Focal stacks and lightfields

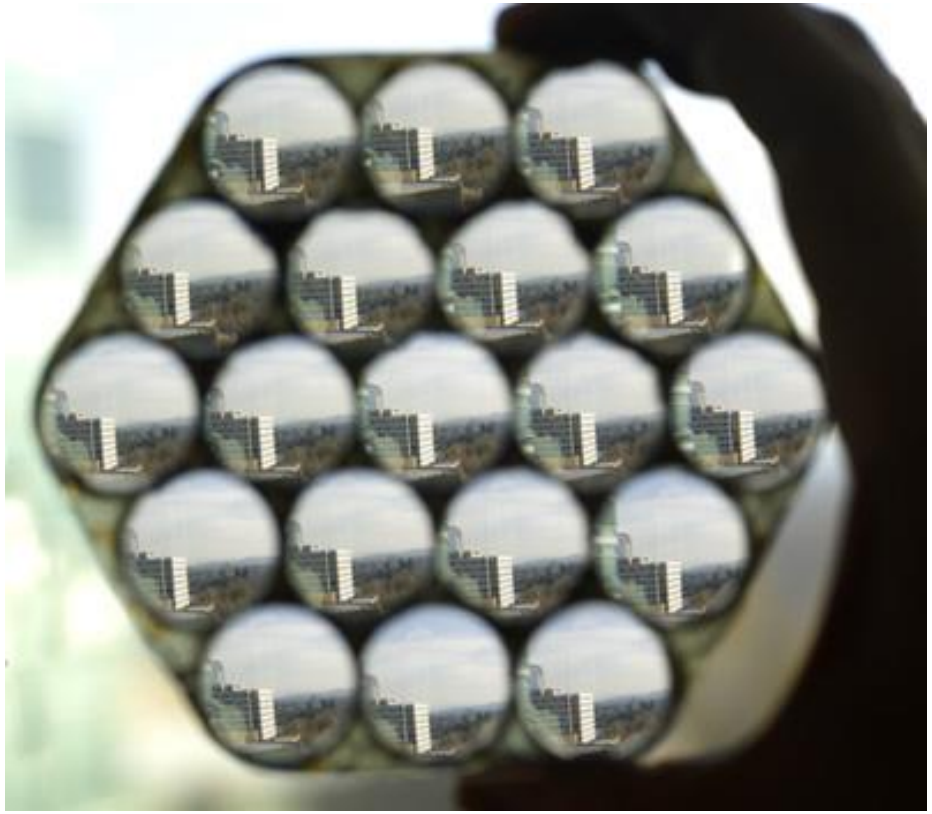

15-463, 15-663, 15-862
Computational Photography
Fall 2017, Lecture 16

http://graphics.cs.cmu.edu/courses/15-463

Course announcements

- Homework 4 is out.
  - Due October 26th.
  - Bilateral filter will take a very long time to run.
  - Drop by Yannis' office to pick up cameras any time.

- Project ideas were due on Piazza on Friday 20th.
  - Responded to most of you.
  - Some still need to post their ideas.

- Project *proposals* are due on Monday 30th.

Overview of today's lecture

- Pinhole vs lens cameras.
- Focal stack.
- Lightfield.
- Measuring lightfields.
- Plenoptic camera.
- Images from lightfields.

Slide credits

Most of these slides were adapted from:

- Fredo Durand (MIT).
- Gordon Wetzstein (Stanford).

Pinhole vs lens cameras

Pinhole camera

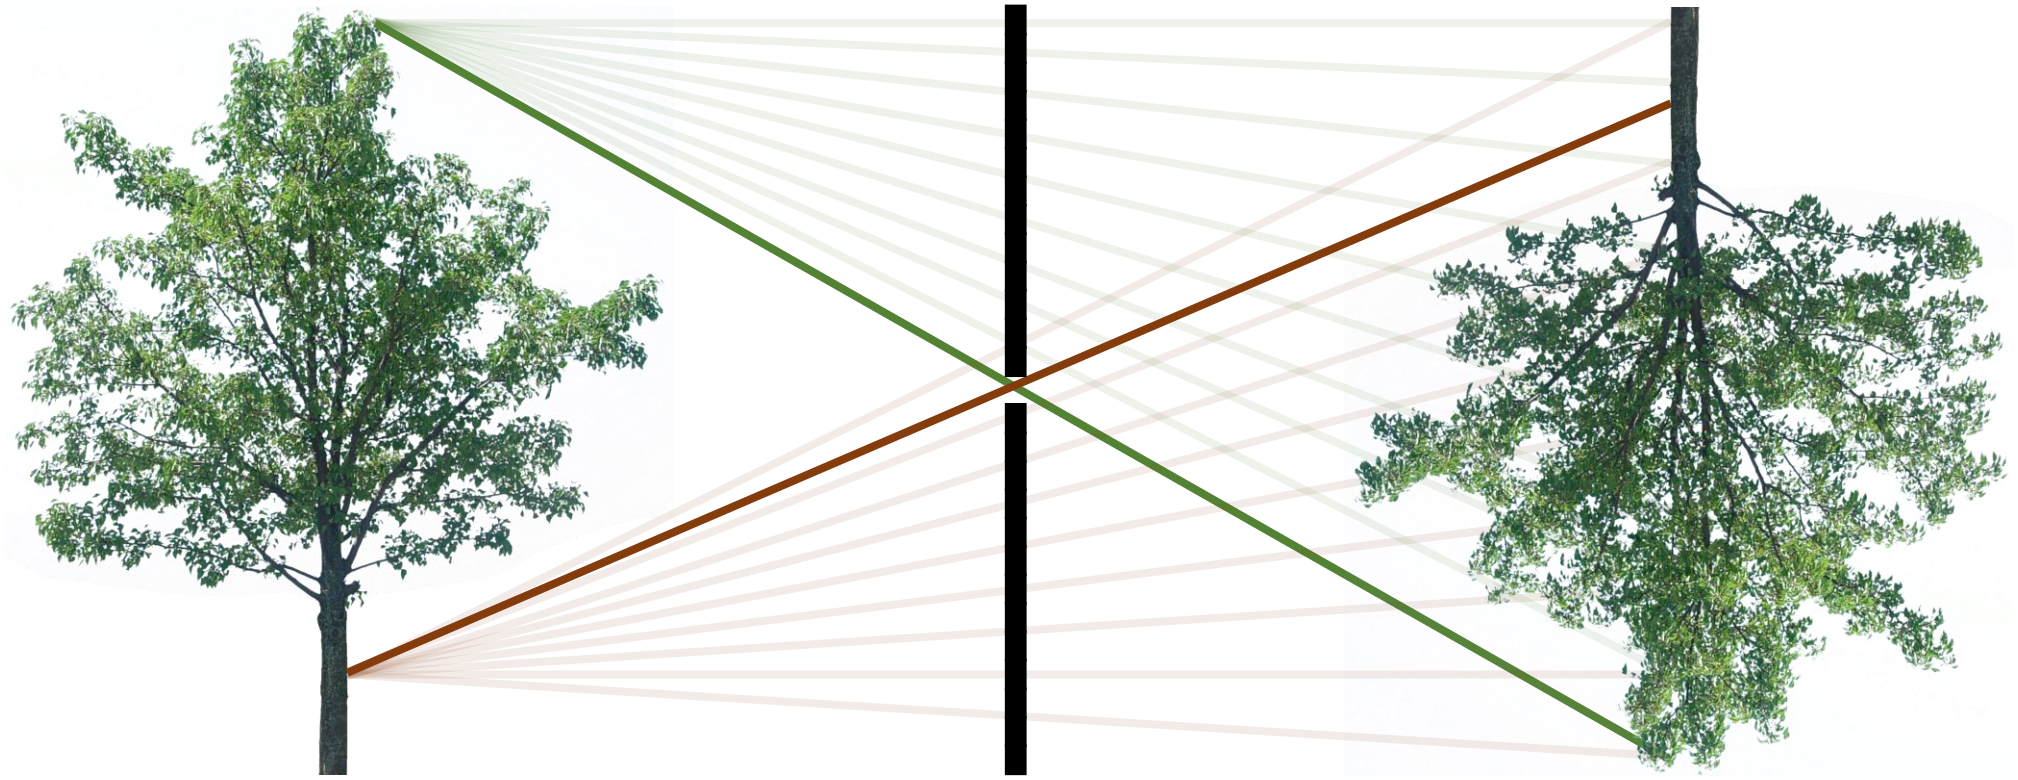

- Everything is in focus.
- Very light inefficient.

Lens camera

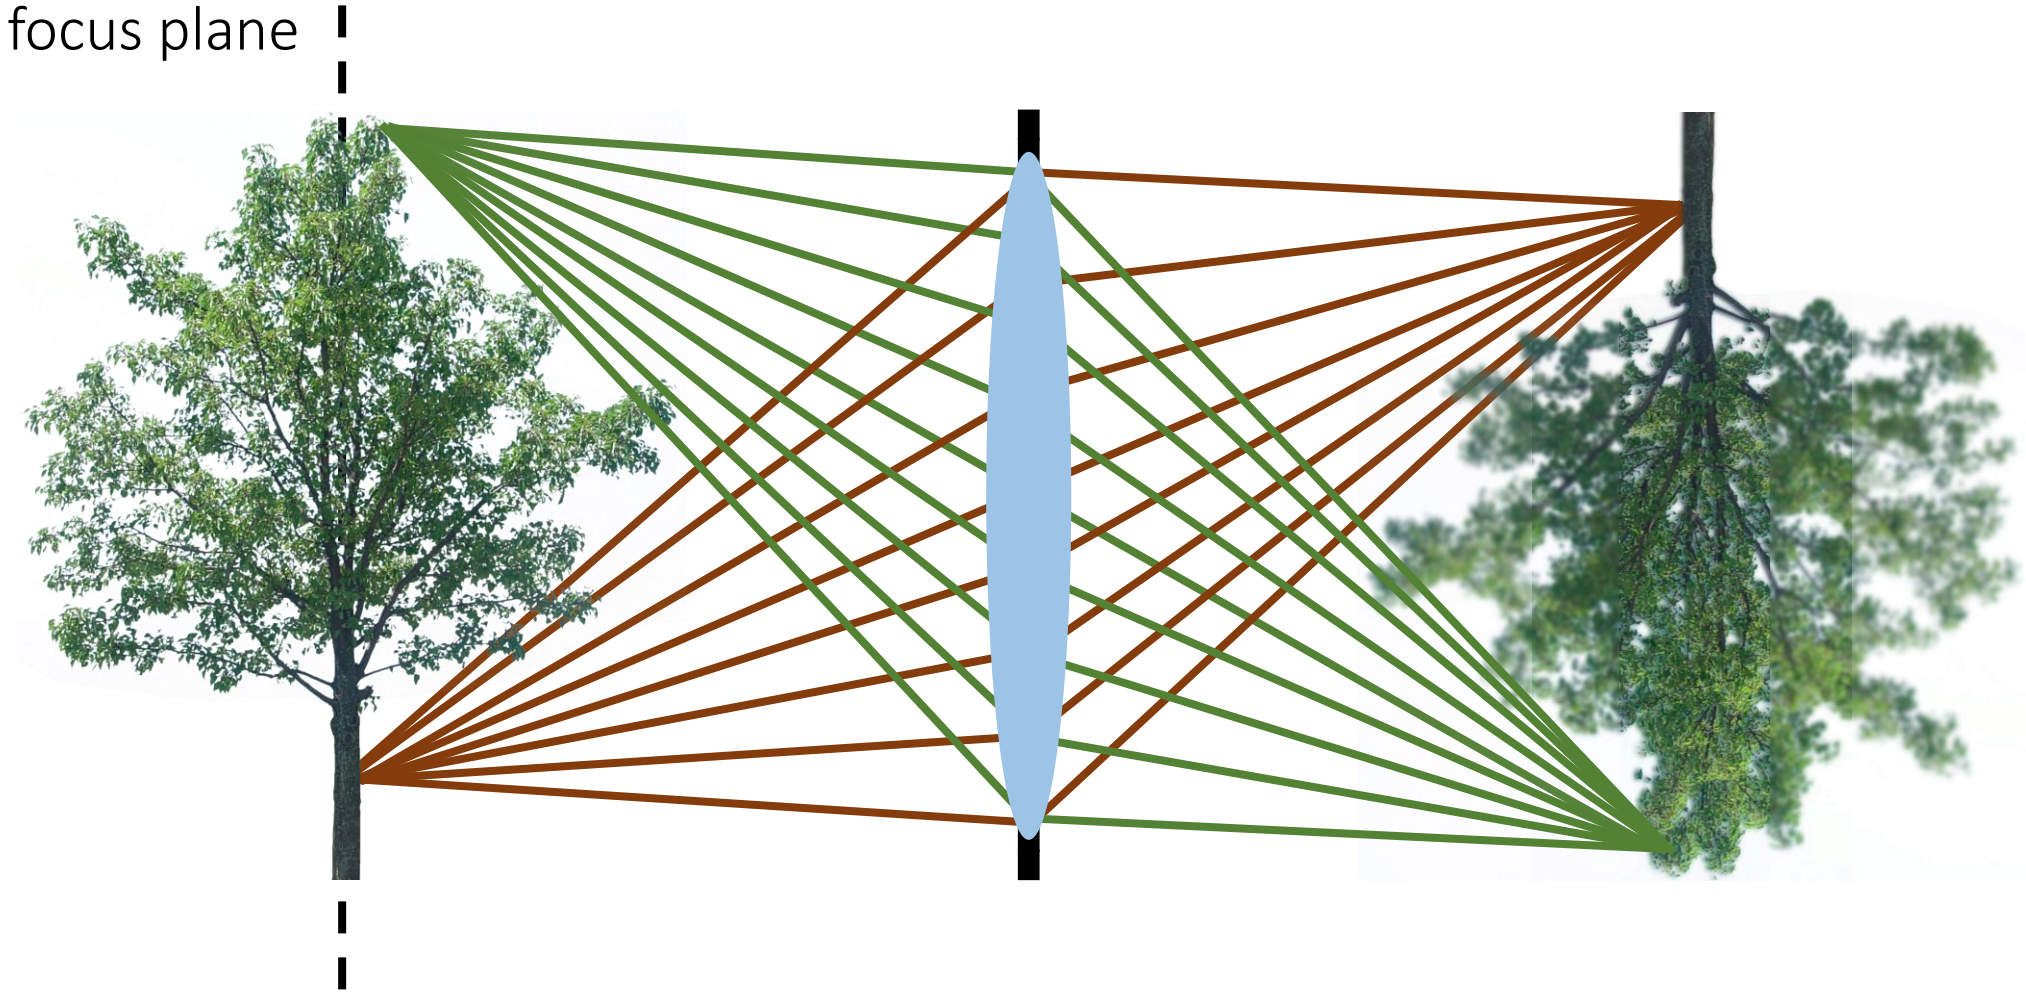

- Only one plane is in focus.
- Very light efficient.

How can we get an all in-focus image?

Focal stack

Focal stack

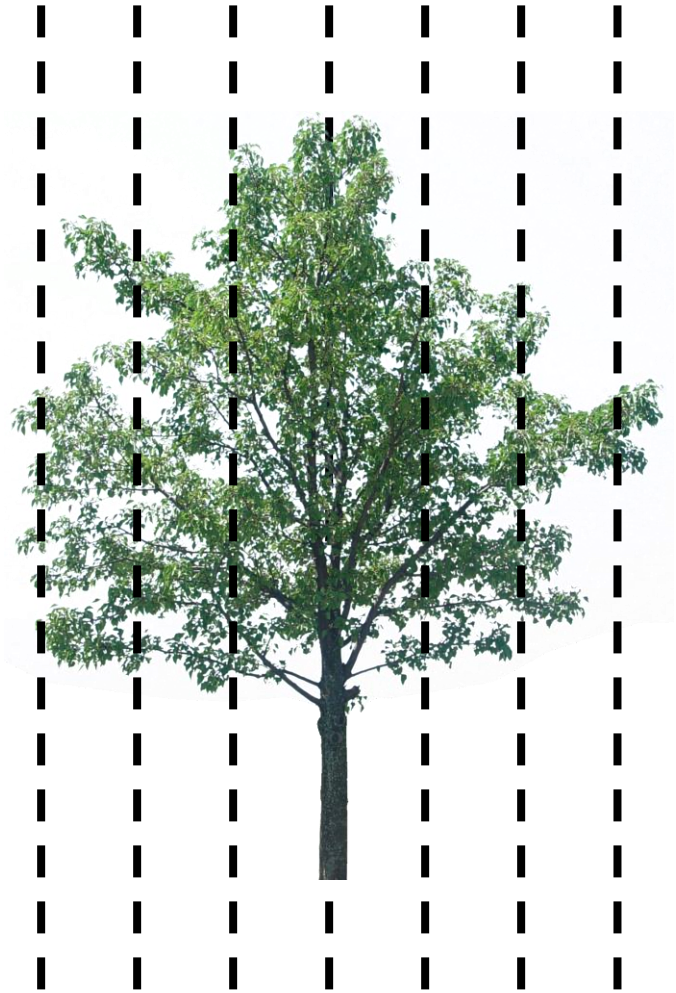

- Capture images focused at multiple planes.
- Merge them into a single all in-focus image.

Analogous to what we did in HDR

- Focal stack instead of exposure stack.

