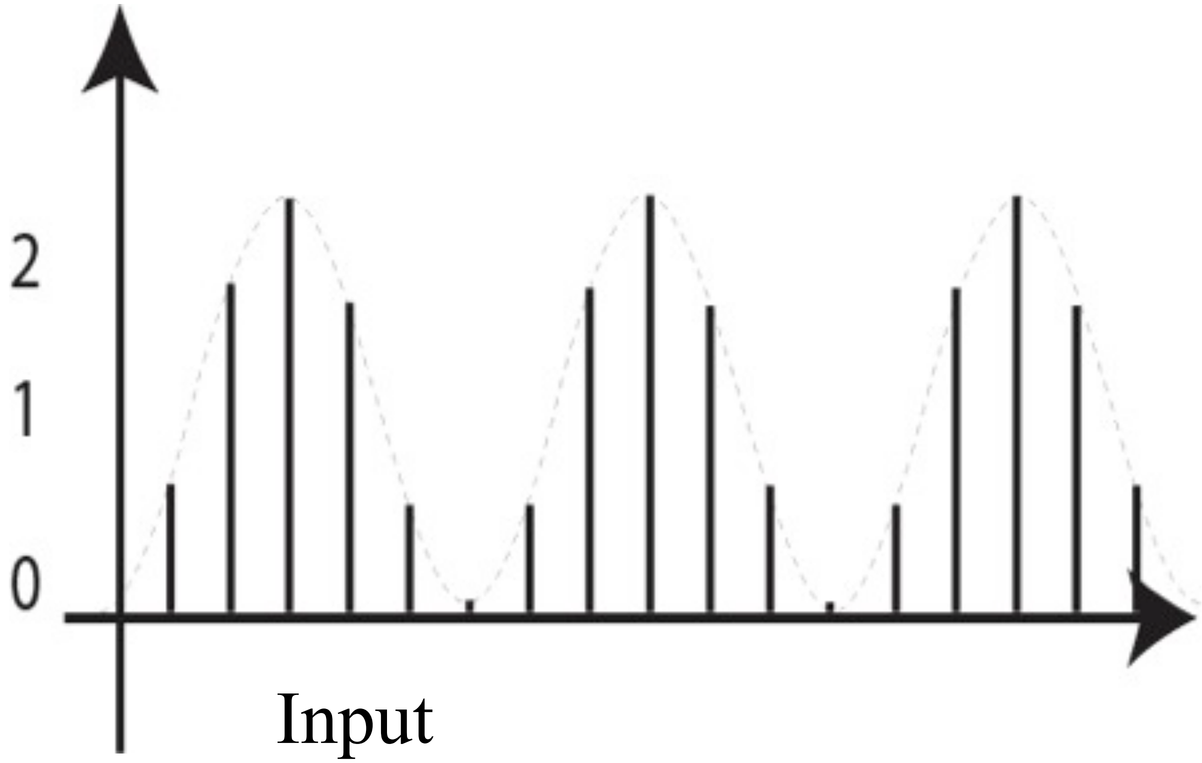

Input

Convolution sign

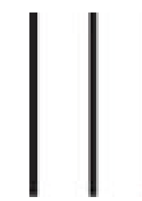

Kernel

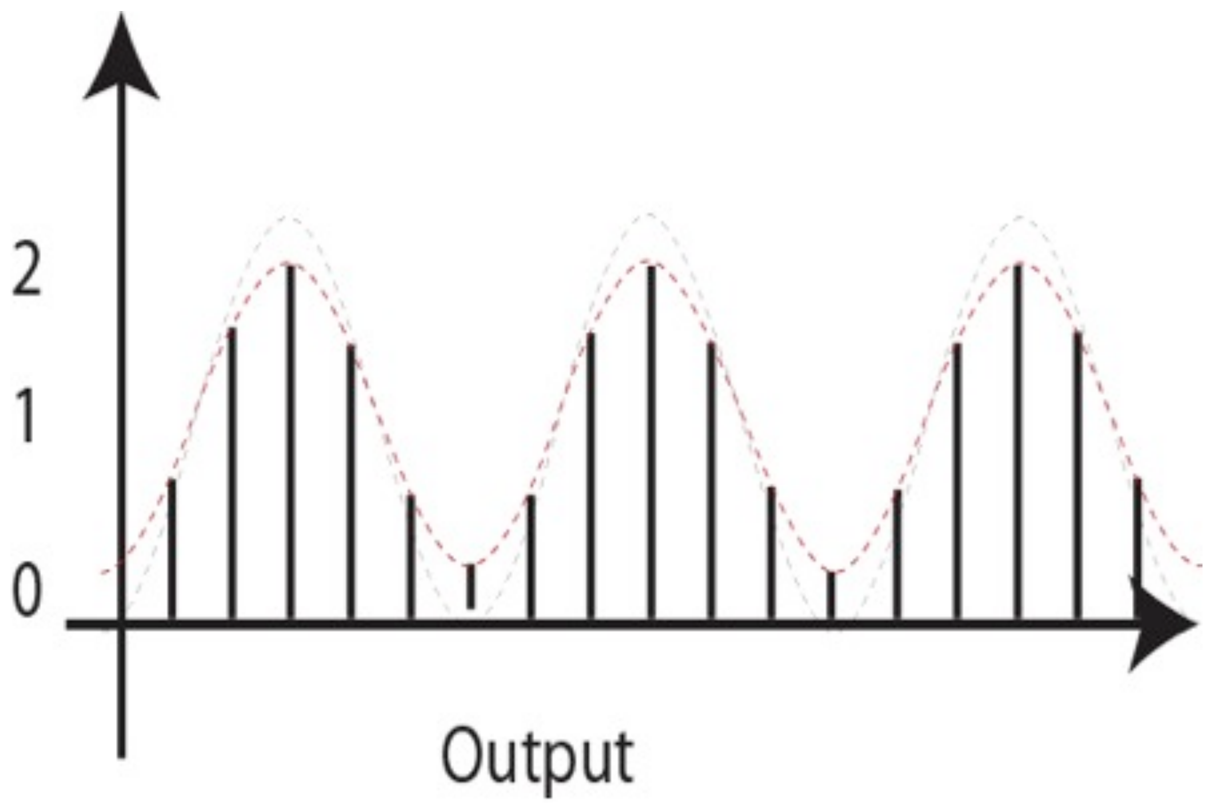

Same shape, just reduced contrast!!!

This is an eigenvector (output is the input multiplied by a constant)

Convolution theorem

A convolution in the primal is a multiplication in Fourier

Primal		Fourier
$f \otimes g$	$\Longleftrightarrow$	FG

i.e. Fourier bases are eigenvectors of convolution
(convolution is diagonal in the Fourier domain)

Big Motivation for Fourier analysis

- (Complex) sine waves are eigenvectors of the convolution operator
  - They diagonalize convolution
  - Convolution theorem

Second motivation for Fourier analysis: sampling

- The sampling grid is a periodic structure
  - Fourier is pretty good at handling that
- Sampling is a linear process
  - (but not shift-invariant)

Sampling

- If we're lucky, sampling density is enough

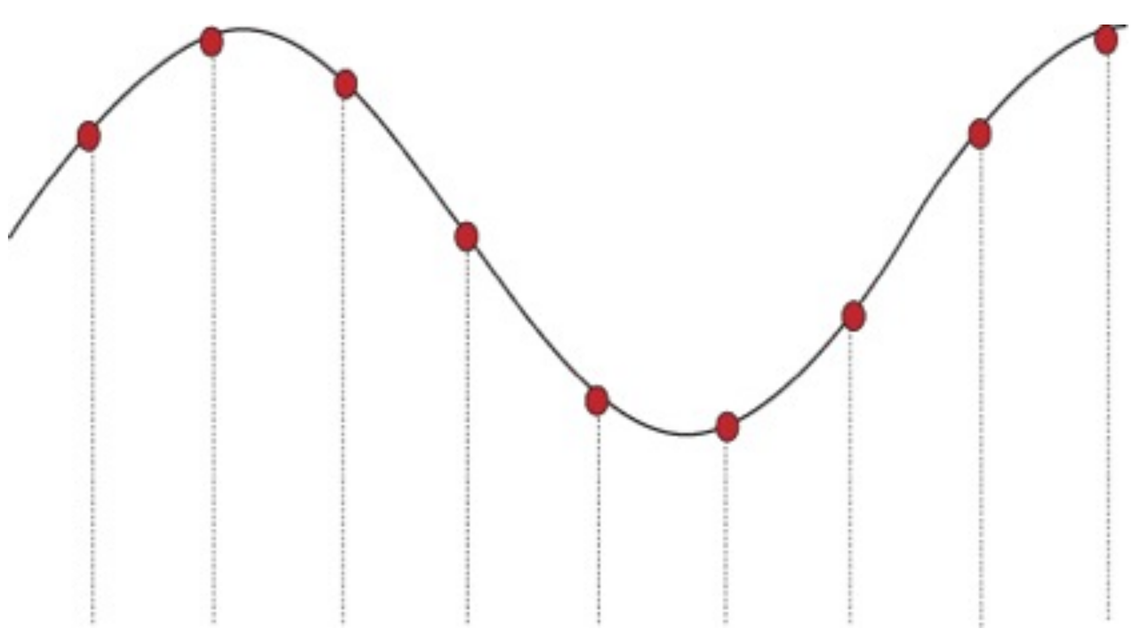

Input

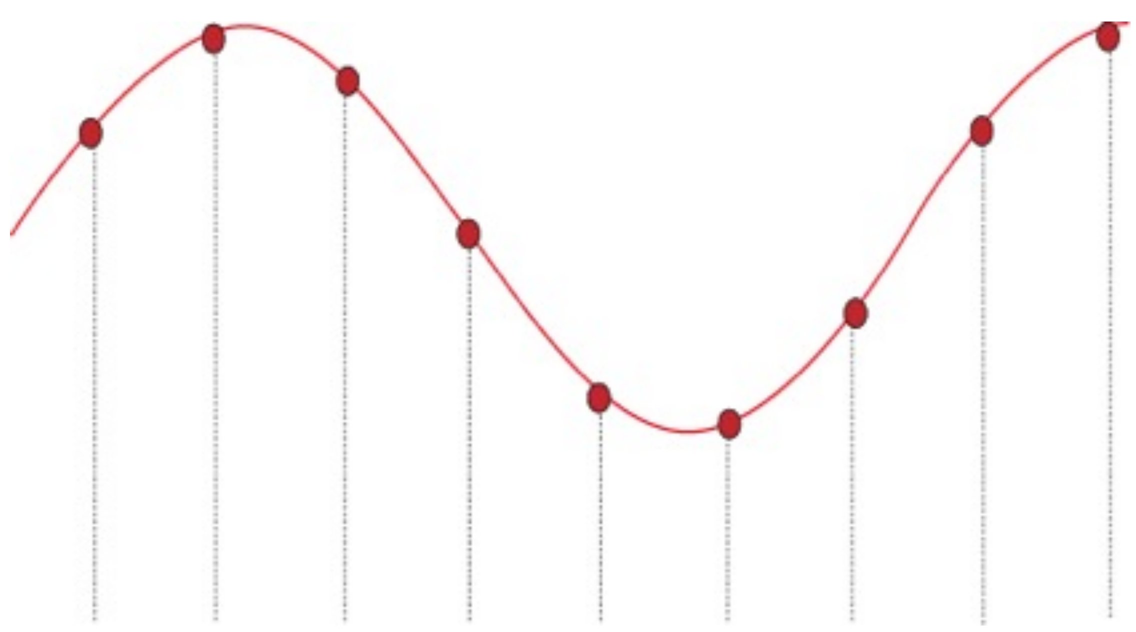

Reconstructed

Sampling

- If we insufficiently sample the signal, it may be mistaken for something simpler during reconstruction (that's aliasing!)

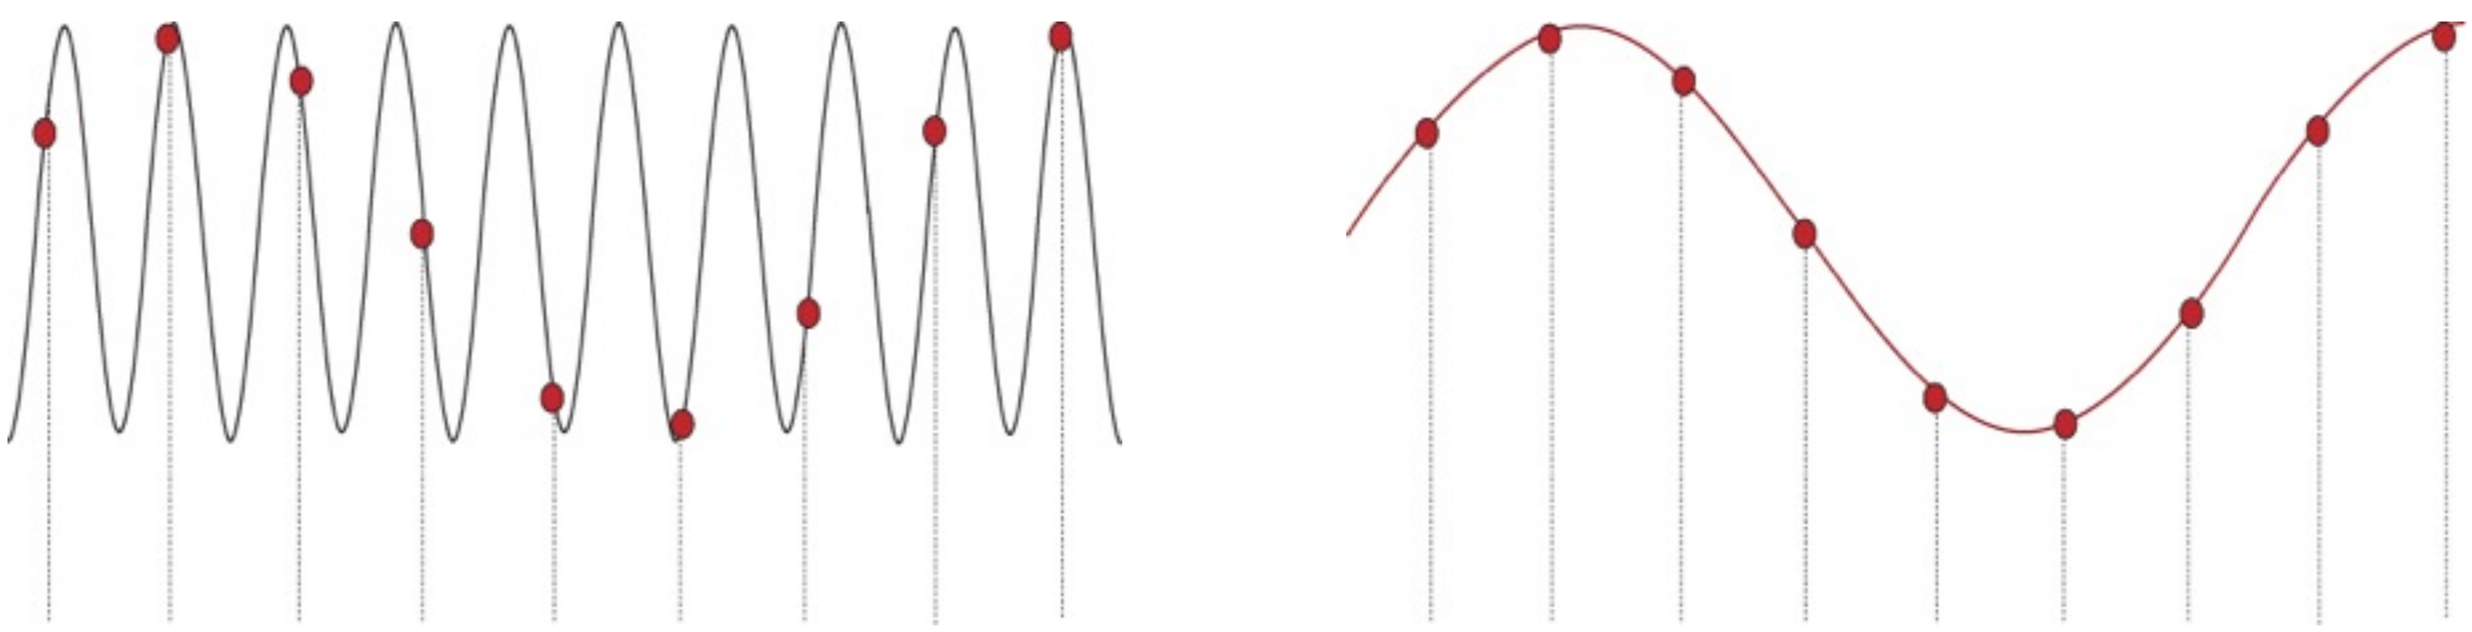

Recap: motivation for sine waves

- Blurring sine waves is simple
  - You get the same sine wave, just scaled down
  - The sine functions are the eigenvectors of the convolution operator
- Sampling sine waves is interesting
  - Get another sine wave
  - Not necessarily the same one! (aliasing)

If we represent functions (or images) with a sum of sine waves, convolution and sampling are easy to study

Questions?

Fourier as change of basis

- Shuffle the data to reveal other information
- E.g., take average & difference: matrix

