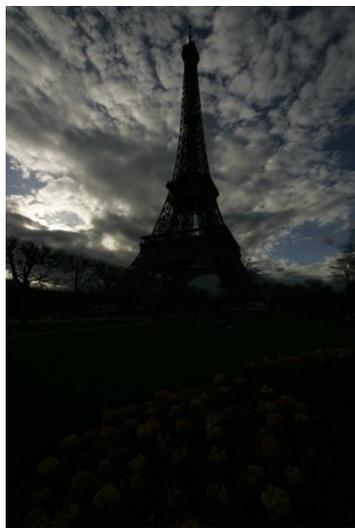
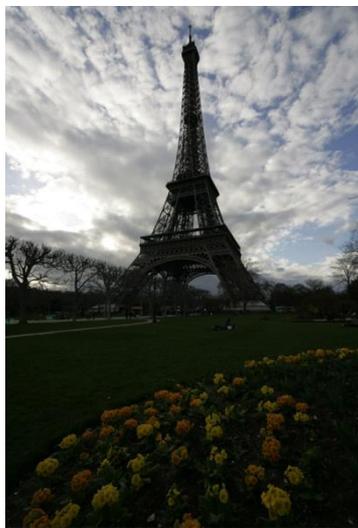


# CS6670: Computer Vision

Noah Snavely

## Lecture 23: Computational photography, Part 2



# Announcements

- Final project midterm reports due today to CMS by 11:59pm
- Final project presentation times

# The ultimate camera

---

Infinite resolution

Infinite zoom control

Desired object(s) are in focus

No noise

No motion blur

Infinite dynamic range (can see dark and bright things)

...

# Creating the ultimate camera

---

The “analog” camera has changed very little in >100 yrs

- we’re unlikely to get there following this path

More promising is to combine “analog” optics with computational techniques

- “Computational cameras” or “Computational photography”

This lecture will survey techniques for producing higher quality images by combining optics and computation

Common themes:

- take multiple photos
- modify the camera

# Capture and composite several photos

---

Works for

- field of view
- resolution
- signal to noise
- dynamic range
- Focus

But sometimes you can do better by modifying the camera...

# Focus

---

Suppose we want to produce images where the desired object is ***guaranteed*** to be in focus?

Or suppose we want ***everything*** to be in focus?

# Light field camera [Ng et al., 2005]

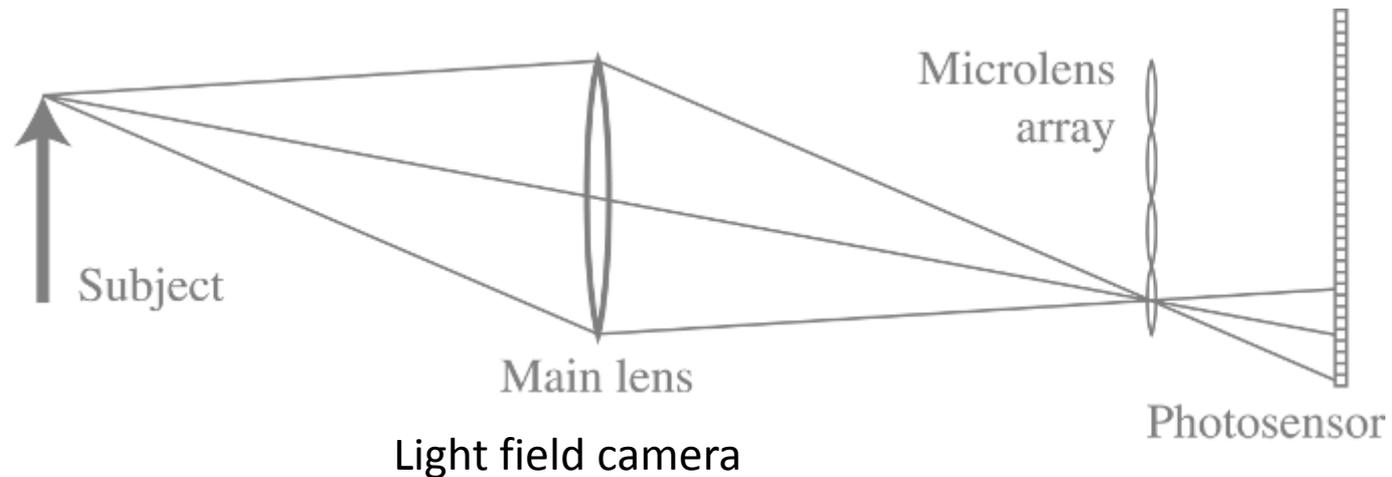
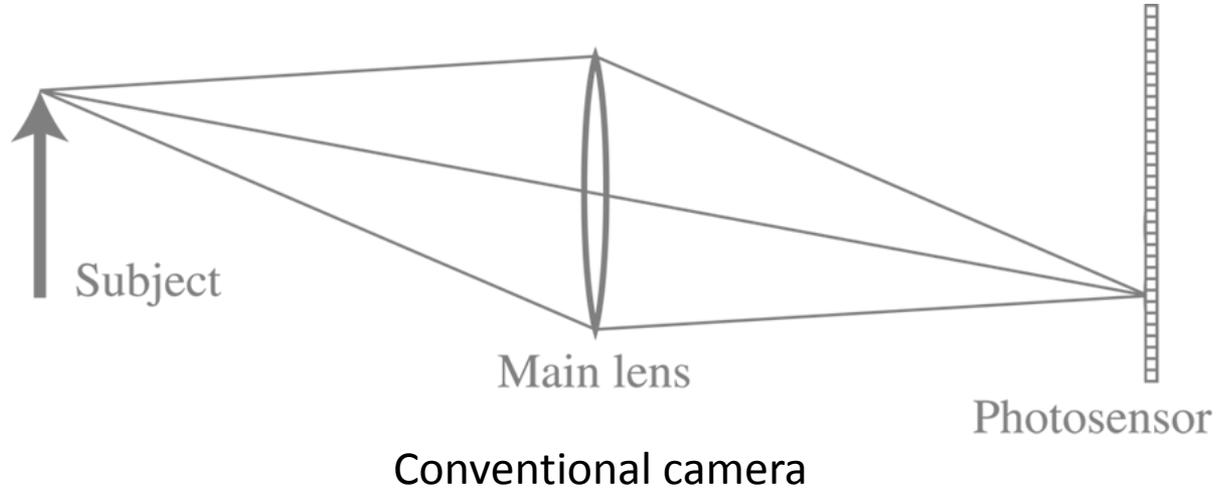
---



<http://www.refocusimaging.com/gallery/>

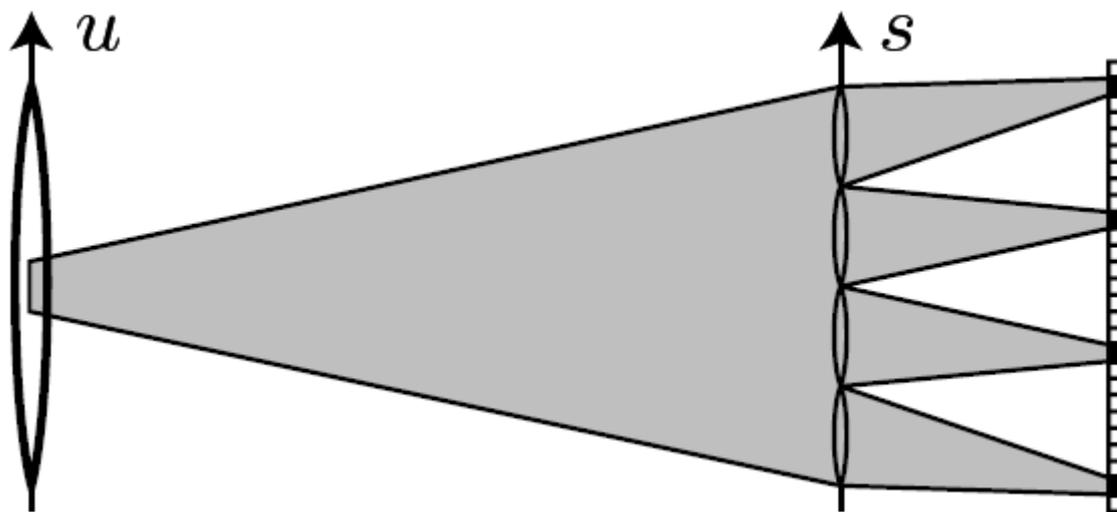
# Conventional vs. light field camera

---



# Light field camera

---



Rays are reorganized into many smaller images corresponding to subapertures of the main lens