Homework 5: focal stack imaging

Focal stack imaging

1. Capture a focal stack




2. Merge into an all in-focus image

Focal stack imaging

1. Capture a focal stack

2. Merge into an all in-focus image

How do you capture a focal stack?

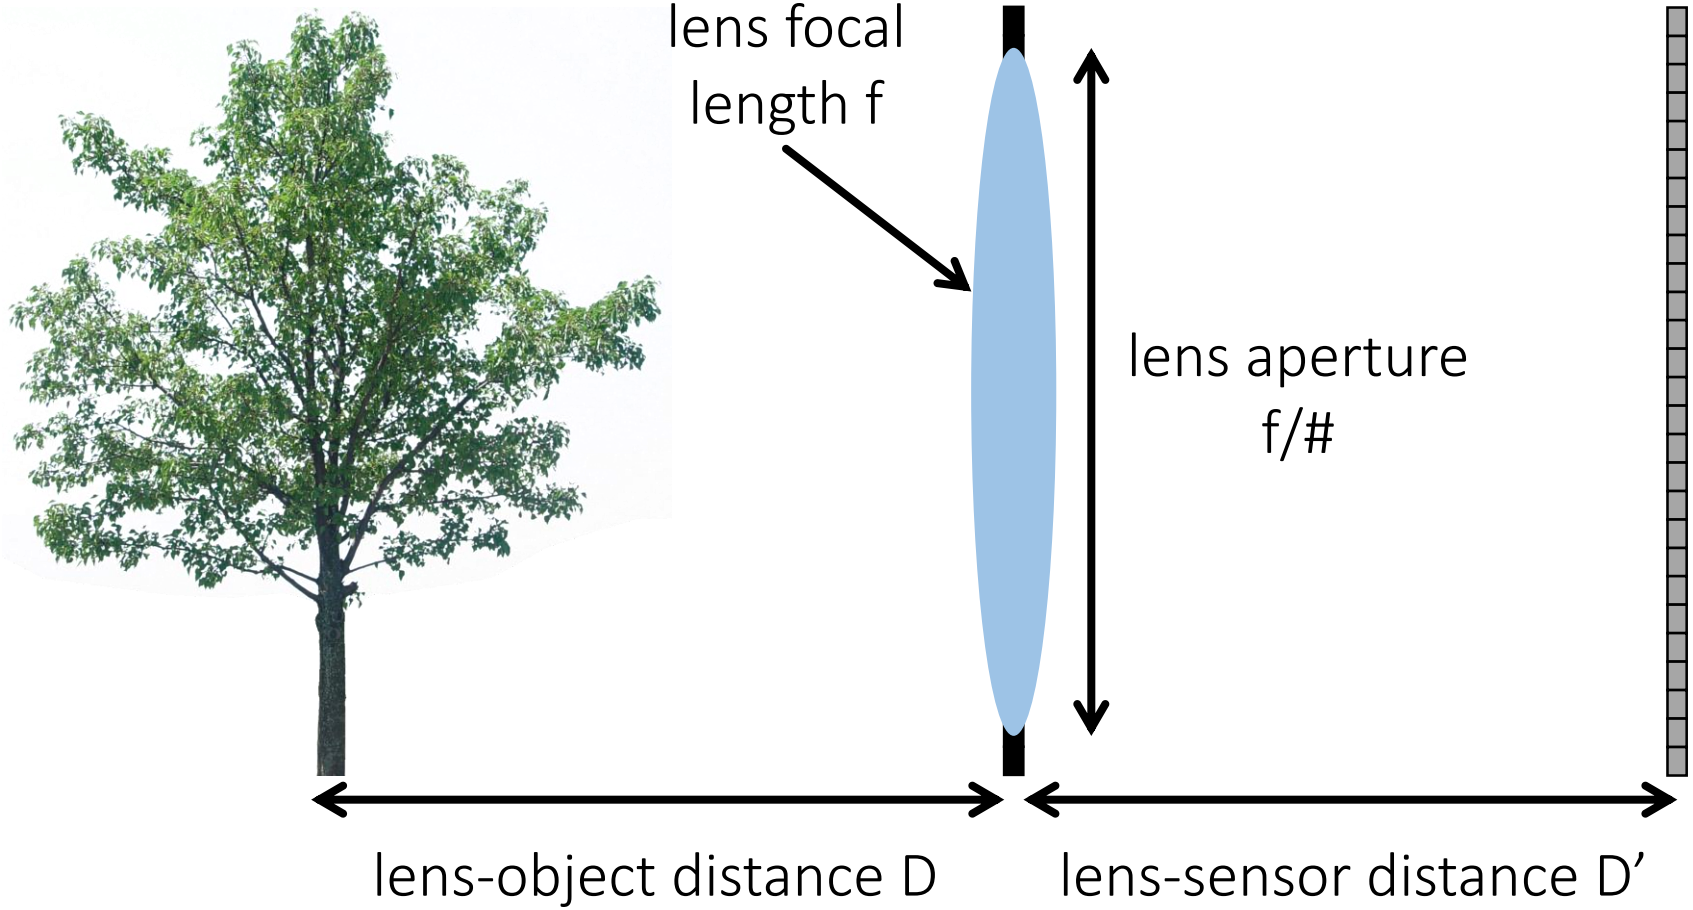

Which of these parameters would you change (and how)?

How do you capture a focal stack?

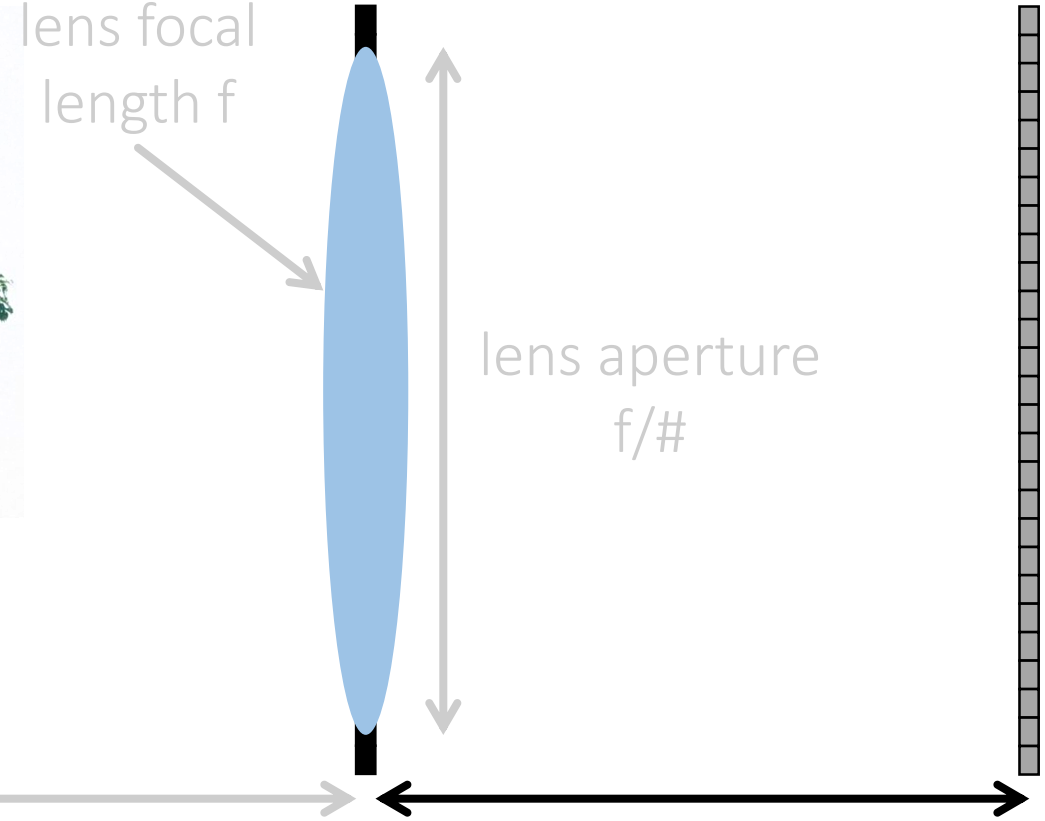

rotate lens focus ring (not zoom!)

Which of these parameters would you change (and how would you achieve that)?

How do you capture a focal stack?

Which of these parameters would you change (and how would you achieve that)?

Capturing a focal stack

In-focus plane in each stack image


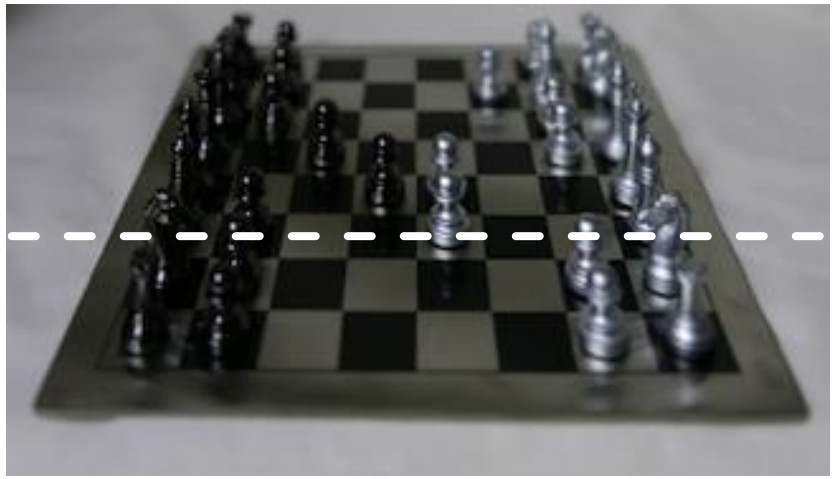


Focal stack imaging

1. Capture a focal stack



2. Merge into an all in-focus image

How do you merge a focal stack?

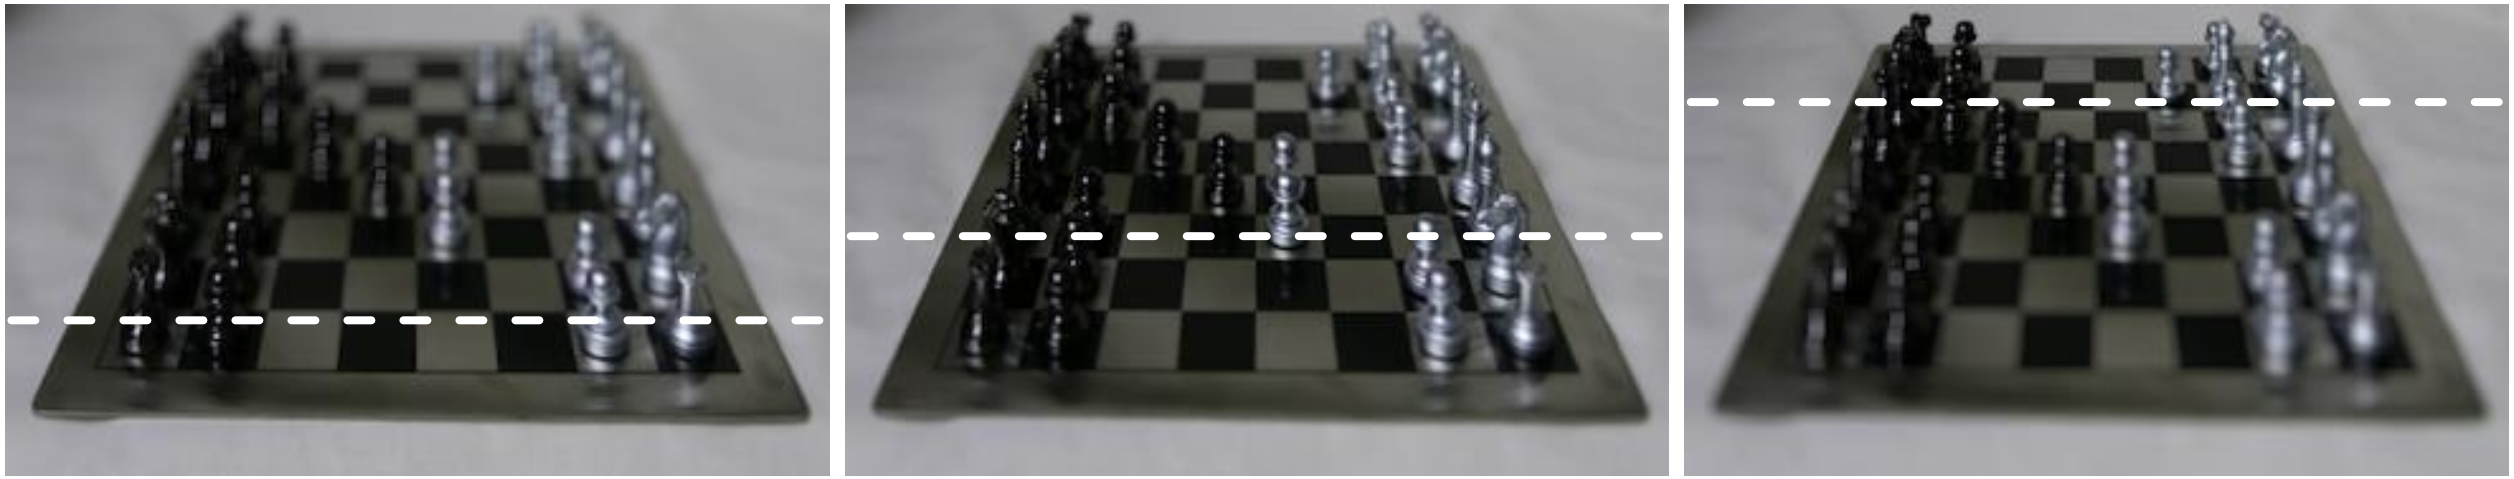

How do you merge a focal stack?

1. Align images
2. Assign per-pixel weights representing “in-focus”-ness
3. Compute image average

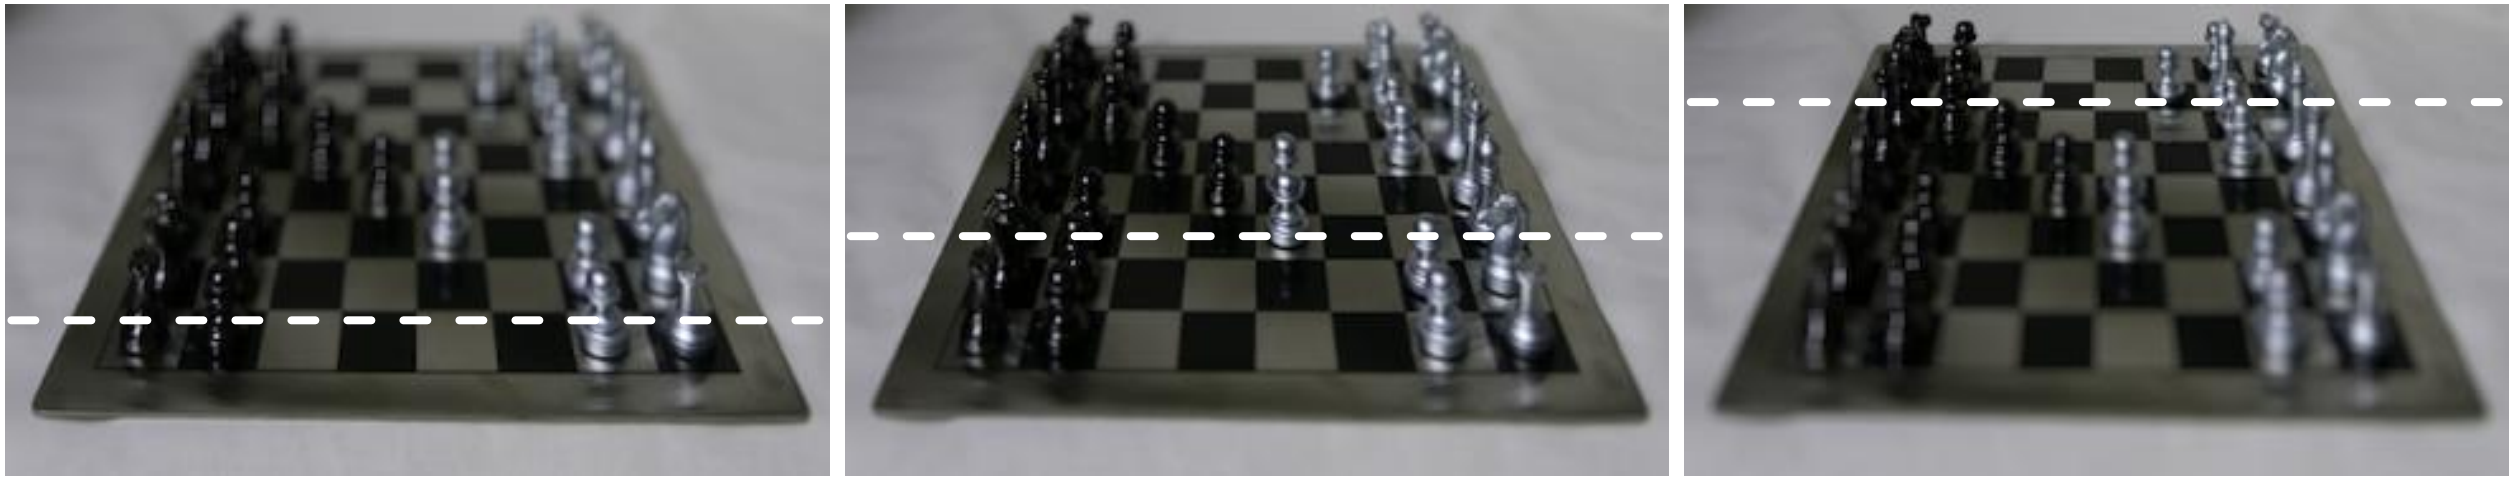

How do you merge a focal stack?