$$\begin{bmatrix} 0.5 & 1 \\ 0.5 & -1 \end{bmatrix}$$

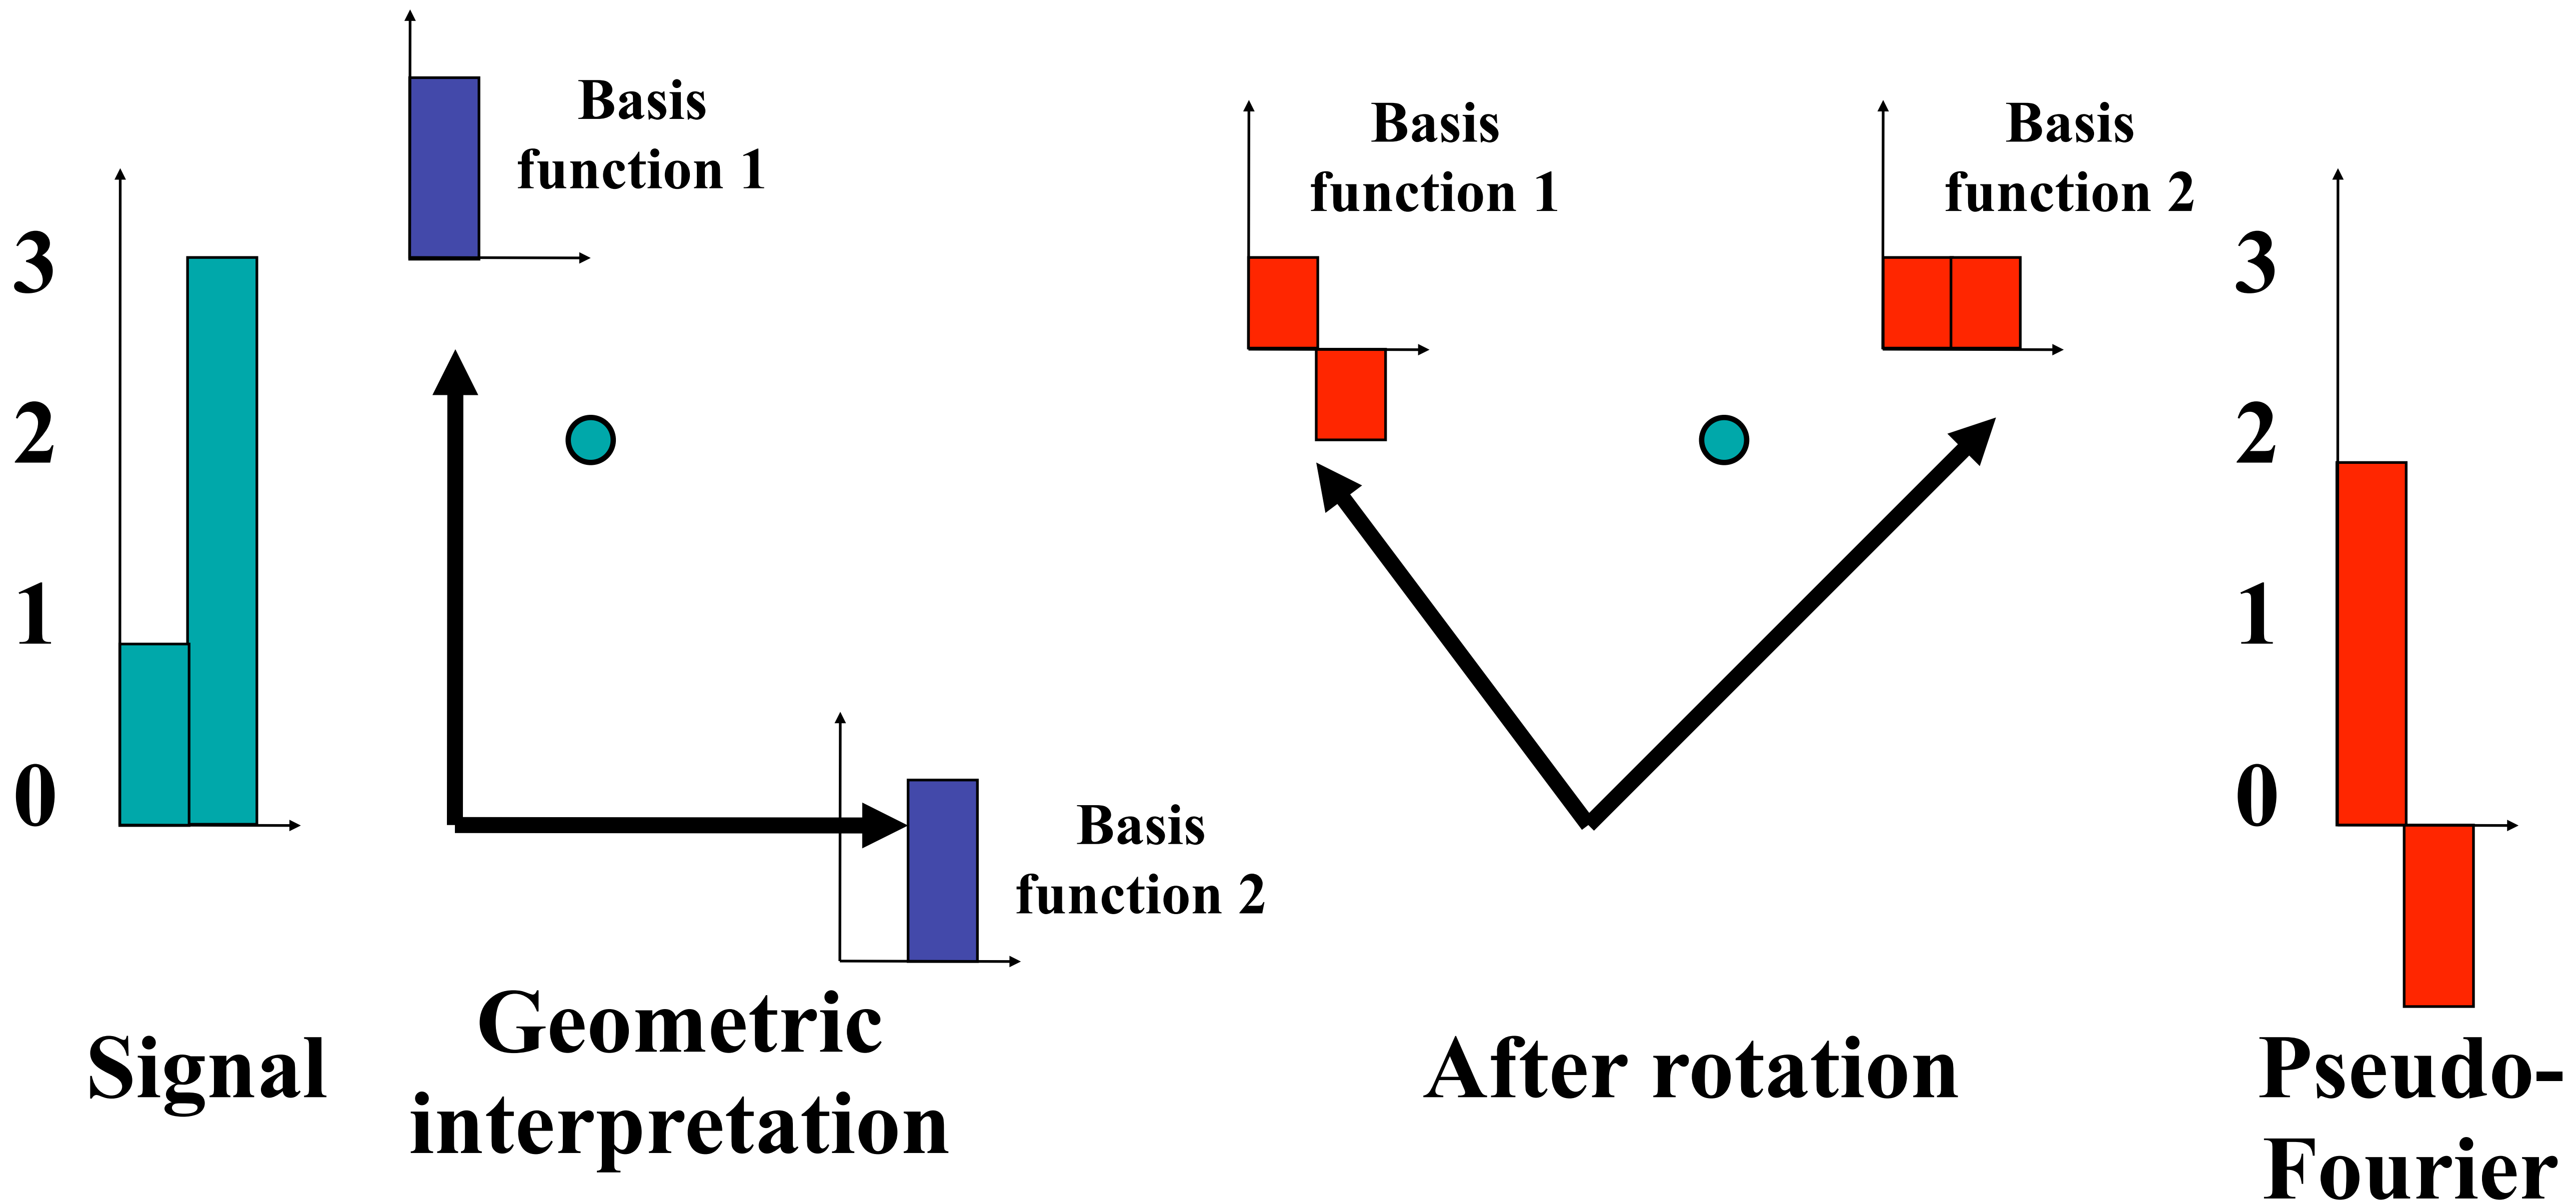

Fourier as change of basis

- Same thing with infinite-dimensional vectors

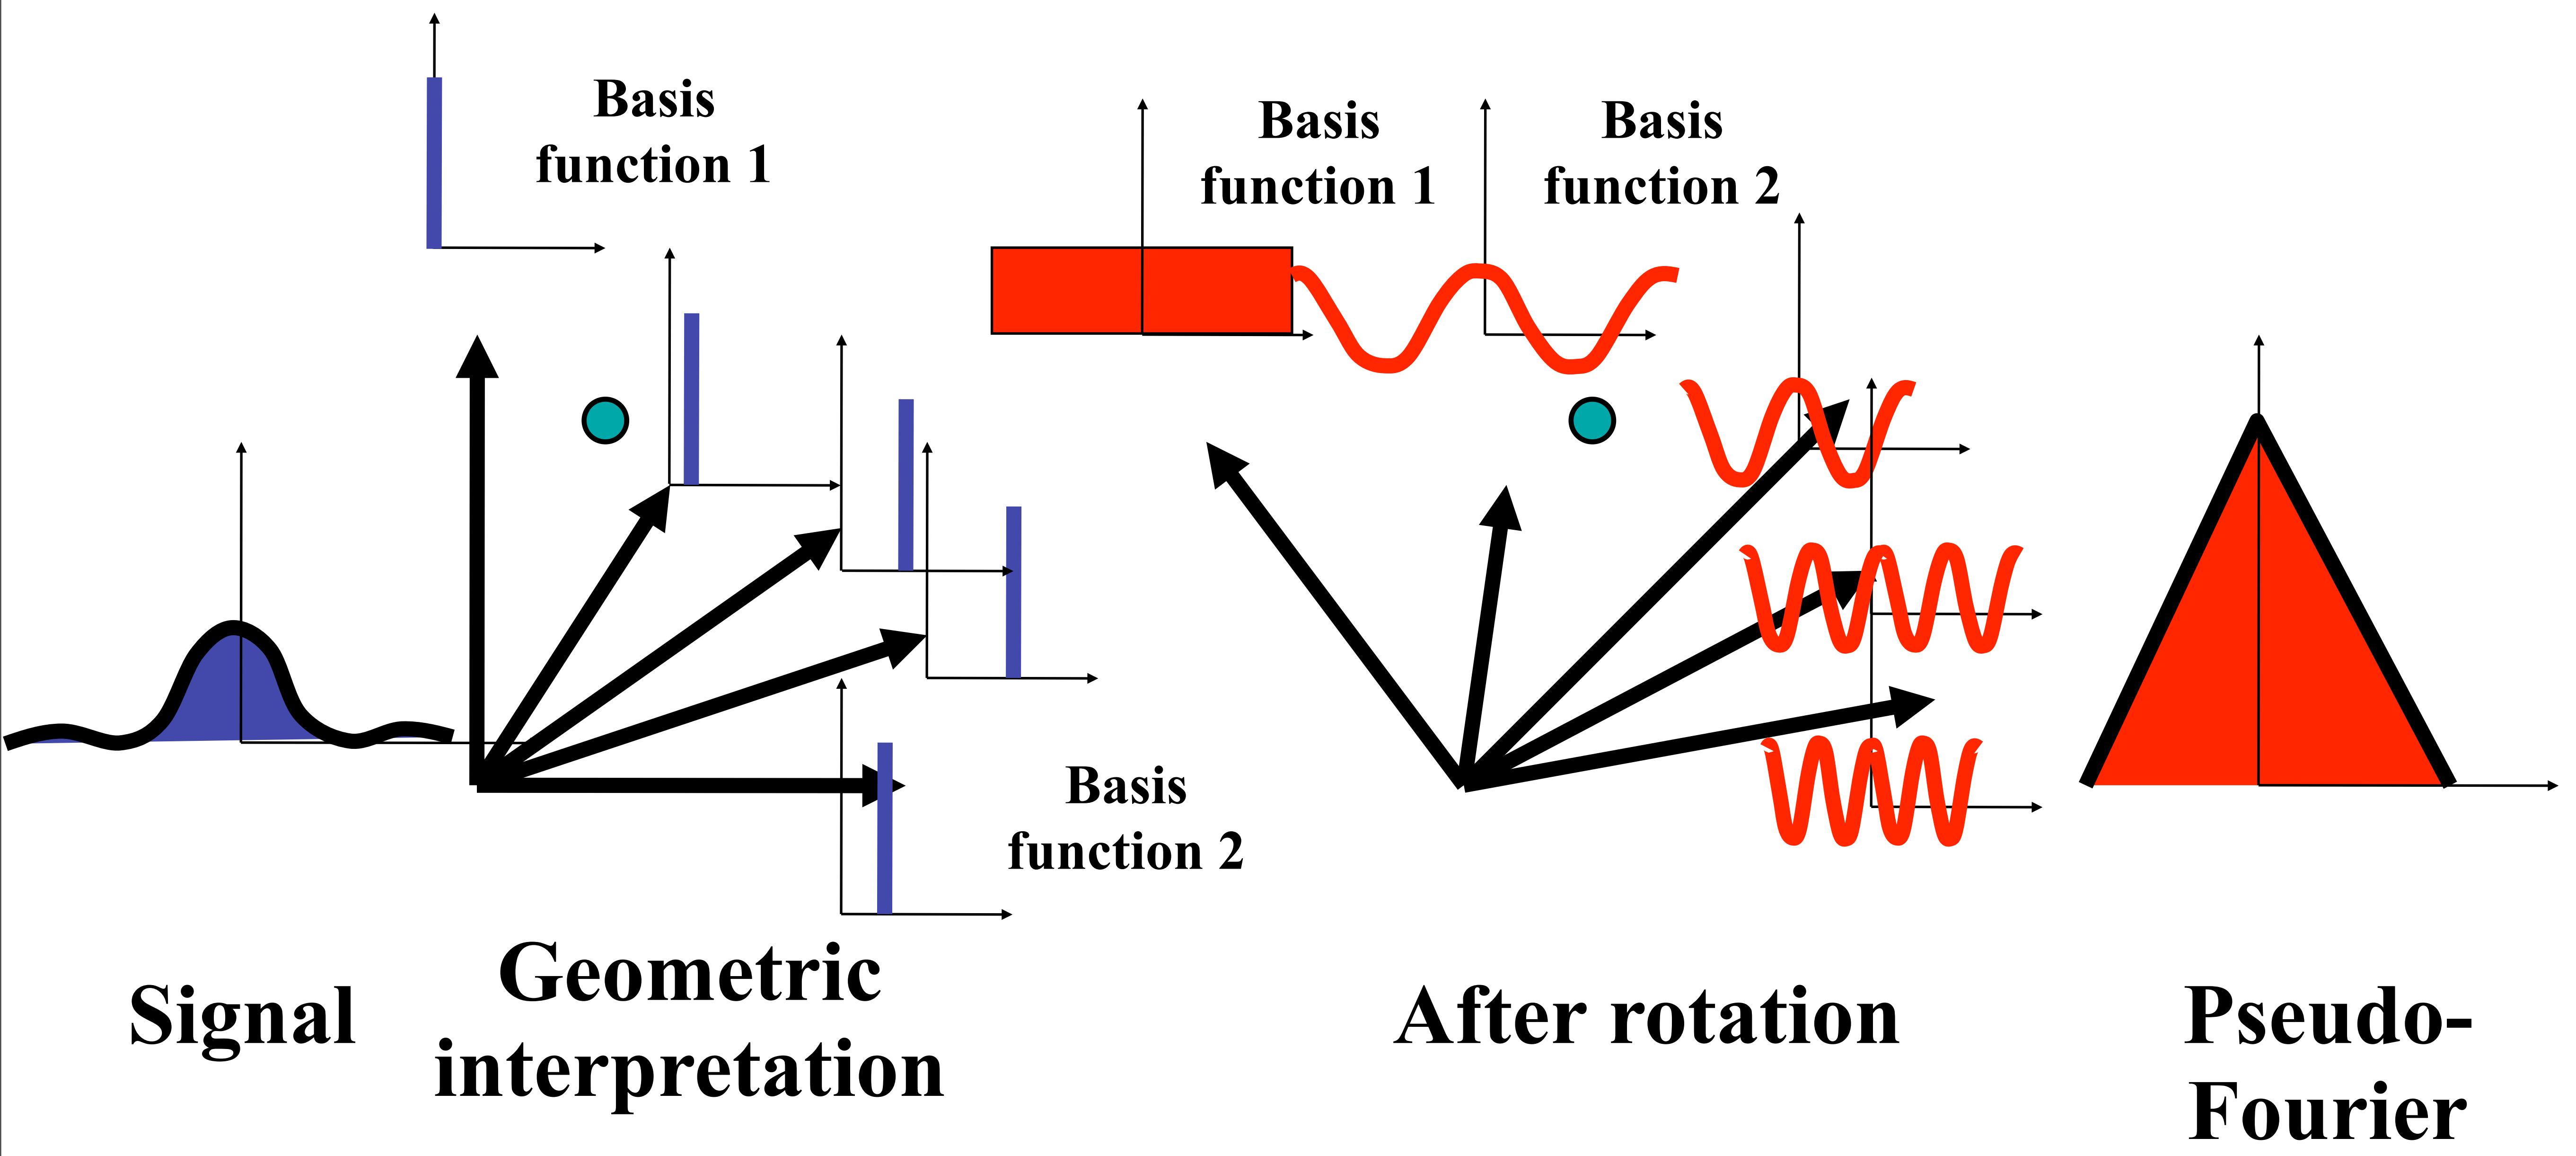

Questions?

Other presentations of Fourier

- Start with Fourier series with periodic signal
- Heat equation
  - –more or less special case of convolution
  - –iterate -> exponential on eigenvalues

Motivations

- Insights & mathematical beauty
- Sampling rate and filtering bandwidth
- Computation bases
  - FFT: faster convolution
  - E.g. finite elements, fast filtering, heat equation, vibration modes
- Optics: wave nature of light & diffraction

Primal vs. dual

- Often, we use the Fourier domain only for analysis
  - –convergence, well-posedness
- Computation is performed in the primal
- In other cases, computation is better in Fourier
  - –faster because diagonal

Questions?

Can we undo blur?

original

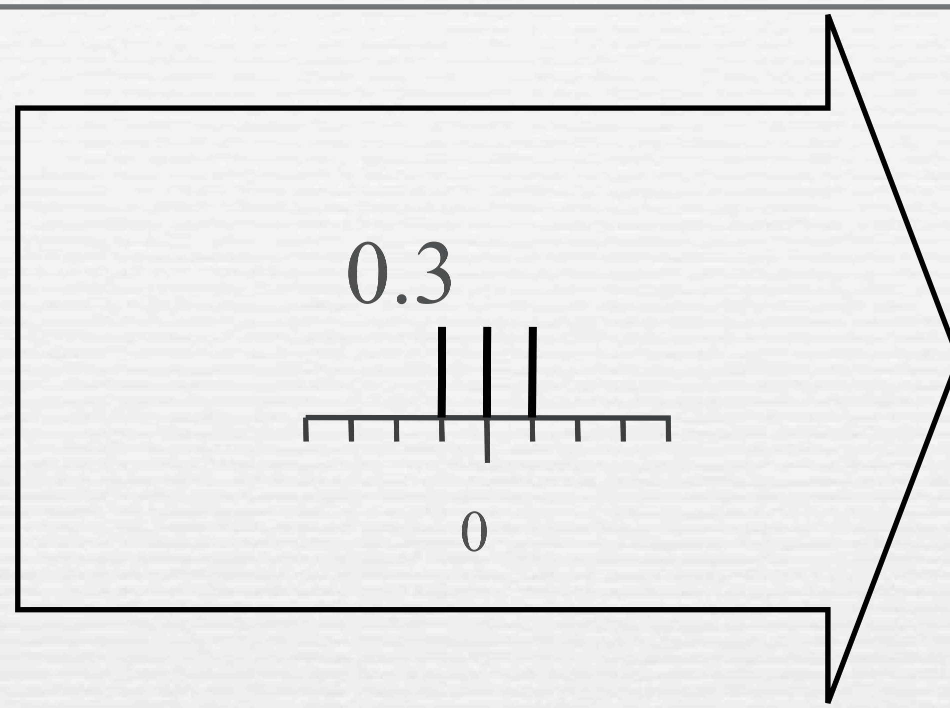

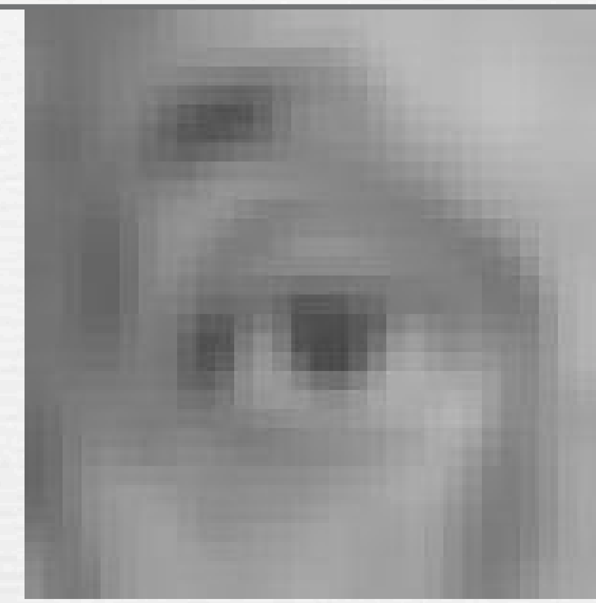

Blurred

Blurred

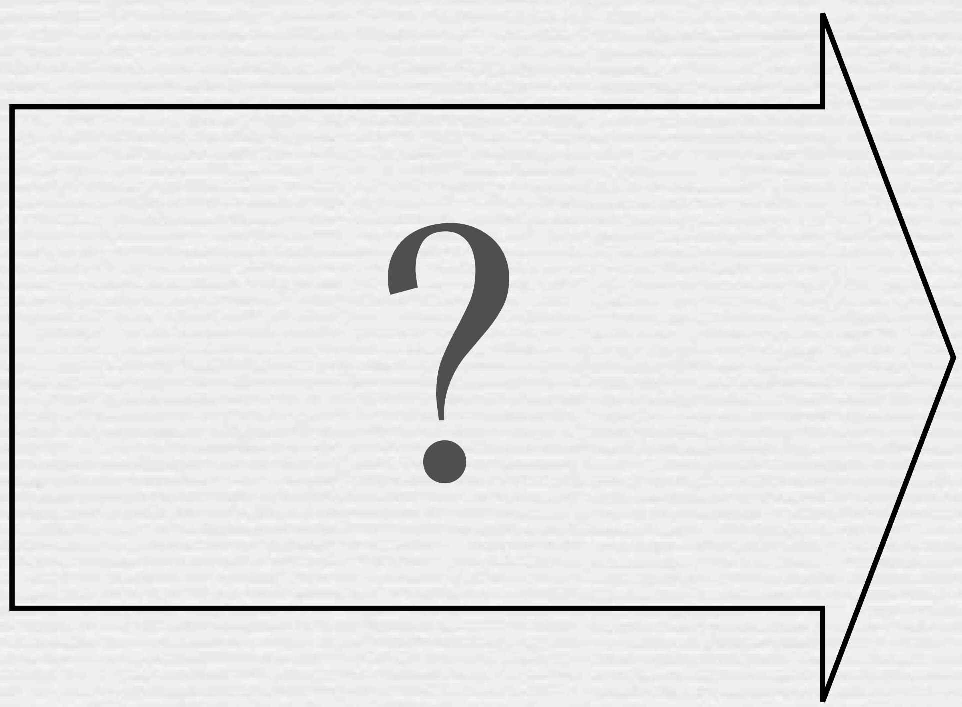

original

Not easy, even when we know the kernel

Recall convolution theorem

- Convolution in space is a multiplication in Fourier
- Note y the observed blurry image and x the original sharp one
- $y = g \otimes x$ in the spatial domain
- $Y = GX$ in the Fourier domain
  - A frequency does not depend on the other ones

Invert the convolution theorem

- Given $y = g \otimes x$ and g, we seek an estimate x' of x
- How do you invert a multiplication?
  - Division!
- $X'(\omega) = Y(\omega)/G(\omega)$
- DECONVOLUTION IS A DIVISION IN THE FOURIER DOMAIN !
- Which means it is also a convolution in the spatial domain, by the inverse Fourier transform of 1/G

Questions?

- Given $y = g \otimes x$ and g, we seek an estimate x' of x
- How do you invert a multiplication?
  - Division!
- $X'(\omega) = Y(\omega)/G(\omega)$
- DECONVOLUTION IS A DIVISION IN THE FOURIER DOMAIN !
- Which means it is also a convolution in the spatial domain, by the inverse Fourier transform of 1/G

Potential problem?

- Deconvolution is a division in the Fourier domain
- Division by zero is bad!
  - Information is lost at the zeros of the kernel spectrum G

Noise problem

- Even when there is no zero, noise is a big problem
- If G has small number, division amplifies noise
- if $y = g \otimes x + v$ where v is additive noise
- $Y = GX + V$
- $X' = (GX+V)/G$
  $= X + V/G$
- V is amplified by 1/G. This is why you typically get more high-frequency noise with deconvolution

Noise problem

blurry, no noise

deconvolved

blurry, with noise

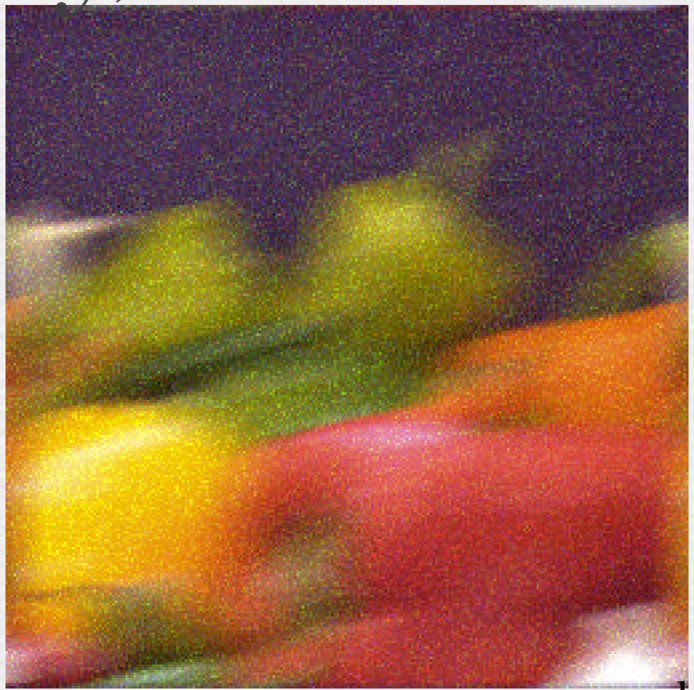

deconvolved

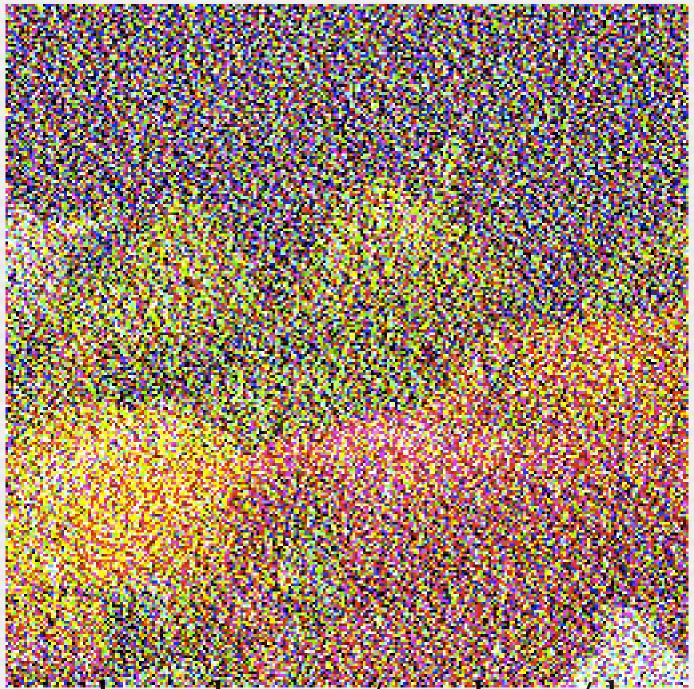

http://www.mathworks.com/products/demos/image/deblur_wiener/deblur.html

Questions?

Noise: extreme case

- If $G(\omega) = 0 \Rightarrow Y(\omega)=V(\omega)$
- what is the best estimate of $X(\omega)$?

  $X'(\omega) = 0$

- Even if $G(\omega)$=tiny, dividing by tiny is a bad idea and something much closer to zero is better

- The strategy should depend on the relative noise
  - low noise: just divide
  - high noise: under-estimate, closer to zero

Noise, even without convolution

- y=x+v
- or in Fourier Y=X+V

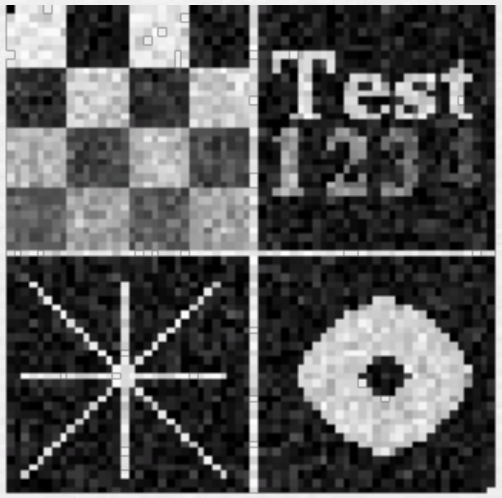

Test image with AWGN (additive White Gaussian Noise)
$\sigma_n^2 = 400$
but centered inside a
256x256 empty image

Noise, even without convolution

- Say we know that $E(X^2)=1$, $E(V^2)=16$
  - Pretty bad signal noise ratio!
- We observe $Y=X+V=5$
- It's more likely to be $X=0.5 + V=4.5$ than $X=4+V=1$
- How can we optimize our bet?

http://www.norbertwiener.umd.edu/NW/index.html

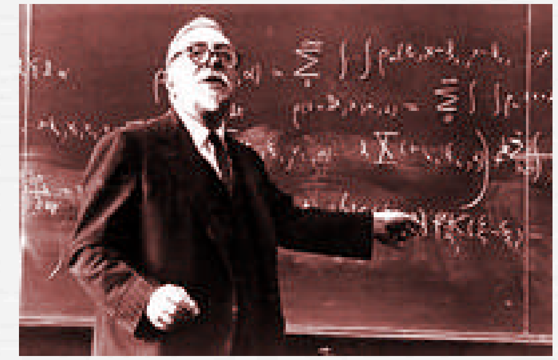

Wiener denoising

- Optimal estimation
  given known noise and signal powers
- Derive in Fourier domain because SNR is best known per frequency
  - Spectrum of images usually falls off as $1/\omega^2$
  - Noise is often white (flat spectrum) or high frequency

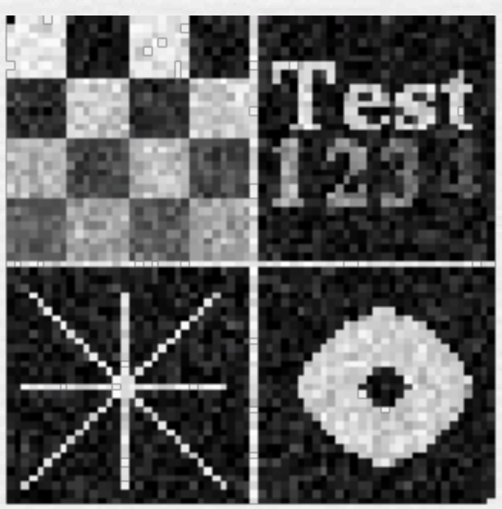

Test image with AWGN
$\sigma_n^2 = 400$
but centered inside a
256x256 empty image

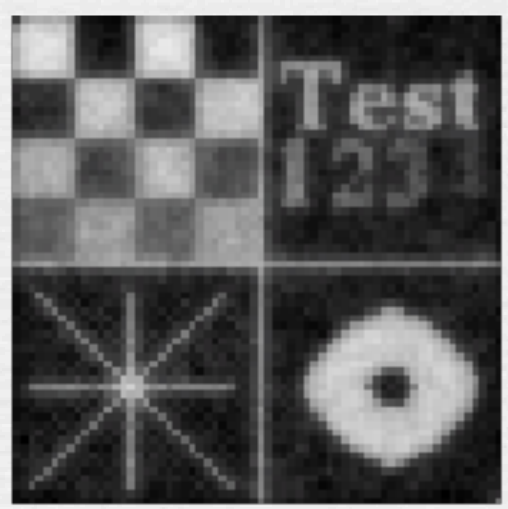

After Wiener filtering
MSE=121 (256x256 image)
MSE=1232 (portion shown)

Wiener denoising (in Fourier)

- Given Y=X+V, find X'=HY to minimize $E(||X-X'||^2)$
- argmin $E(||X-HY||^2)$ => argmin $E(||X-H(X+V)||^2)$
- argmin $E(||(1-H)X-HV||^2)$
- X and N are assumed independent, E(XV)=0
  expand and ignore cross terms
- argmin $||1-H||^2E(||X||^2)+||H||^2E(||V||^2)$
- derive wrt H, set to zero
- $(2H-2)\ E(||X||^2)+2H\ E(||V||^2)=0$
- $H=E(||X||^2)\ /\ (E(||X||^2+E(||V||^2)$
- divide by $E(||X||^2)$ to get a function of SNR

$$H = \frac{1}{1 + 1/SNR}$$

Wiener denoising

$$X' = Y \frac{1}{1 + 1/SNR}$$

- When SNR is high, gain goes towards 1
- When SNR is low, gain goes to zero

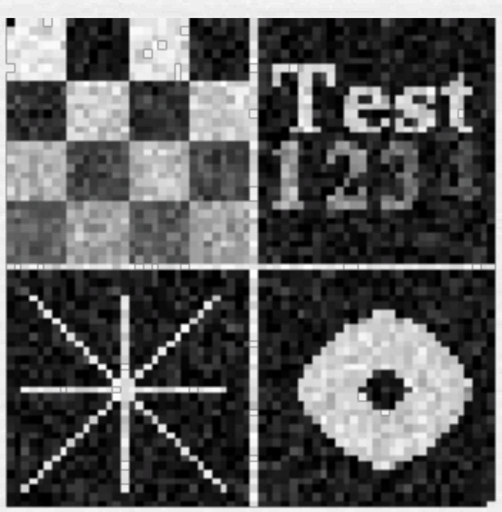

Test image with AWGN
$\sigma_n^2 = 400$
but centered inside a
256x256 empty image

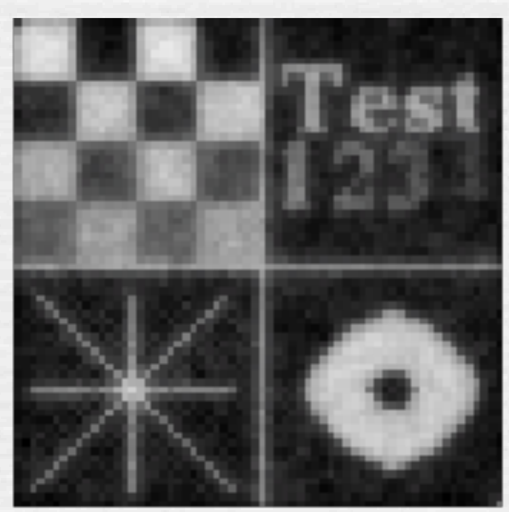

After Wiener filtering
MSE=121 (256x256 image)
MSE=1232 (portion shown)

Questions?