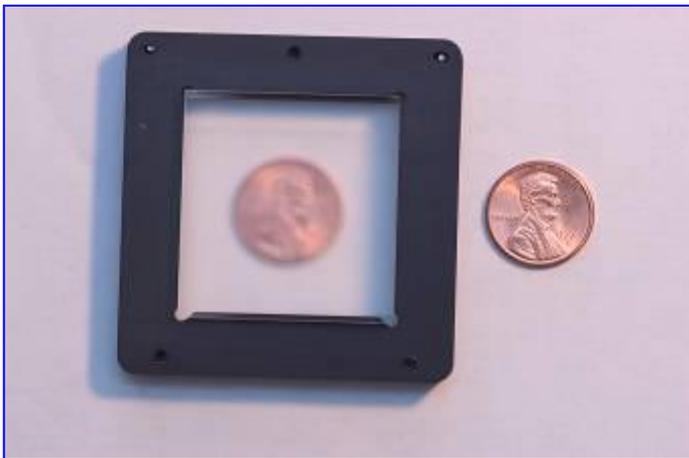
# Prototype camera



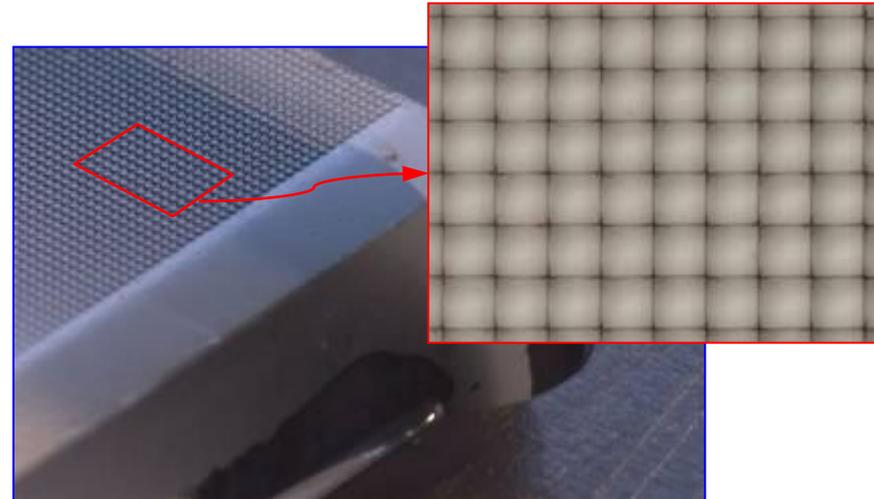
Contax medium format camera



Kodak 16-megapixel sensor

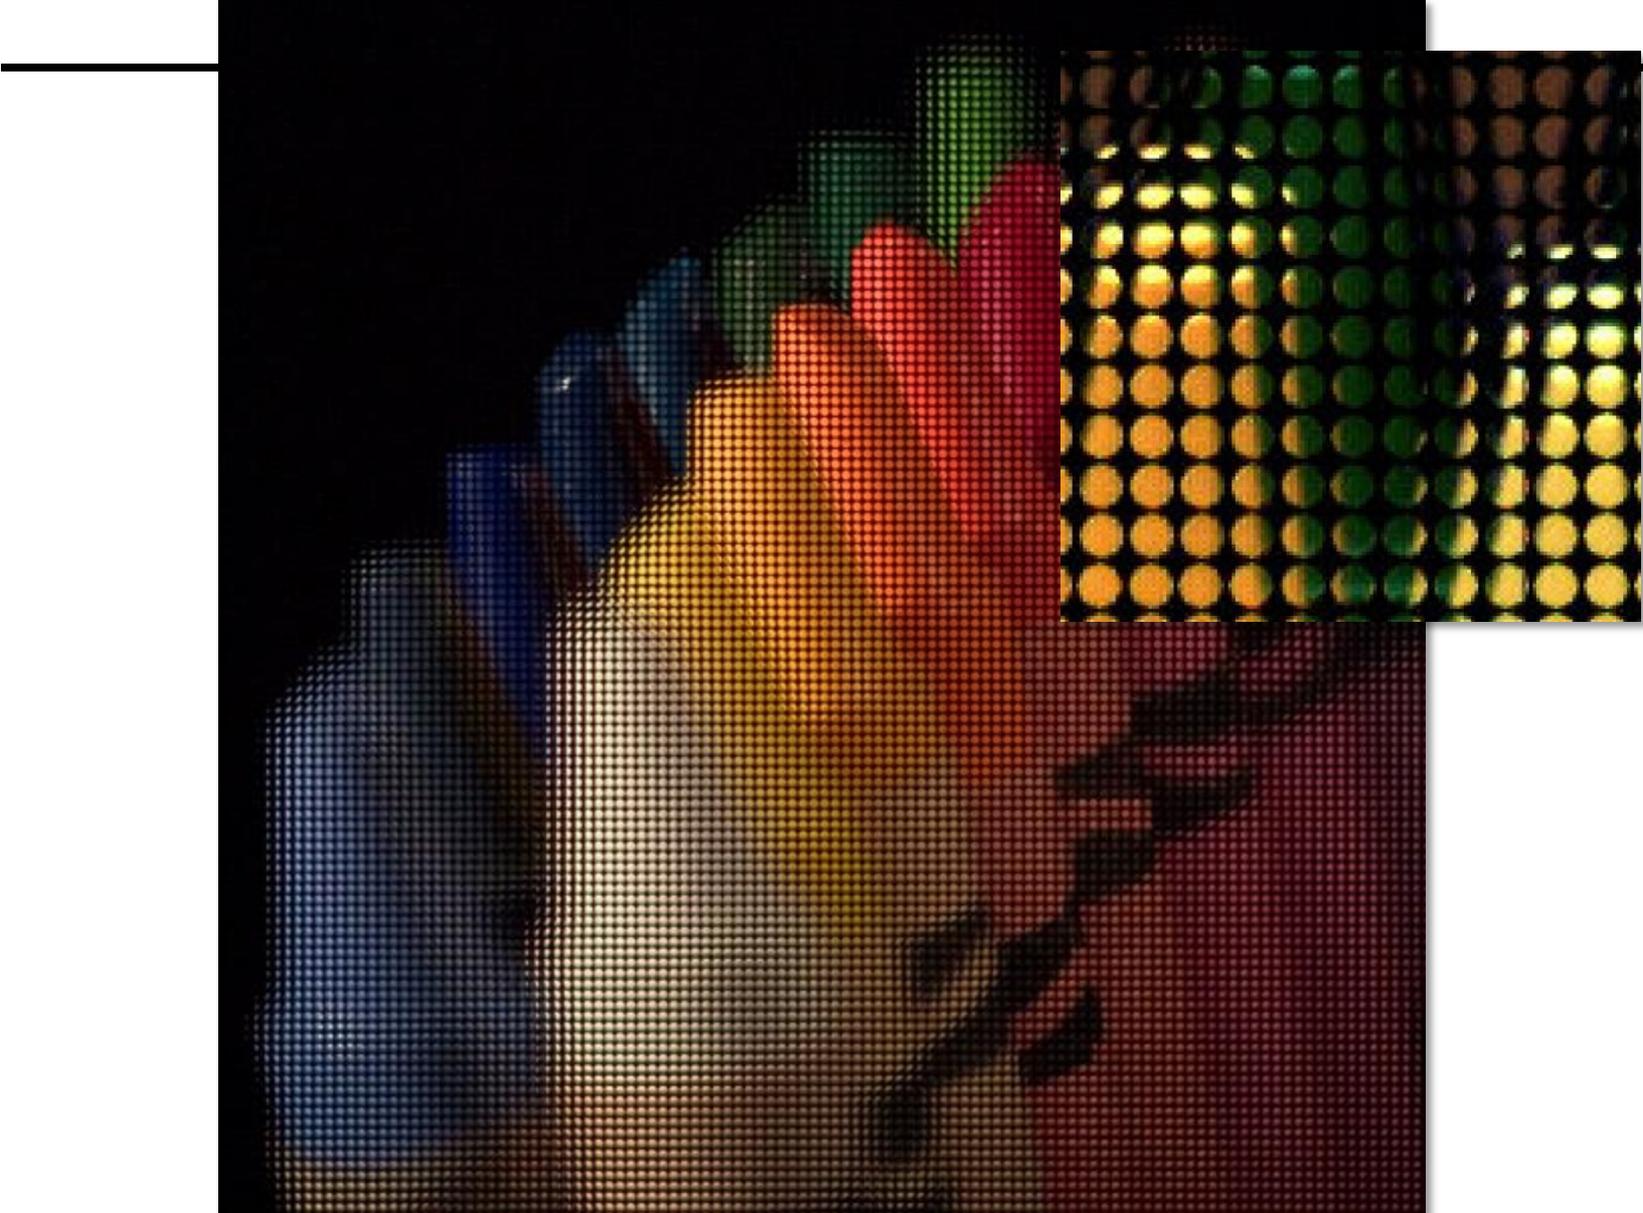


Adaptive Optics microlens array



125 $\mu$  square-sided microlenses

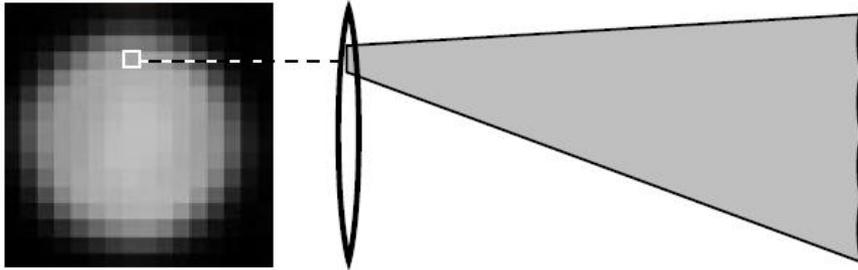
$$4000 \times 4000 \text{ pixels} \div 292 \times 292 \text{ lenses} = 14 \times 14 \text{ pixels per lens}$$



# What can we do with the captured rays?

---

Change viewpoint



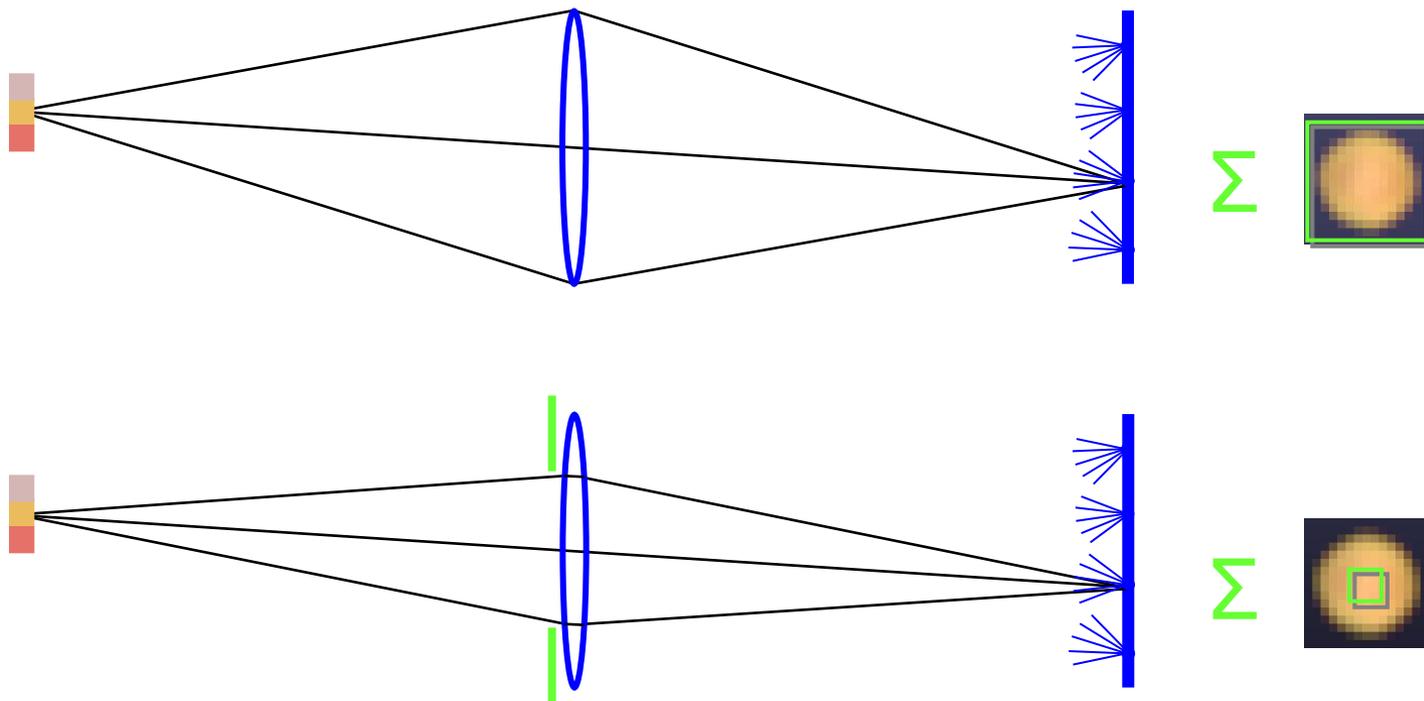
[ray-traced edge](#)