1. Align images
2. Assign per-pixel weights representing “in-focus”-ness
3. Compute image average


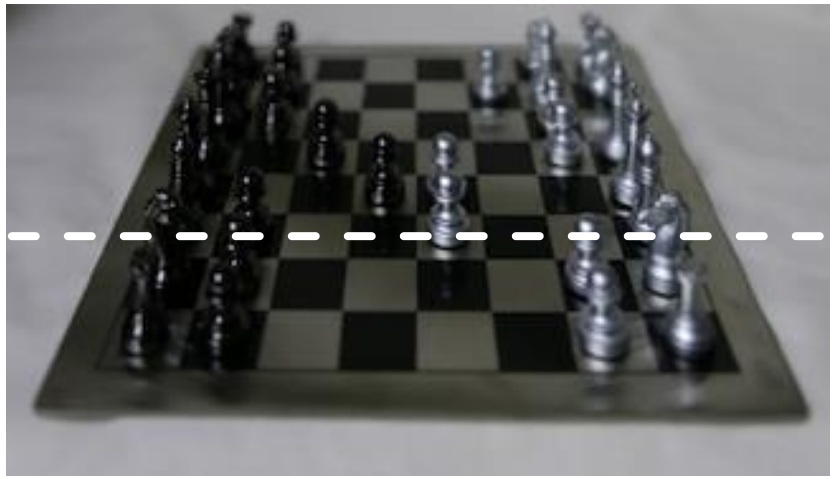


Image alignment

Why do we need to align the images?

Image alignment

Why do we need to align the images?

- When we change focus distance, we also change field of view (magnification).
- Also, scene may not be static (but we will be ignoring this for now).

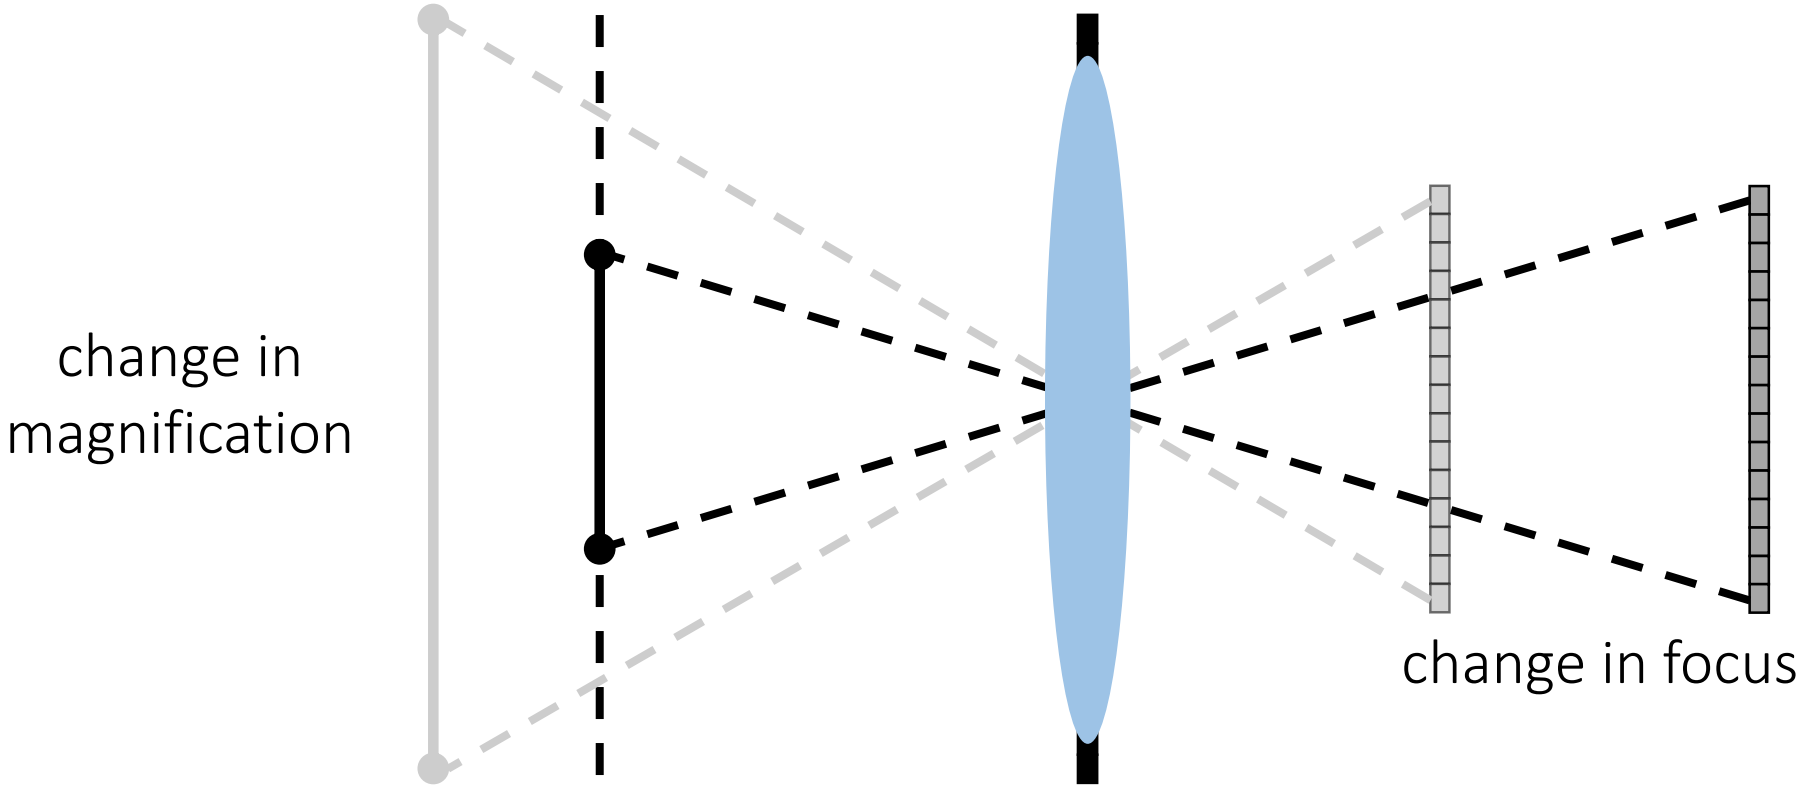

Image alignment

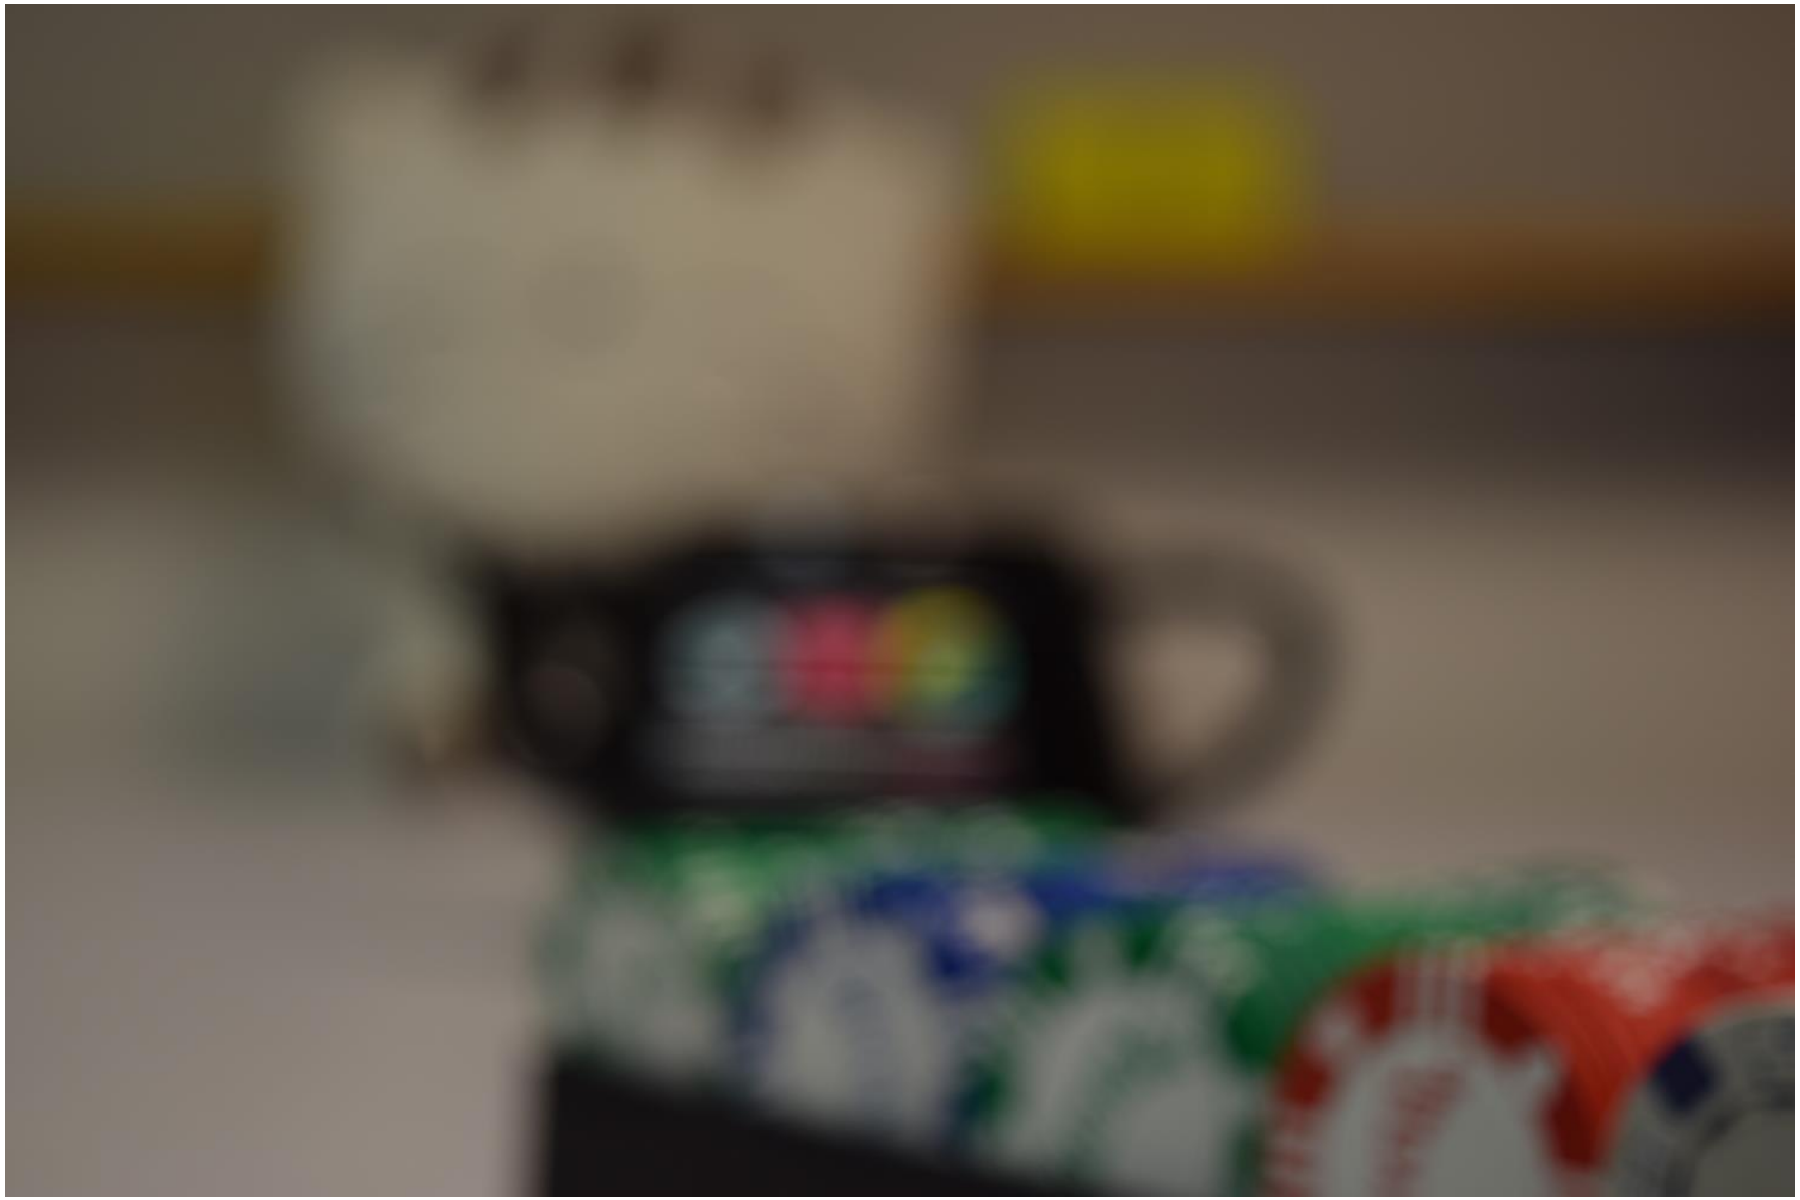

Image alignment

Image alignment

Why do we need to align the images?

- When we change focus distance, we also change field of view (magnification).
- Also, scene may not be static (but we will be ignoring this for now).

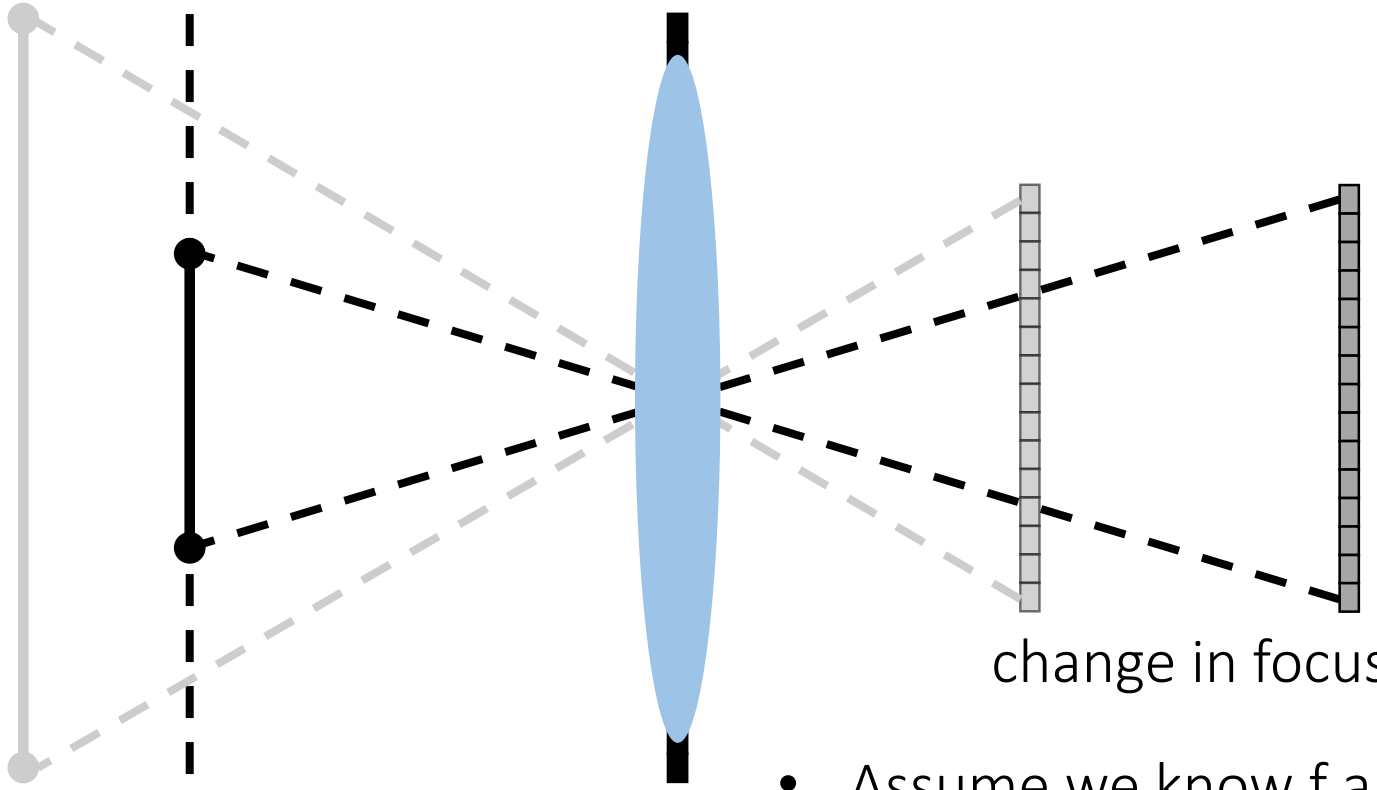

- Assume we know f and all D’ values.
- How do can we align the images?

Image alignment

Why do we need to align the images?