Back to convolution

- Assume we know the expect noise power spectrum and expected signal power spectrum

- Can we tweak 1/G to reduce output noise ?
- Maybe if we use something smaller than 1/G
  - we won't amplify noise as much
  - but the inversion won't be as correct

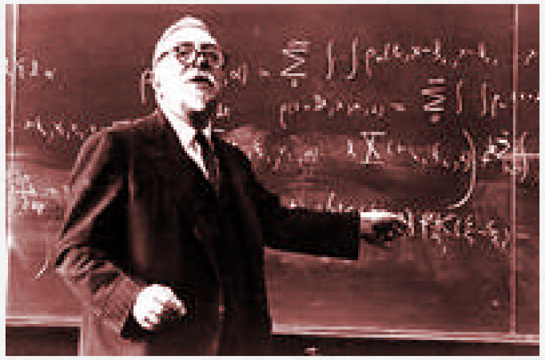

Wiener deconvolution

- Find the gain H that minimize $||X'(\omega)-X(\omega)||^2$ where X'=HY
- We need to know the signal noise ratio $SNR(\omega)=E(|X(\omega)|^2)/E(|V(\omega)|^2)$
- Optimal filter

$$\frac{1}{G(\omega)}\left[\frac{|G(\omega)|^2}{|G(\omega)|^2+1/SNR(\omega)}\right]$$

- See http://en.wikipedia.org/wiki/Wiener_deconvolution
  - careful, their notations are different from mine

Wiener deconvolution derivation

- Find H to minimize $||X'(\omega)-X(\omega)||^2$ where $X'=HY$
- argmin $E(||HY-X||^2)$=>argmin $E(||H(XG+V)-X||^2)$
- argmin $E(||(HG-1)X+HV||^2)$
- X and N are assumed independent:$E(XV)=0$
  Expand and ignore cross terms
- argmin $||HG-1||^2E(|X|^2)+||H||^2E(|V|^2)$
- $E(|X|^2)$ and $E(|V|^2)$ given by expected spectrum
- argmin$||HG-1||^2 +1/SNR\ ||H||^2$
- derive wrt H, set to zero, get Wiener

I remove ω for simplicity

Wiener result

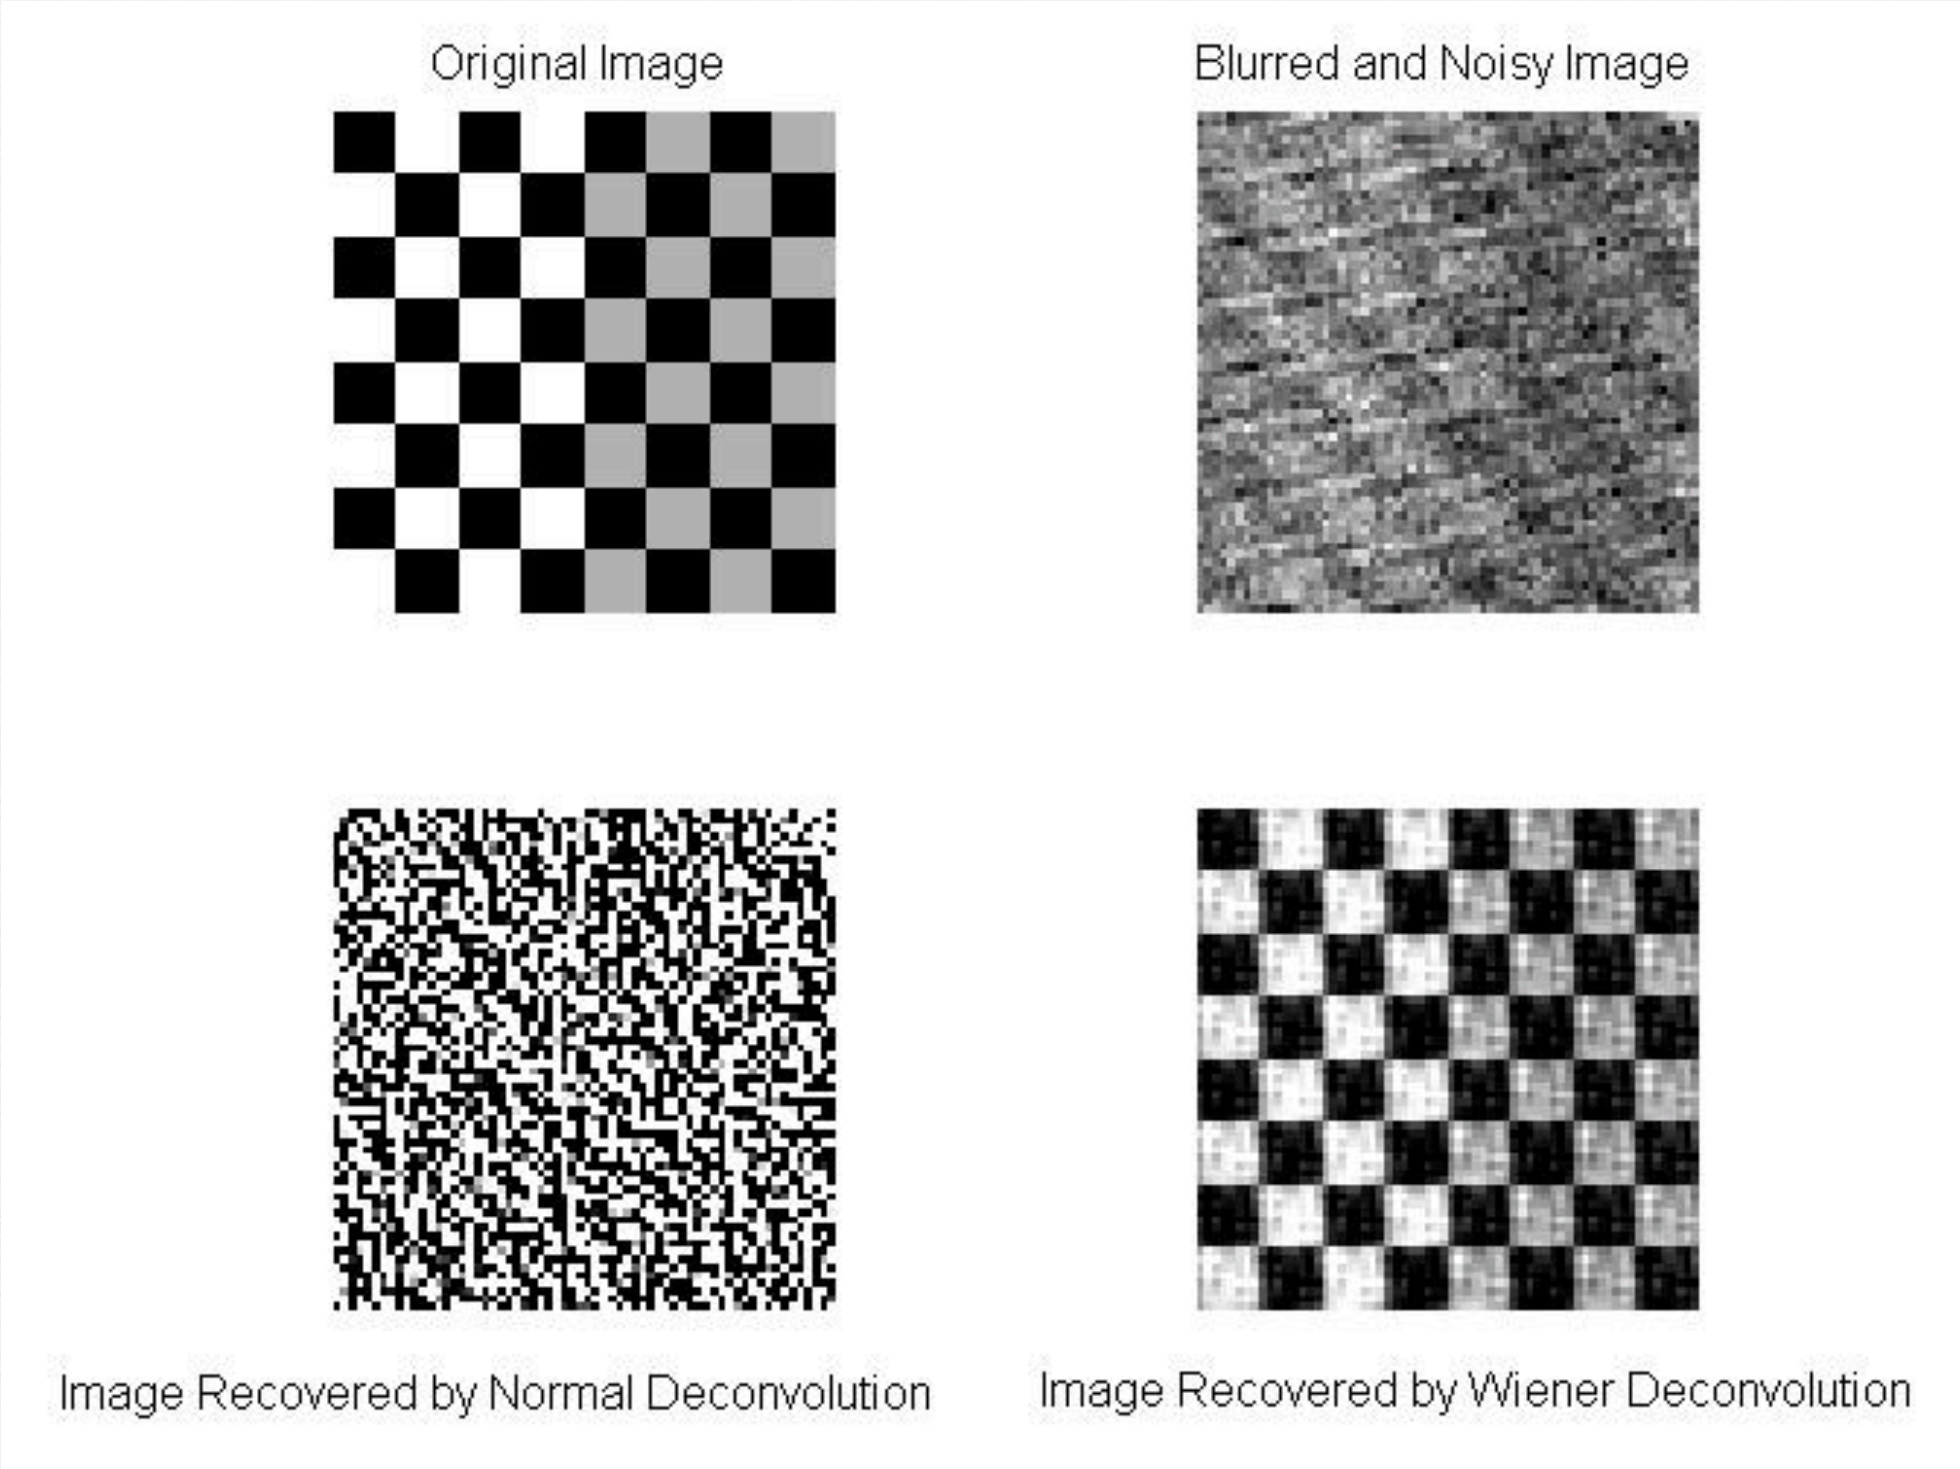

http://cnx.org/content/m15167/latest/

Results from Wiener

blurry with noise

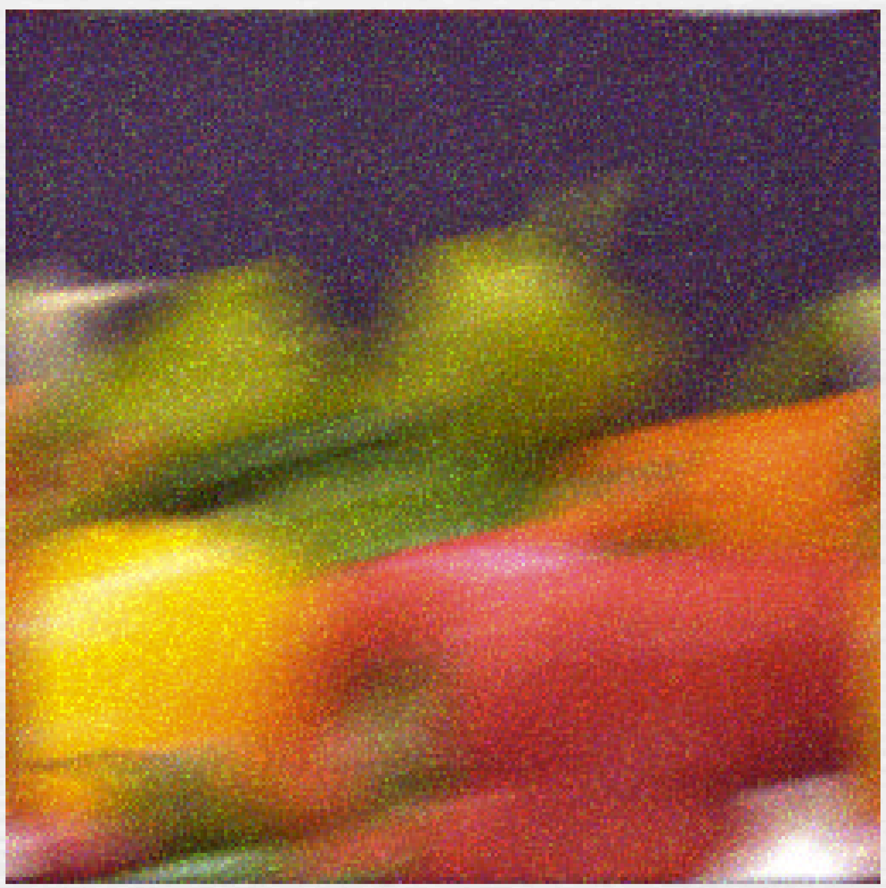

naive deconvolution

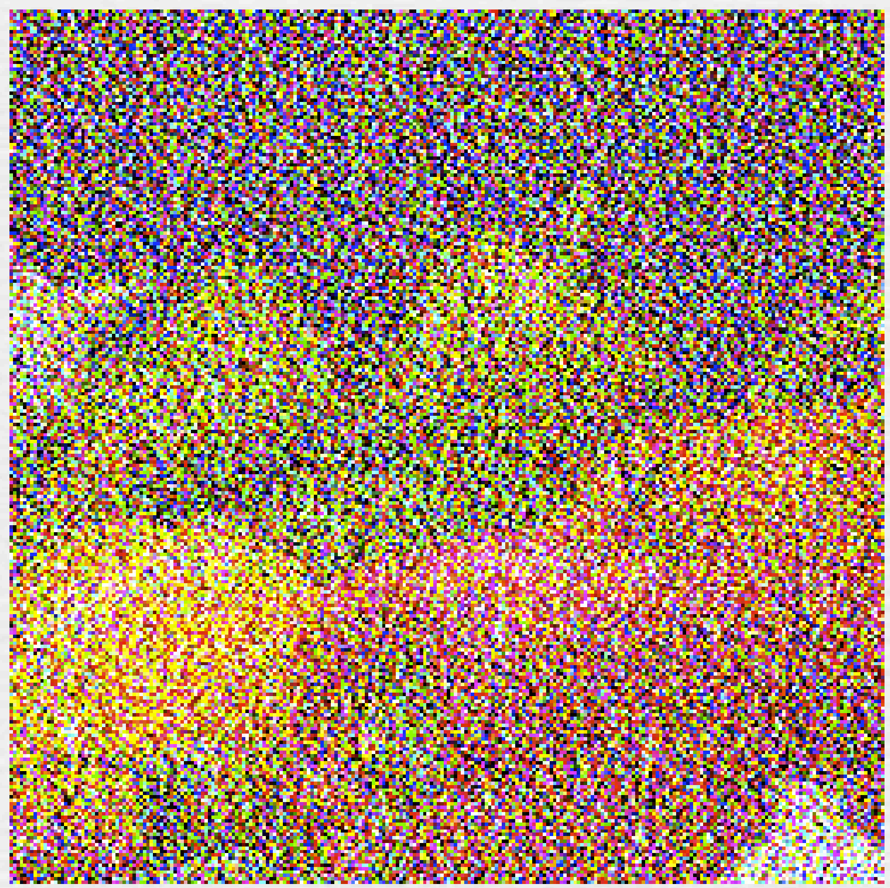

Wiener deconvolution

http://www.mathworks.com/products/demos/image/deblur_wiener/deblur.html

Questions?

Note

- Wiener is derived in Fourier domain
- But in some cases, can be applied directly in primal
- In particular when SNR is $1/\omega^2$
  - Rely on image gradient
  - min $||y-x \otimes g||^2 + k|\nabla x|^2$
    - where k depends on SNR

At a high level

- Wiener can be seen as an example of regularization with a ***prior*** on the signal
- We know the power spectrum
- Other priors are possible
  - Sparsity of the gradient
  - other filters
- Active idea in computer vision

Recap

- Fourier bases diagonalize convolution
  - Convolution > multiplication in Fourier
- Naive deconvolution is division
- Optimal deconvolution takes SNR into account

$$\frac{1}{G(\omega)}\left[\frac{|G(\omega)|^2}{|G(\omega)|^2 + 1/SNR(\omega)}\right]$$

- Deconvolution quality depends on blur frequency response
  - We want a high blur spectrum

Questions?

- Fourier bases diagonalize convolution
  - Convolution > multiplication in Fourier
- Naive deconvolution is division
- Optimal deconvolution takes SNR into account

$$\frac{1}{G(\omega)}\left[\frac{|G(\omega)|^2}{|G(\omega)|^2 + 1/SNR(\omega)}\right]$$

- Deconvolution quality depends on blur frequency response
  - We want a high blur spectrum

Blind deconvolution

- So far we have assumed we know the kernel
- When both x and g are unknown, badly ill-posed
  - It is called blind deconvolution
- See e.g.
  - http://www.wisdom.weizmann.ac.il/~levina/papers/deconvLevinEtal09-MIT-TR.pdf
  - http://cs.nyu.edu/~fergus/research/deblur.html

Questions?

Wavefront coding

Is depth of field a blur?

- Depth of field is NOT a convolution of the image
- The circle of confusion varies with depth
- There are interesting occlusion effects
- (If you really want a convolution, there is one, but in the light field...)

From Macro Photography

Wavefront coding

- **CDM-Optics, U of Colorado, Boulder**
- **Improve depth of field using weird optics & deconvolution**
- **http://www.cdm-optics.com/site/publications.php**
  - The worst title ever: "A New Paradigm for Imaging Systems", Cathey and Dowski, Appl. Optics, 2002

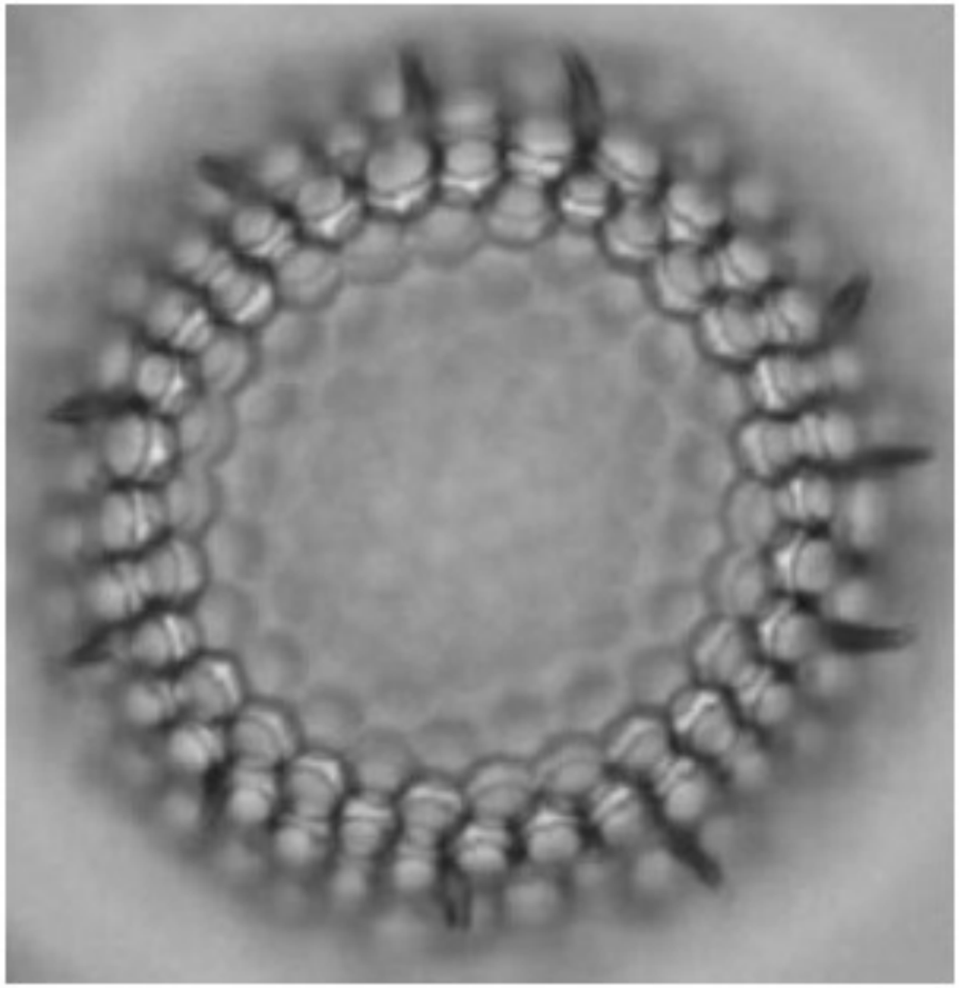

Single-cell algae imaged without wavefront coding.

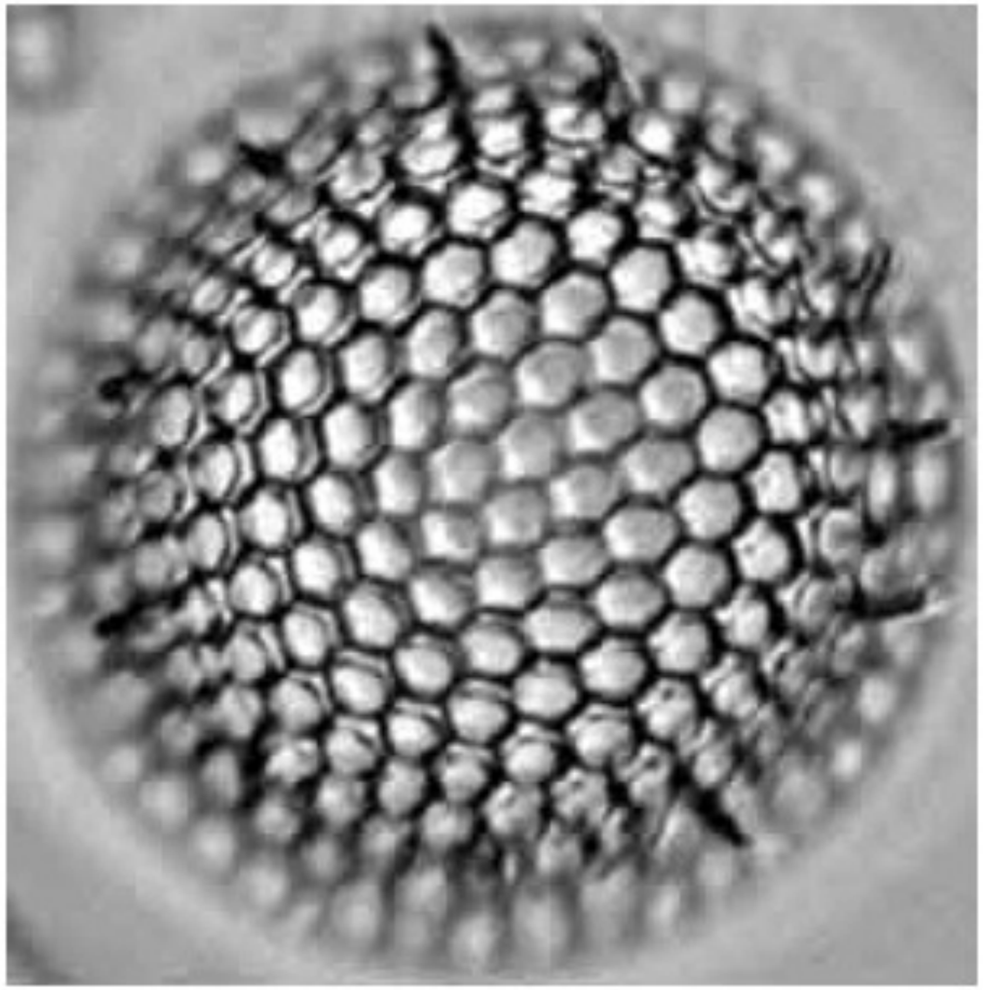

Single-cell algae imaged with wavefront coding.