# What can we do with the captured rays?

---

Simulate depth of field

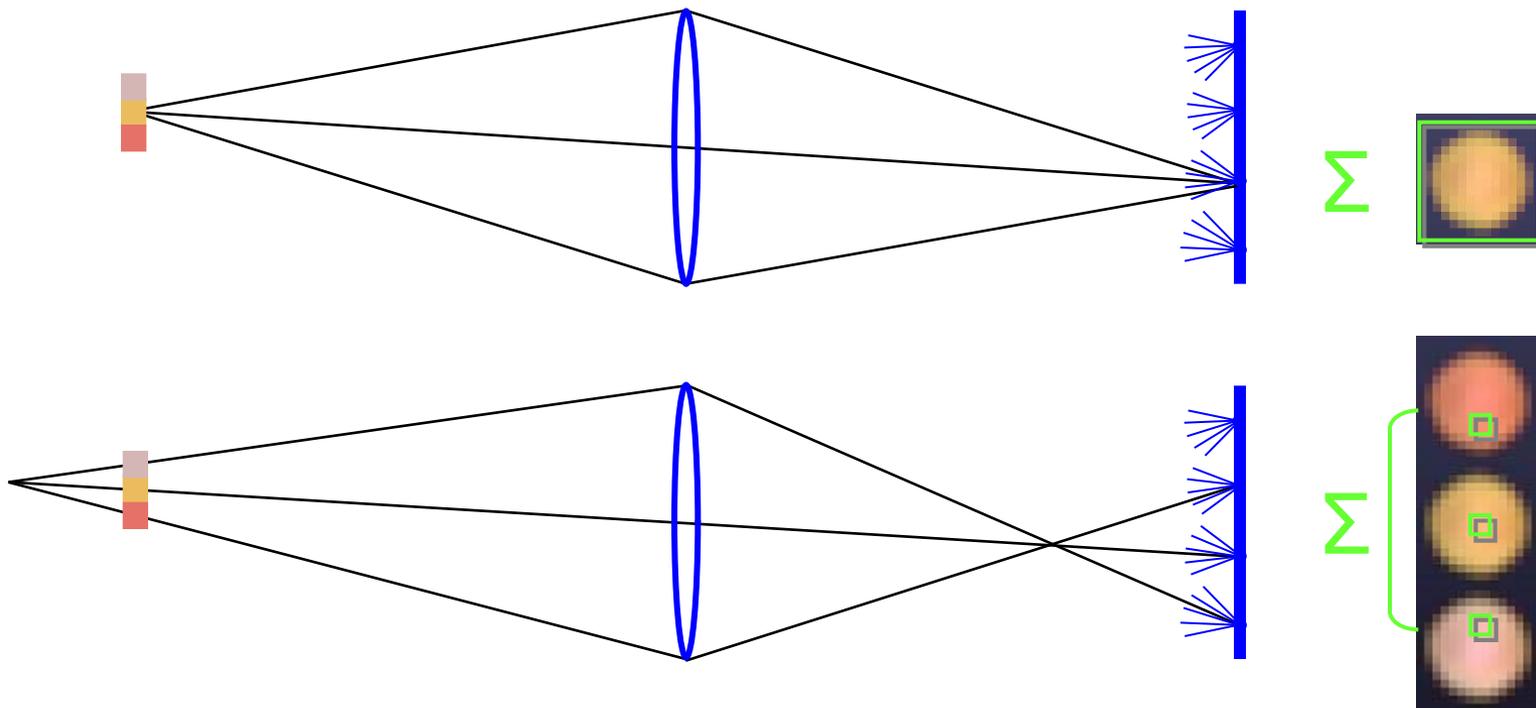


stopping down aperture = summing only  
the central portion of each microlens

# What can we do with the captured rays?

---

## Digital refocusing



refocusing = summing windows extracted  
from several microlenses

# Example of digital refocusing

---



# Seeing beyond occlusion

---

Photo taken through zoo bars  
Telephoto at full aperture  
The bars are so blurry  
that they are invisible

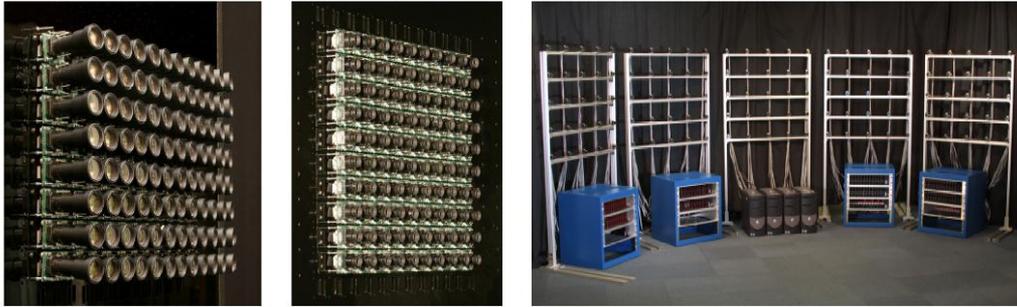


# Synthetic aperture

---

Stanford Camera array (Willburn et al.

<http://graphics.stanford.edu/papers/CameraArray/>)



(a)



(b)



(c)

Figure 11: Matted synthetic aperture photography. (a) A sample image from one of 90 cameras used for this experiment. (b) The synthetic aperture image focused on the plane of the people, computed by aligning and averaging images from all 90 cameras as described in the text. (c) Suppressing contributions from static pixels in each camera yields a more vivid view of the scene behind the occluder. The person and stuffed toy are more clearly visible.

# All-in-focus images

---

Combines sharpest parts of all of the individual refocused images

Using single pixel from each subimage



# All-in-focus

---

If you only want to produce an all-focus image, there are simpler alternatives

E.g.,

- Wavefront coding [Dowsky 1995]
- Coded aperture [Levin SIGGRAPH 2007], [Raskar SIGGRAPH 2007]
  - can also produce change in focus (ala Ng's light field camera)

# Why are images blurry?

---



Depth of field



Camera focused at wrong distance



Motion blur

How can we remove the blur?