- When we change focus distance, we also change field of view (magnification).
- Also, scene may not be static (but we will be ignoring this for now).

change in magnification

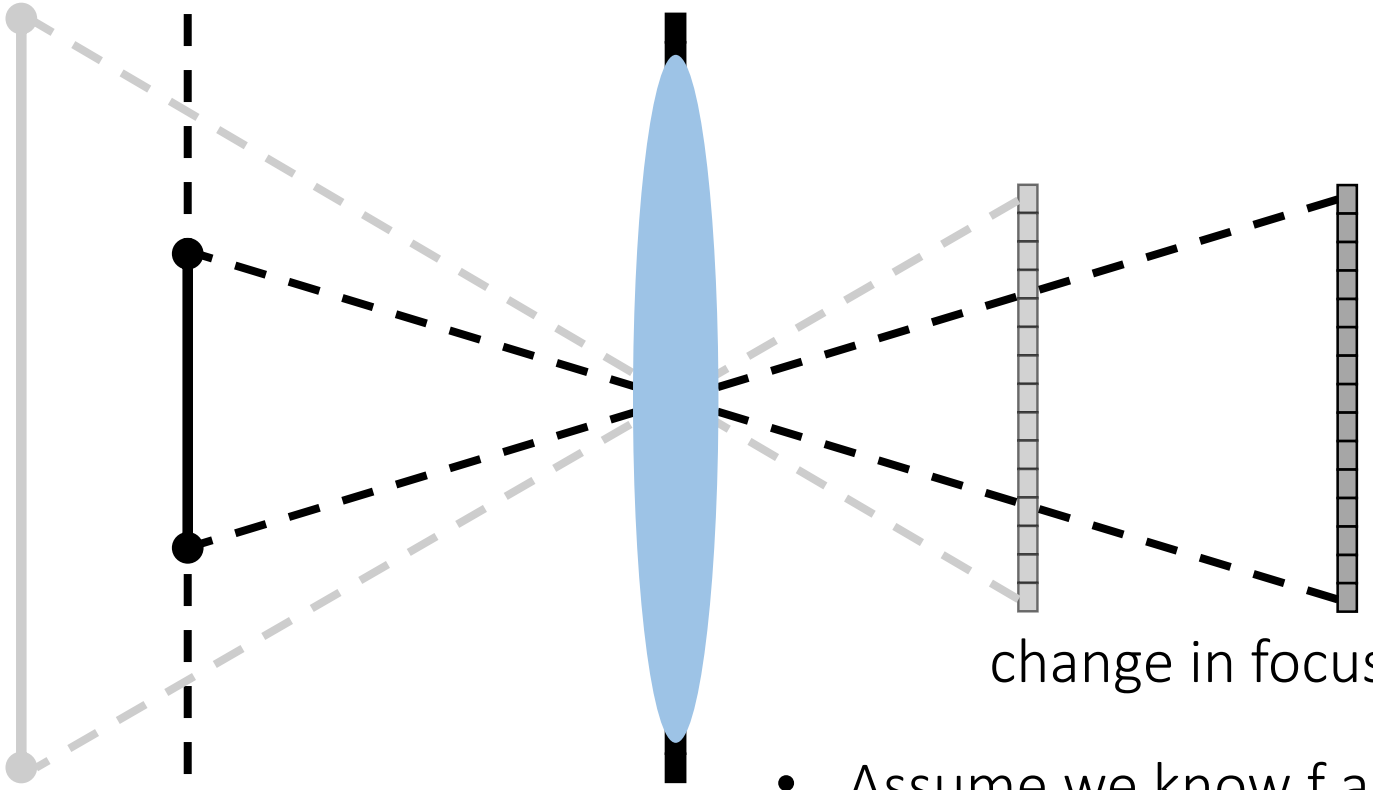

change in focus

$$m = \frac{f}{D' - f}$$

$$\frac{1}{D'} + \frac{1}{D} = \frac{1}{f}$$

resize using these equations

- Assume we know f and all D’ values.
- How do can we align the images?

How do you merge a focal stack?

1. Align images
2. Assign per-pixel weights representing “in-focus”-ness
3. Compute image average


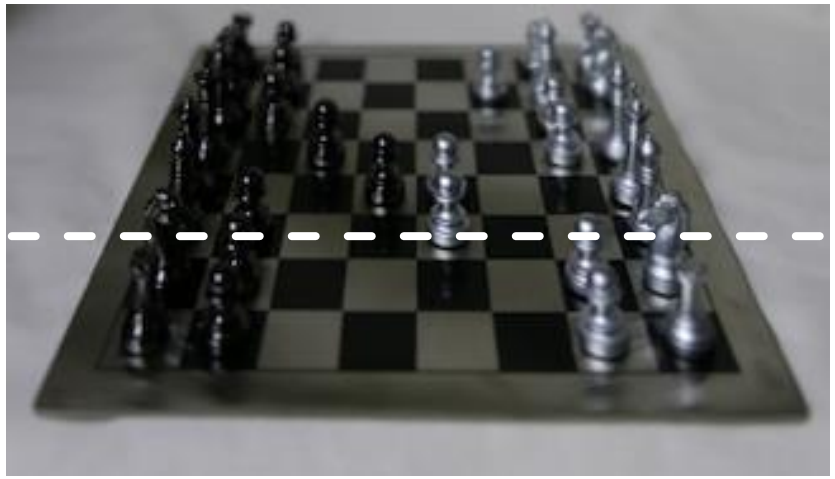


Weight assignment

How do we measure how much “in-focus” each pixel is?

Weight assignment

How do we measure how much “in-focus” each pixel is?

- Measure local sharpness.
- This is also how auto-focus works.

How do we measure local sharpness?

Weight assignment

How do we measure how much “in-focus” each pixel is?

- Measure local sharpness.
- This is also how auto-focus works.

run Laplacian operator

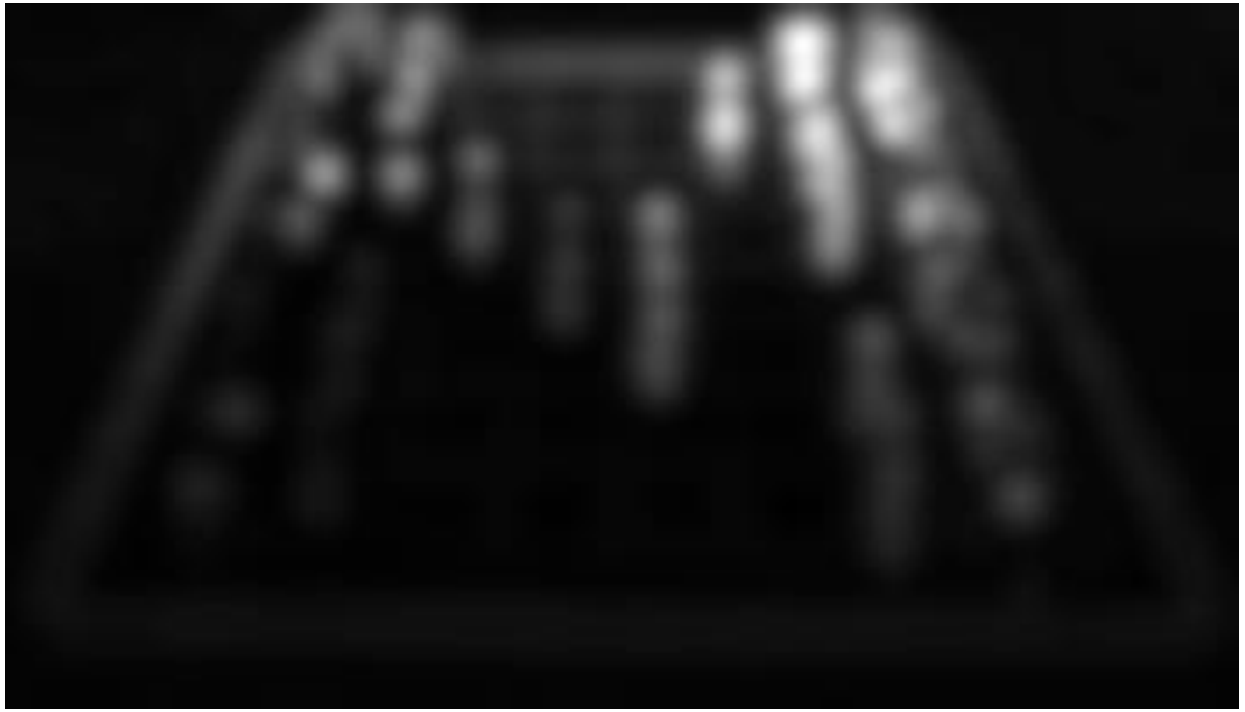

do some Gaussian blurring (why?)

Weight assignment

How do we measure how much “in-focus” each pixel is?

- Measure local sharpness.
- This is also how auto-focus works.

run Laplacian operator

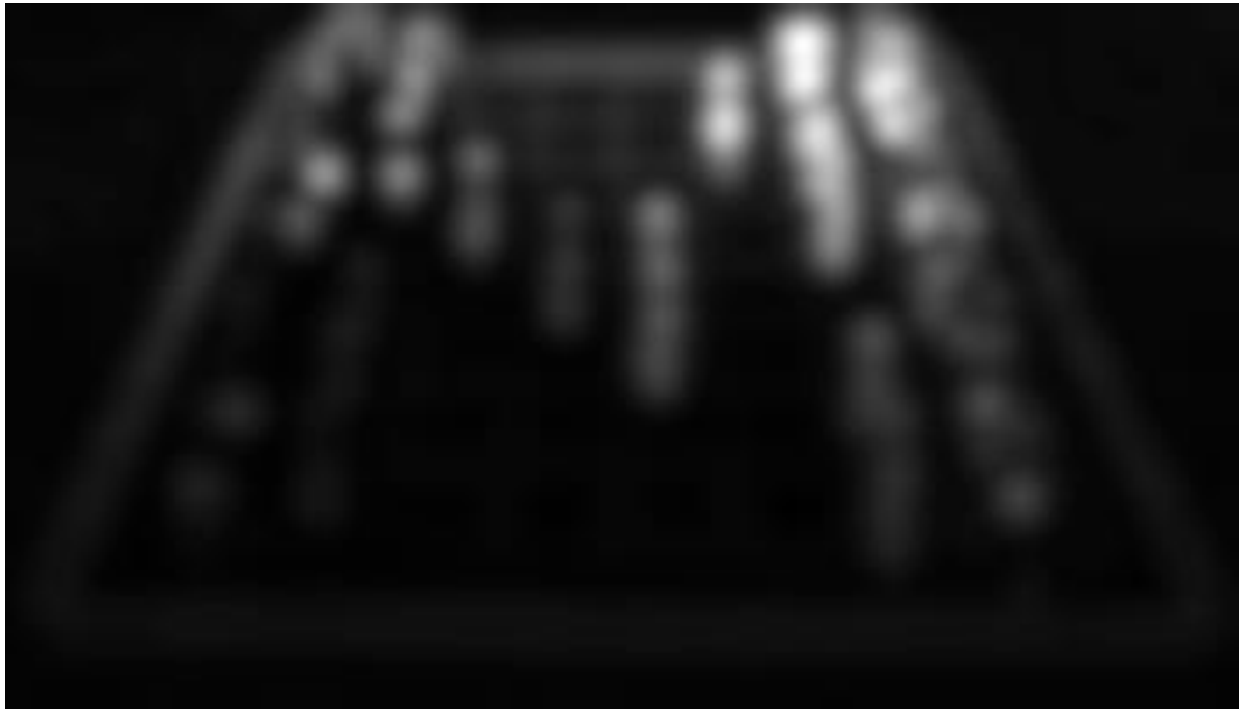

do some Gaussian blurring (so that nearby pixels have similar weights)

Just one example, many alternatives possible.

How do you merge a focal stack?