Wavefront coding

- **Idea: deconvolution to deblur out of focus regions**
- **Problem 1: depth of field blur is not shift-invariant**
  - Depends on depth
    - ➔If depth of field is not a convolution, it's harder to use deconvolution ;-(
- **Problem 2: Depth of field blur "kills information"**
  - Fourier transform of blurring kernel has low frequency response

Wavefront coding

- **Idea: deconvolution to deblur out of focus regions**
  - Problem 1: depth of field blur is not shift-invariant
  - Problem 2: Depth of field blur "kills information"
- **Solution: change optical system so that**
  - Rays don't converge anymore
  - Image blur is the same for all depth
  - Blur spectrum is higher

Wavefront coding

- **Idea: deconvolution to deblur out of focus regions**
  – Problem 1: depth of field blur is not shift-invariant
  – Problem 2: Depth of field blur "kills information"
- **Solution: change optical system so that**
  – Rays don't converge anymore
  – Image blur is the same for all depth
  – Blur spectrum is higher
- **How it's done**
  – Phase plate (cubic lens $z=y^3+x^3$)
  – Will do things similar to spherical aberrations

Ray version

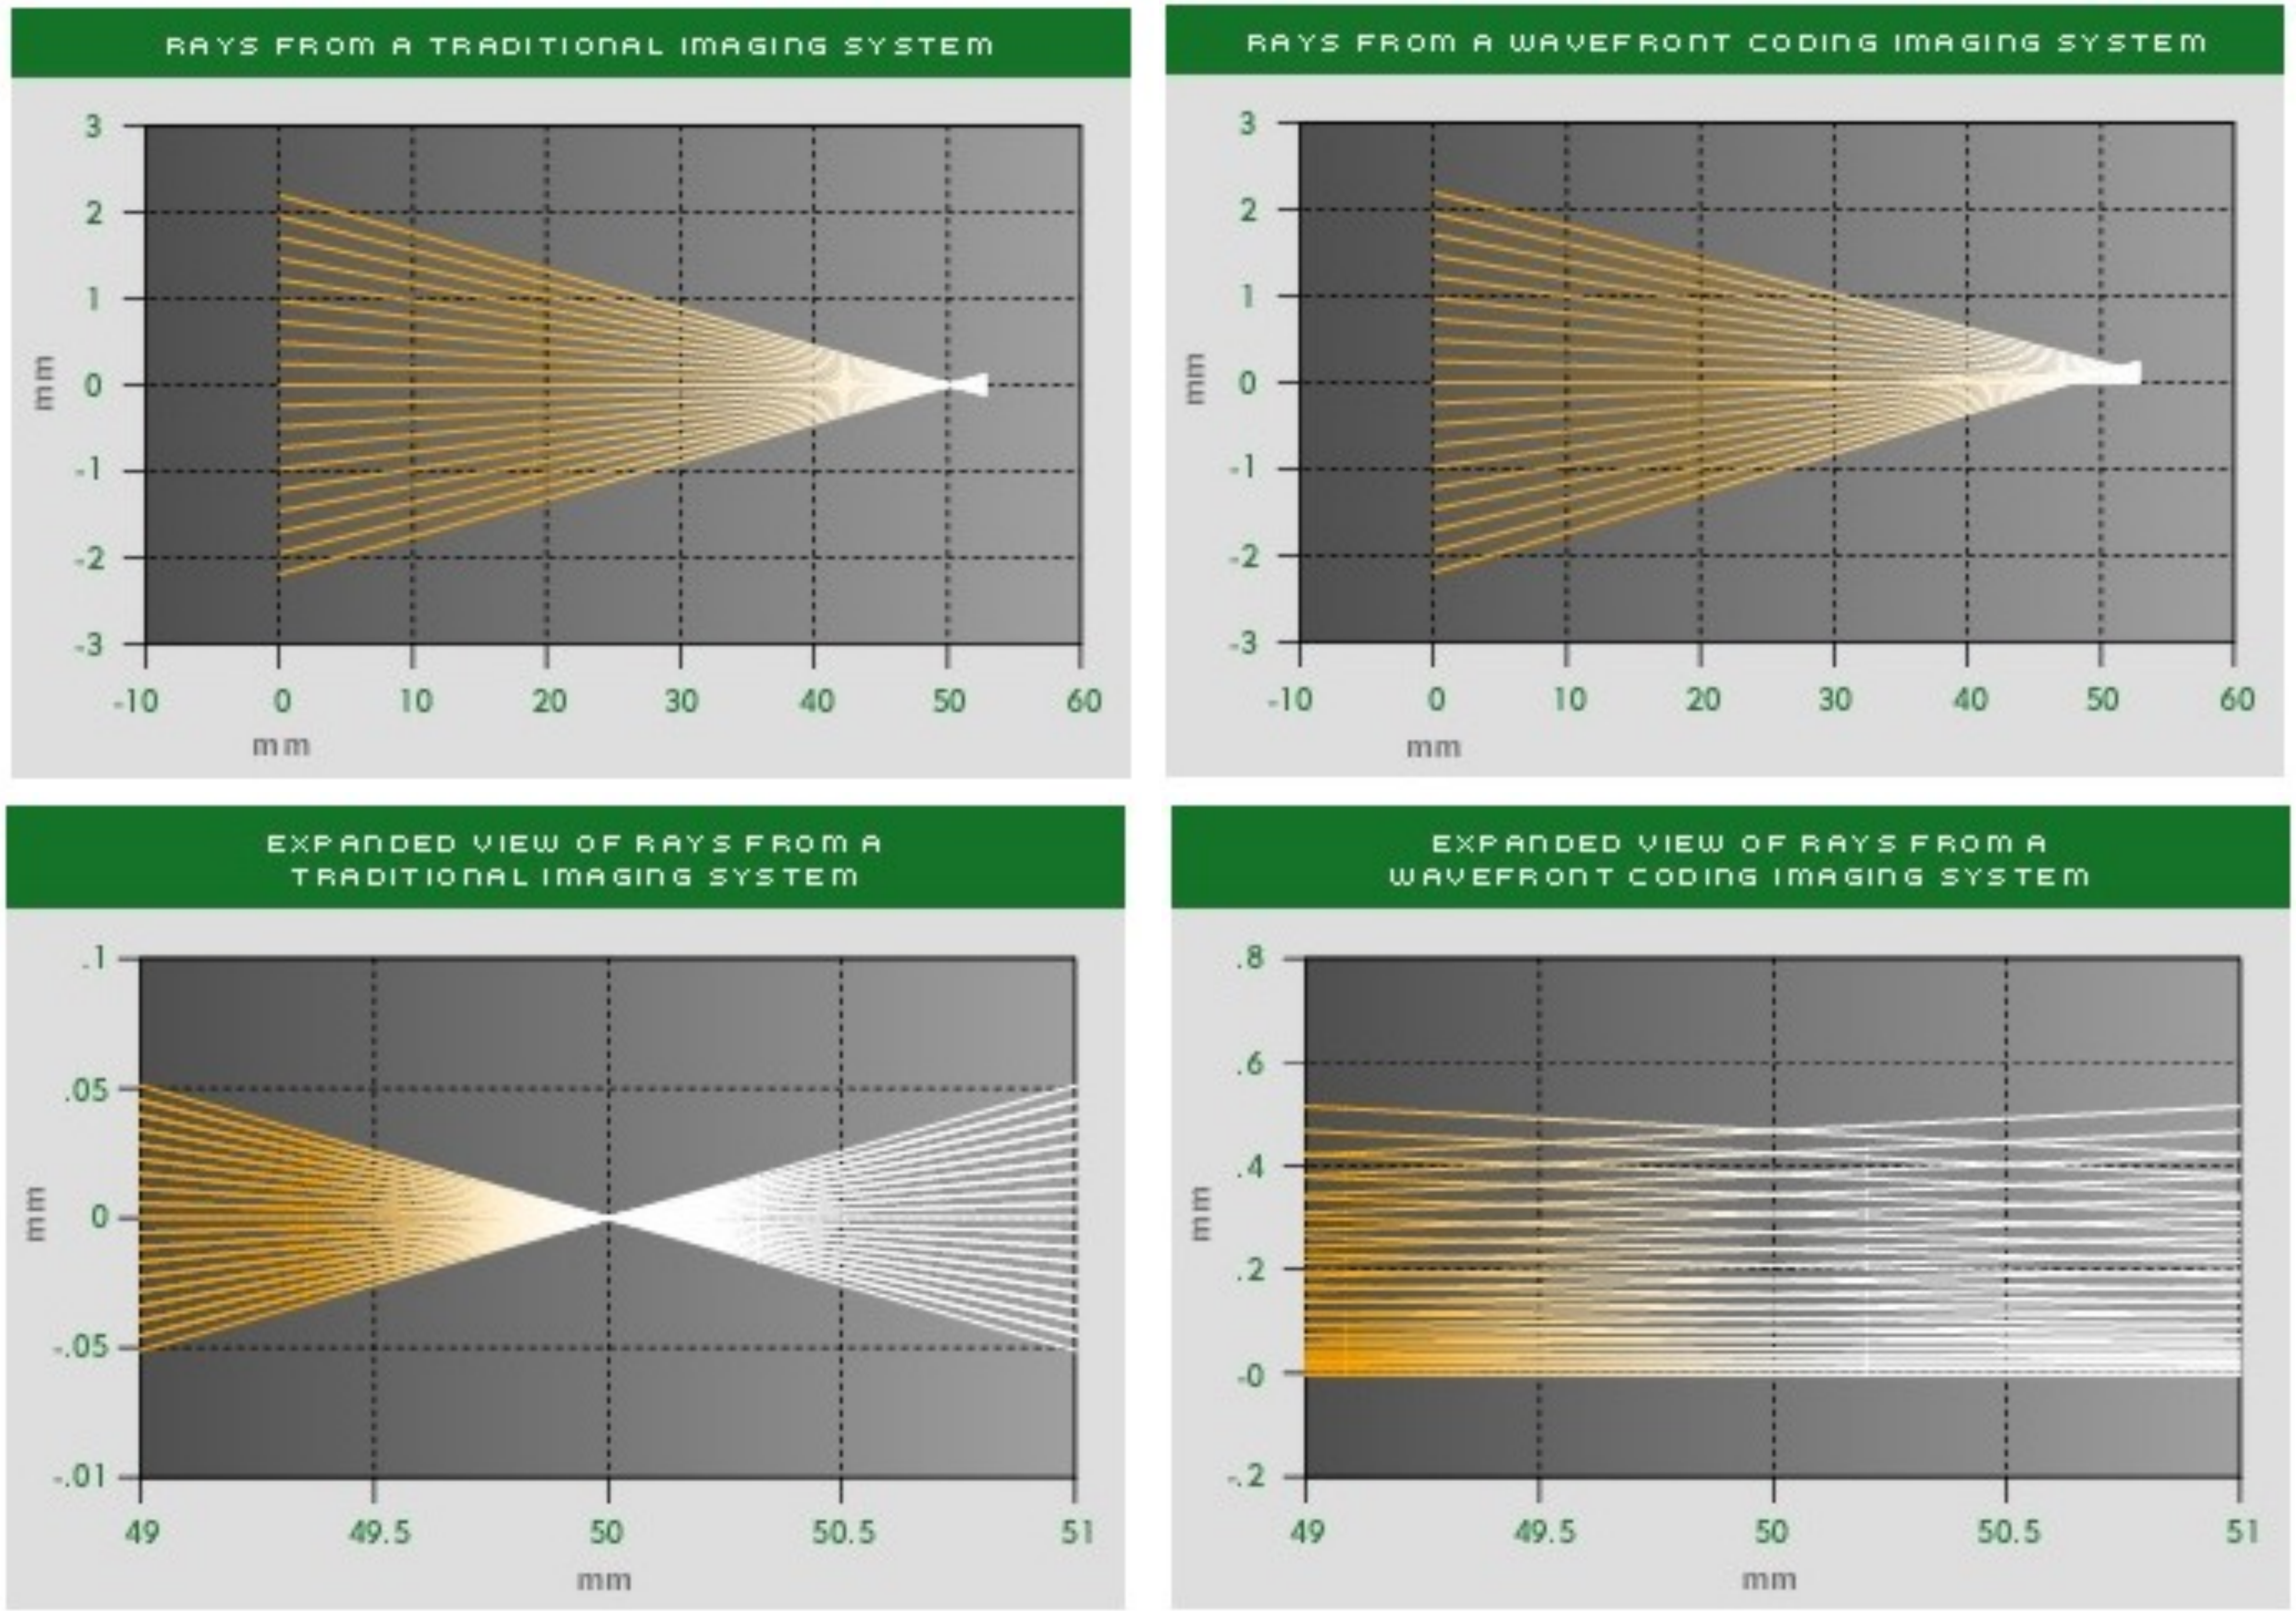
RAYS FROM A TRADITIONAL IMAGING SYSTEM
RAYS FROM A WAVEFRONT CODING IMAGING SYSTEM
EXPANDED VIEW OF RAYS FROM A TRADITIONAL IMAGING SYSTEM
EXPANDED VIEW OF RAYS FROM A WAVEFRONT CODING IMAGING SYSTEM
mm
3
2
1
0
-1
-2
-3
-10
10
20
30
40
50
60
.1
.05
-.05
-.01
49
49.5
50.5
51
.8
.6
.4
.2
-0
-.2

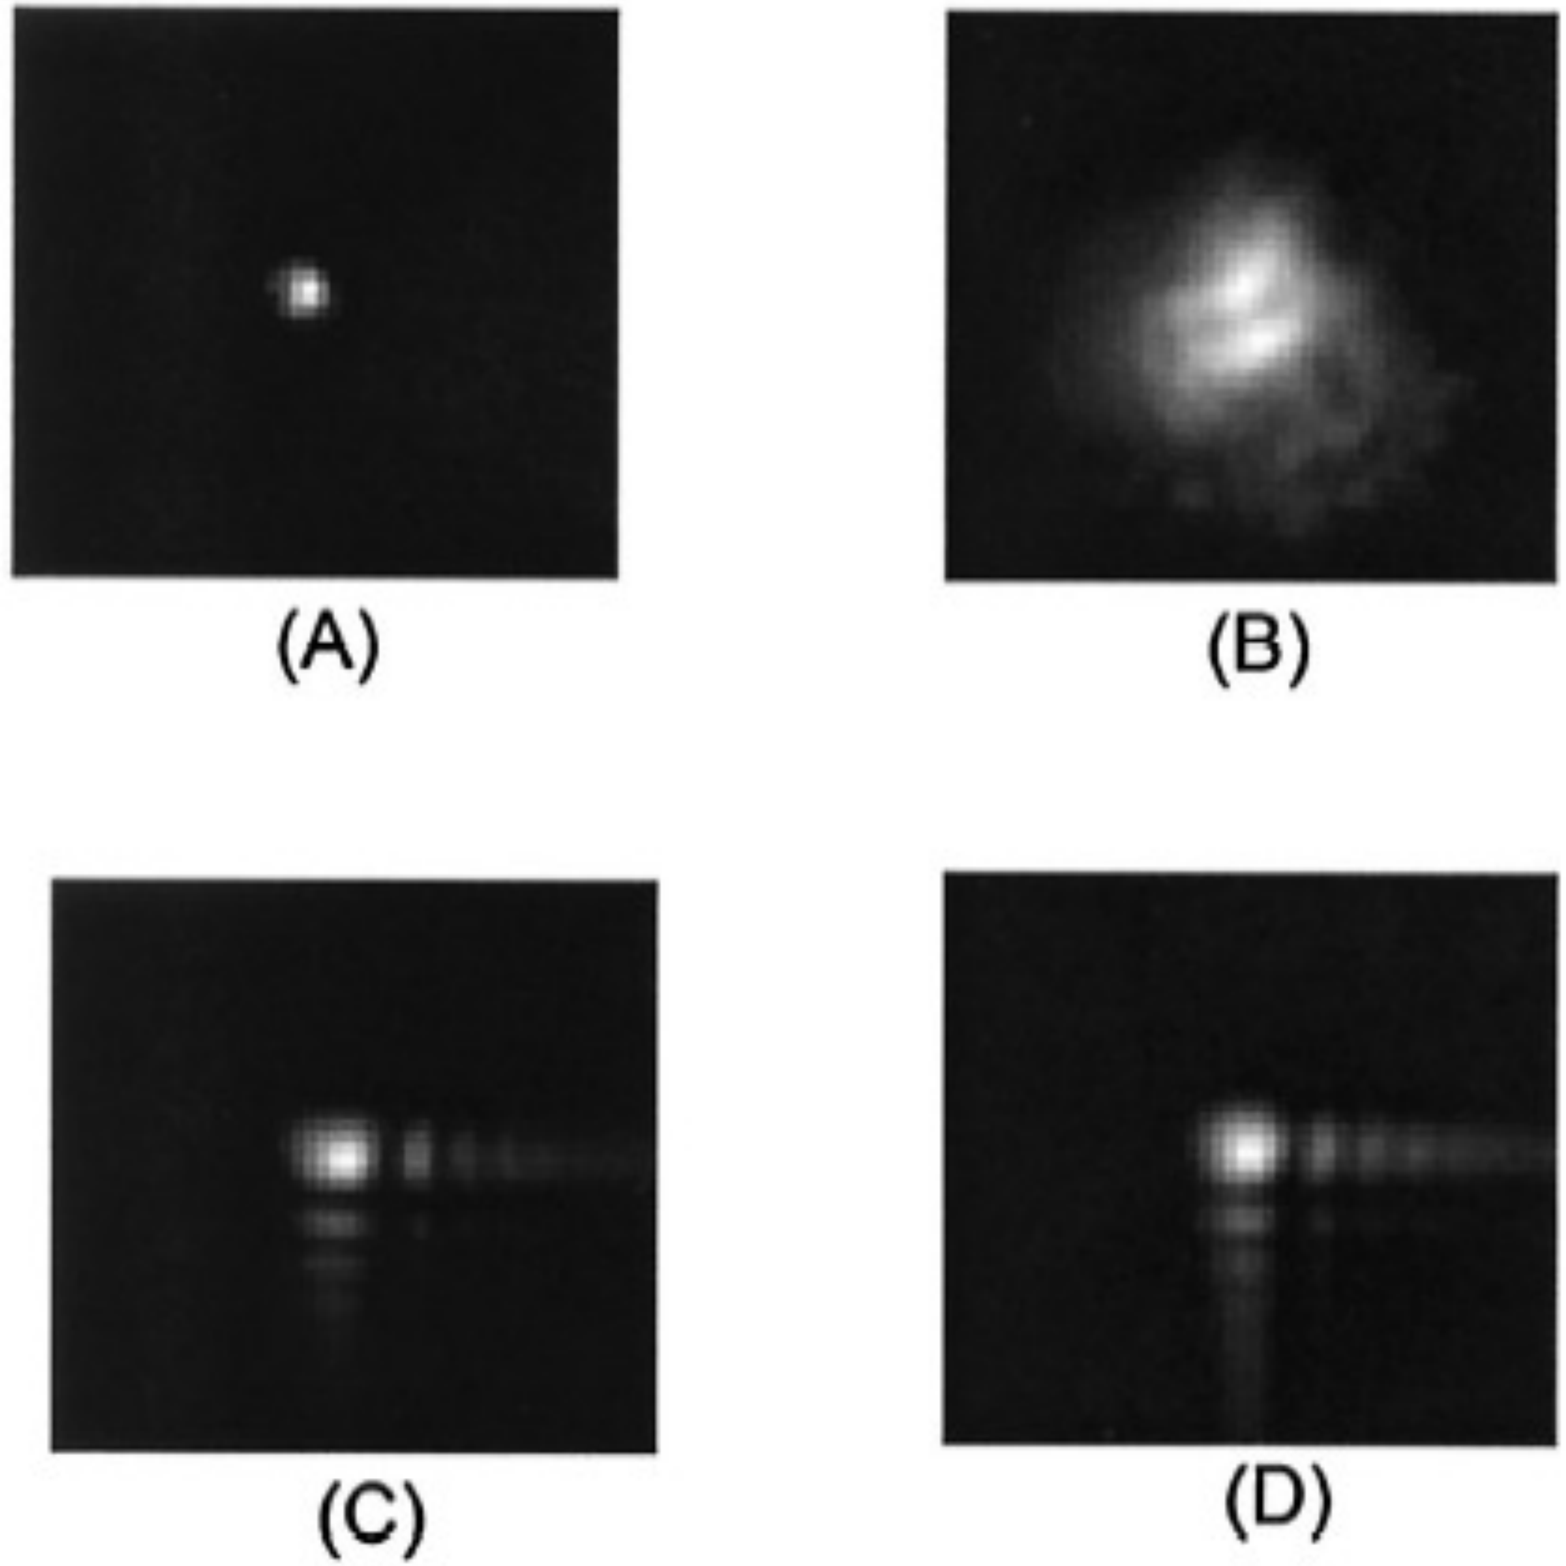

Fig. 3. PSFs associated with the rays of Fig. 2. The PSFs for a normal system are shown for (A) in focus and (B) out of focus. The PSFs for a coded system are shown (C) in the normal region of focus and (D) in the out-of-focus region.

Frequency response (MTF)

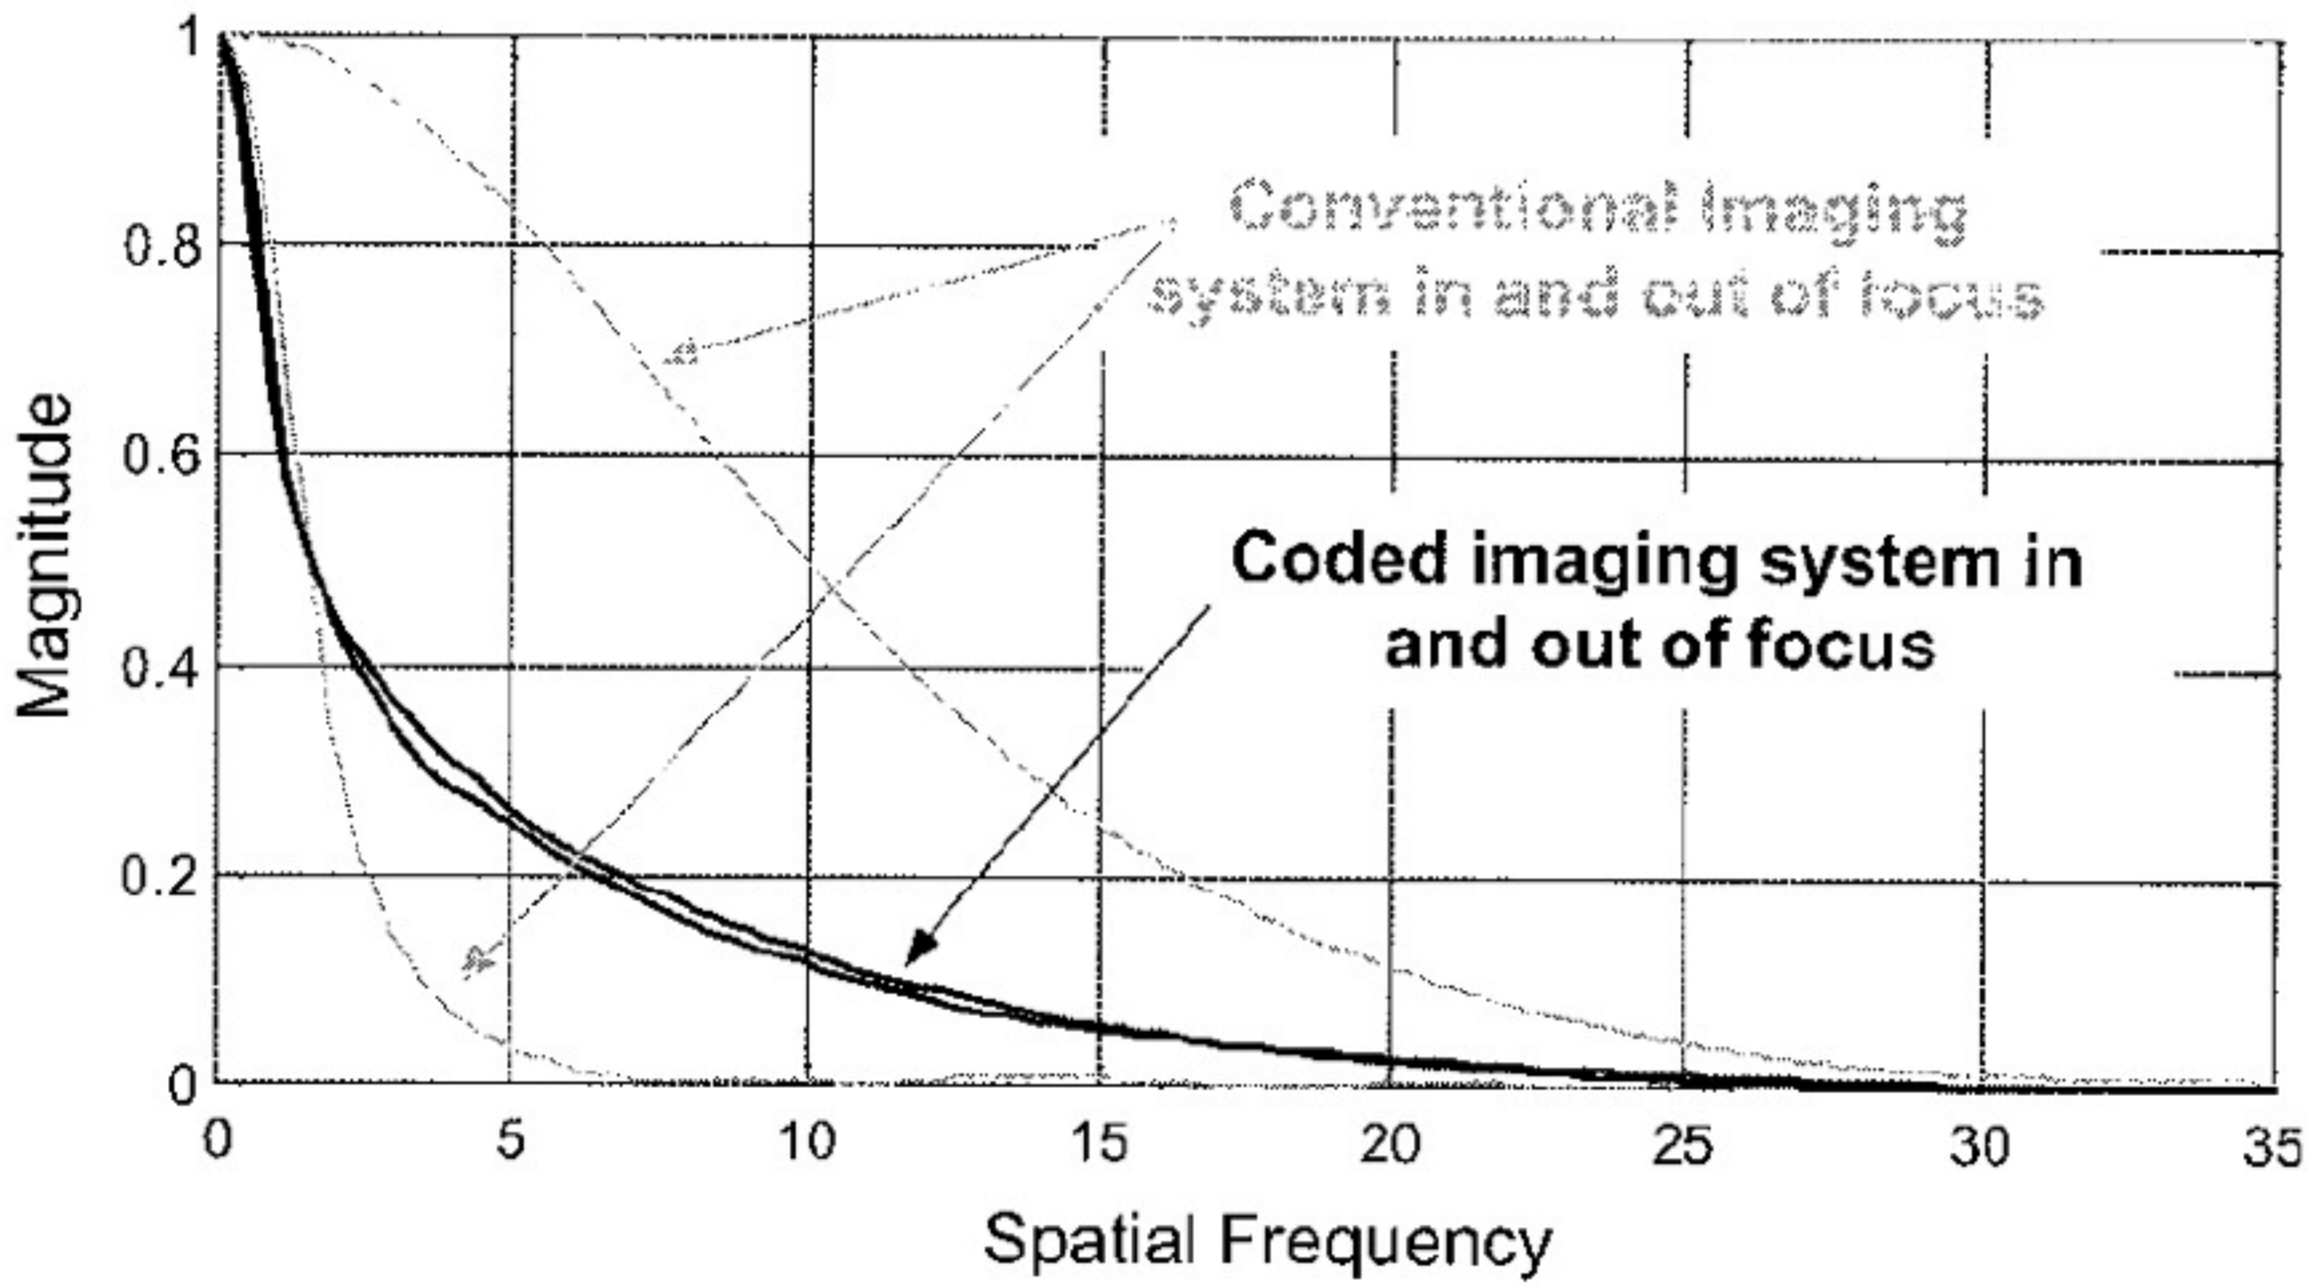

Fig. 5. MTFs corresponding with the PSFs of Fig. 3 for a conventional image in and out of focus and a coded image for the same misfocus values.

Results

Traditional Optical System Image

Intermediate Extended Depth of Field Image

Stopped Down Traditional System Image

Final Wavefront Coded™ Image

Questions?