# Motion blur

---

Especially difficult to remove, because the blur kernel is unknown



=



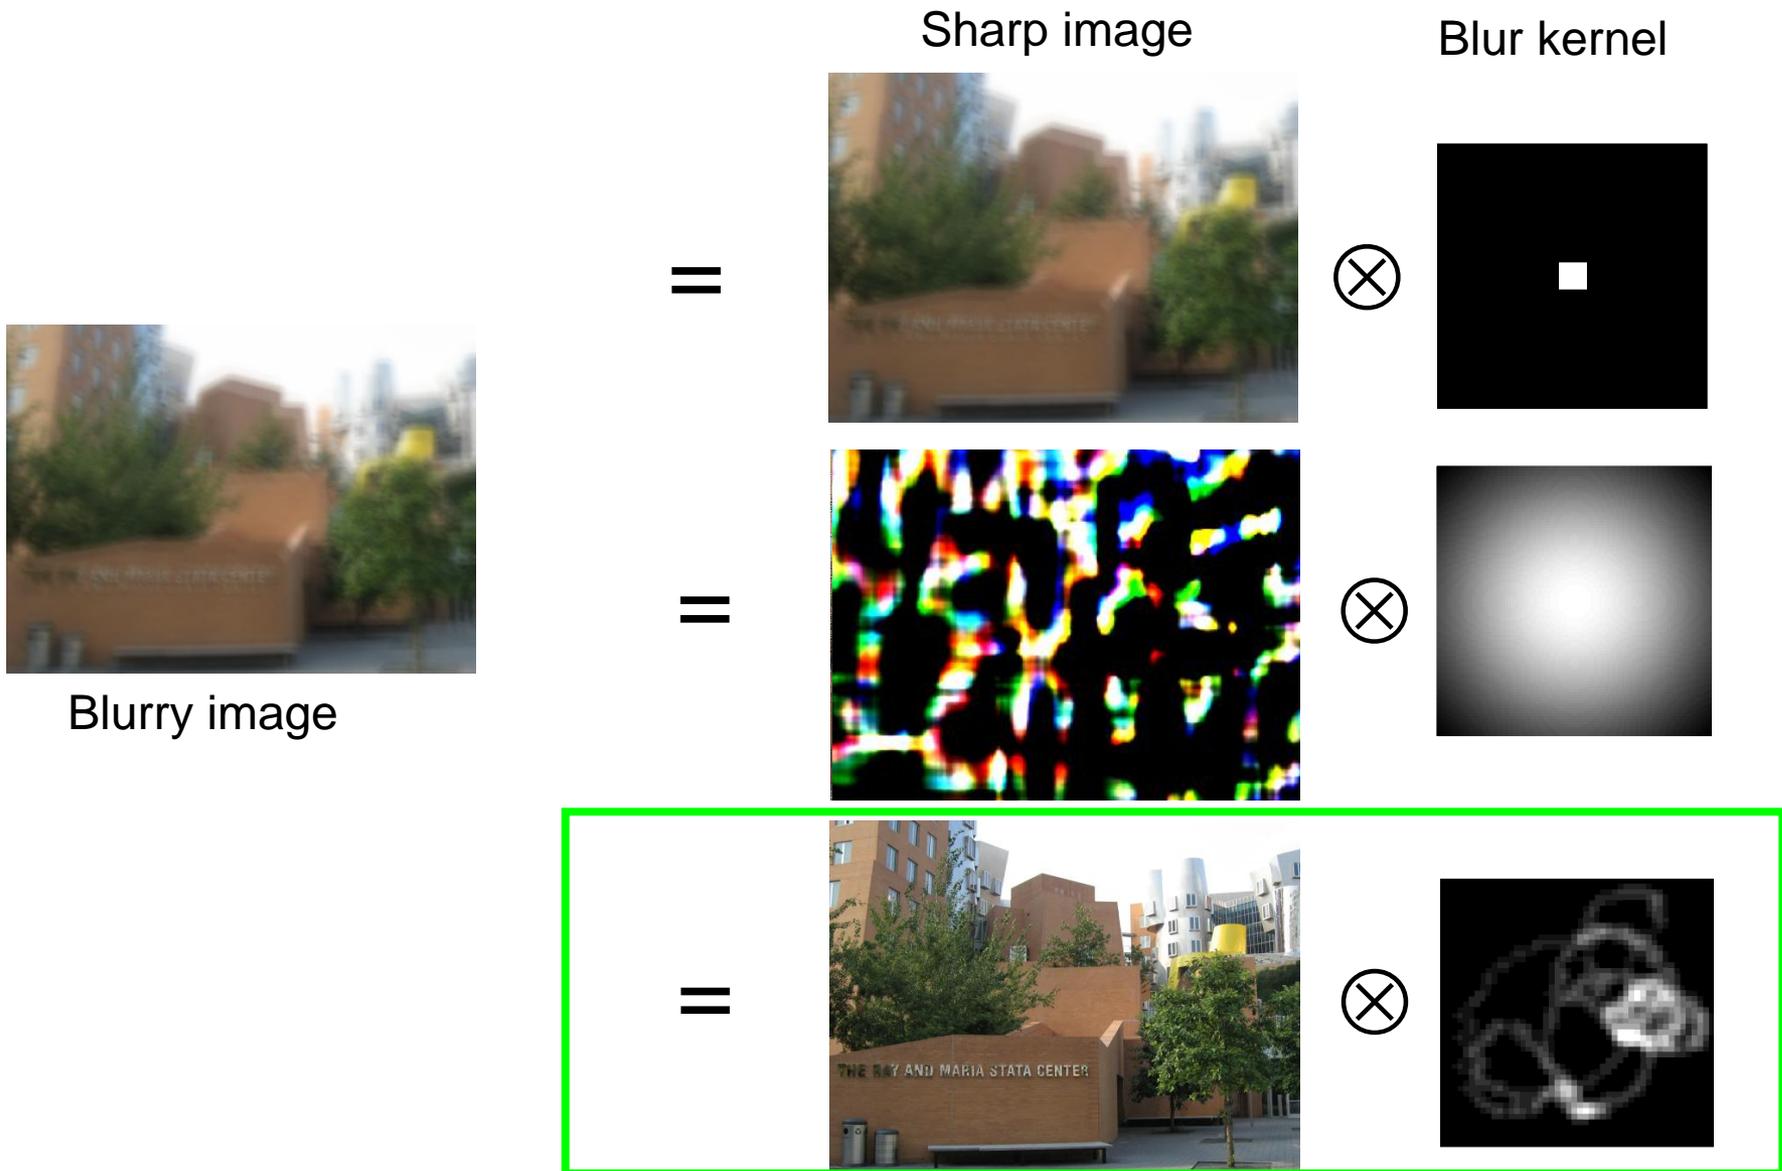
⊗



both unknown

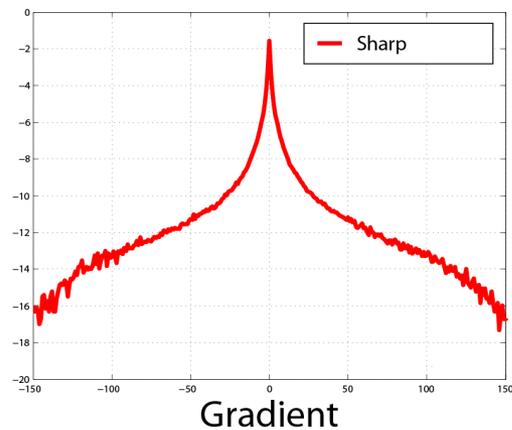
# Multiple possible solutions

---



# Priors can help

---



Priors on natural images

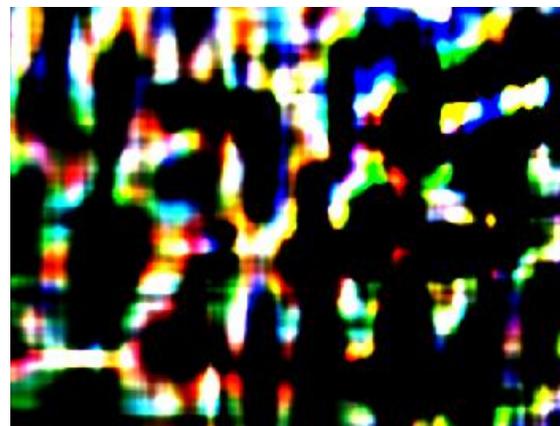


Image A is more “natural” than image B

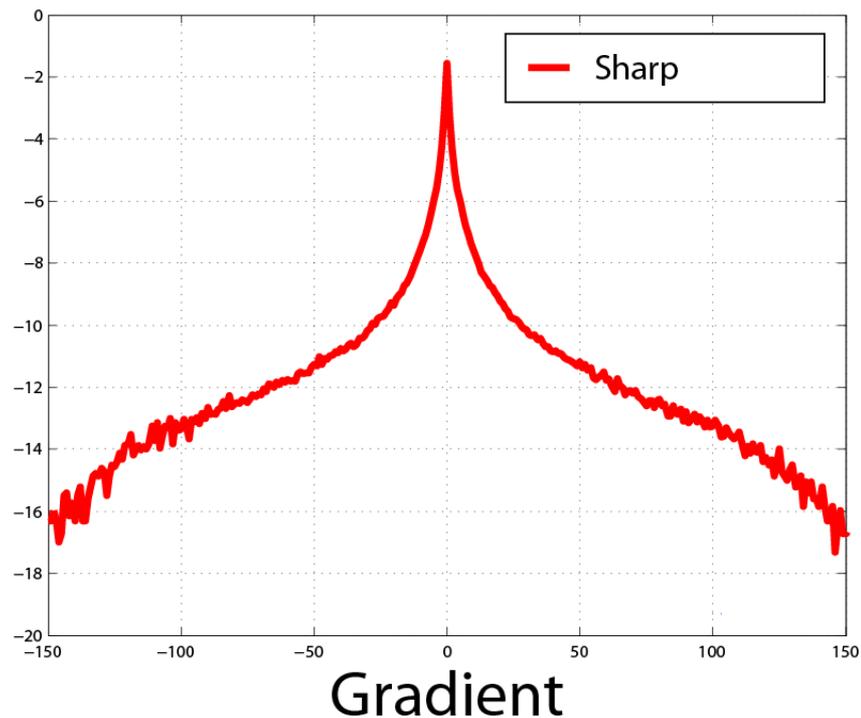
# Natural image statistics

---

Characteristic distribution with heavy tails



Histogram of image gradients

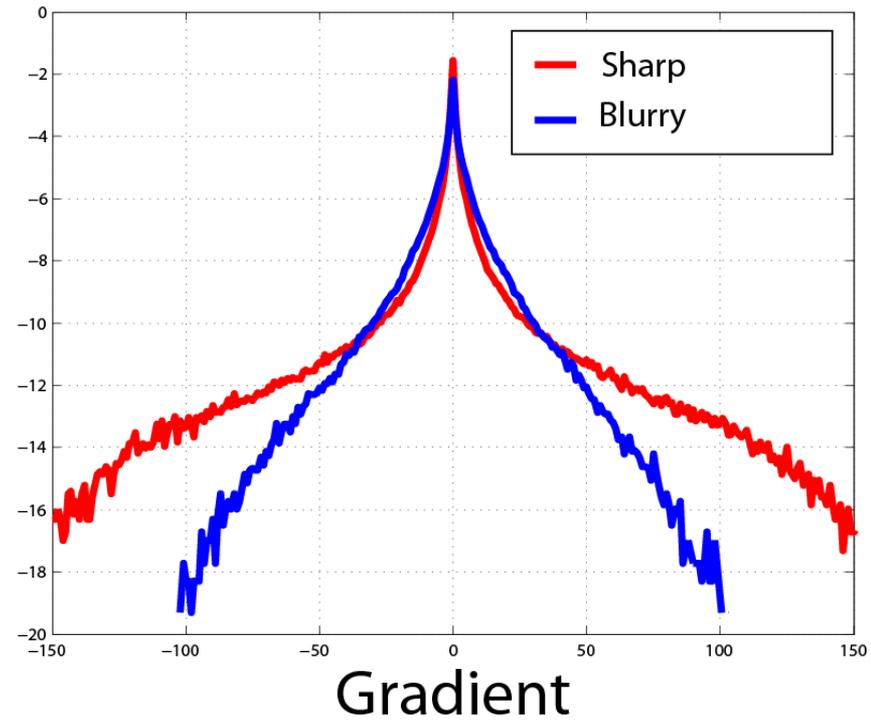


# Blurry images have different statistics

---

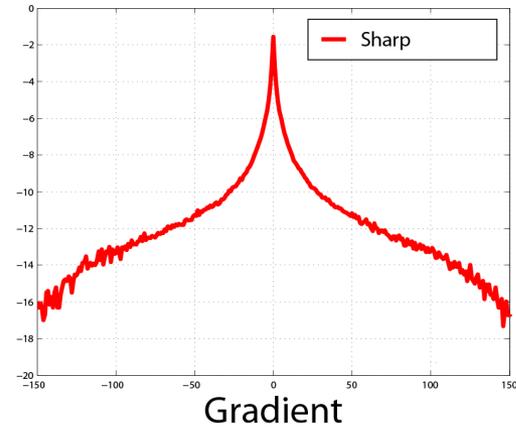


Histogram of image gradients

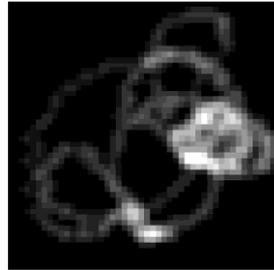


# Priors can help

---



Priors on natural images



Prior on blur kernels (positive and sparse)

# Three sources of information

---

## 1. Reconstruction constraint:



Estimated sharp image

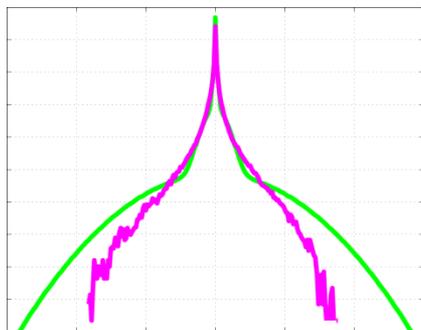


Estimated blur kernel



Input blurry image

## 2. Image prior:



Distribution of gradients

## 3. Blur prior:



Positive  
&  
Sparse

# Results [Fergus, *et al*, 2006]

---

Original



Algorithm of Fergus



# Close-up

---

Original



Naïve Sharpening

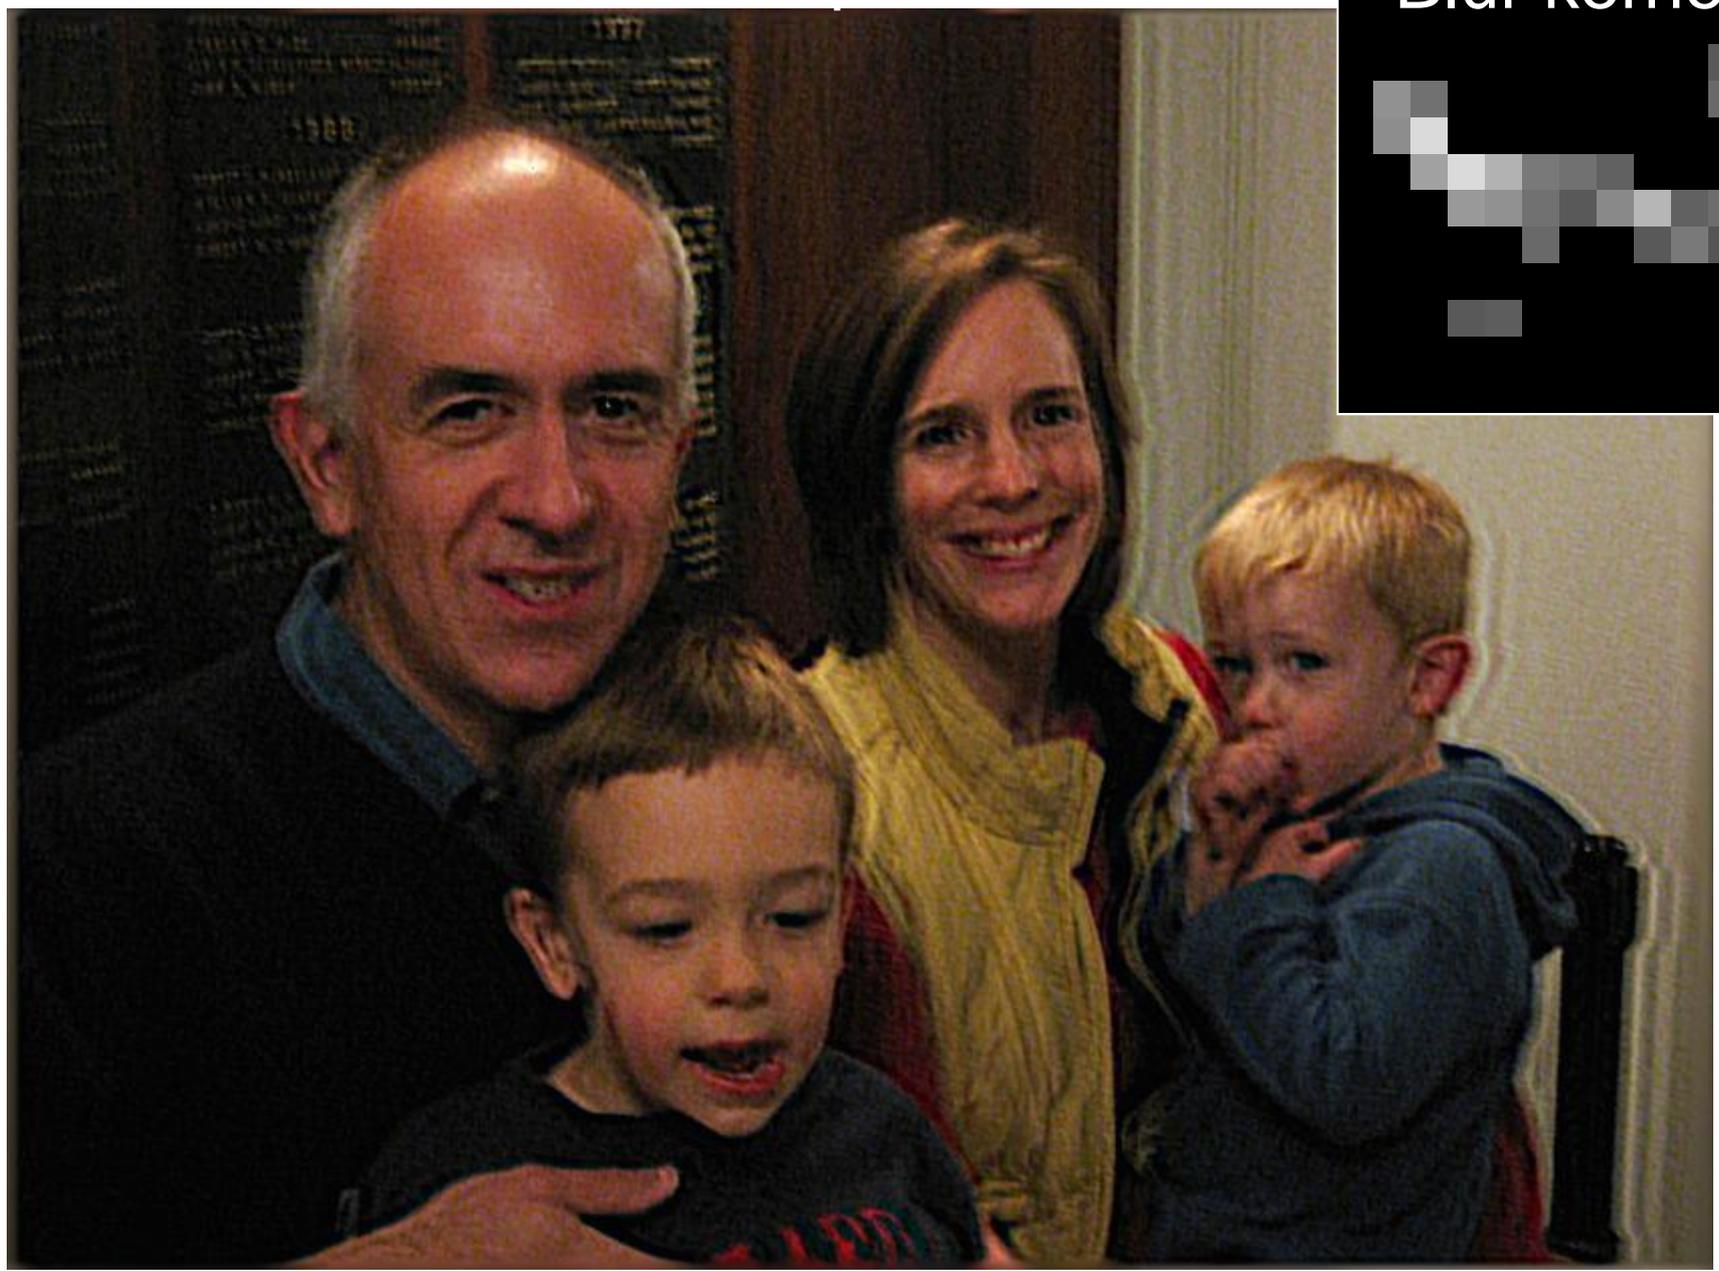
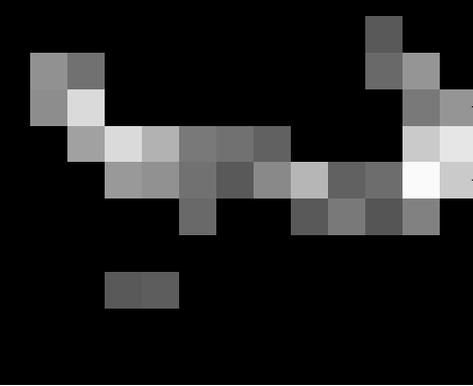