1. Align images
2. Assign per-pixel weights representing “in-focus”-ness
3. Compute image average


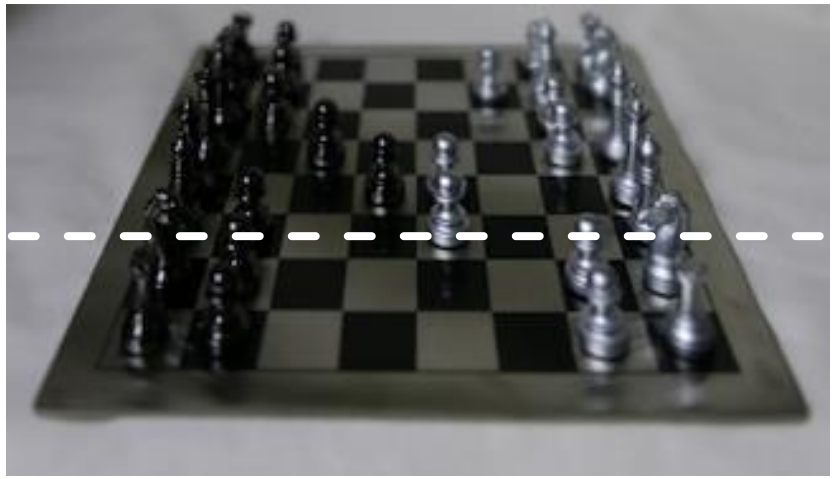


Focal stack merging

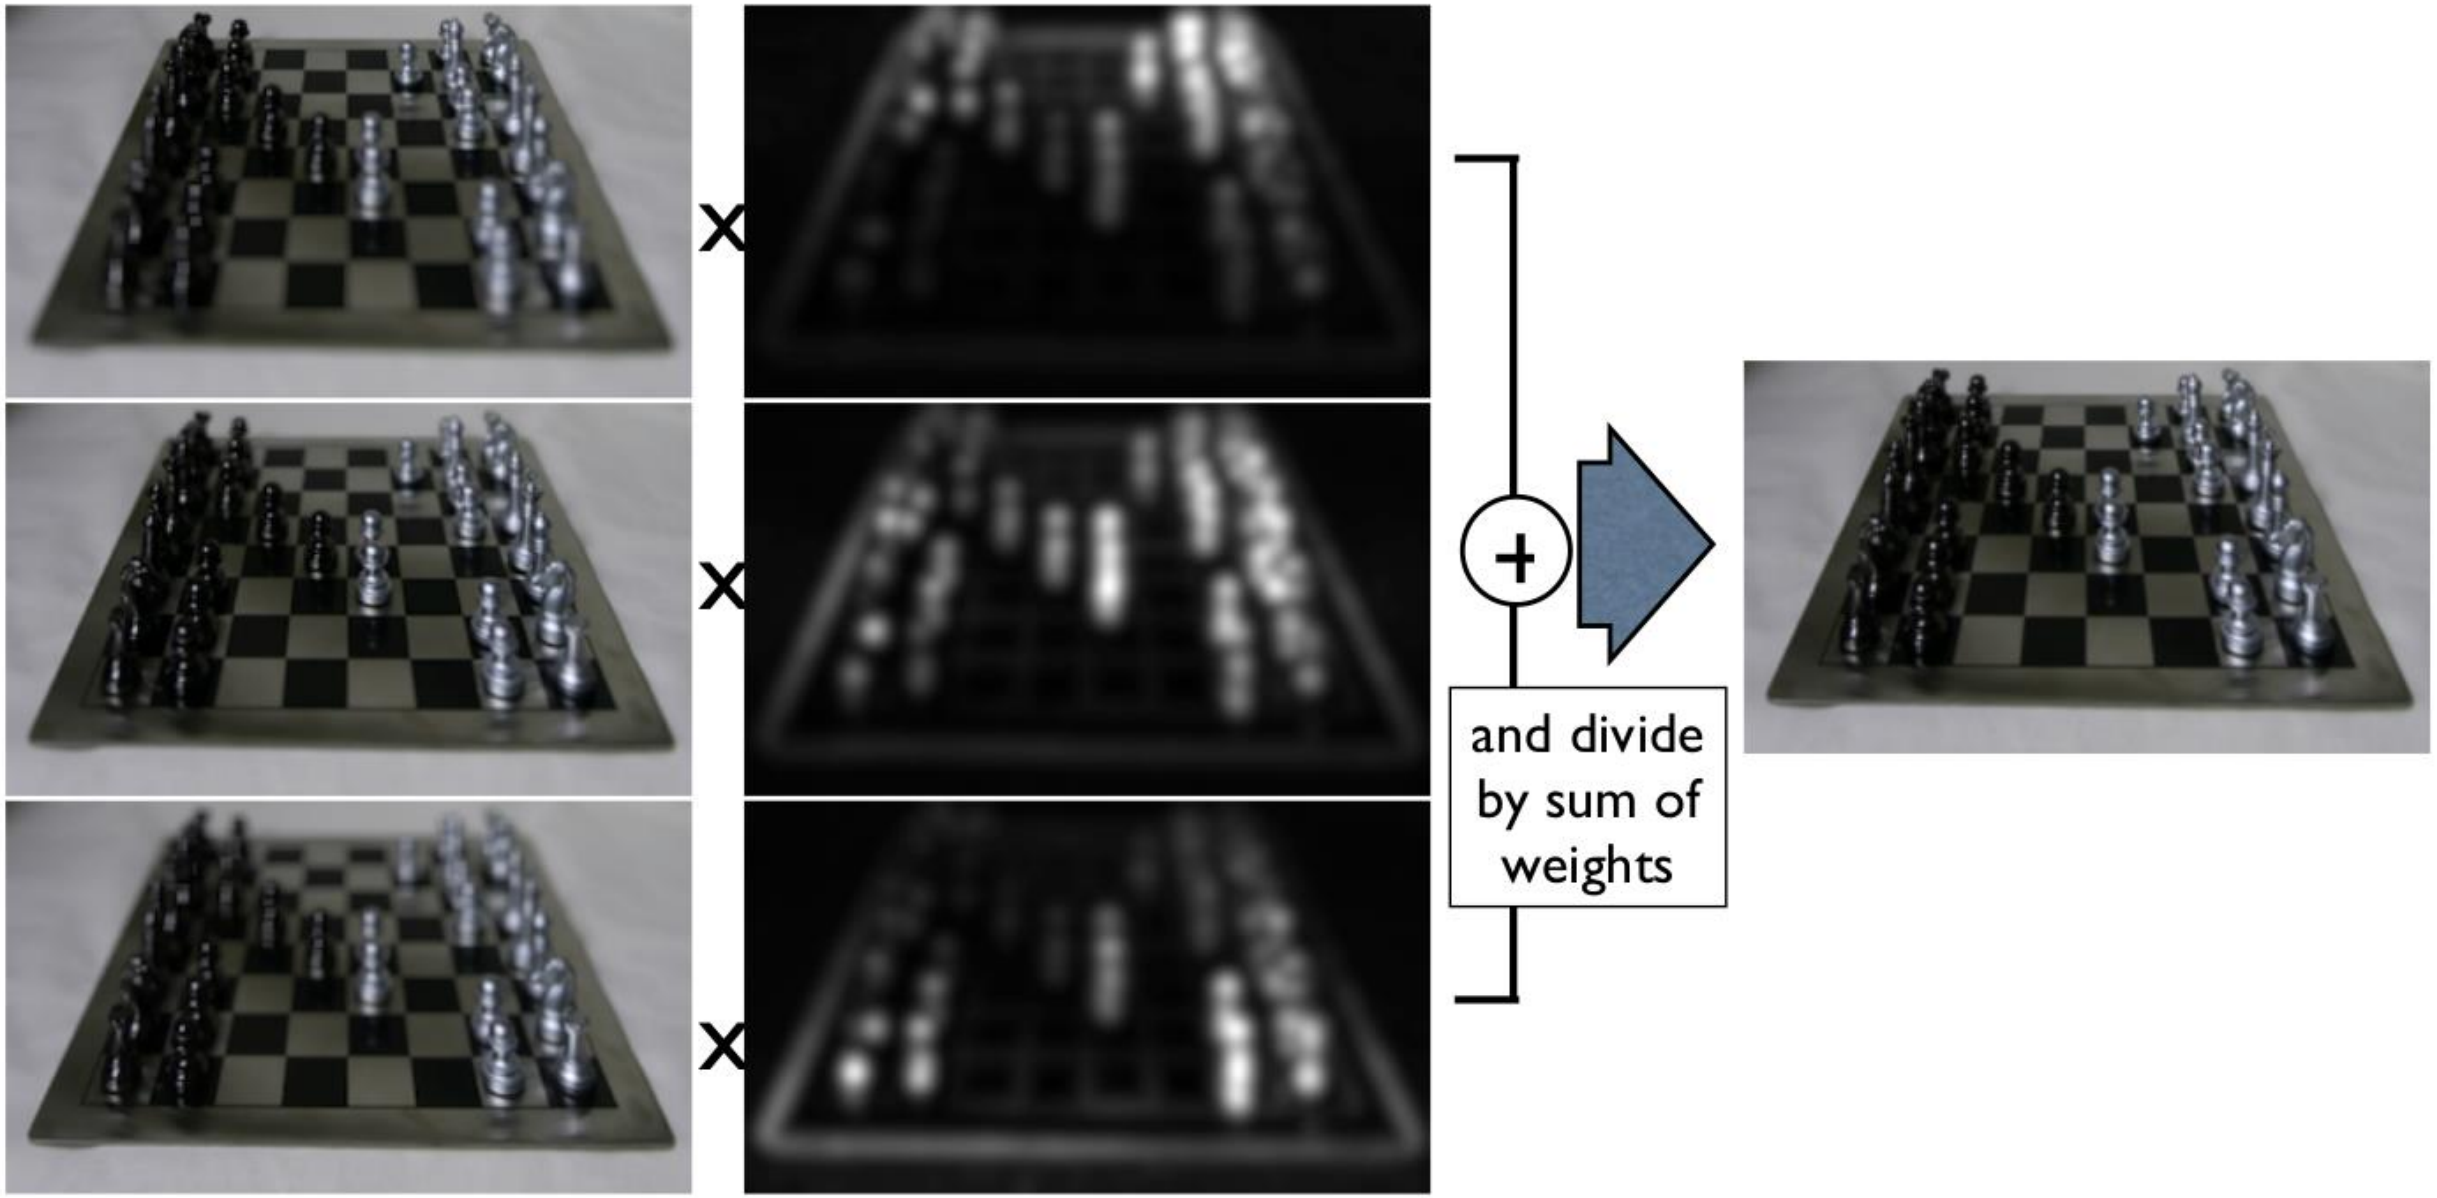

Some results

example image from stack

all in-focus image

Another example

Focal stacking is very useful in macrophotography, where depths of field are very shallow

Another example

middle image from stack

all in-focus image

Another look at the mixing weights

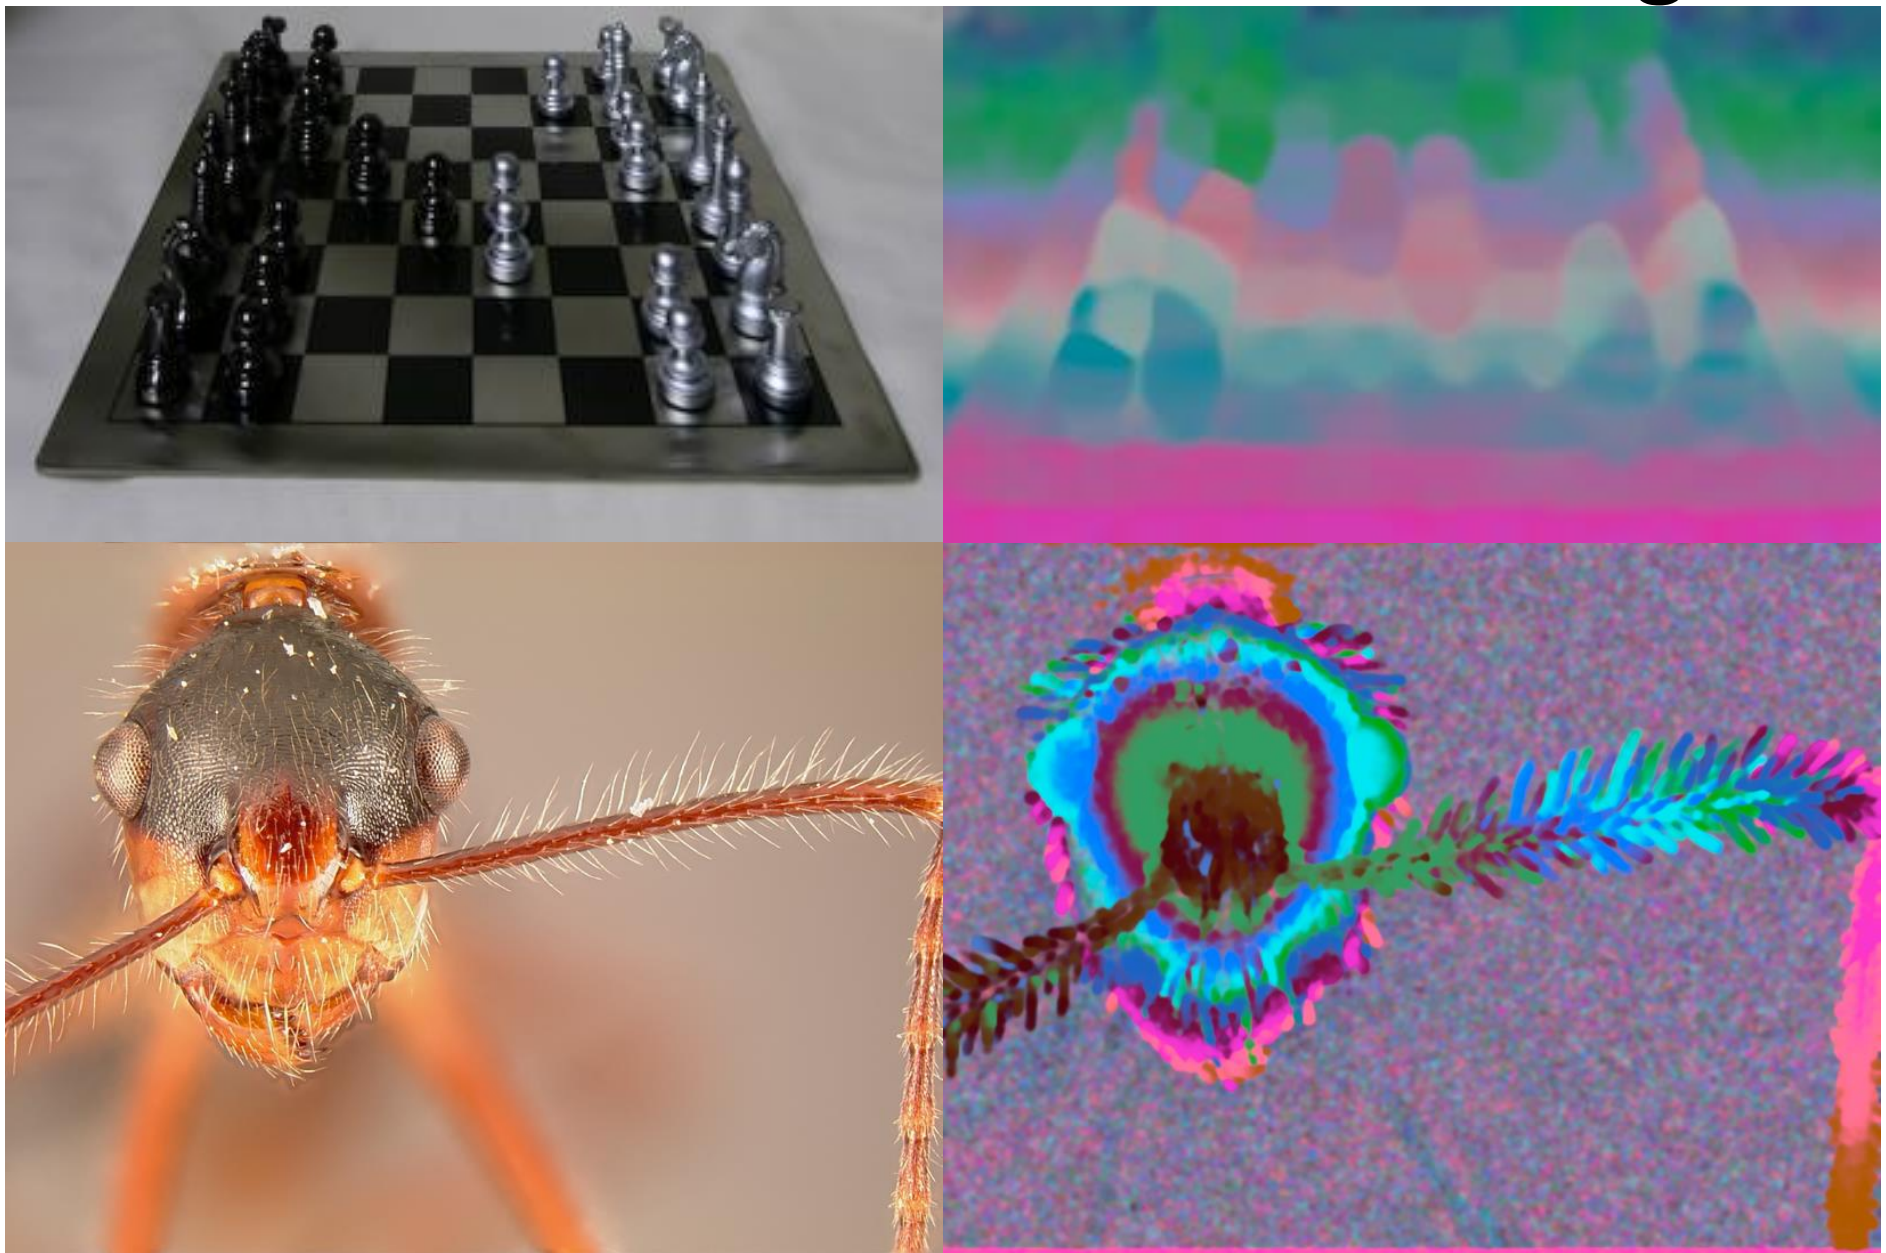

What do the mixing weights look like?