Philosophy: Image capture

- A sensor placed alone in the middle of the visual world does not record an image

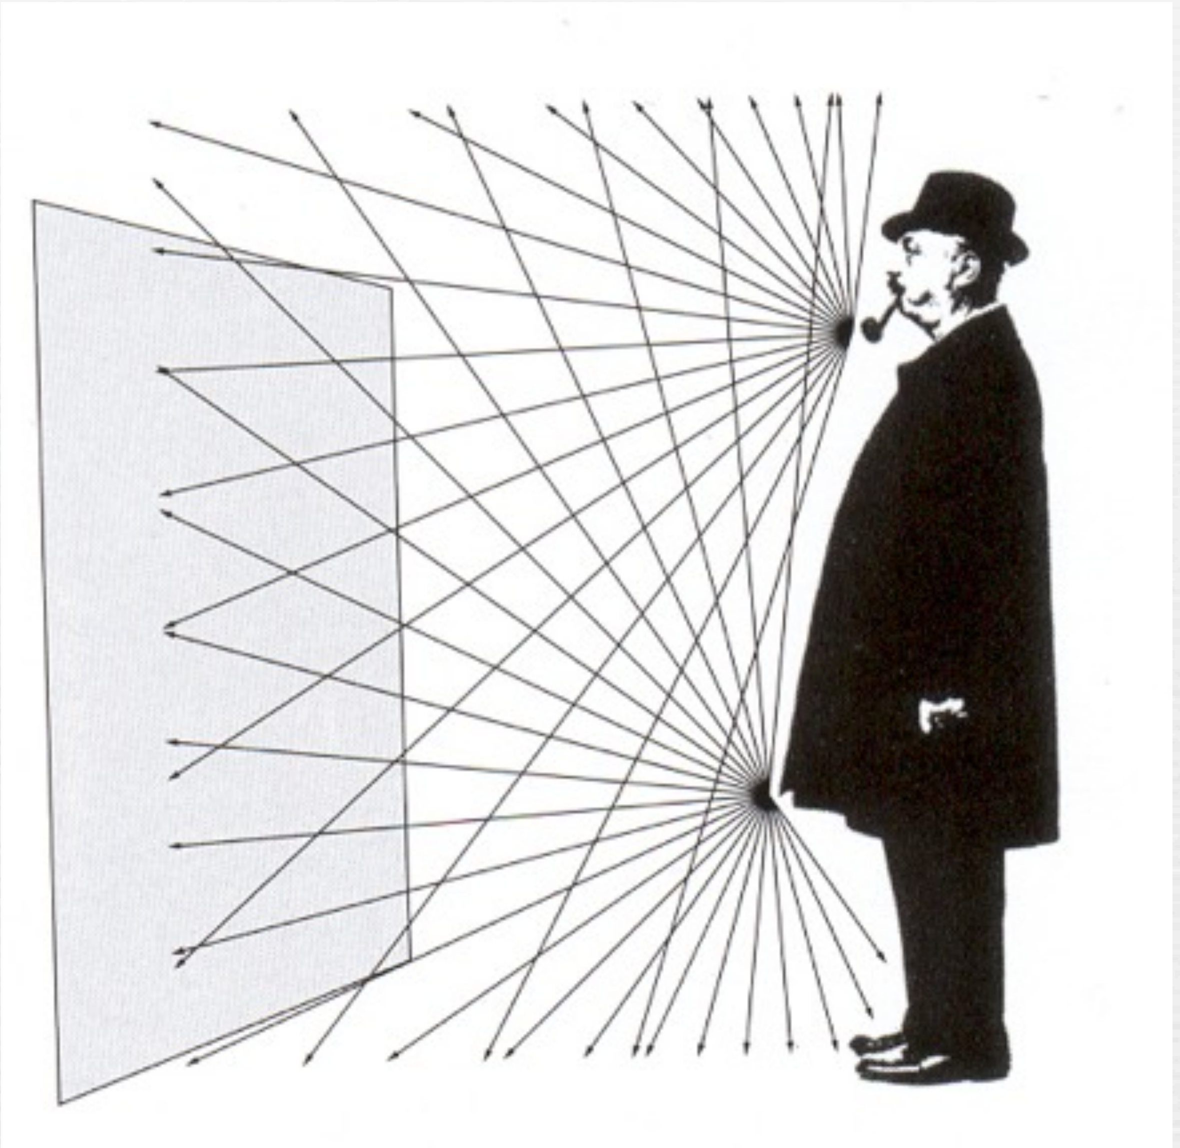

Image capture

- Pinhole allows you to select light rays

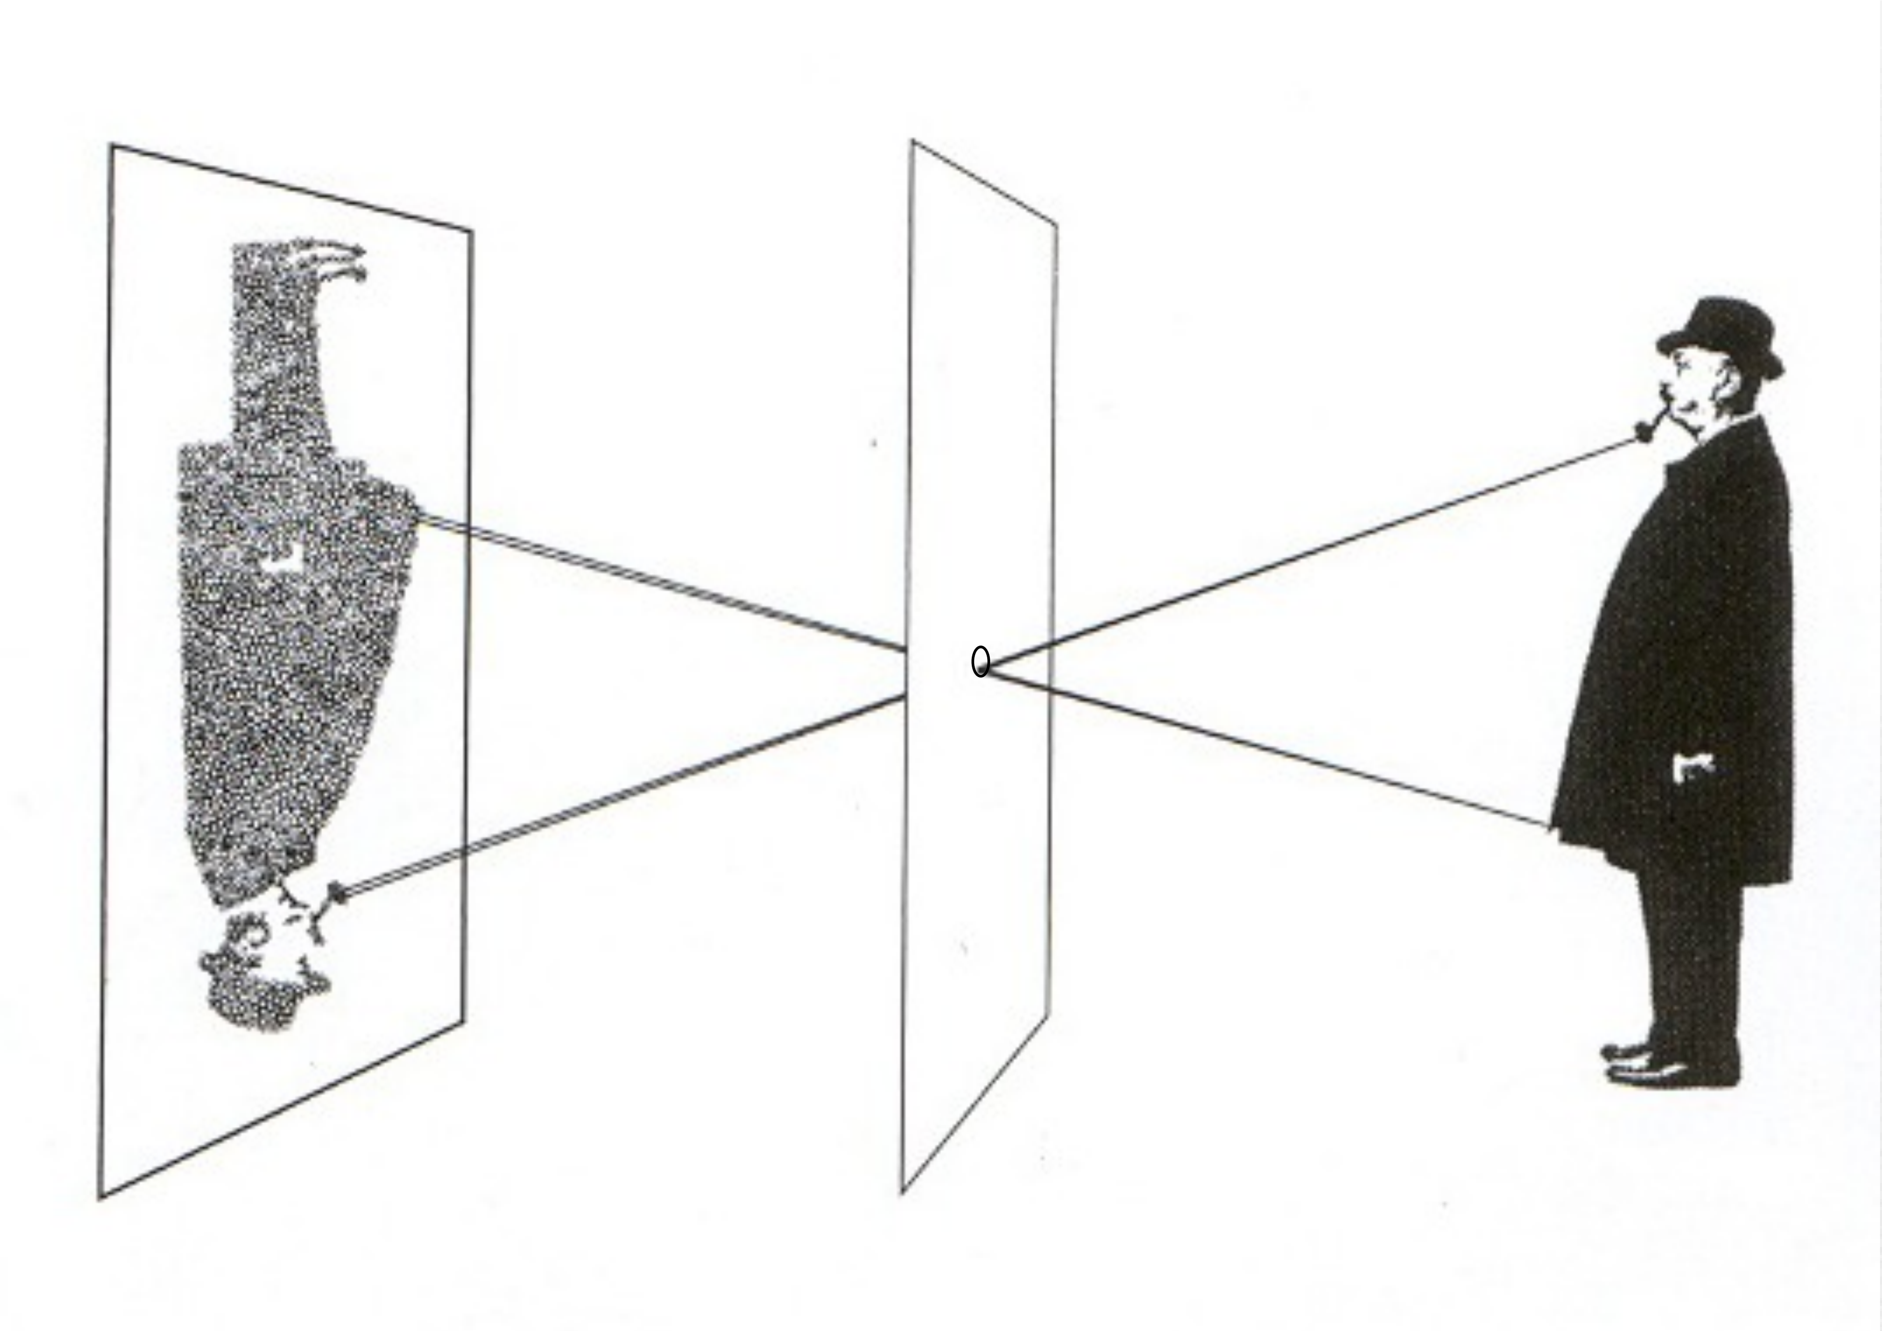

Image formation: optics

- Optics forms an image: selects and integrates light rays

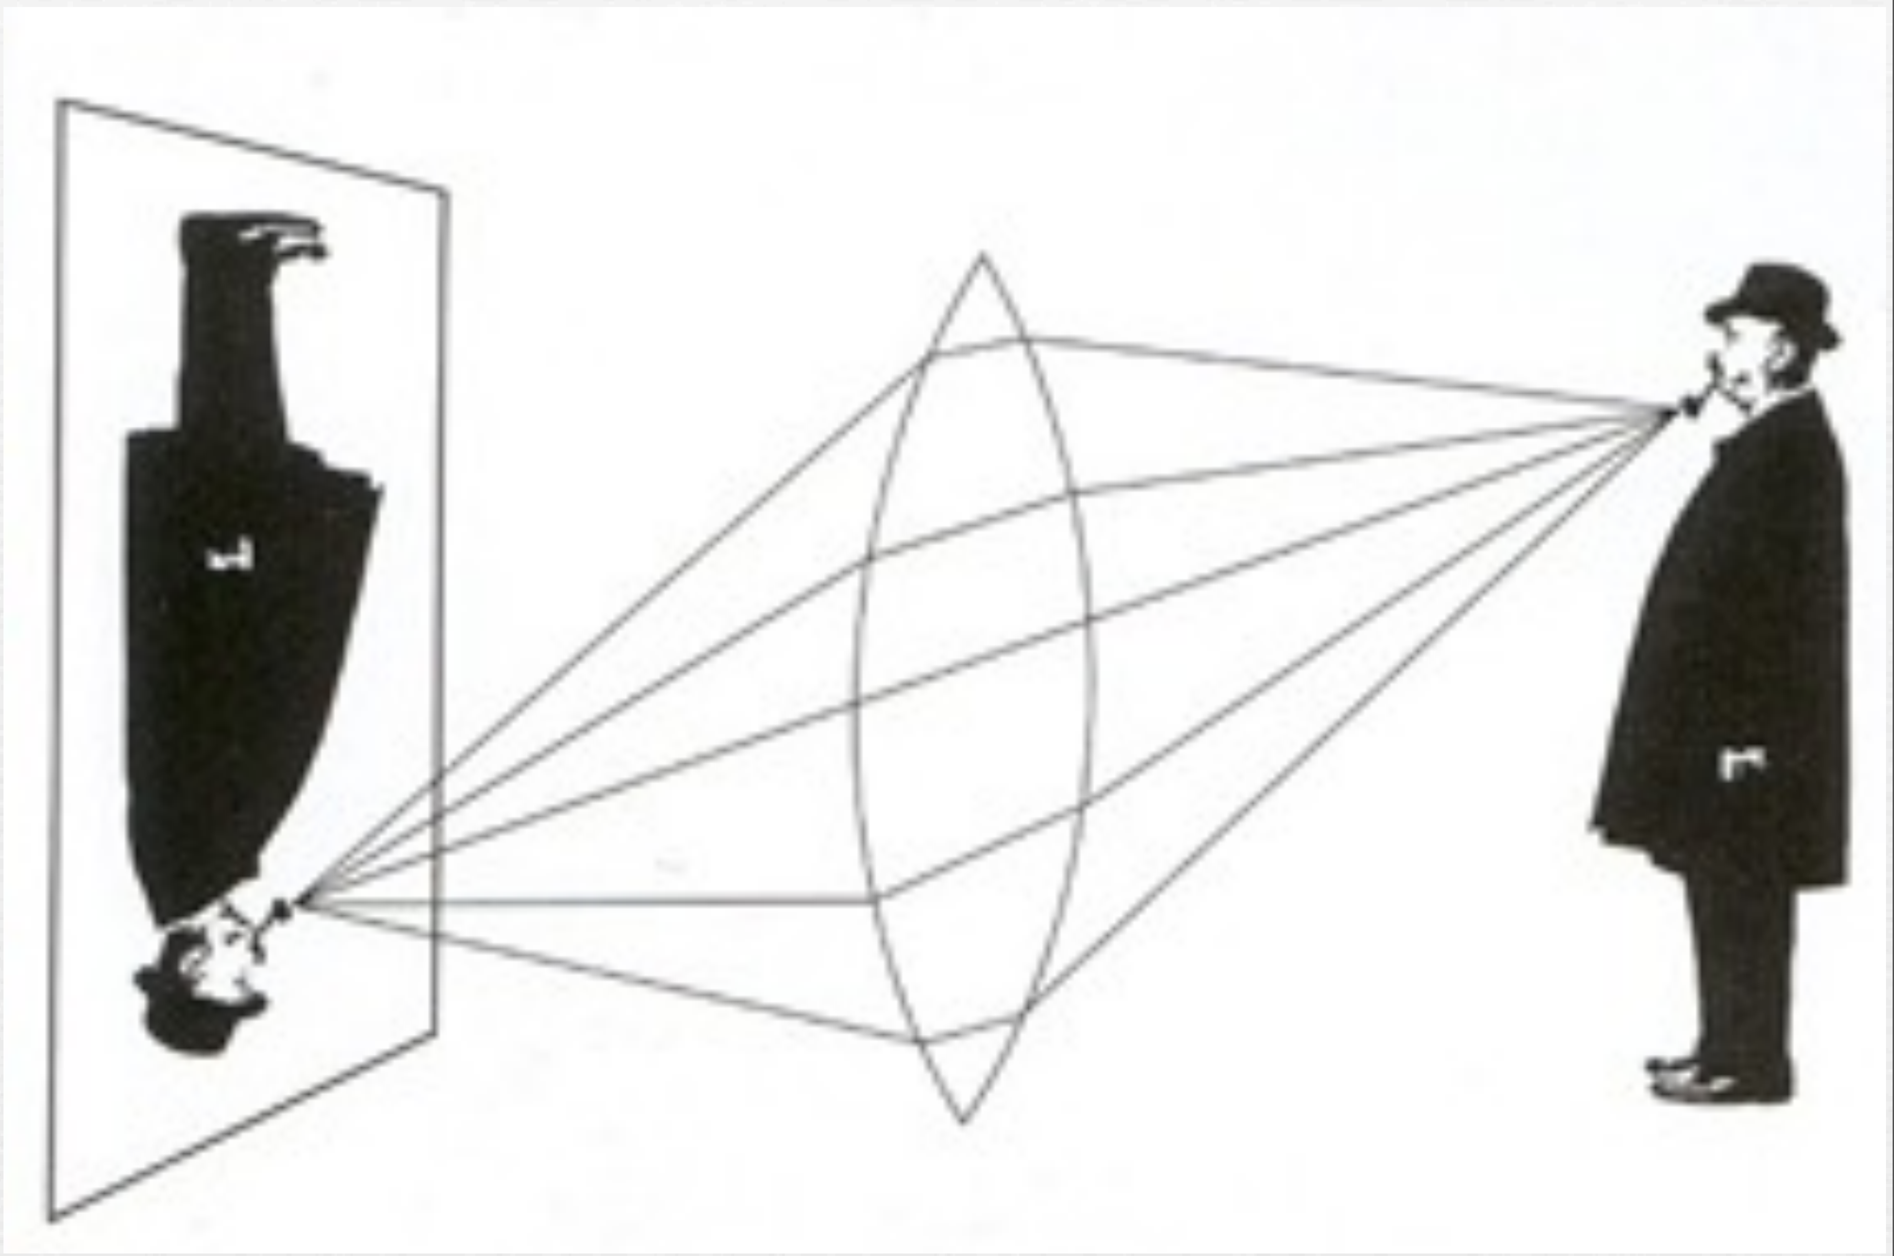

Image formation: computation

- The combination of optics & computation forms the image: selects and combines rays