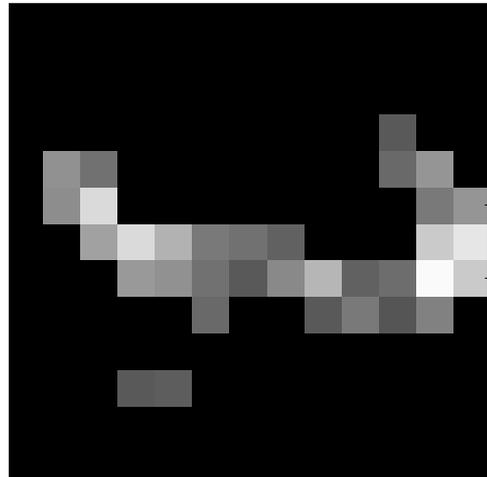
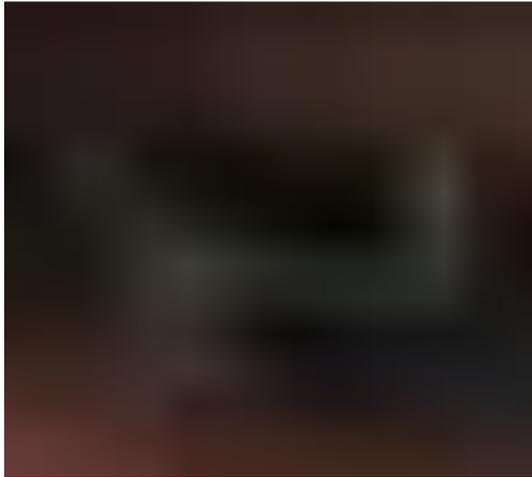
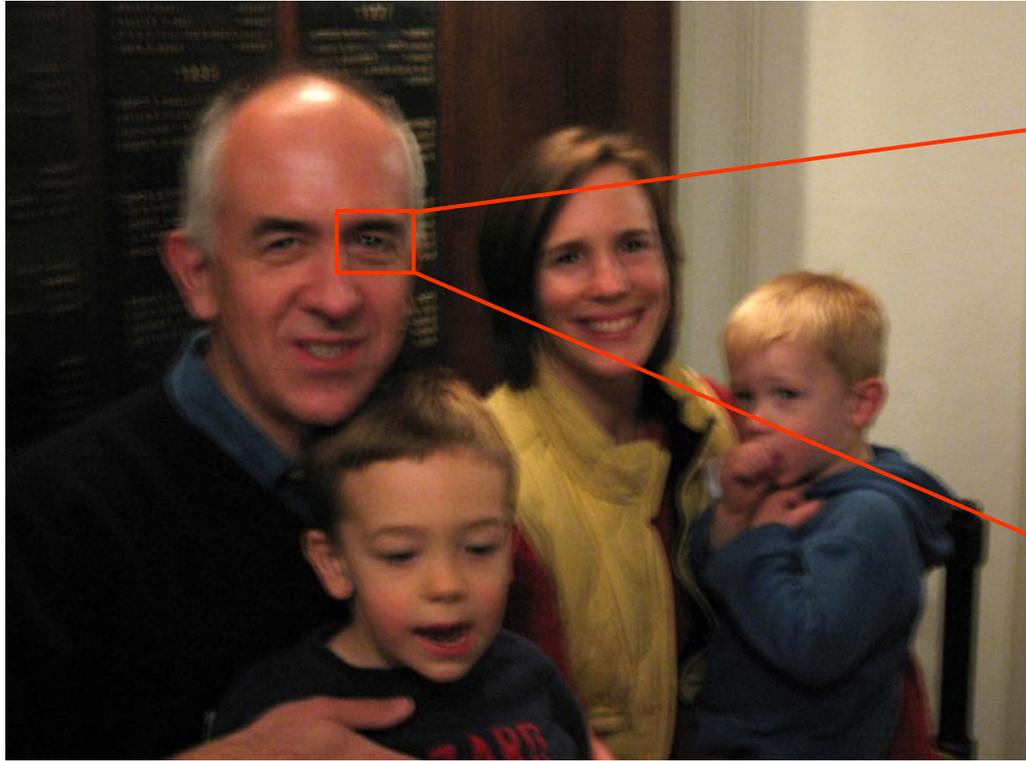
Algorithm of Fergus