Another look at the mixing weights

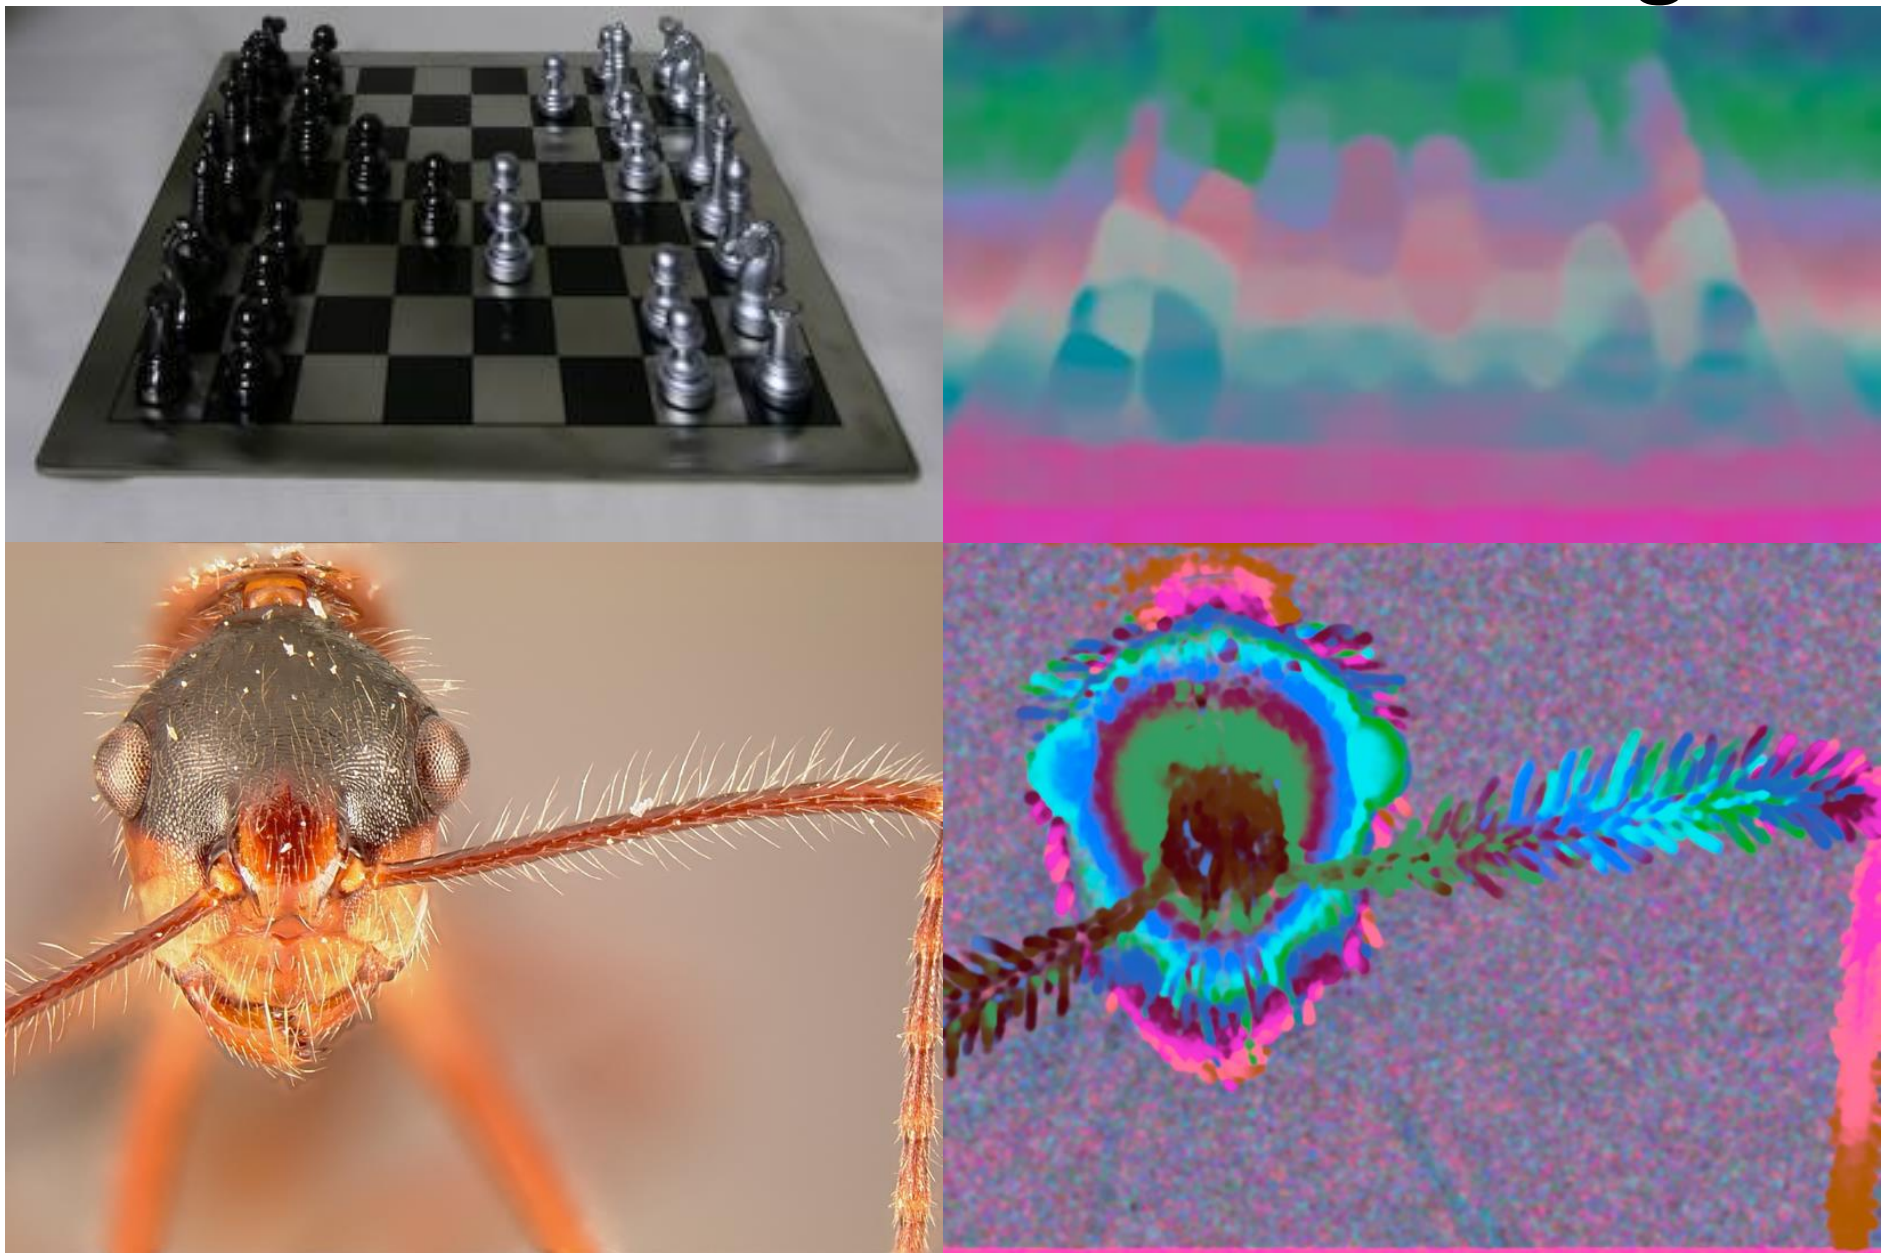

Depth from defocus = determining sharpest pixel in focal stack

Depth from defocus on a mobile phone

Use focal stack from autofocus

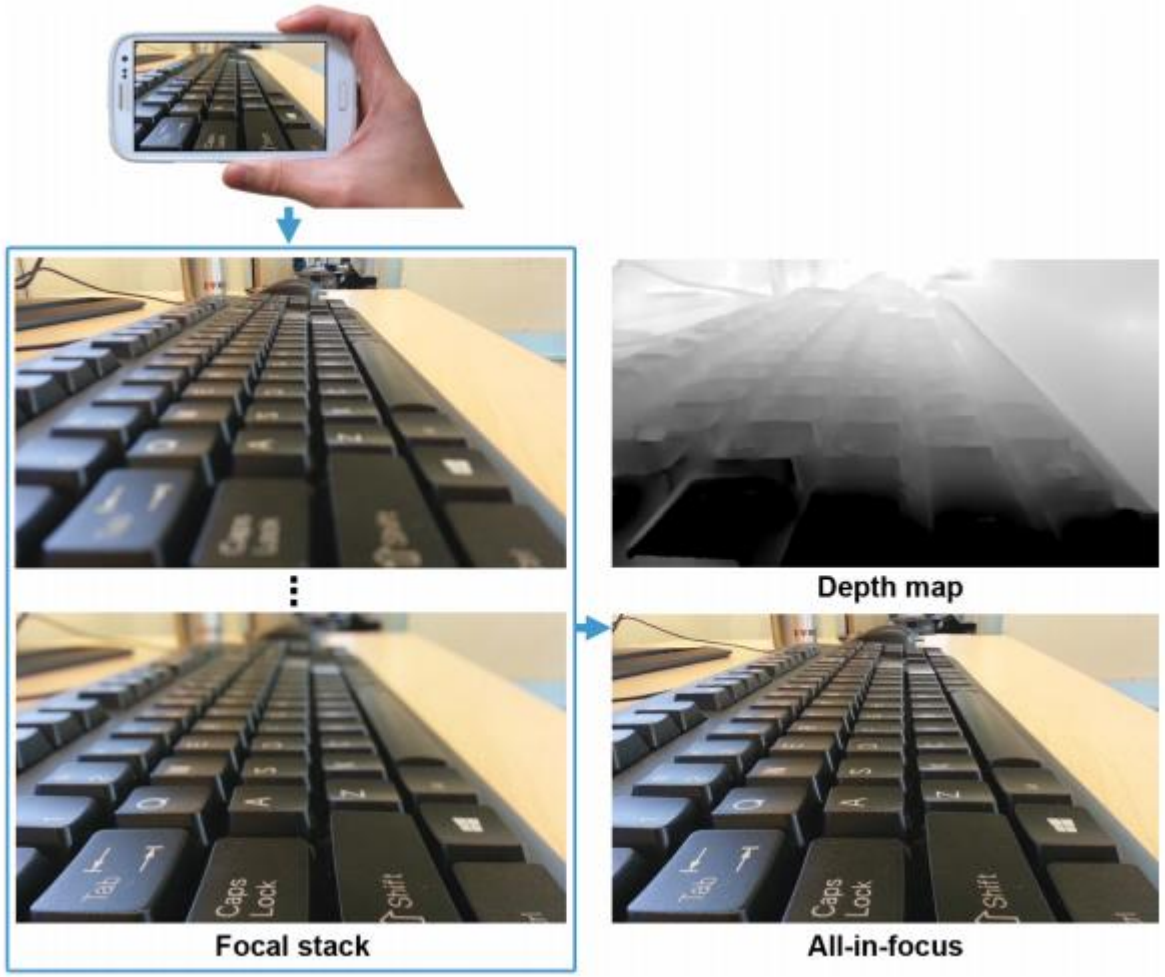

[Suwajanakorn et al., CVPR 2015]

Can we use both focus and aperture?

Confocal stereo

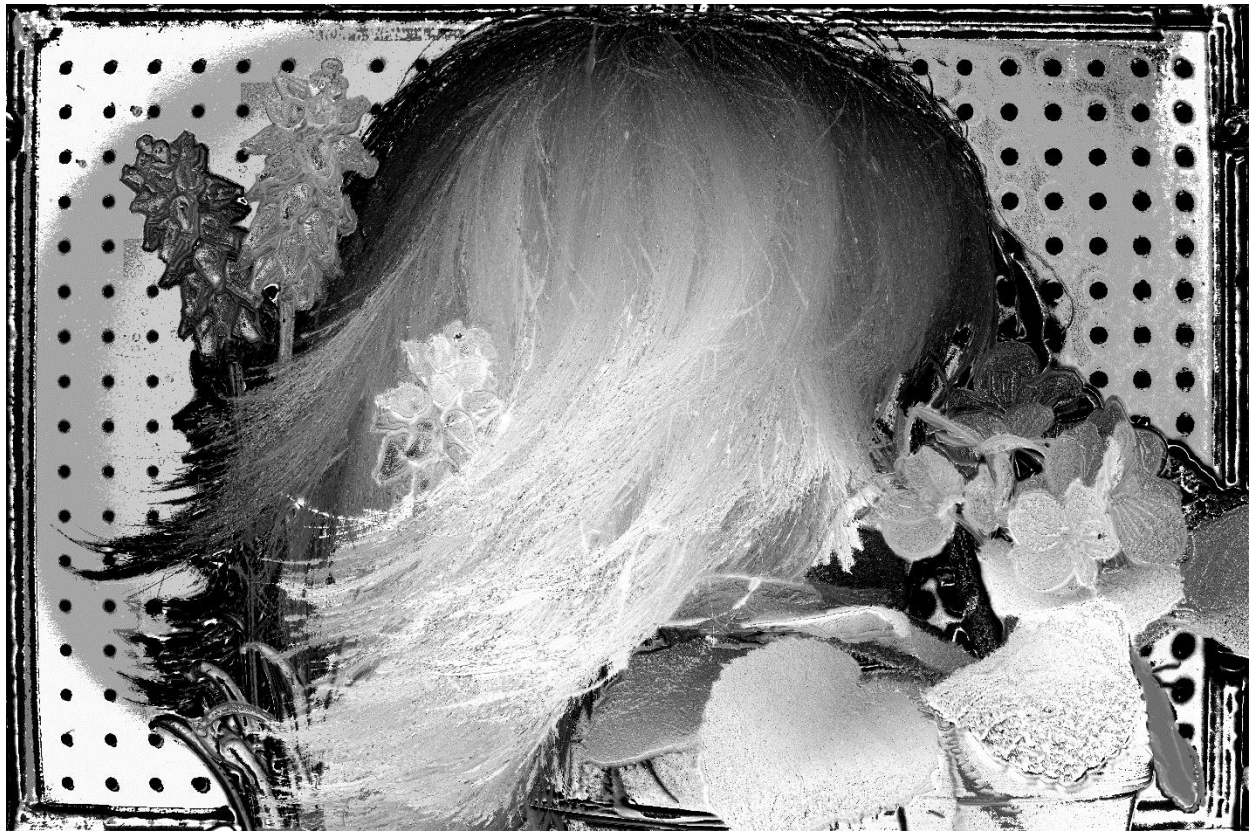

[Hassinof and Kutulakos, ECCV 2006]

Lightfield

Measuring rays

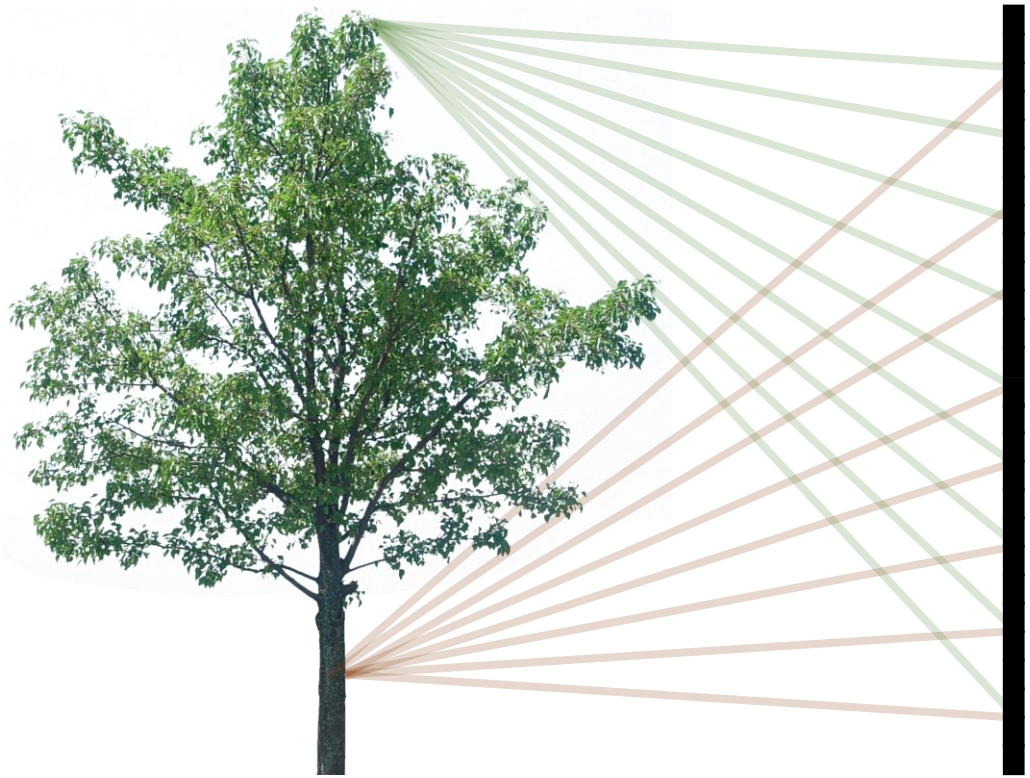

We can capture the same set of rays by using a pinhole camera from multiple viewpoints

Measuring rays

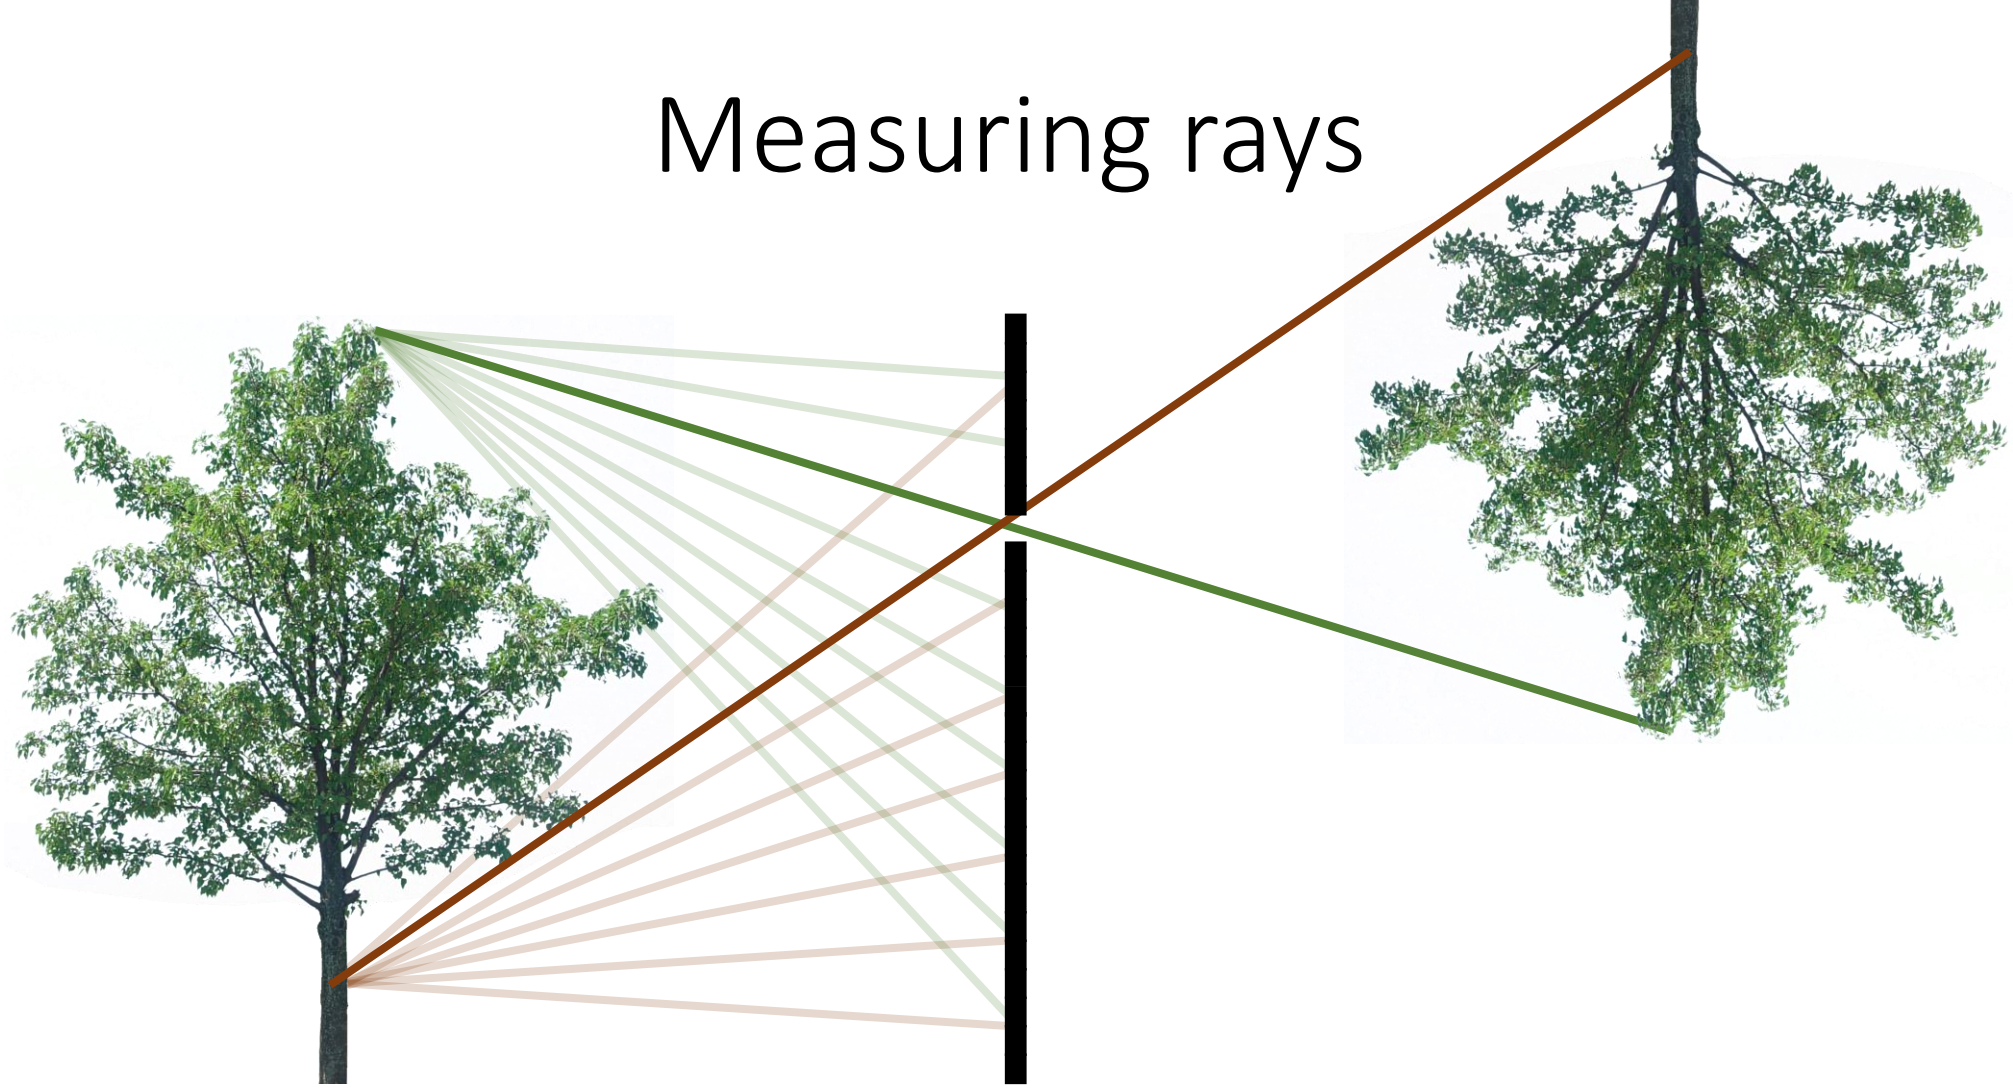

We can capture the same set of rays by using a pinhole camera from multiple viewpoints