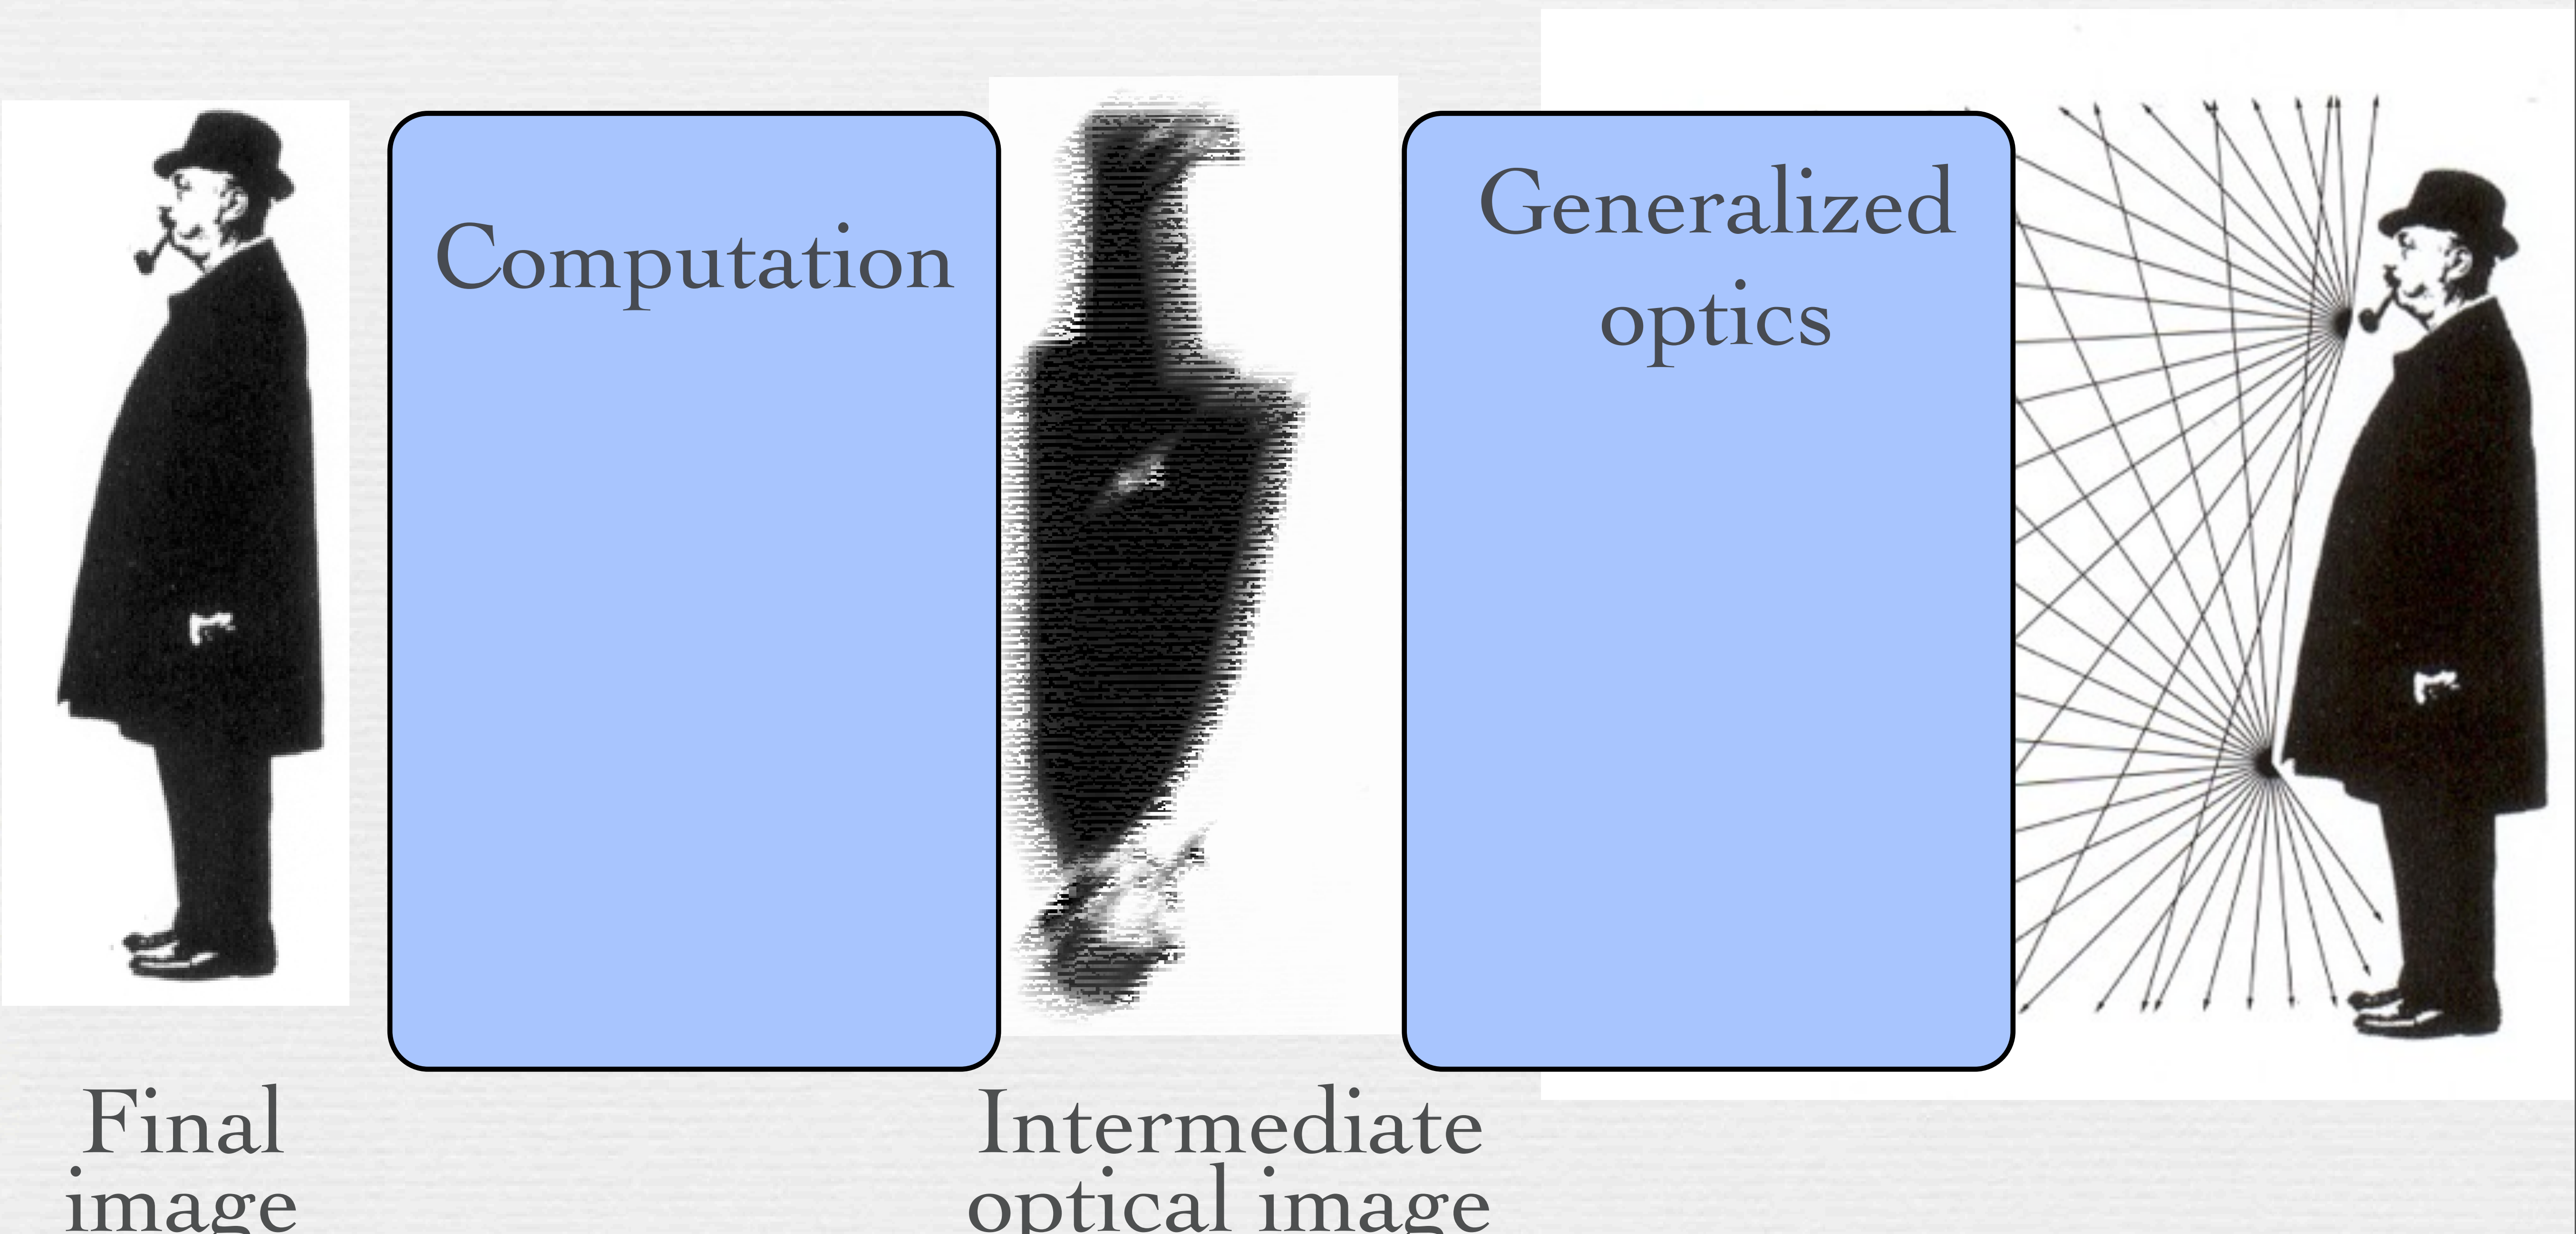

Computational imaging goals

- Better capture information
- Form image as a post-process

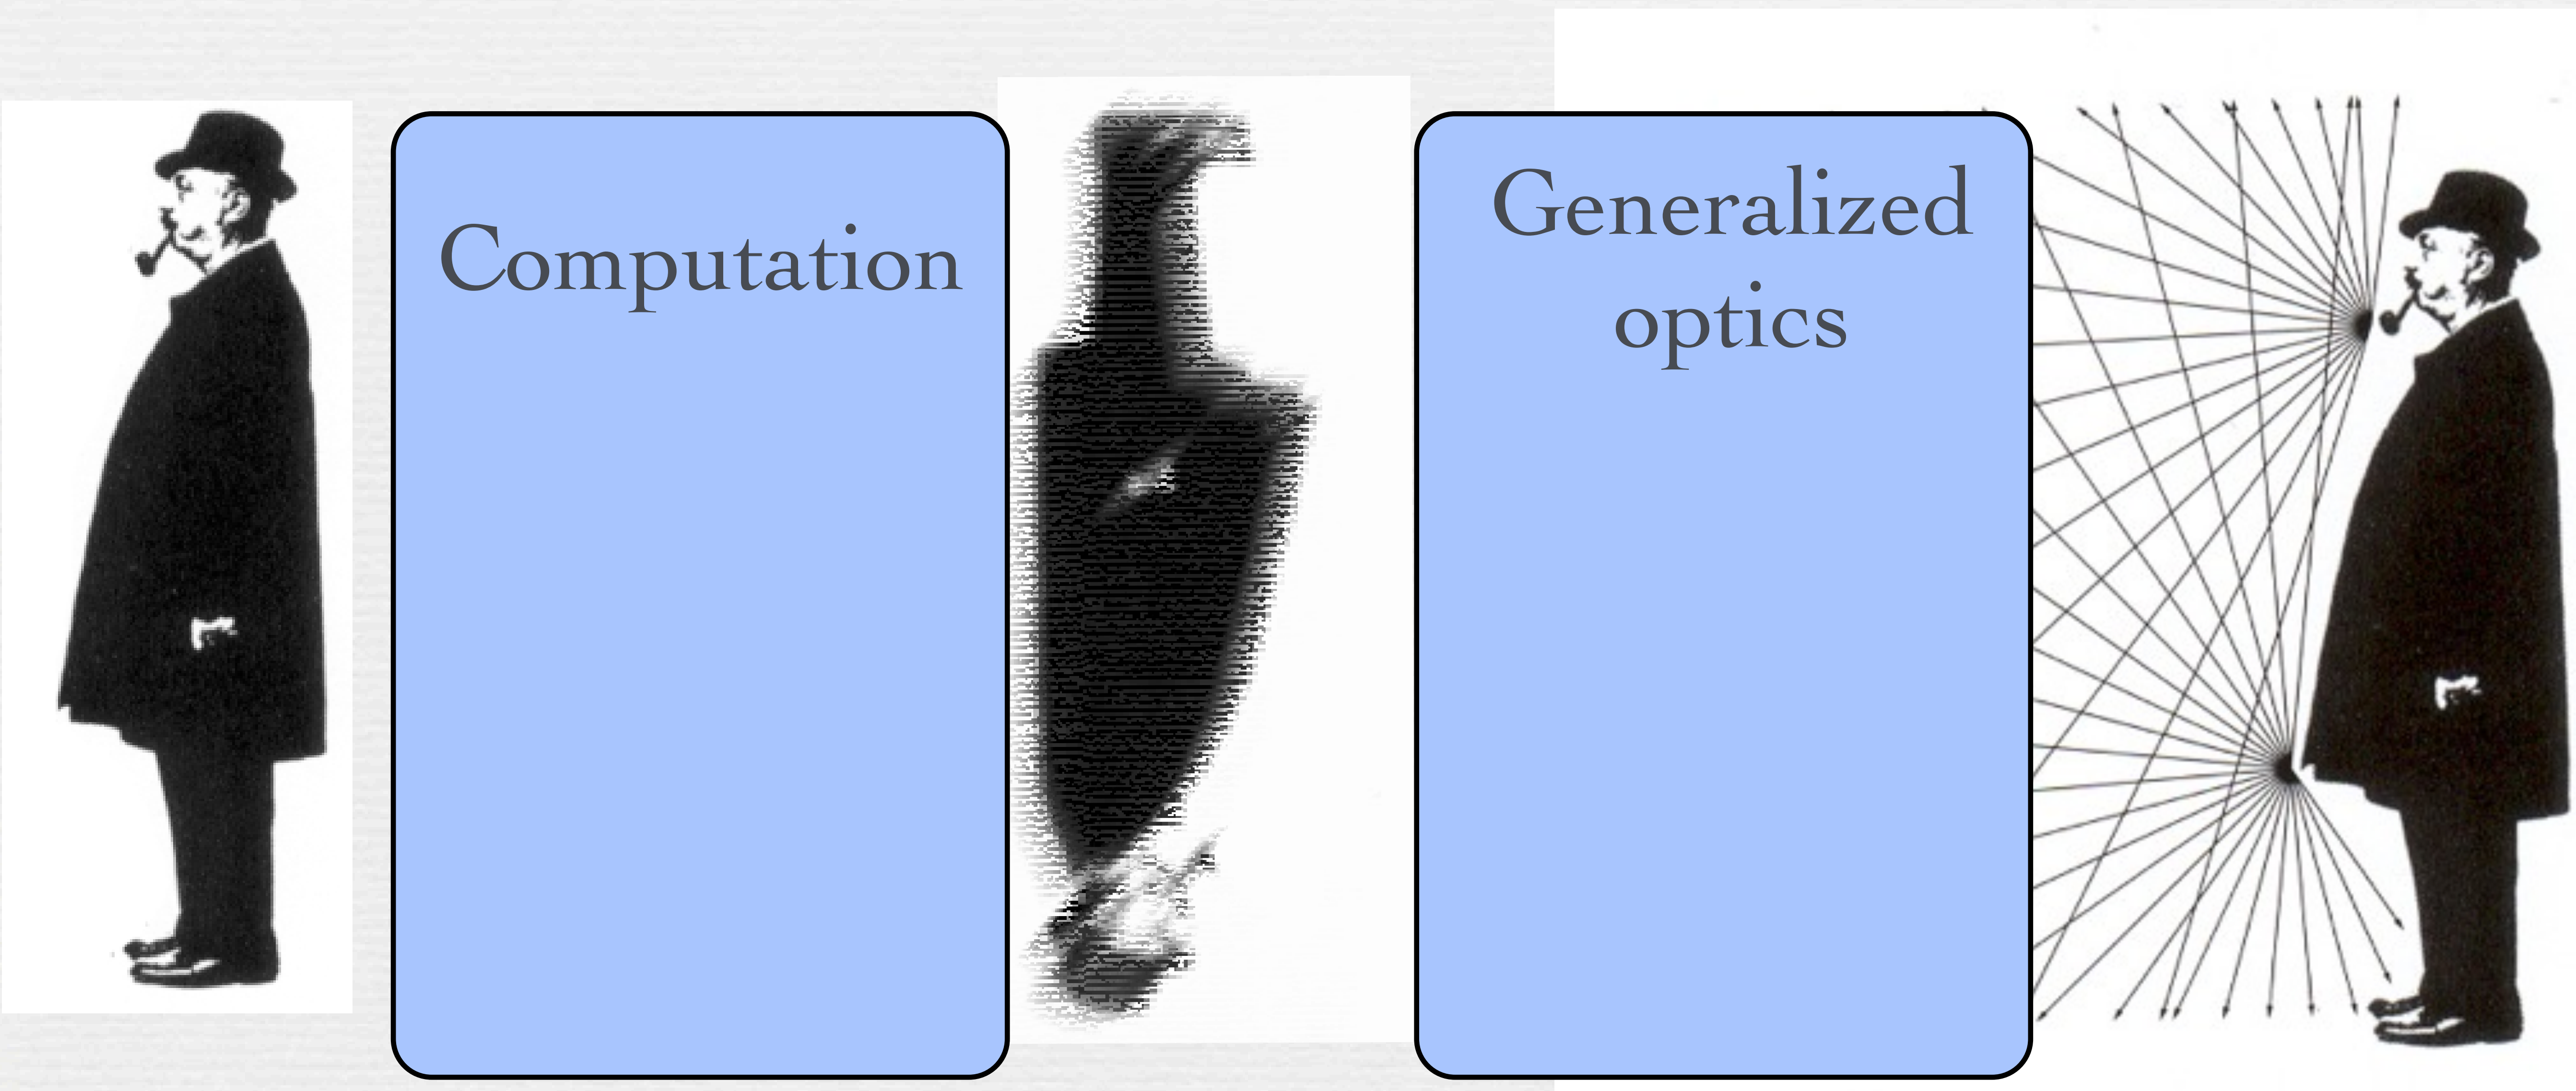

Final image

Intermediate optical image

Better capture information

- Same as communication theory:
  optics encodes , computation decodes
- Code seeks to minimize distortion

Computation

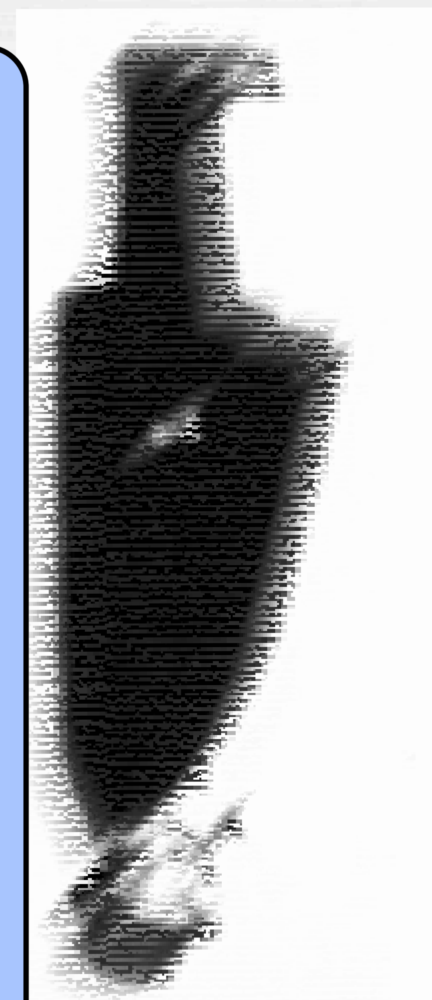

Generalized optics

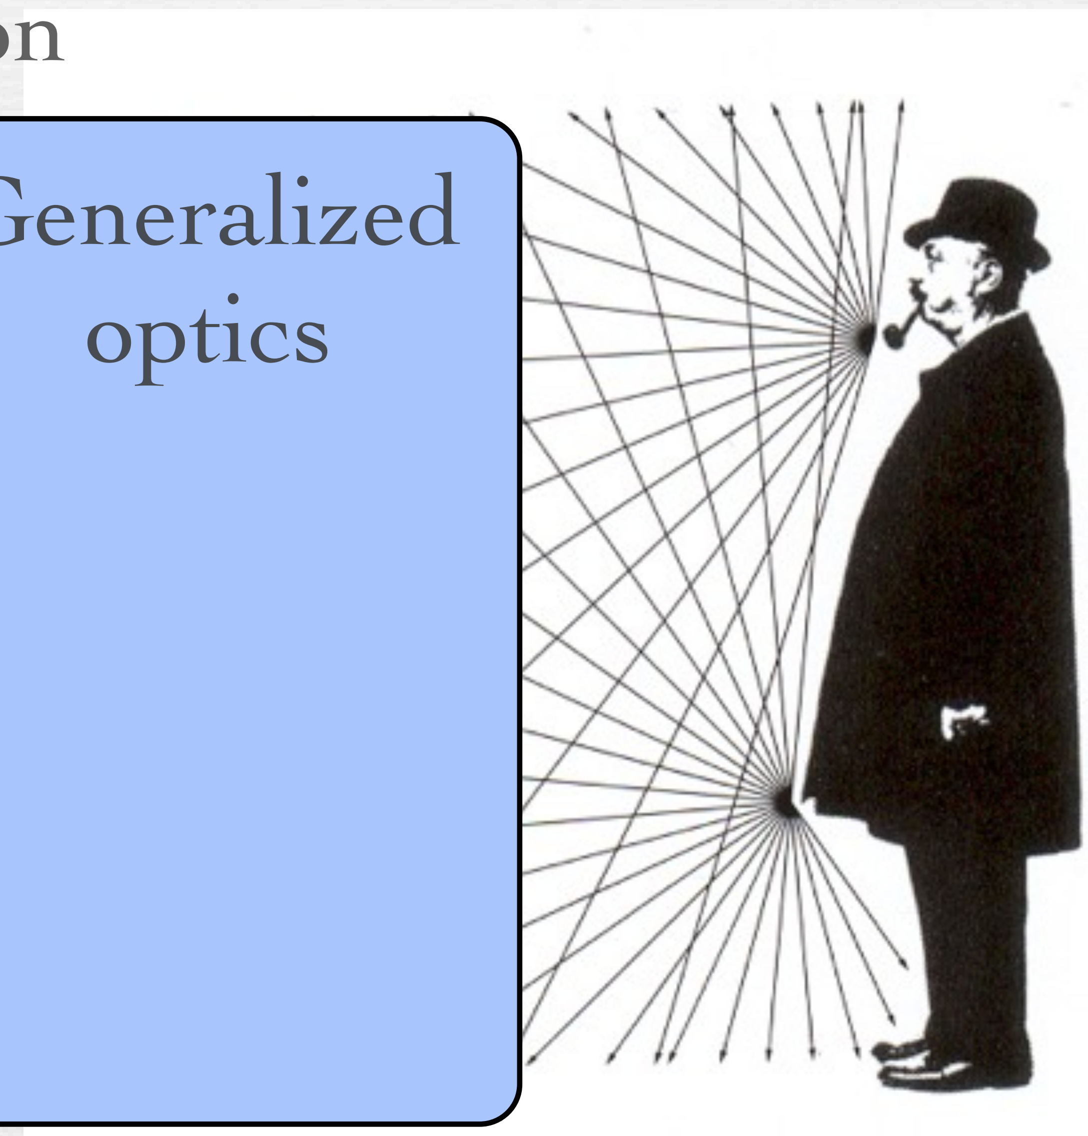

Final image

Intermediate optical image

Form images as a post-process

- The computational part of formation can be done later and multiple times
- e.g., enable refocusing

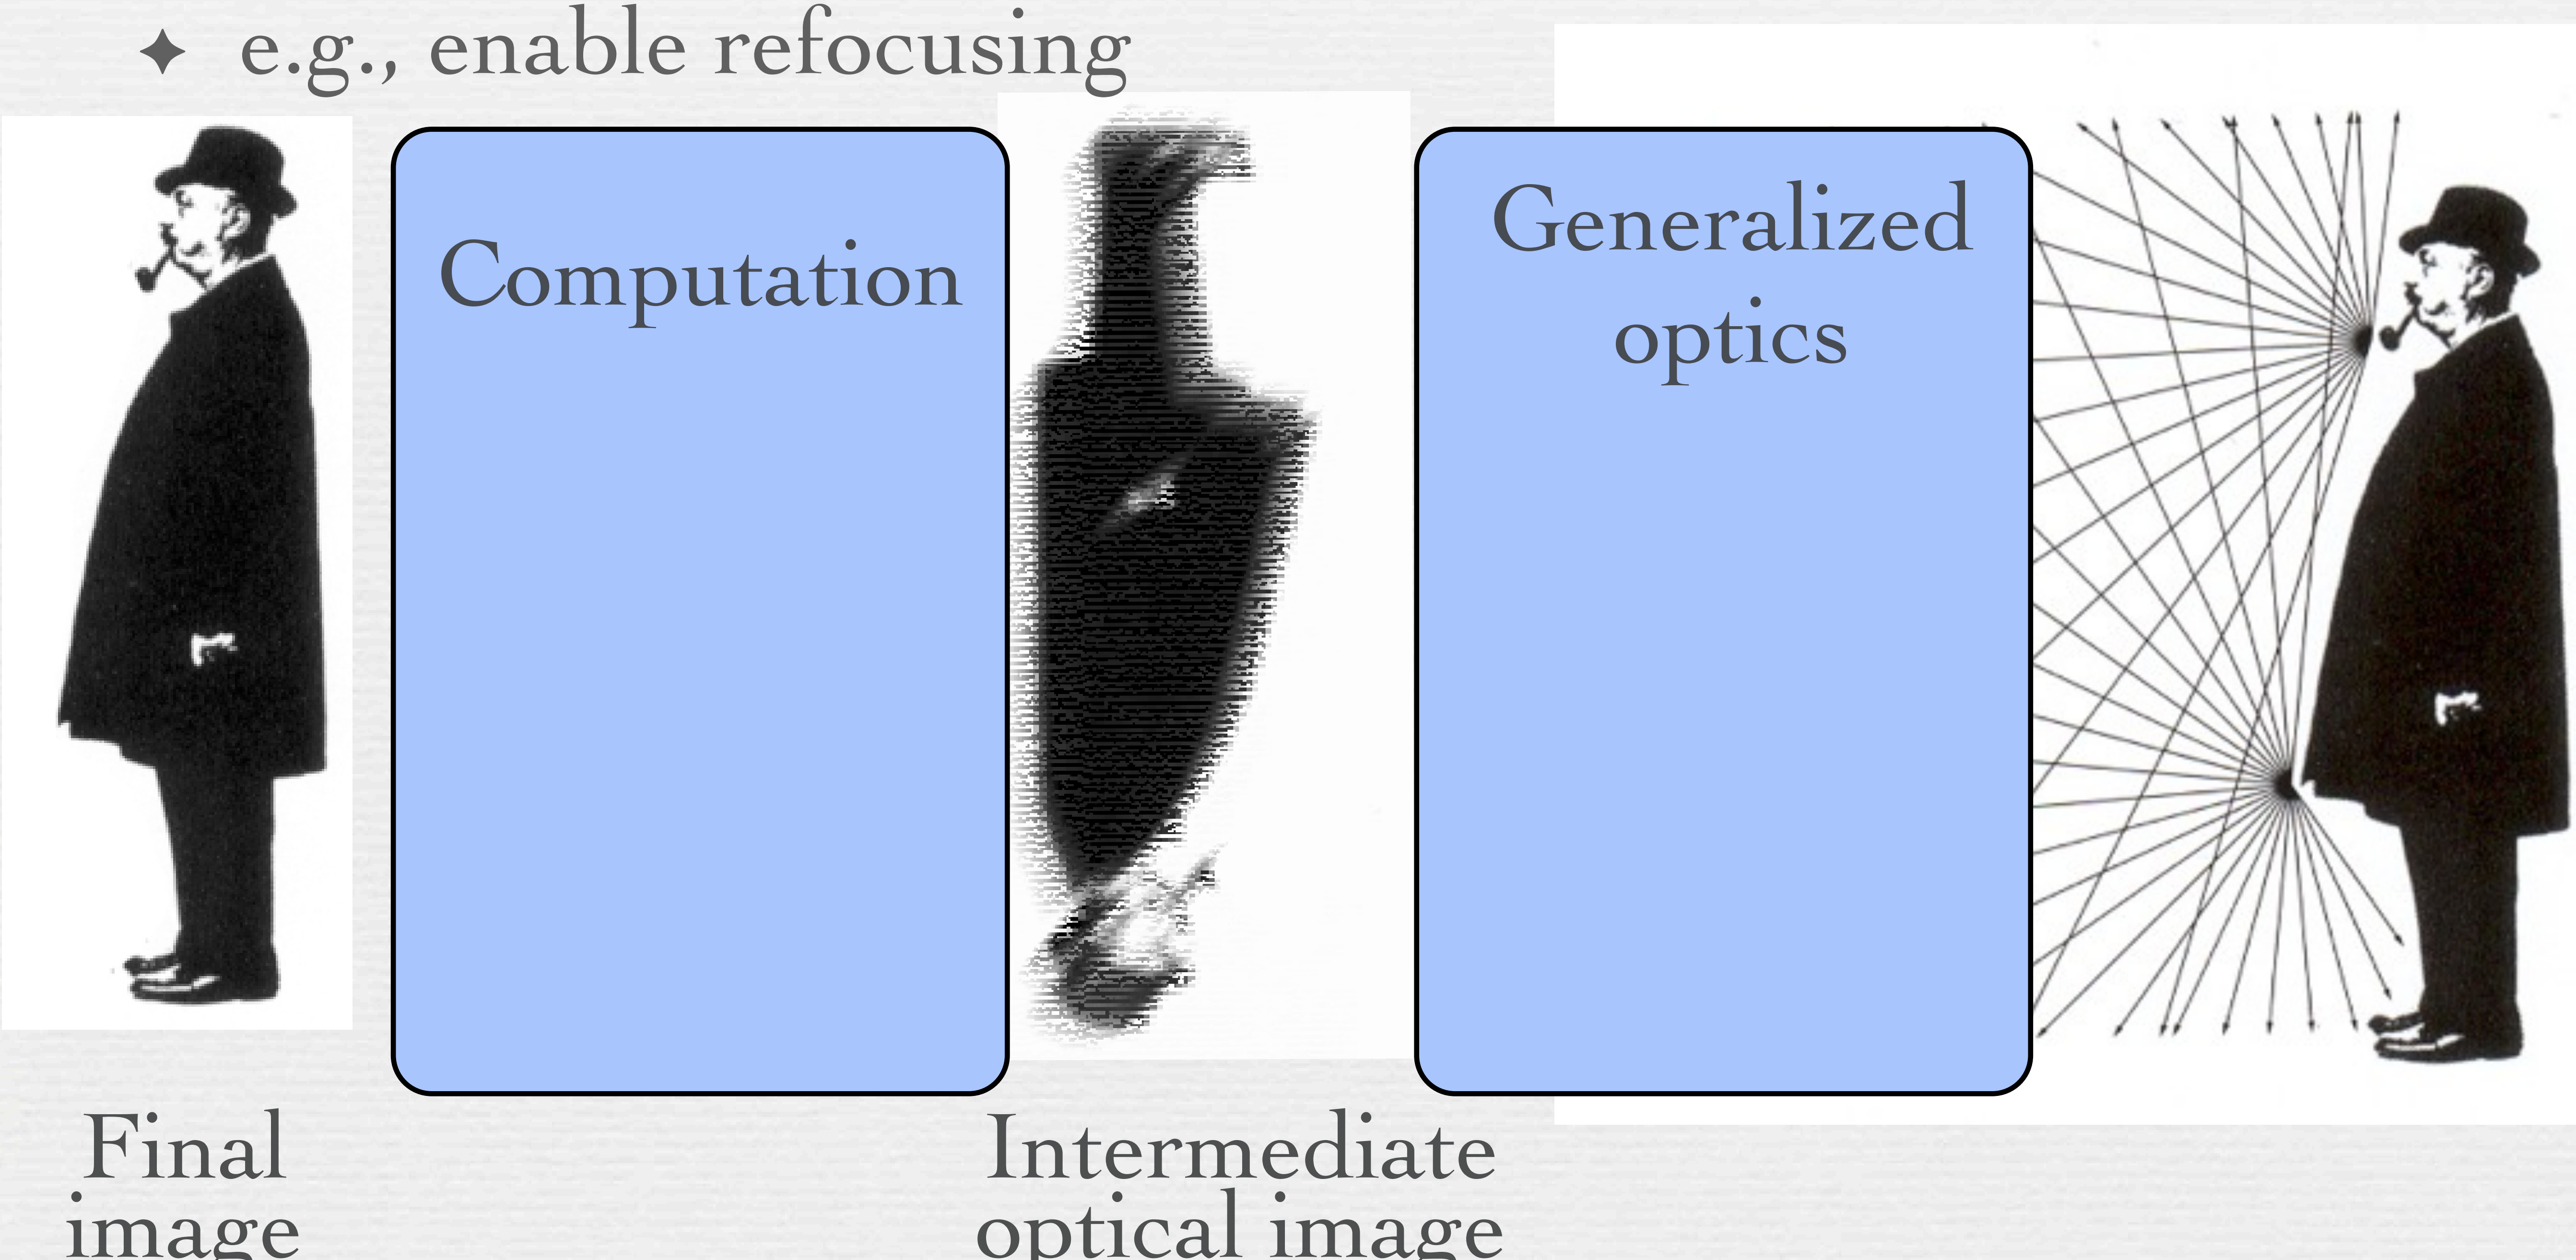

Questions?