Blur kernel

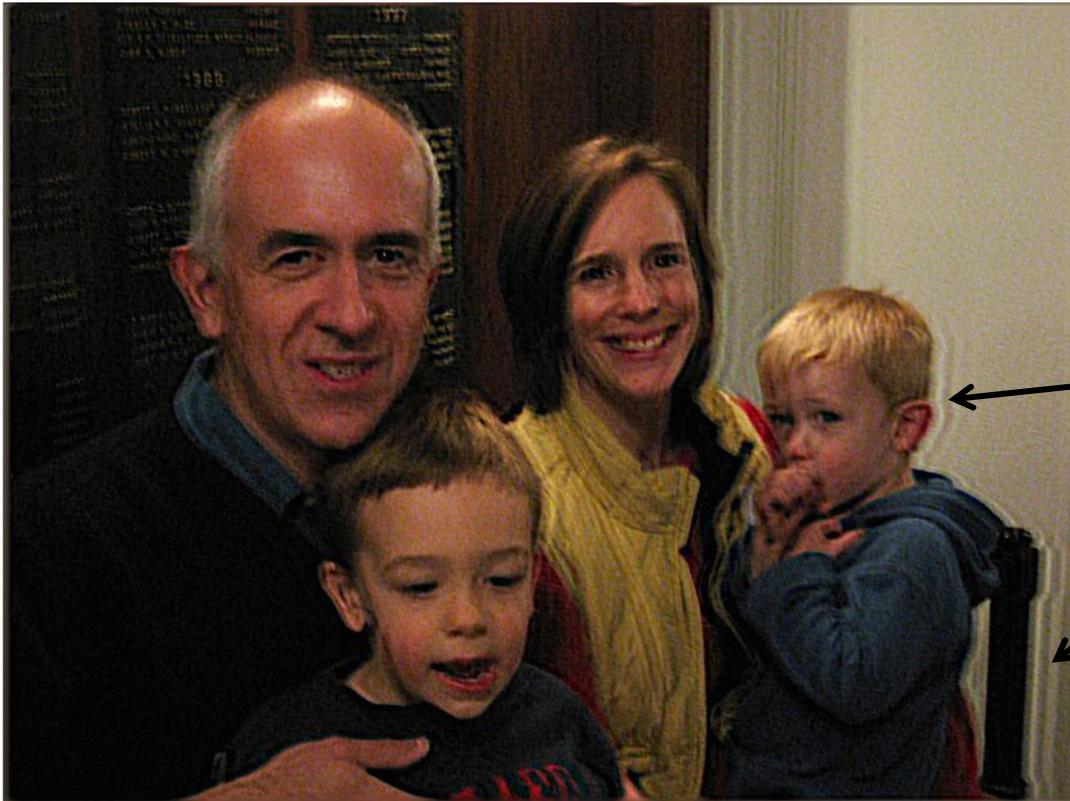




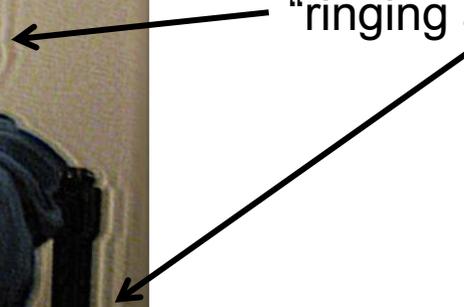
# Deconvolution

---

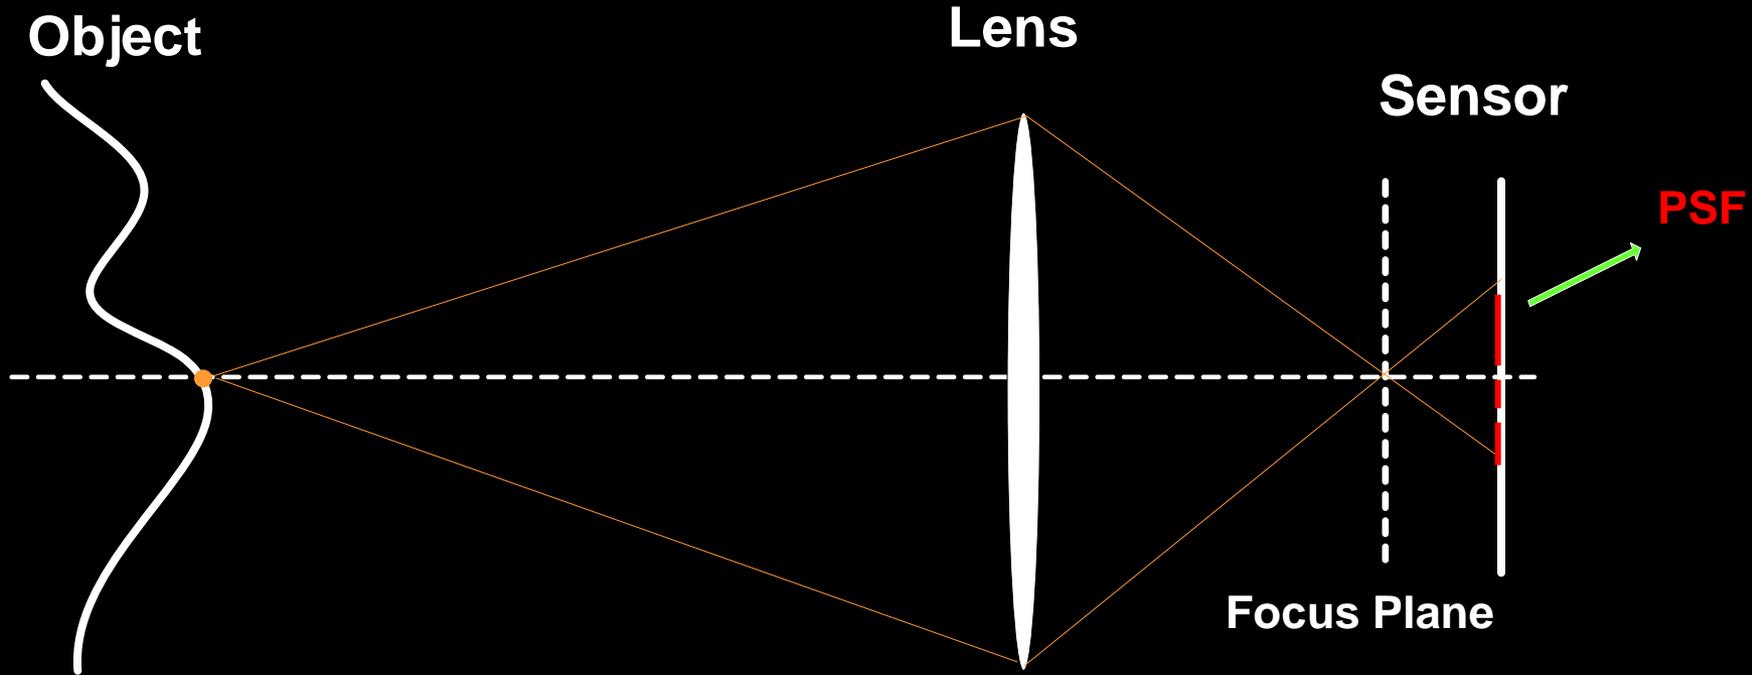
Even if blur kernel is known, recovering a good blurry image can be challenging



“ringing artifacts”

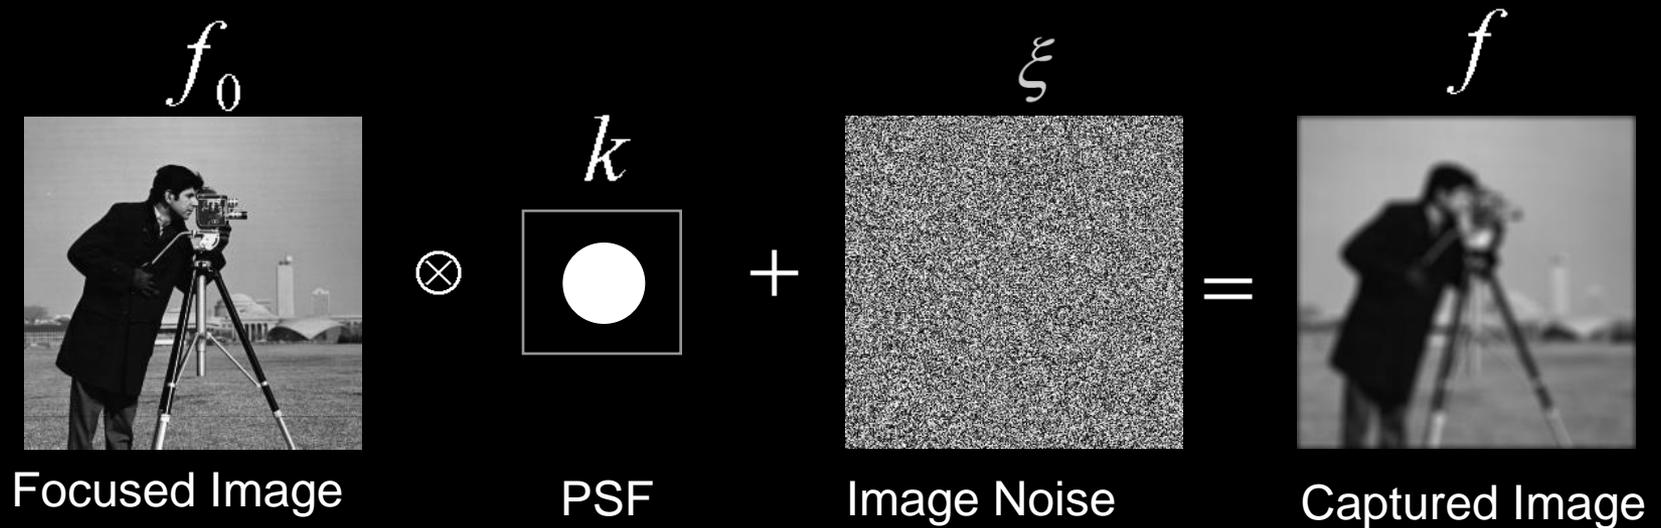


# Geometry of Defocus



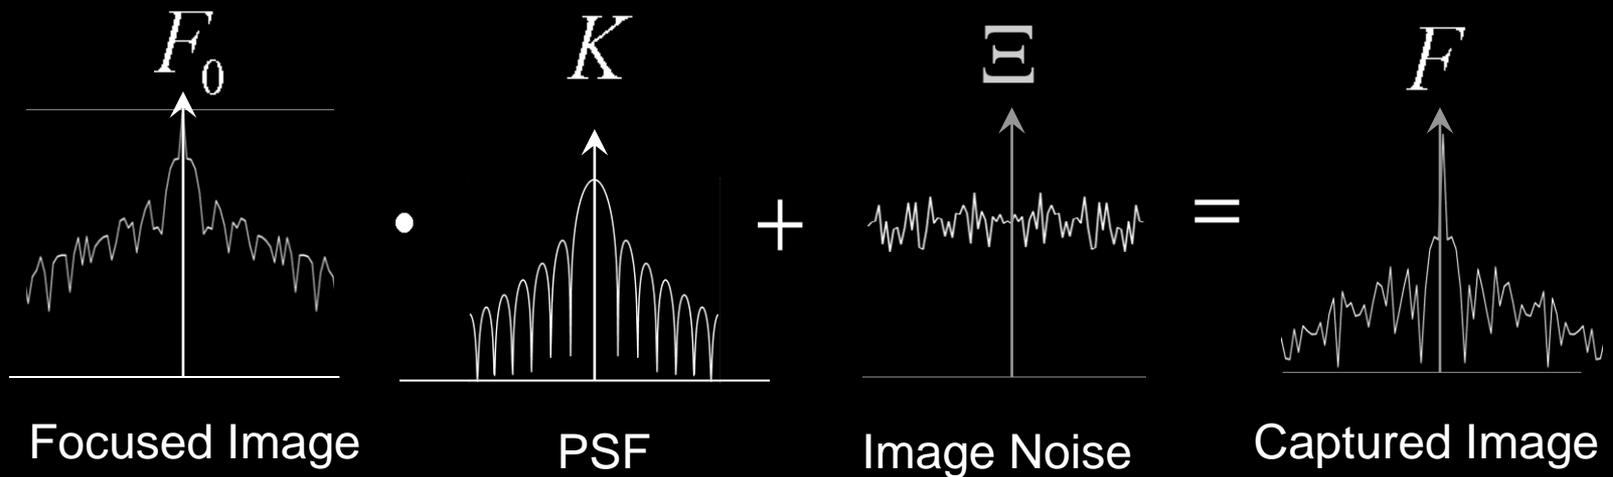
# Formulation of Defocus

In the Spatial Domain



# Formulation of Defocus

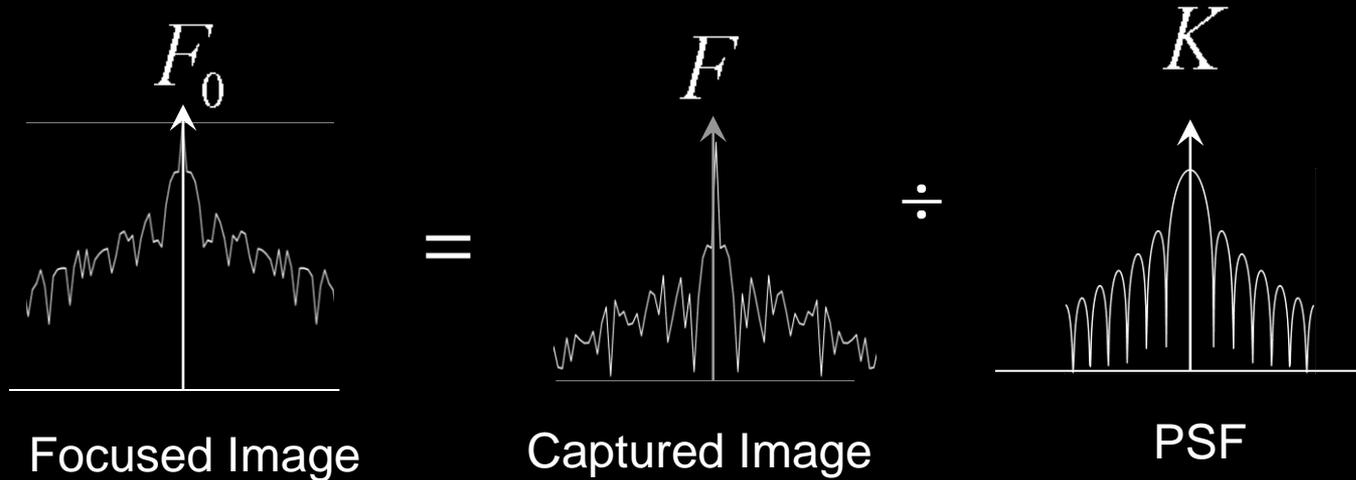
In the Fourier Domain



← Many values close to zero!