Measuring rays

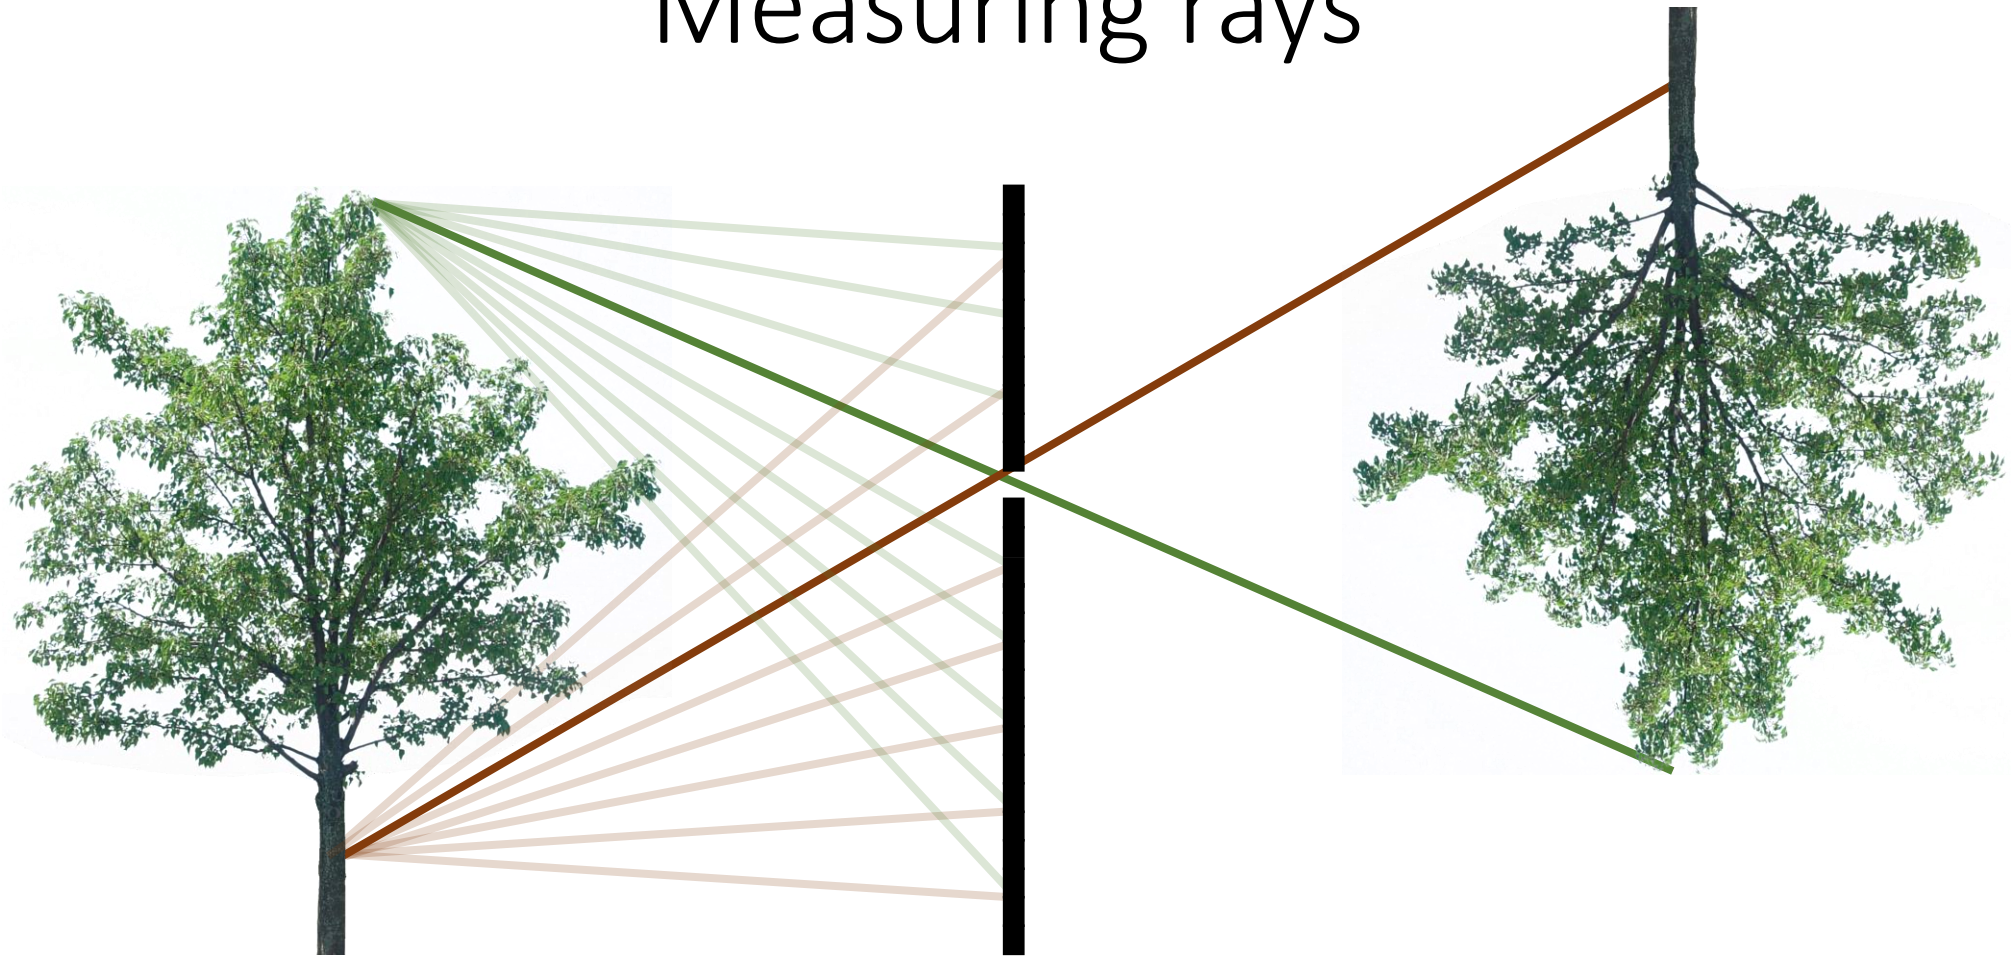

We can capture the same set of rays by using a pinhole camera from multiple viewpoints

Measuring rays

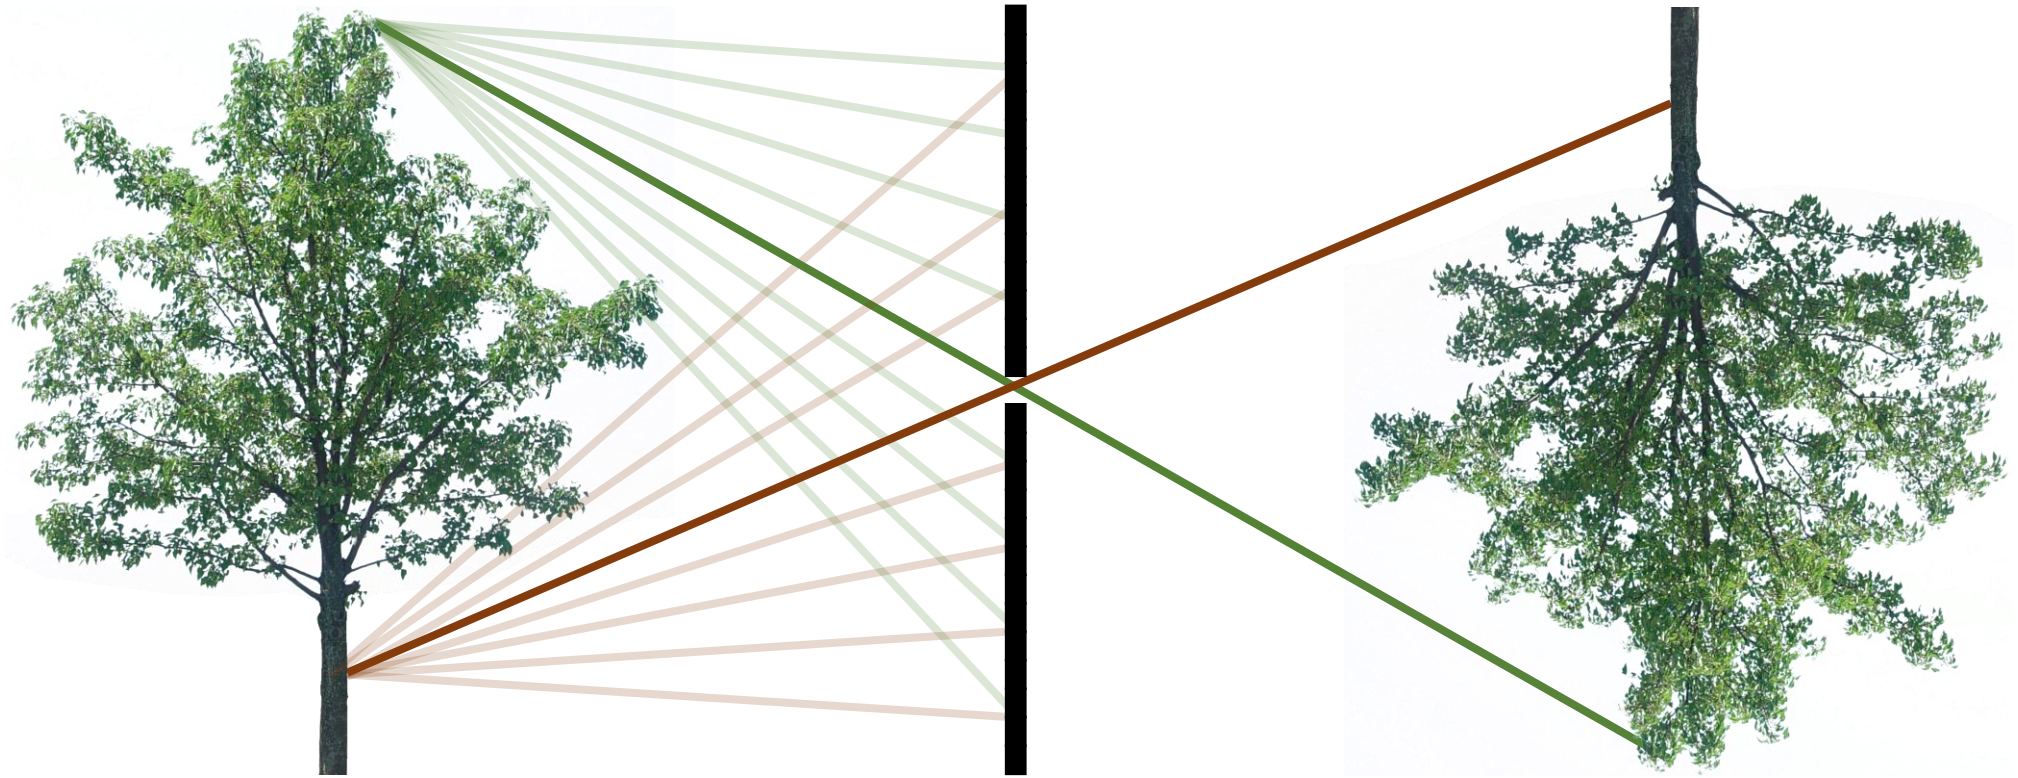

We can capture the same set of rays by using a pinhole camera from multiple viewpoints

Measuring rays

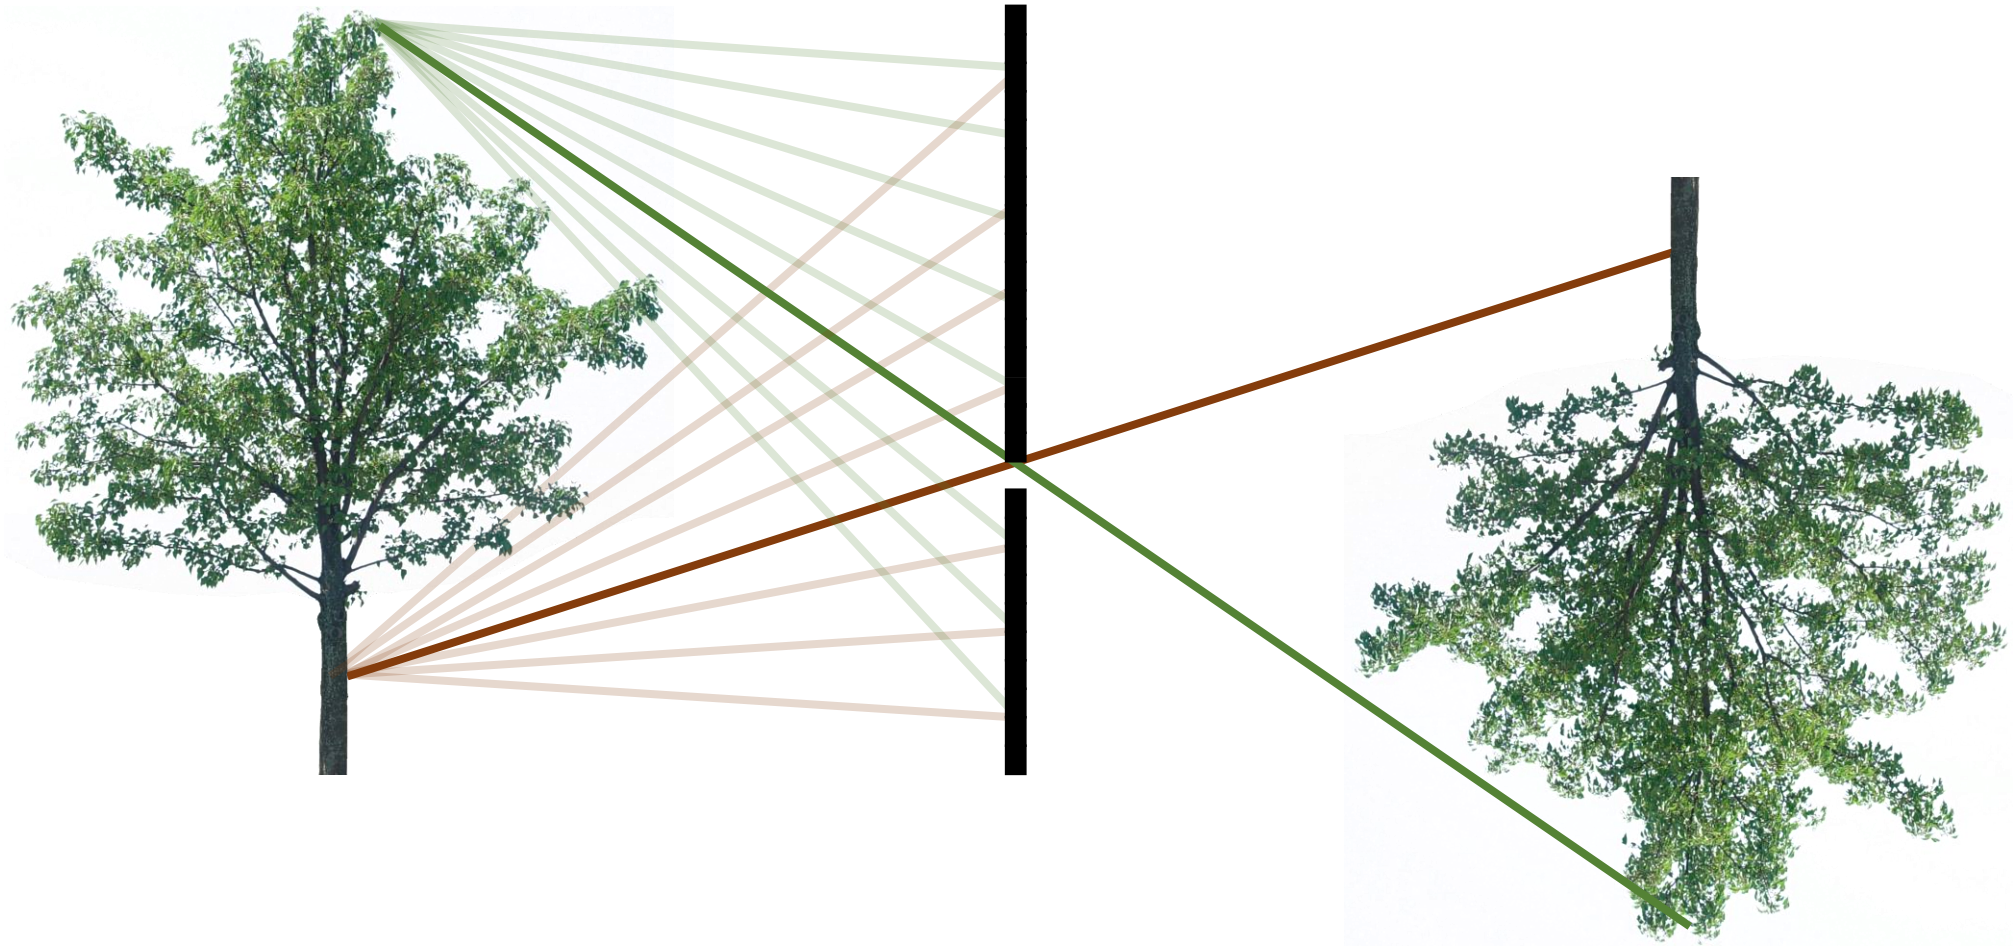

We can capture the same set of rays by using a pinhole camera from multiple viewpoints