Other forms of coded imaging

- **Tomography**
  - e.g. http://en.wikipedia.org/wiki/Computed_axial_tomography
  - Lots of cool Fourier transforms there
- **X-ray telescopes & coded aperture**
  - e.g. http://universe.gsfc.nasa.gov/cai/coded_intr.html
- **Radar, synthetic aperture**
- **Raskar's motion blur**
- **and to some extend, Bayer mosaics**

See Berthold Horn's course at MIT

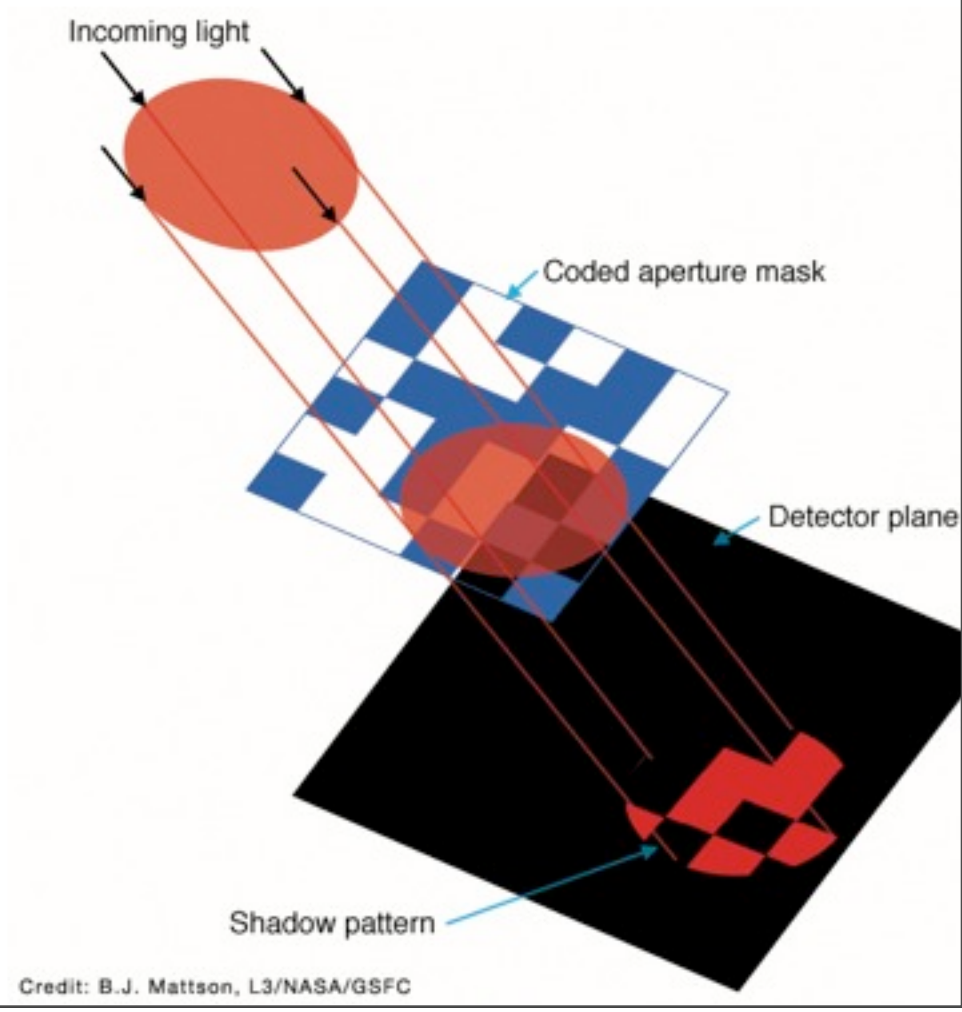

Other computational depth of field extension

Focus sweep
Hausler 72
Nagahara et al. 08

Wavefront coding
Dowski & Cathey 94

Coded aperture
Levin et al. 07
Veeraraghavan et al. 07

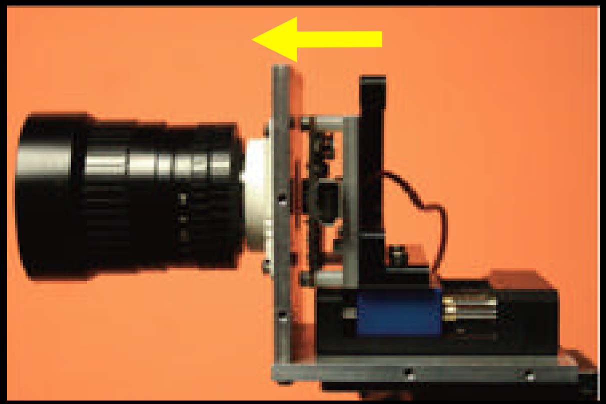

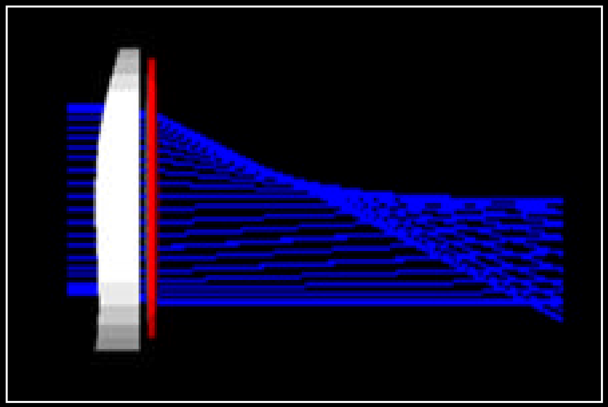

Depth-invariant blur

Depth-varying blur
Depth estimation required

**And in next slide,
lattice focal lens**

Levin et al. 09

New solution: The lattice-focal lens

Levin et al. 09: assembly of subsquares with different focal powers

➡ **each element focuses on a different depth**

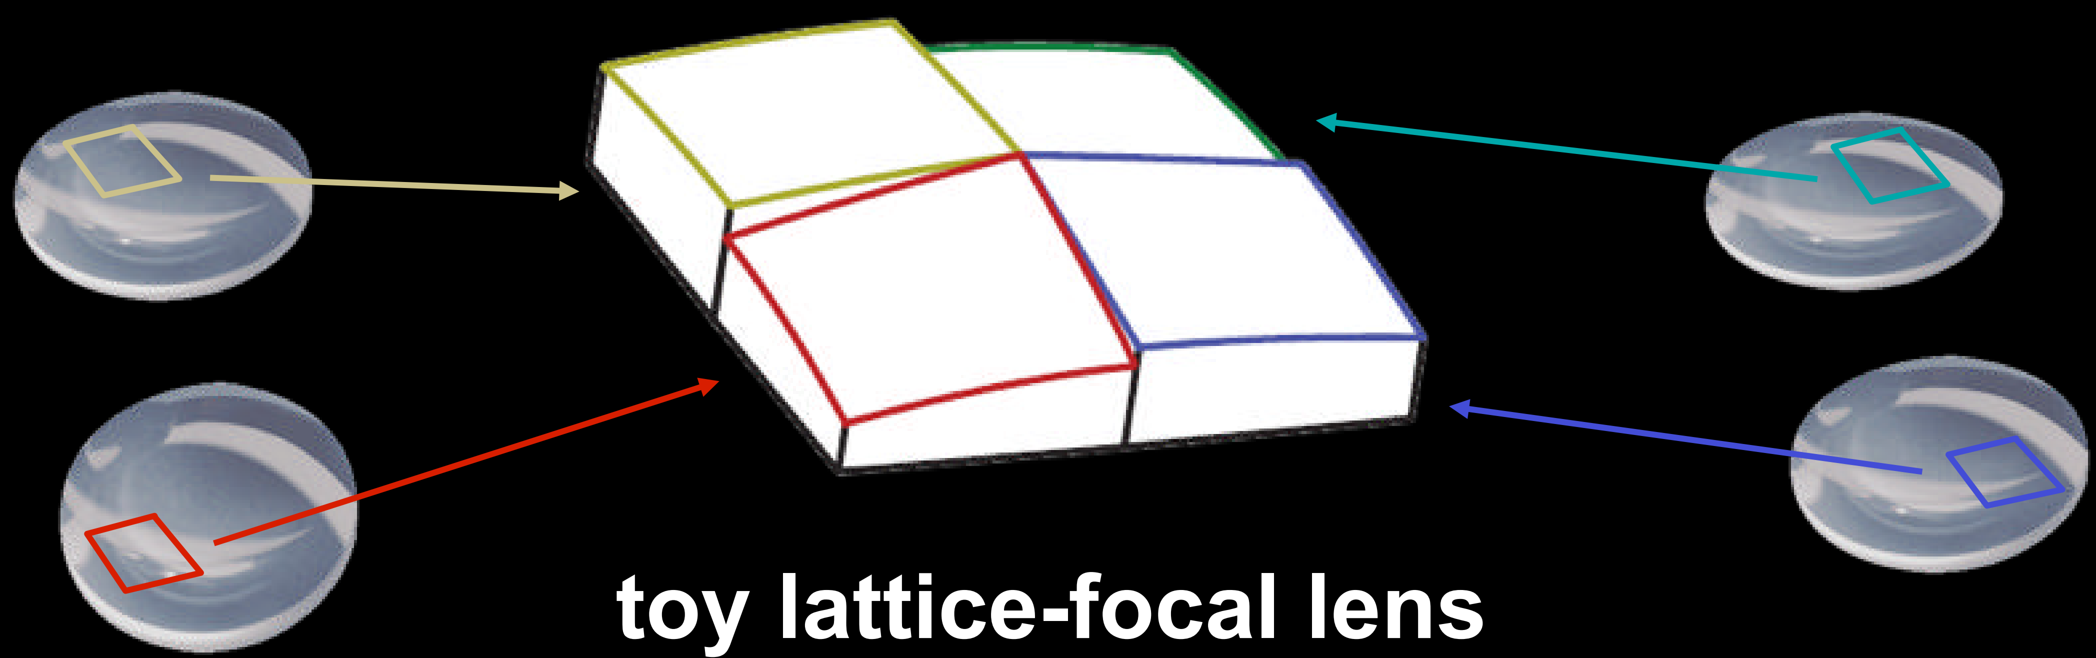

toy lattice-focal lens with 4 elements

$$E\left[\left|\hat{\phi}_s(\omega_{x,y})\right|^2\right] \approx \frac{A^{8/3}}{S^{4/3}\Omega^{1/3}\,|\omega_{x,y}|}$$

Hardware construction

Proof of concept

- 12 subsquares cut from plano-convex spherical lenses
- Attached to main lens

  ➡ extra focal power needed very low
- Modest DOF extension with only 12 subsquares

Depth estimation

- **Defocus kernels vary with depth**

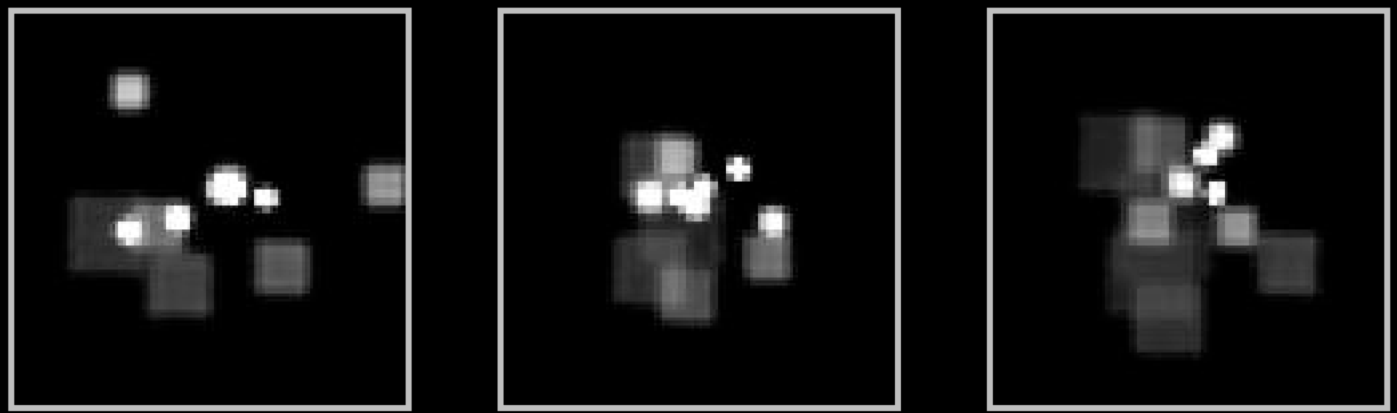

defocus kernels at different depths

- **Depth estimation as for the coded aperture camera [Levin et al. 07]**

input

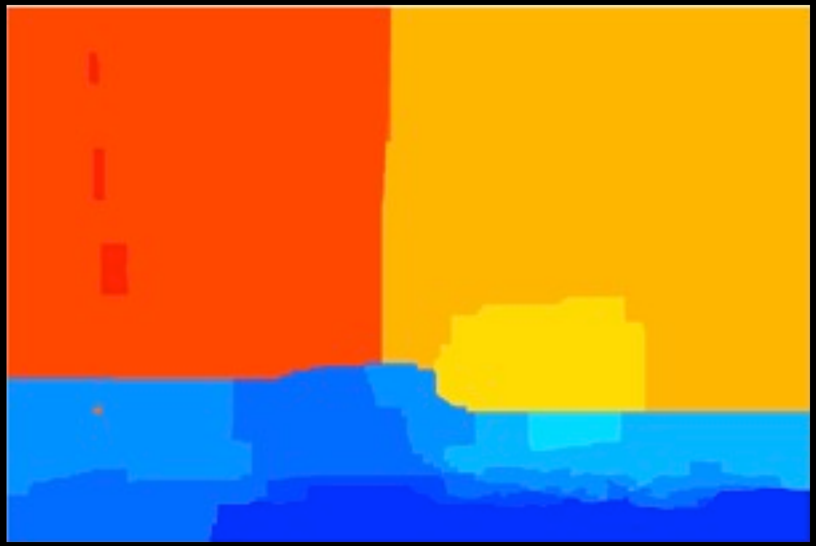

depth map

Standard lens reference

Lattice-focal lens

Lattice focal lens

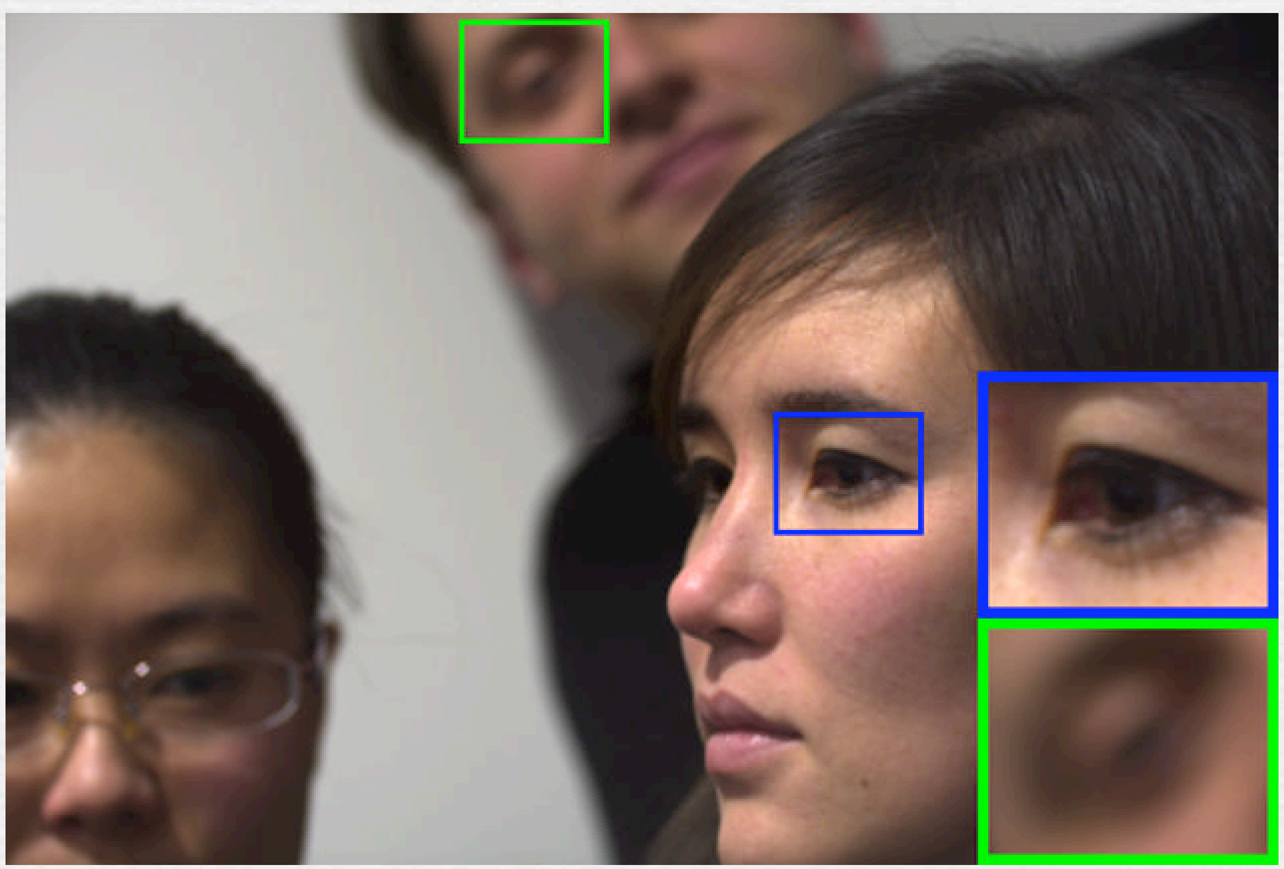

Standard lens image

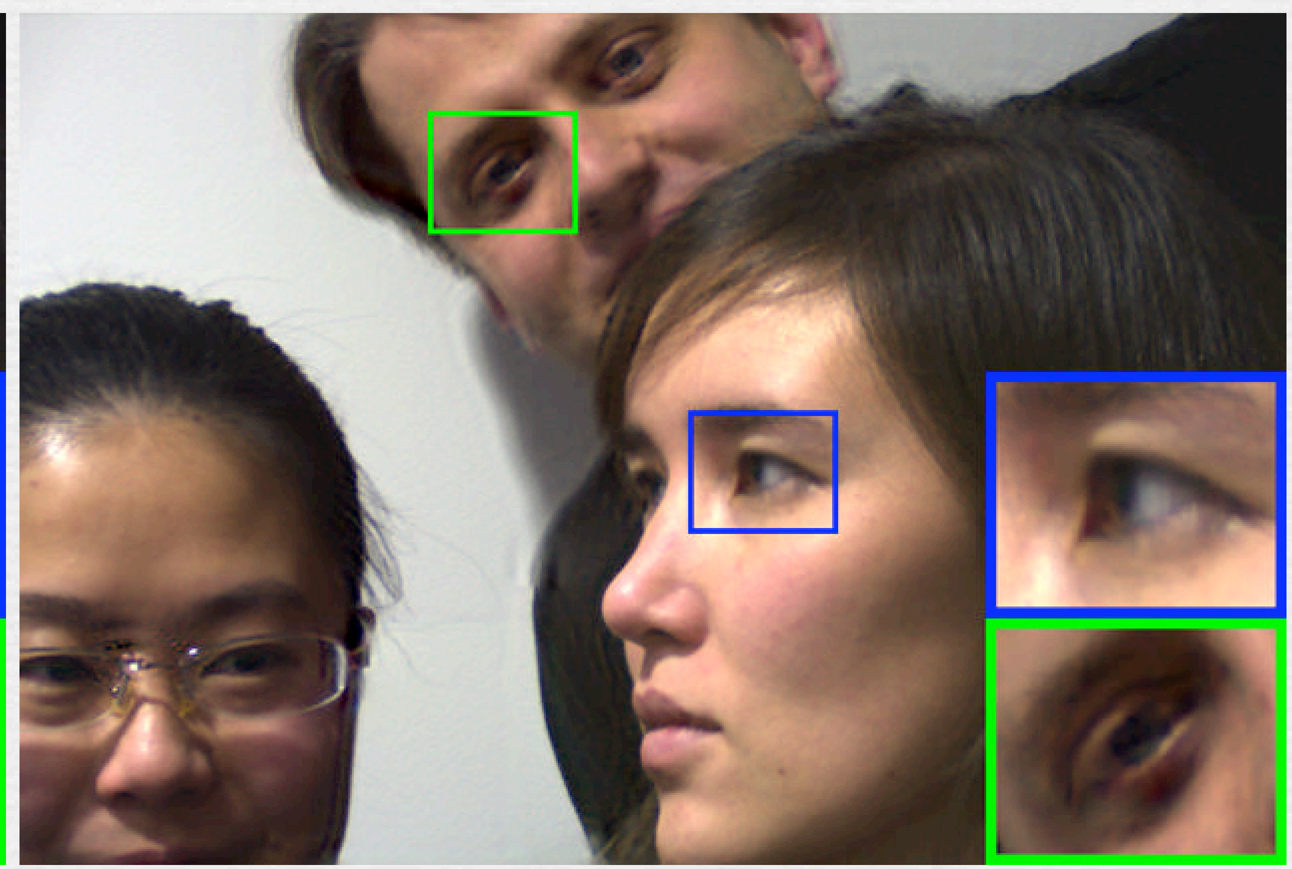

Lattice-focal lens: all-focused output

Application: Refocusing from single captured image

Application: Refocusing from single captured image

Application: Refocusing from single captured image

Depth of field analysis

- How do different cameras compare?
- What is the best that can be done?

Our new theoretical analysis

- http://www.wisdom.weizmann.ac.il/~levina/papers/lattice/
- In the 4D light field
  - Fourier analysis
- Shows that only a 3D subset of the 4D spectrum is useful (dimensionality gap)
- Inspires new lens design: lattice-focal lens