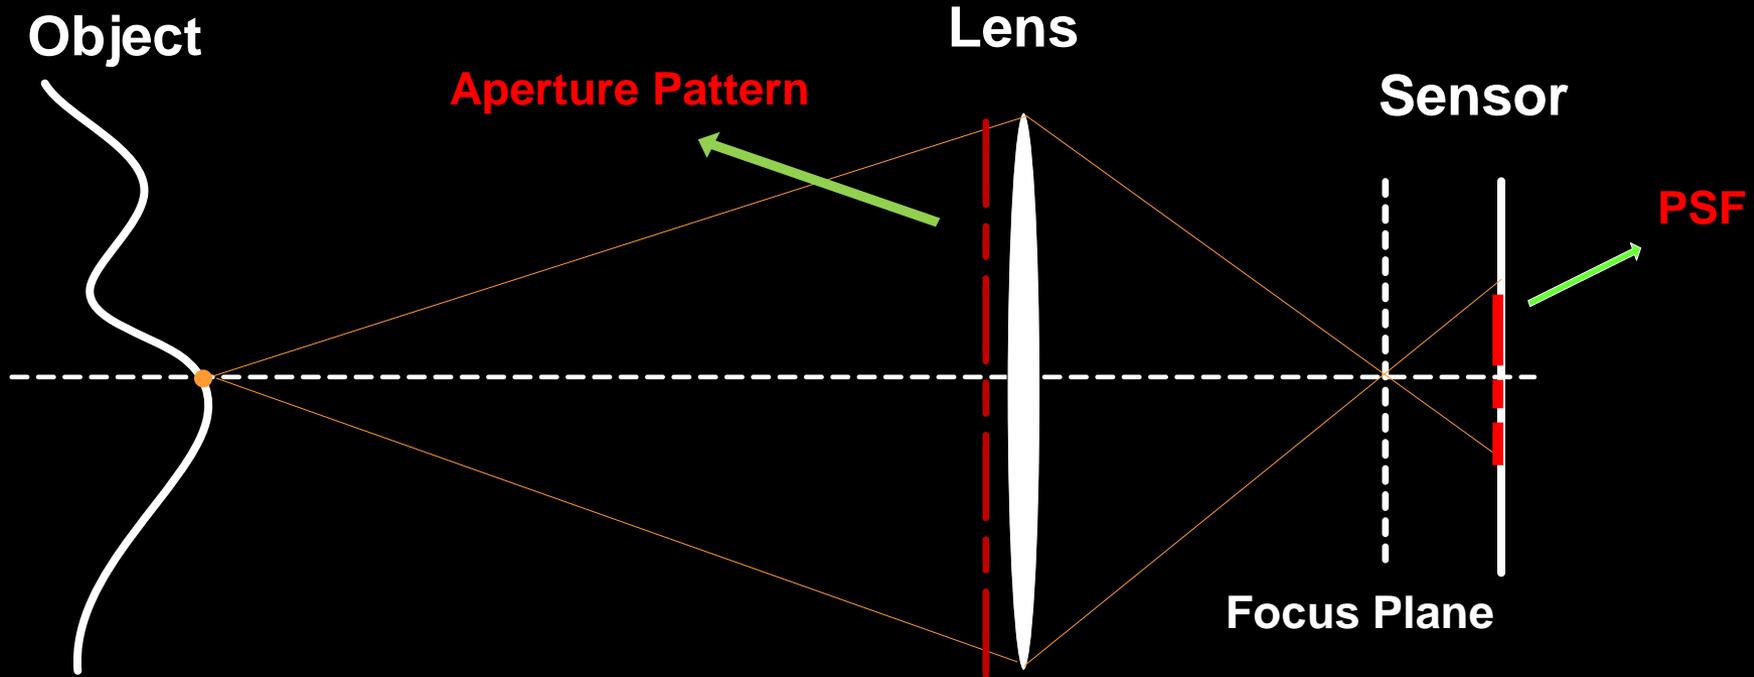
# Deconvolution

In the Fourier Domain

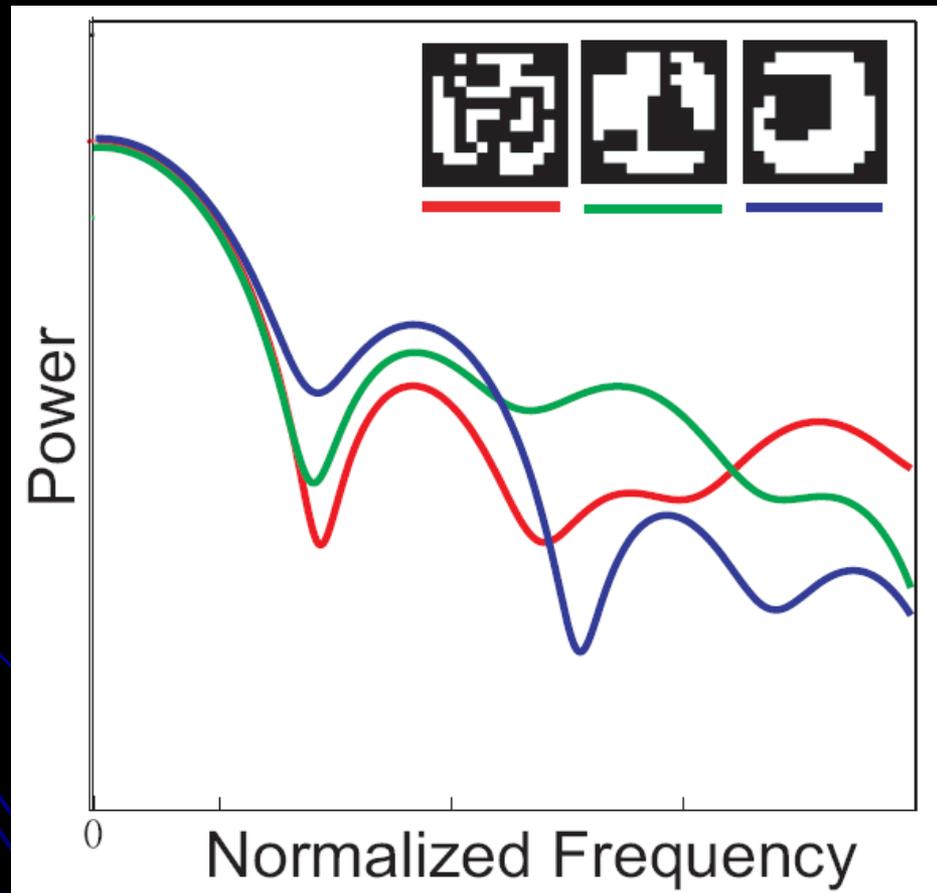


This division can be very unstable

# Geometry of Defocus



# Spectra of optimized kernels



# Implementation

Image Pattern



Veeraraghavan et al.'s Pattern



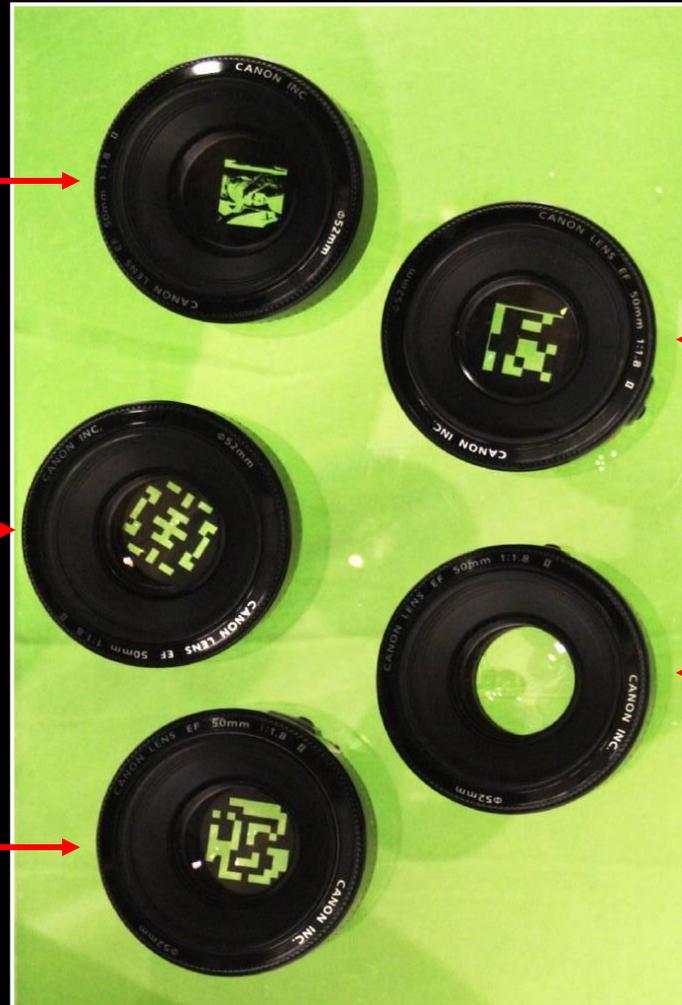
Levin et al.'s Pattern



Circular Pattern (wide open)

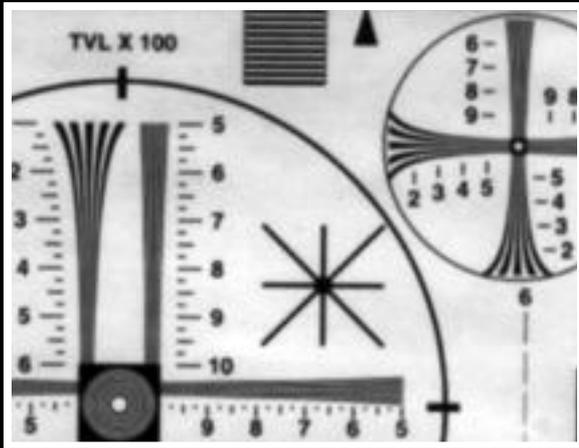


Our Optimized Pattern

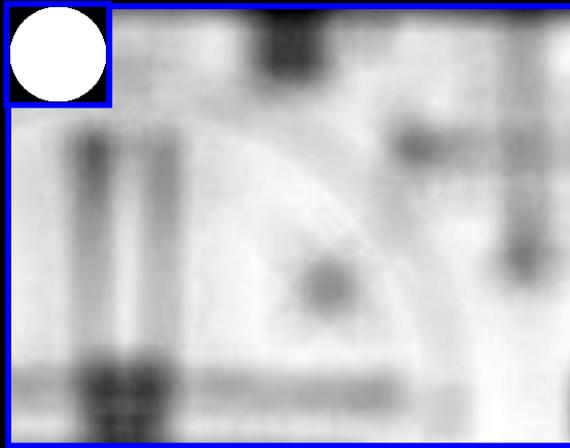


# Comparison Experiments on a CZP Chart

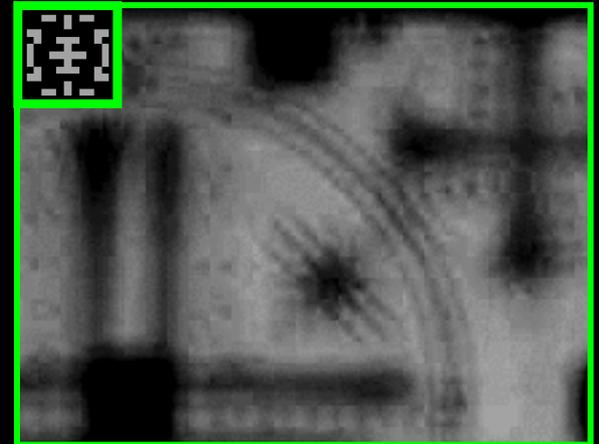
## Captured Images



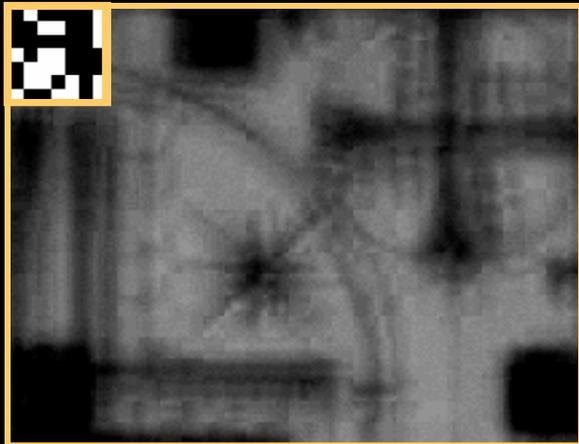
Focused Image



Circular Pattern



Levin et al.'s Pattern



Veeraraghavan et al's Pattern

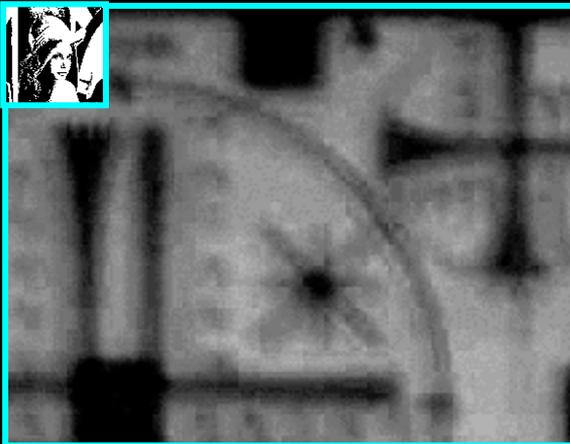
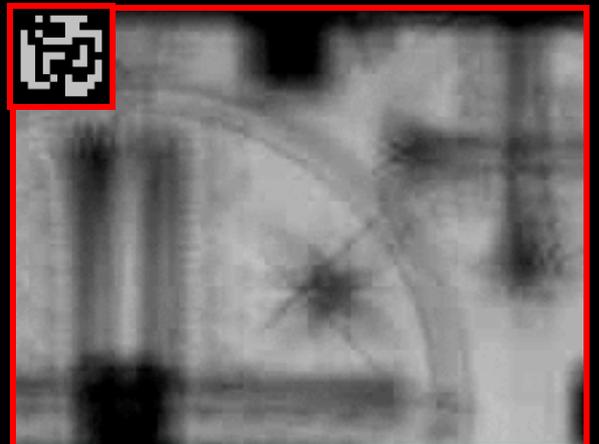