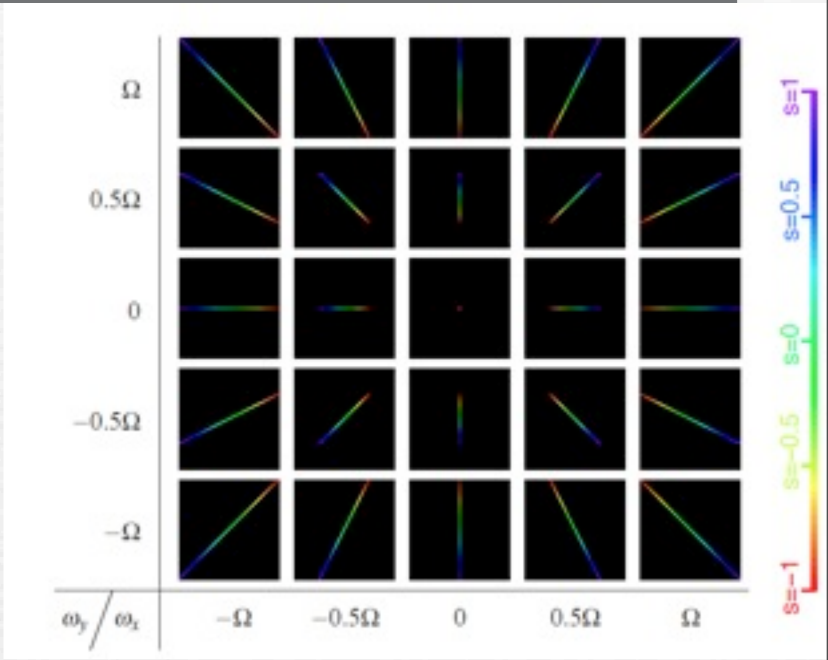

Only a 3D subset of the 4D spectrum is useful

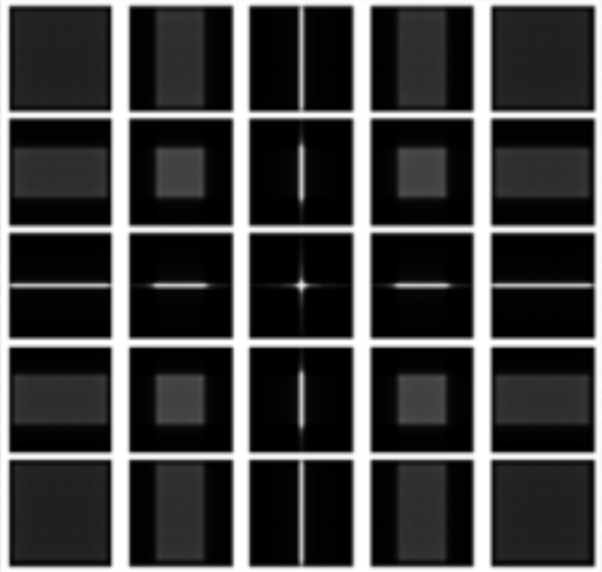

Previous designs spend energy outside the useful subset

Comparison of different cameras

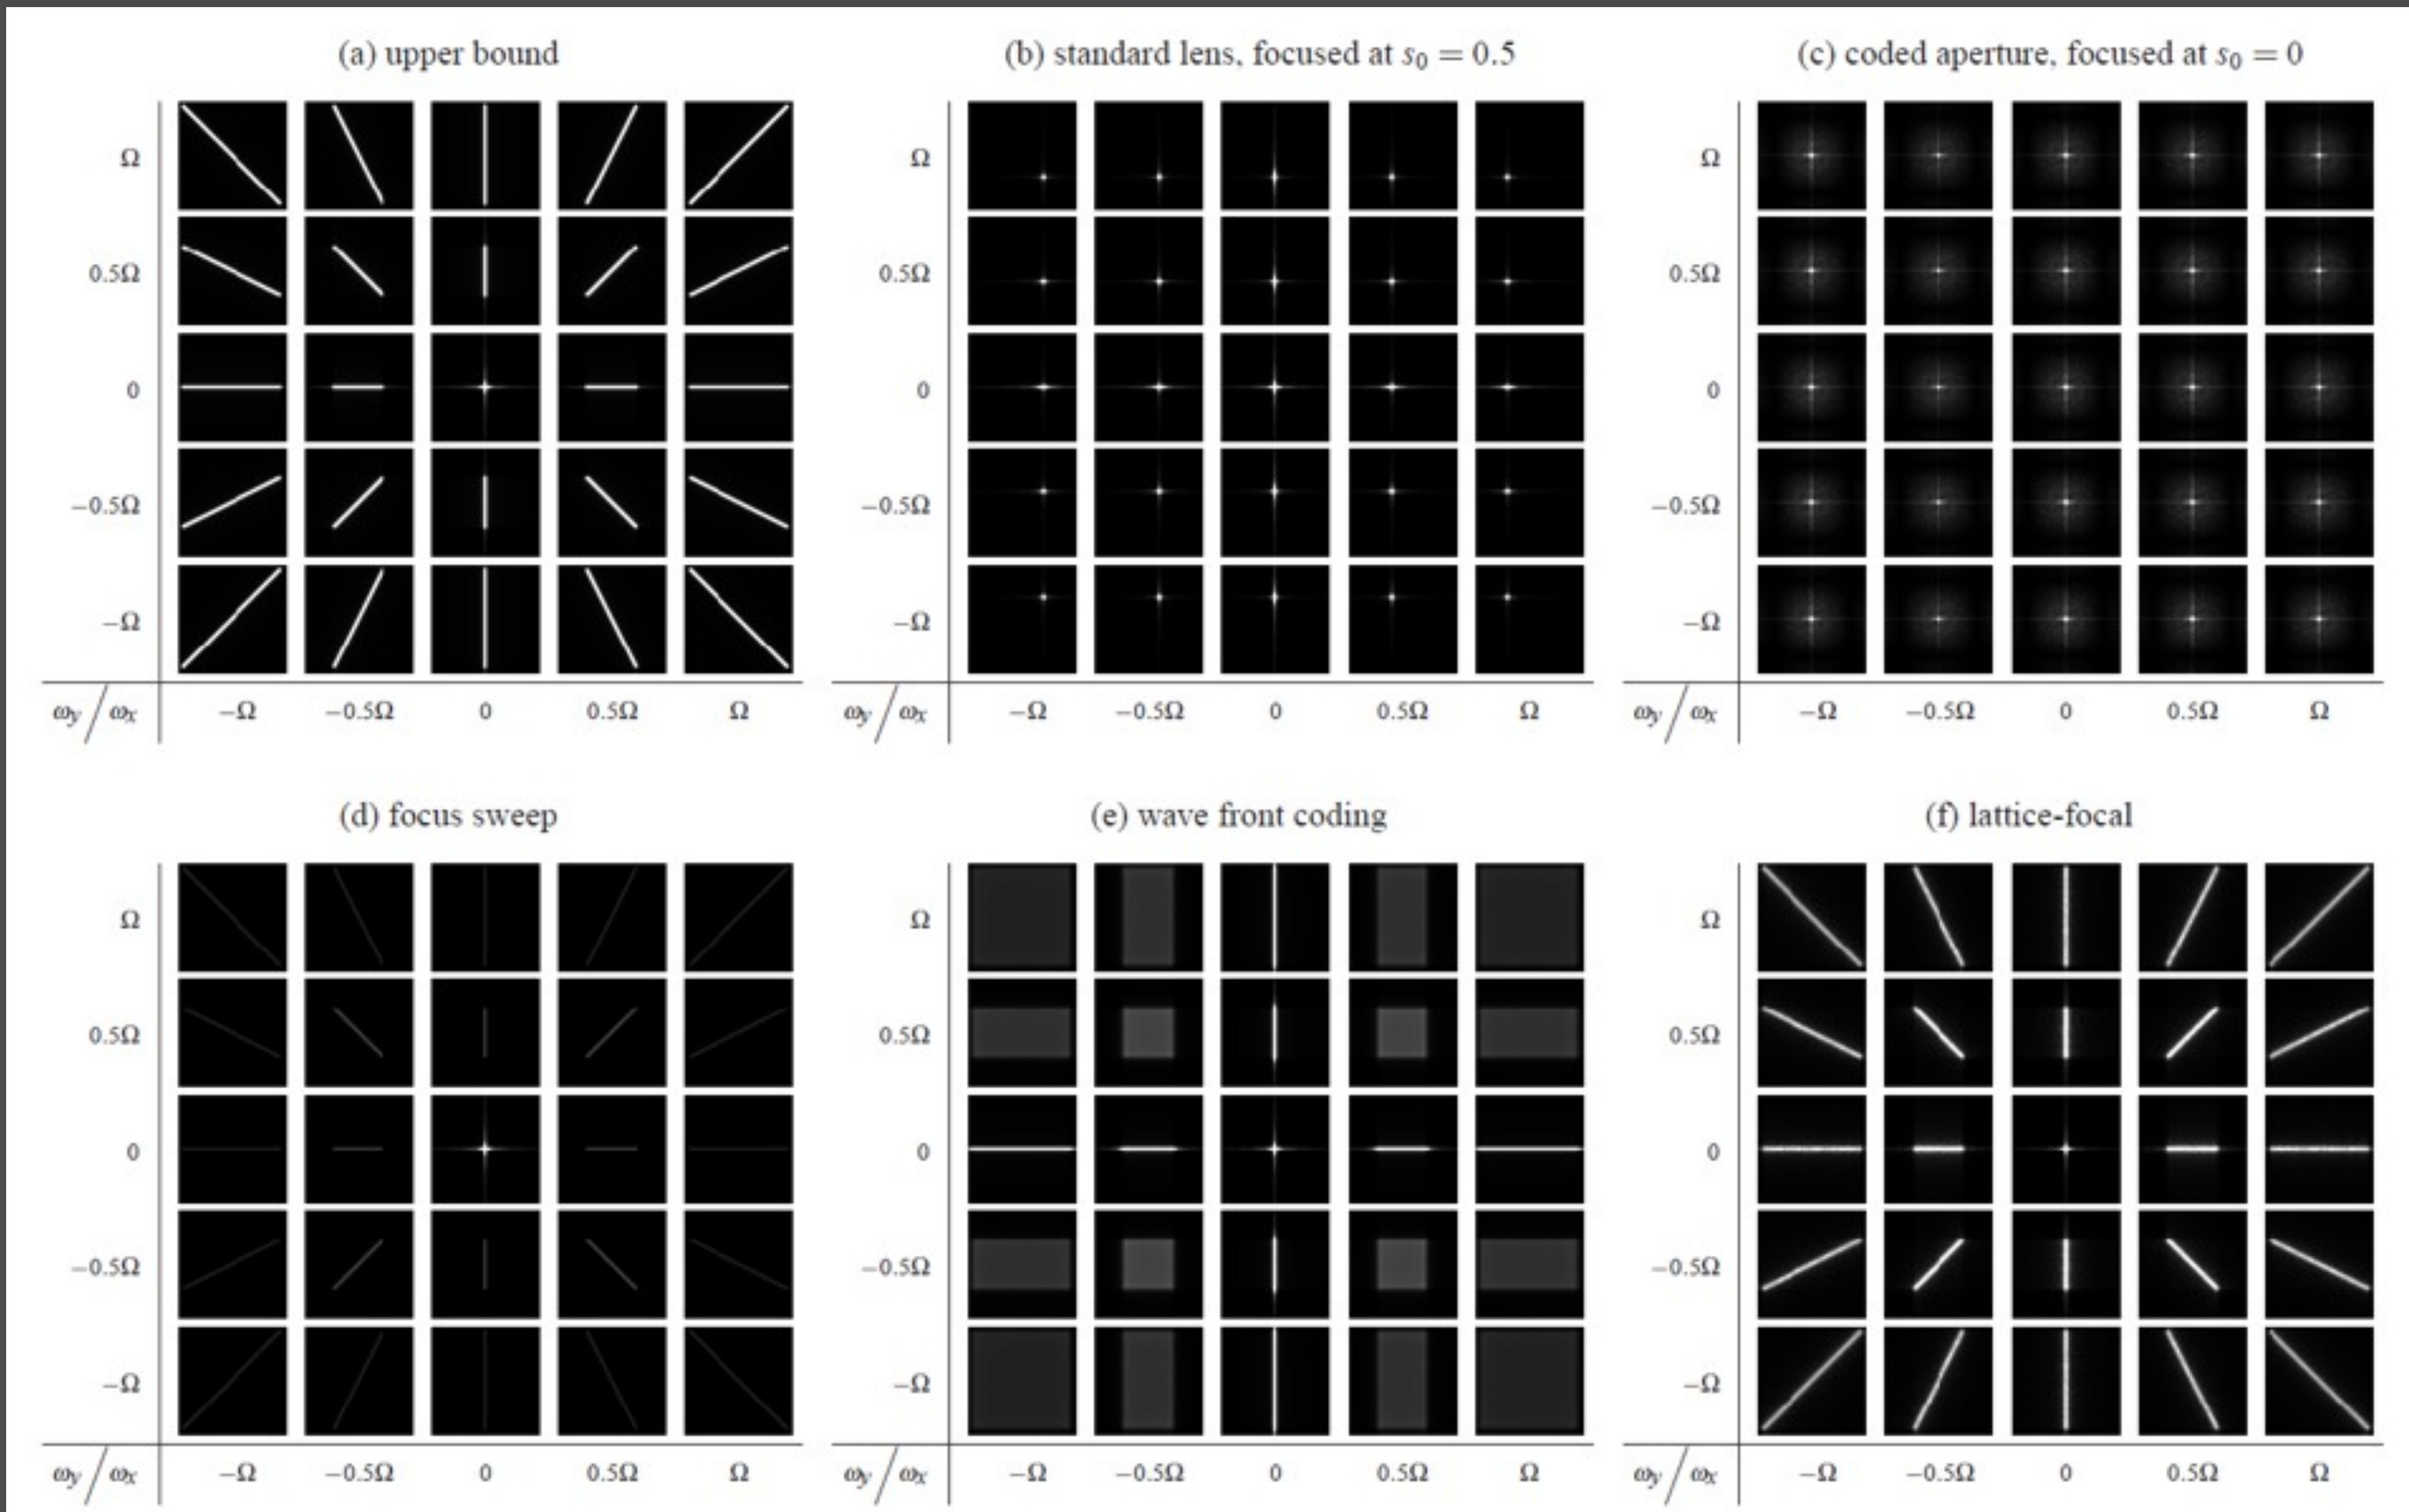

Comparing image reconstruction (simulation)

standard lens

<

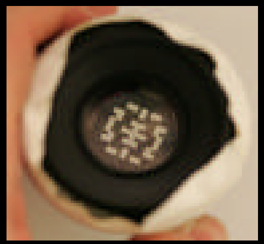

coded aperture

<<

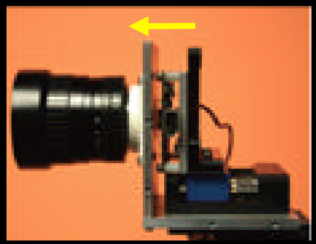

focus sweep

<

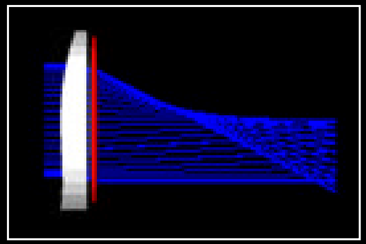

wavefront coding

<

lattice-focal lens

Object at in-focus depth

Object at extreme depth

References

- http://www.cdm-optics.com/site/publications.php
- http://www.wisdom.weizmann.ac.il/~levina/papers/lattice/
- http://www1.cs.columbia.edu/CAVE/publications/pdfs/Nagahara_ECCV08.pdf
- http://groups.csail.mit.edu/graphics/CodedAperture/
- http://web.media.mit.edu/~raskar/Mask/

Focal stacks

Focal stack DoF extensions

- **Capture N images focused at different distances**
- **For each output pixel, choose the sharpest image**
  - e.g. look at local variance, gradient.

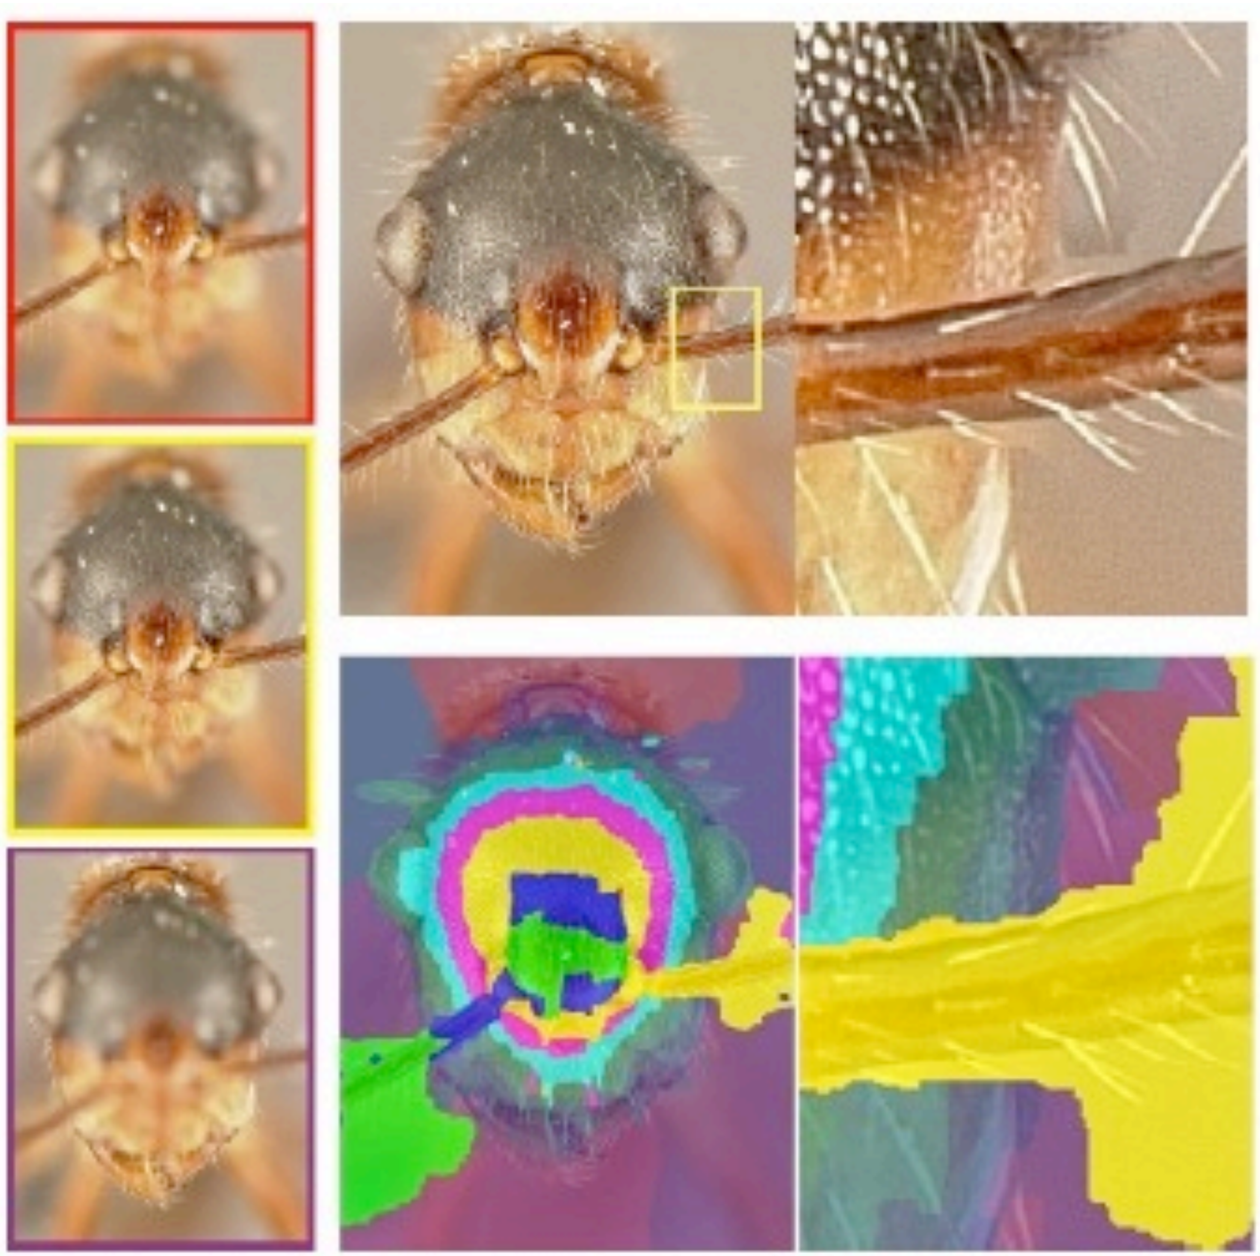

From Agarwala et al.

Focal stack

Montage

Macro montage

- **55 images here**

Focal stack & plenoptic camera

Light Field Photography with a Hand-Held Plenoptic Camera, Ren Ng, Marc Levoy, Mathieu Brédif, Gene Duval, Mark Horowitz, Pat Hanrahan

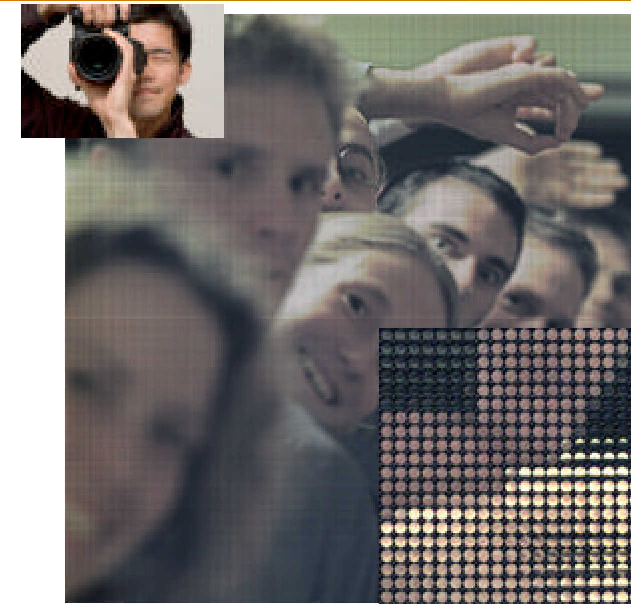

- **Capture light field**

- **Refocus to create focal stack**

- **Use photomontage to generate all-focus image**

Focal stack & plenoptic camera

Figure 15: *Left*: Extended depth of field computed from a stack of photographs focused at different depths. *Right:* A single sub-aperture image, which has equal depth of field but is noisier.

From Ng et al. http://graphics.stanford.edu/papers/lfcamera/

References

- http://www.janrik.net/ptools/ExtendedFocusPano12/index.html
- http://www.outbackphoto.com/workflow/wf_72/essay.html
- http://grail.cs.washington.edu/projects/photomontage/
- http://people.csail.mit.edu/hasinoff/timecon/

- http://graphics.stanford.edu/papers/lfcamera/

Slicing

Scanning: combination in 1 exposure

Macrophotography scanning device. The subject is lit by a fine ray of light with a thickness less than the depth of field; the lens can be used with average apertures that provide maximum sharpness. Mounted on a stand with a headless screw, it is moved forward and backward by a slow and regular movement that is controlled by a motorized micrometer. This device, which can be made by a meticulous handyman, lets you take spectacular shots of large insects with total depth of field.

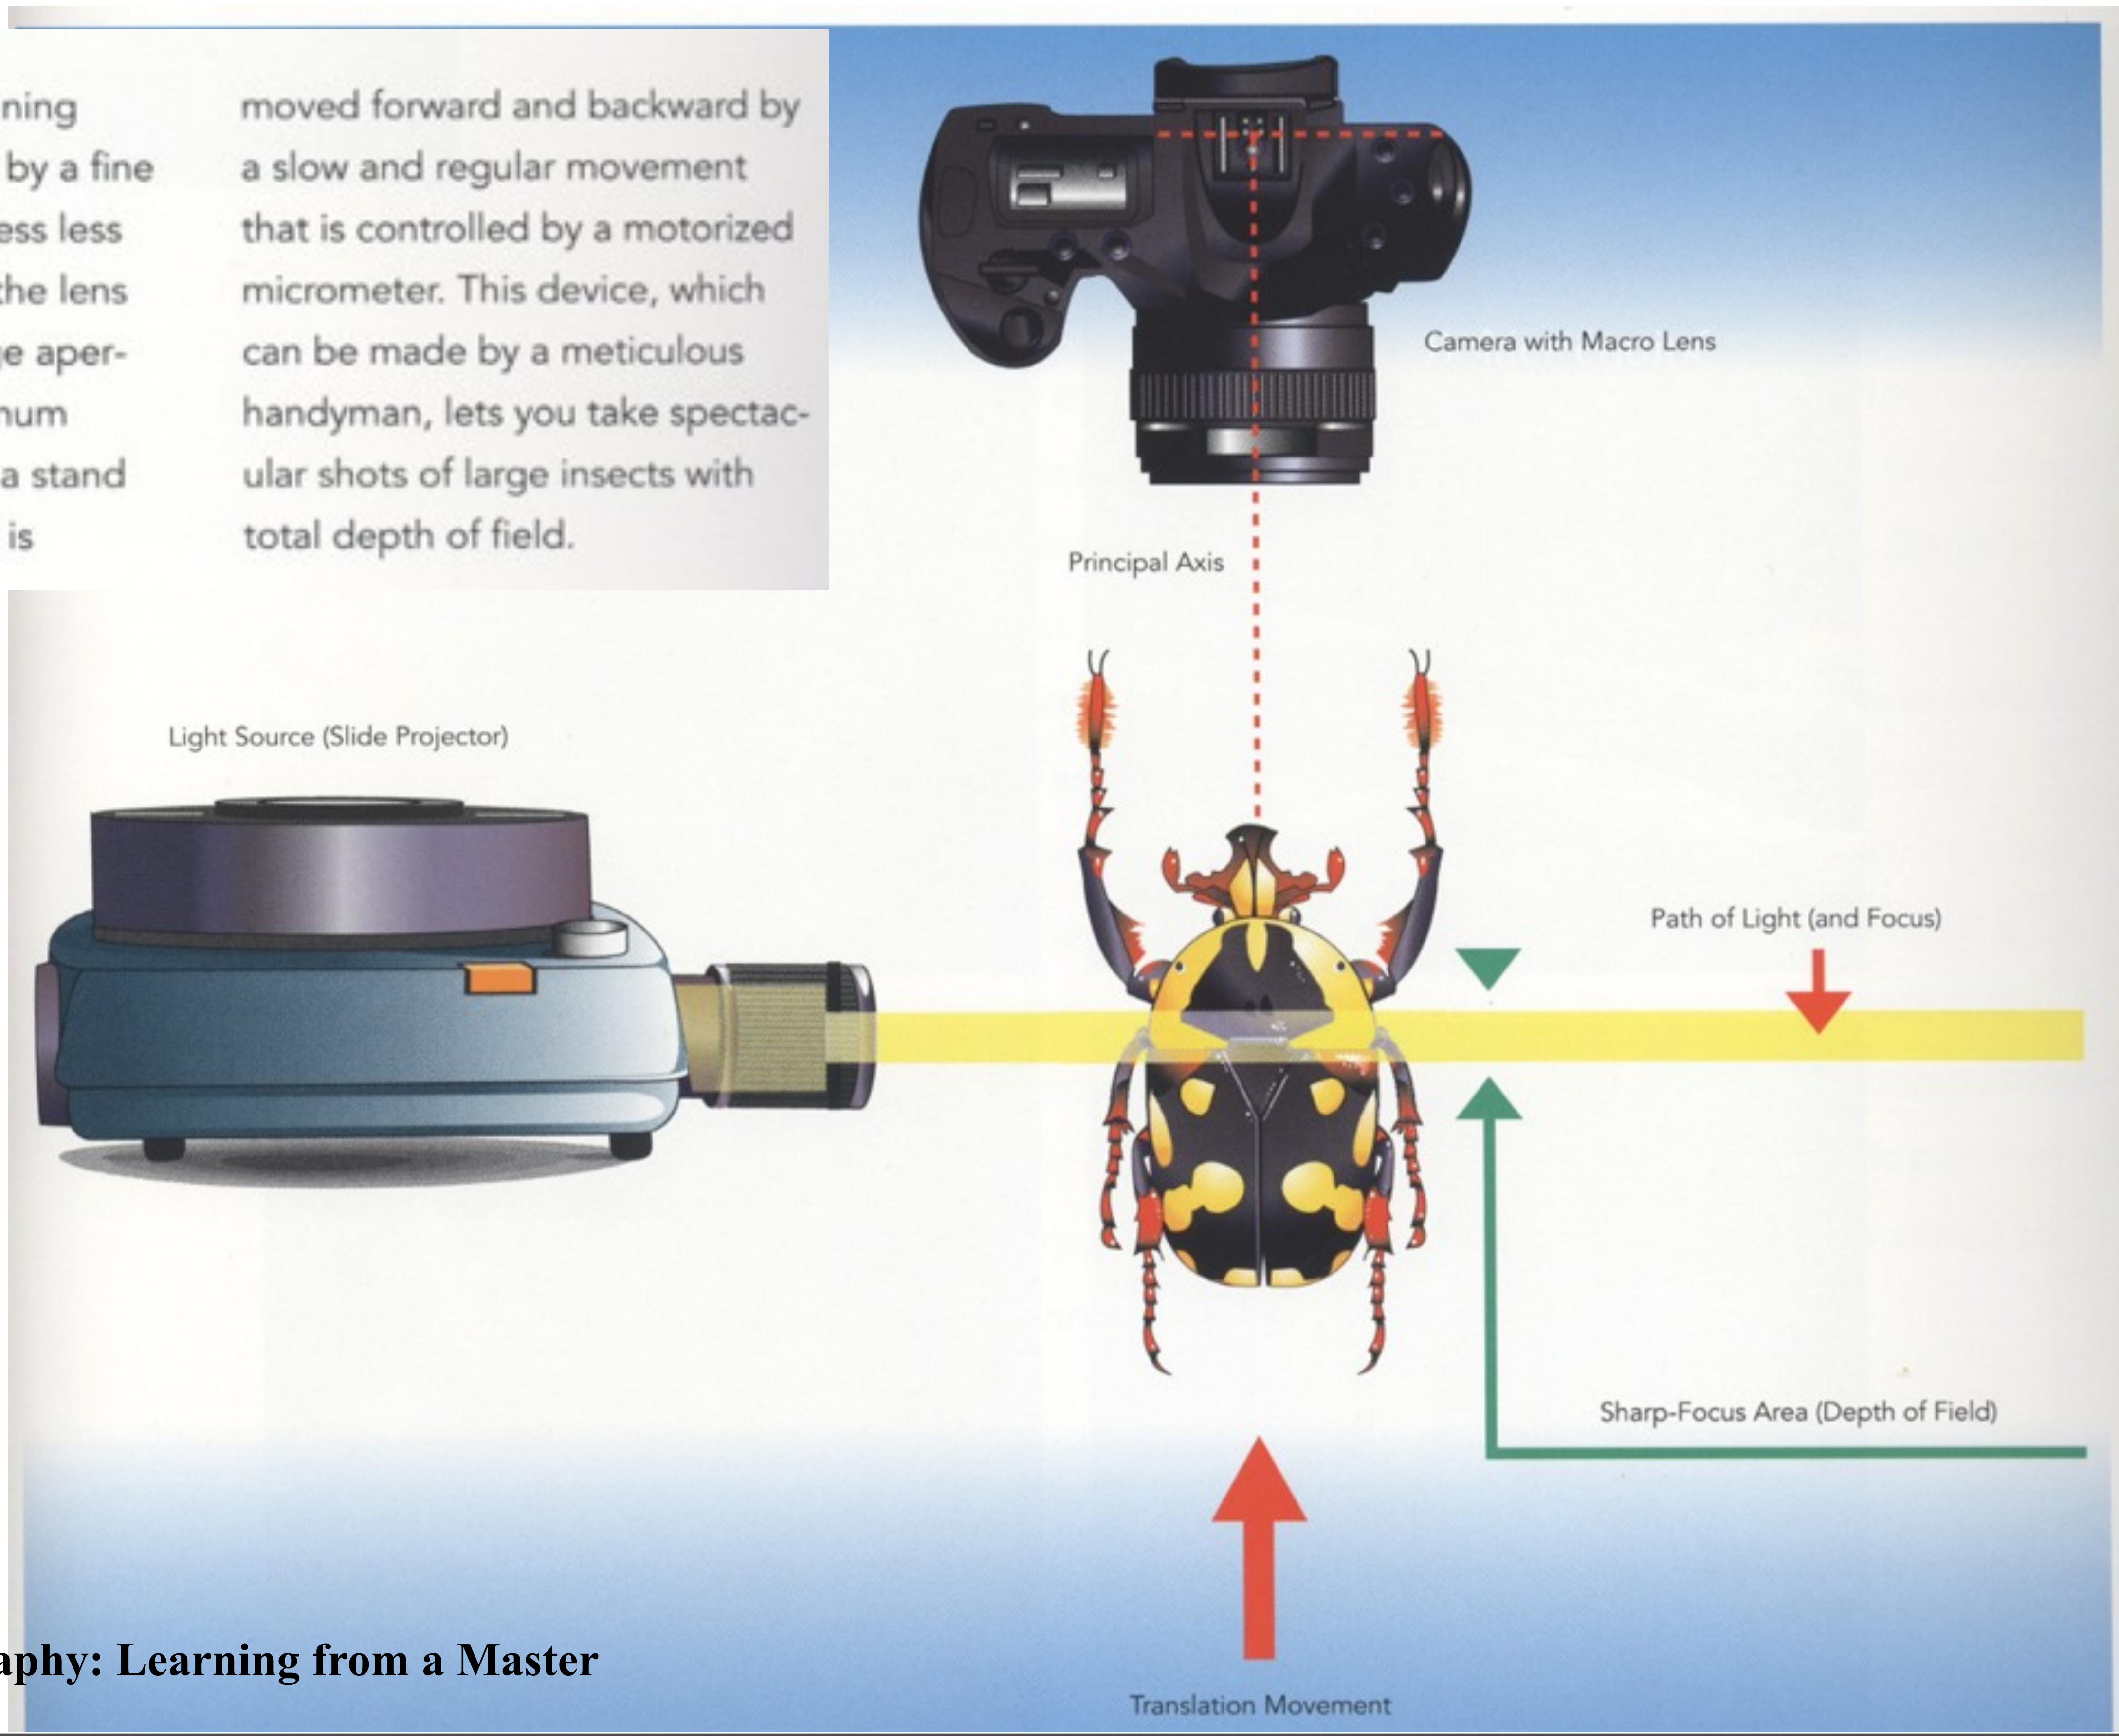

From Macro photography: Learning from a Master