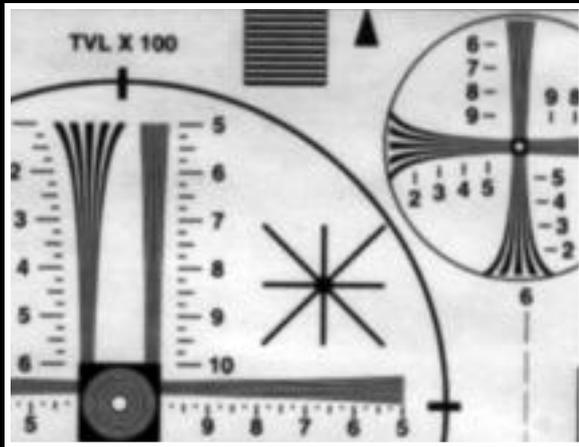
Image Pattern



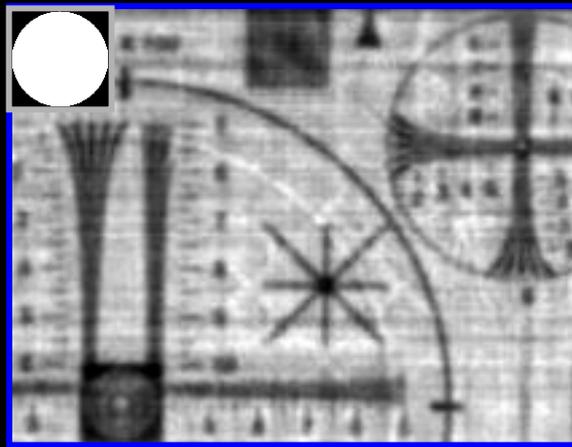
Our Optimized Pattern

# Comparison Experiments on a CZP Chart

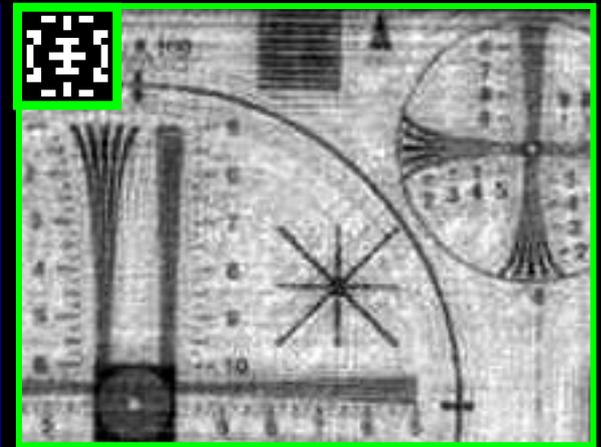
## Deblurred Images



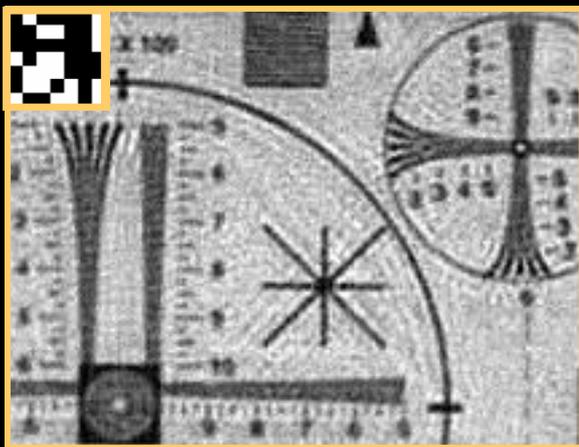
Focused Image



Circular Pattern



Levin et al.'s Pattern



Veeraraghavan et al.'s Pattern

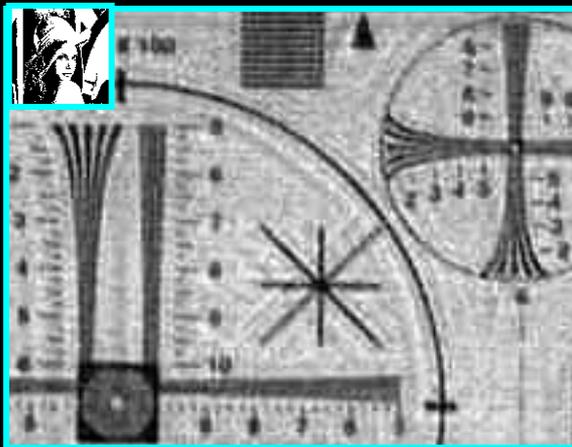
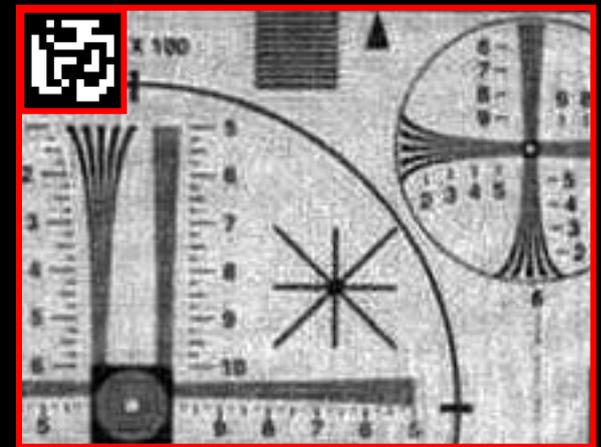


Image Pattern



Our Optimized Pattern

# Build your own coded aperture

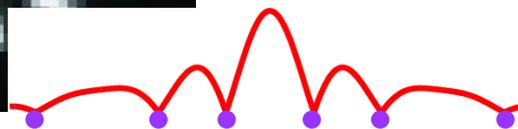
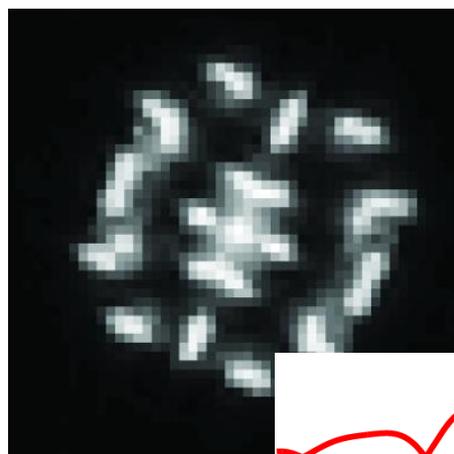
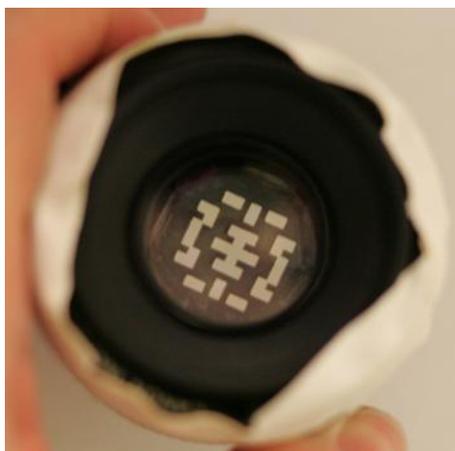
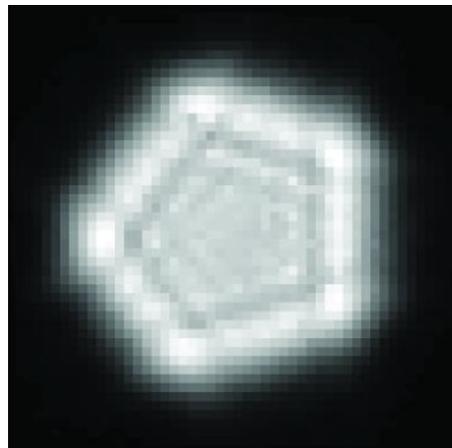


# Voilà!



# Comparison of PSFs

---



Fourier spectrum

# Input



# All-focused (deconvolved)



# Close-up

---

---

Original image



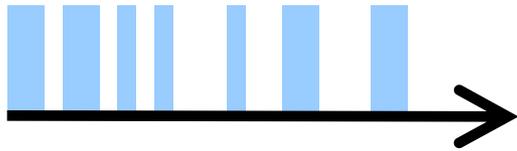
All-focus image



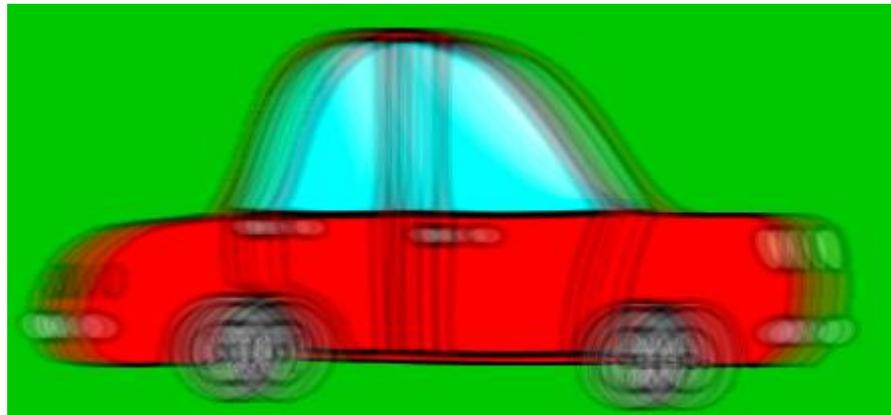
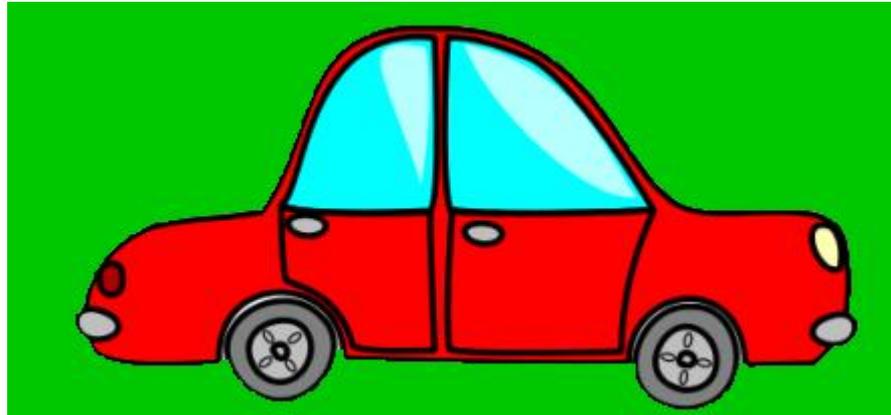
# Motion blur removal

---

Instead of coding the aperture, code the...



Shutter is OPEN and  
CLOSED



Raskar et al.,  
SIGGRAPH 2007



License Plate Retrieval

# Many more possibilities

---

## Seeing through/behind objects

- Using a camera array (“synthetic aperture”)
- [Levoy et al., SIGGRAPH 2004](#)

## Removing interreflections

- [Nayar et al., SIGGRAPH 2006](#)

## Family portraits where everyone’s smiling

- [Photomontage \(Agarwala et al., SIGGRAPH 2004\)](#)

...

# More on computational photography

---

SIGGRAPH course [notes and video](#)

Other courses

- [MIT course](#)
- [CMU course](#)
- [Stanford course](#)
- [Columbia course](#)

[Wikipedia page](#)

[Symposium on Computational Photography](#)

[ICCP 2009 \(conference\)](#)