Measuring rays

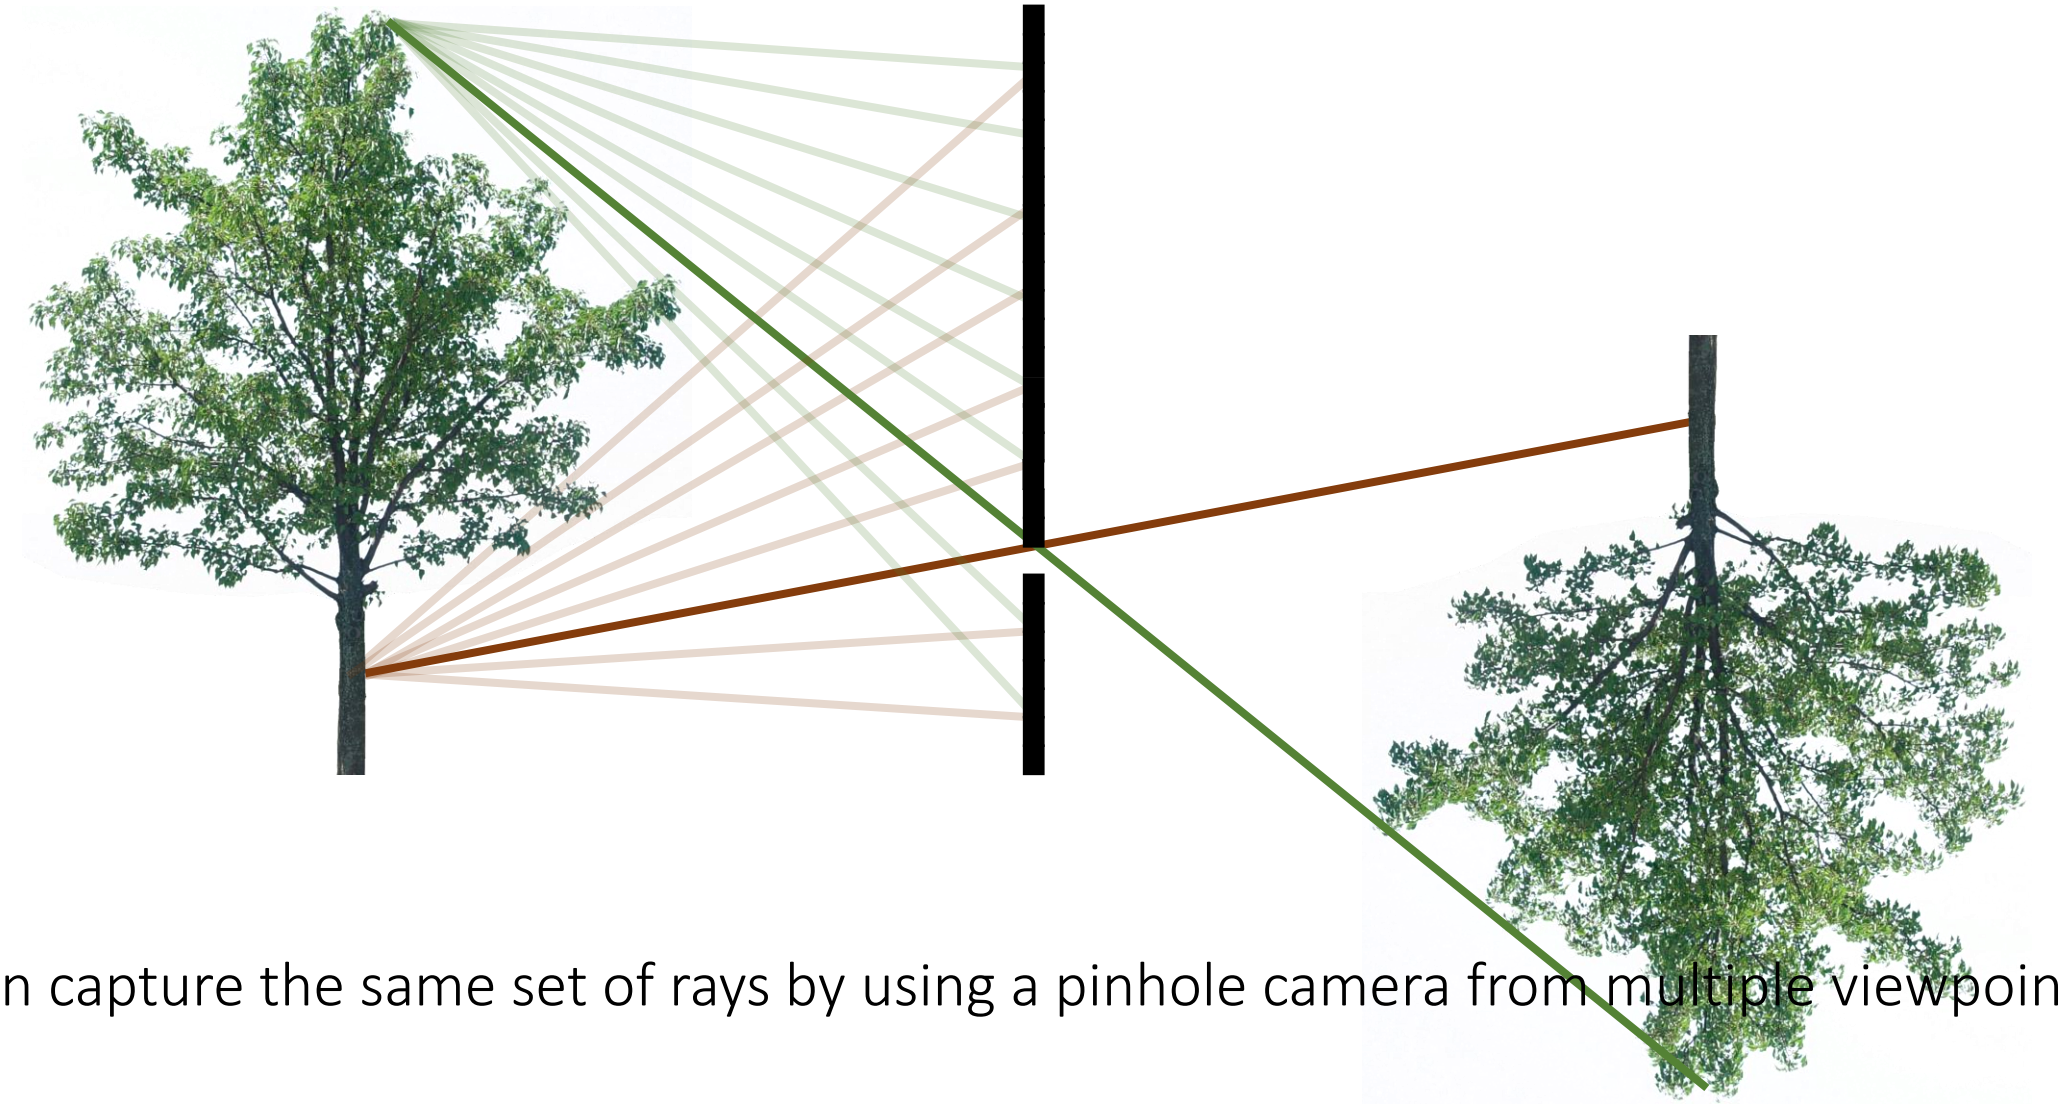

We can capture the same set of rays by using a pinhole camera from multiple viewpoints

Measuring rays

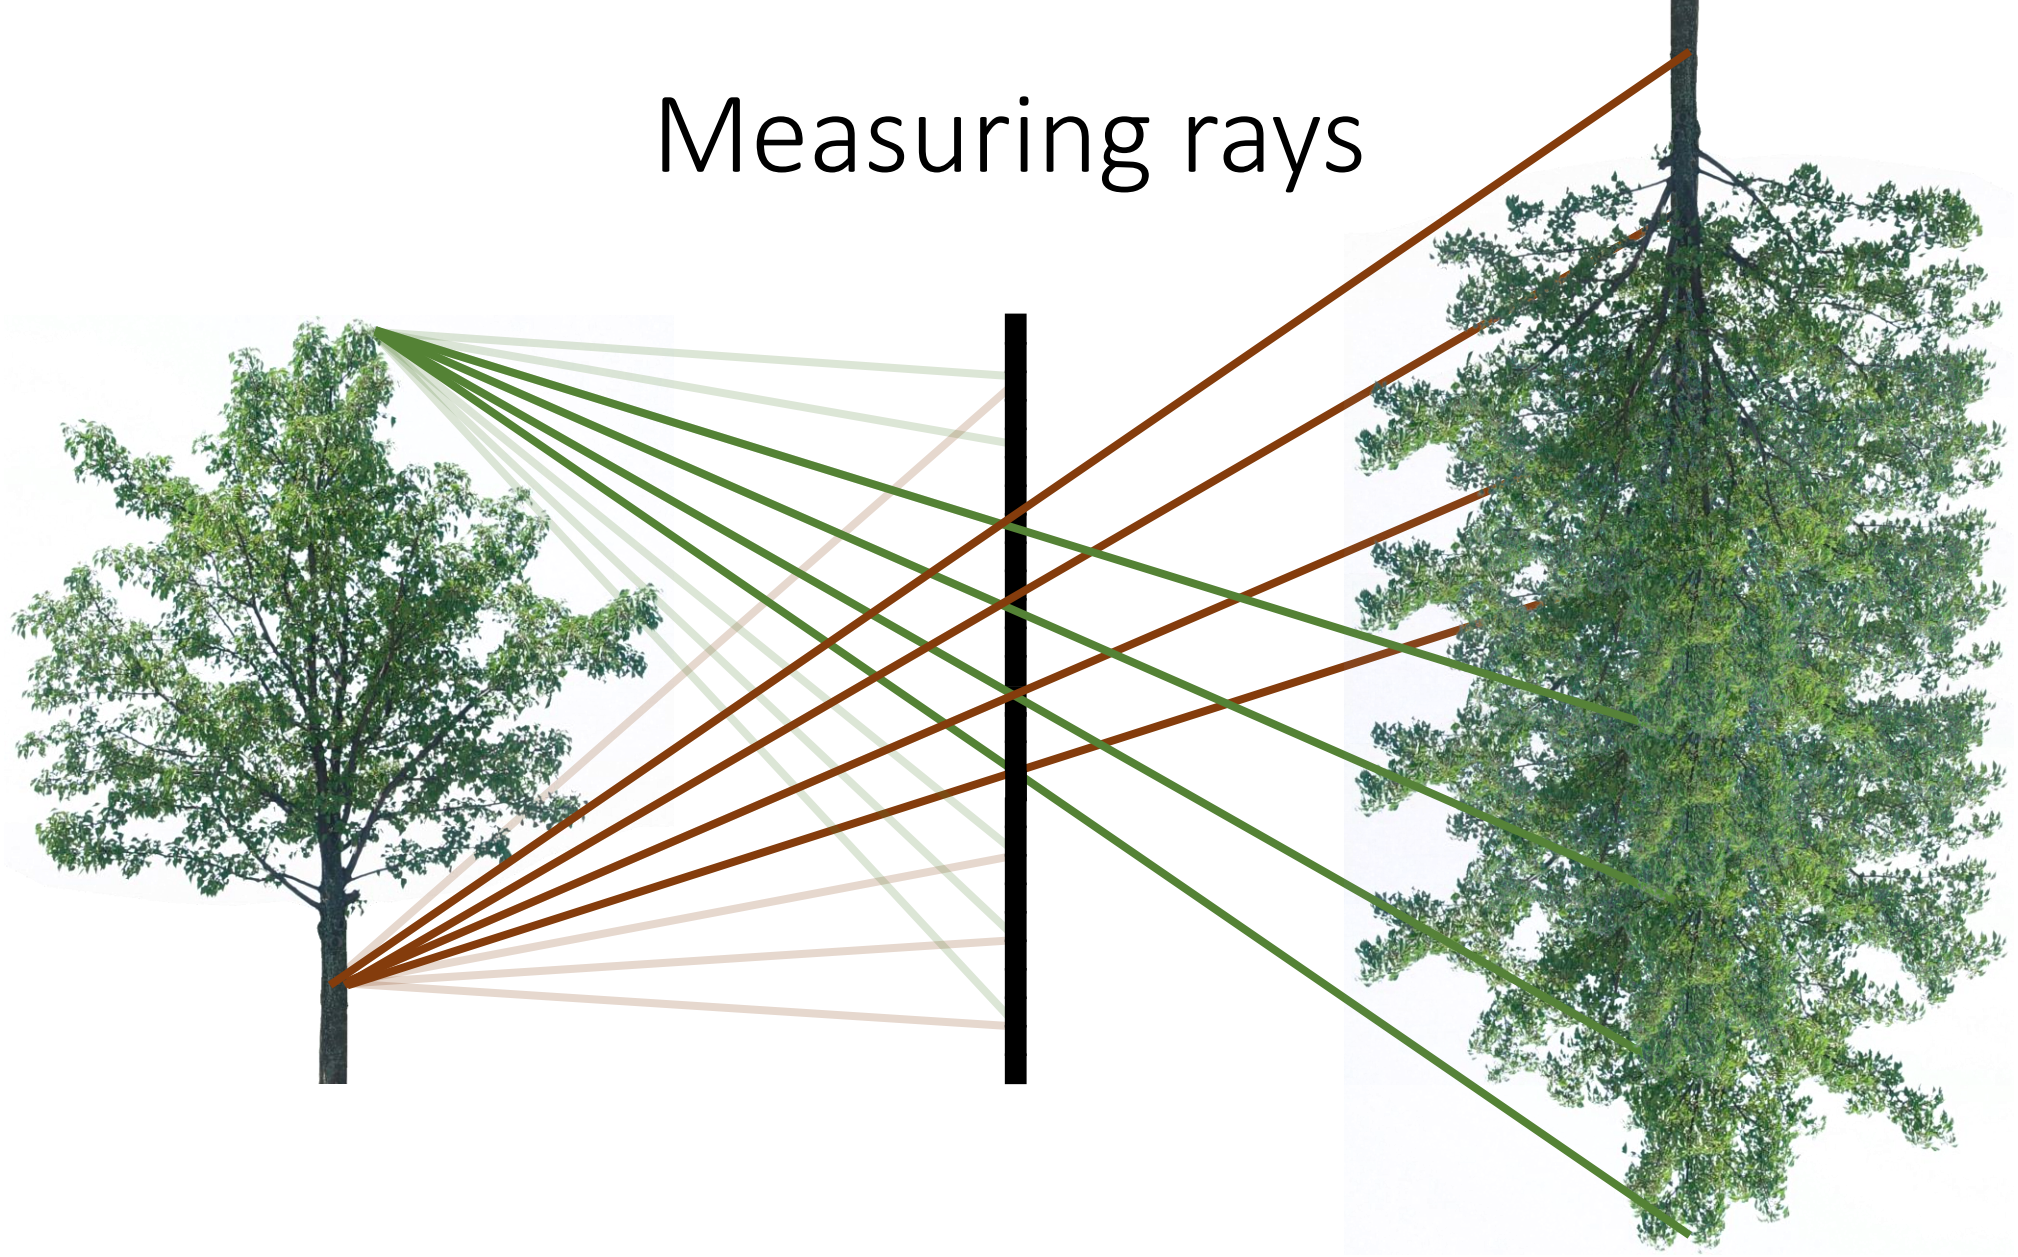

We can capture the same set of rays by using a pinhole camera from multiple viewpoints

- How would you merge these images into a lens, defocused image?

Lightfield: all rays in a scene

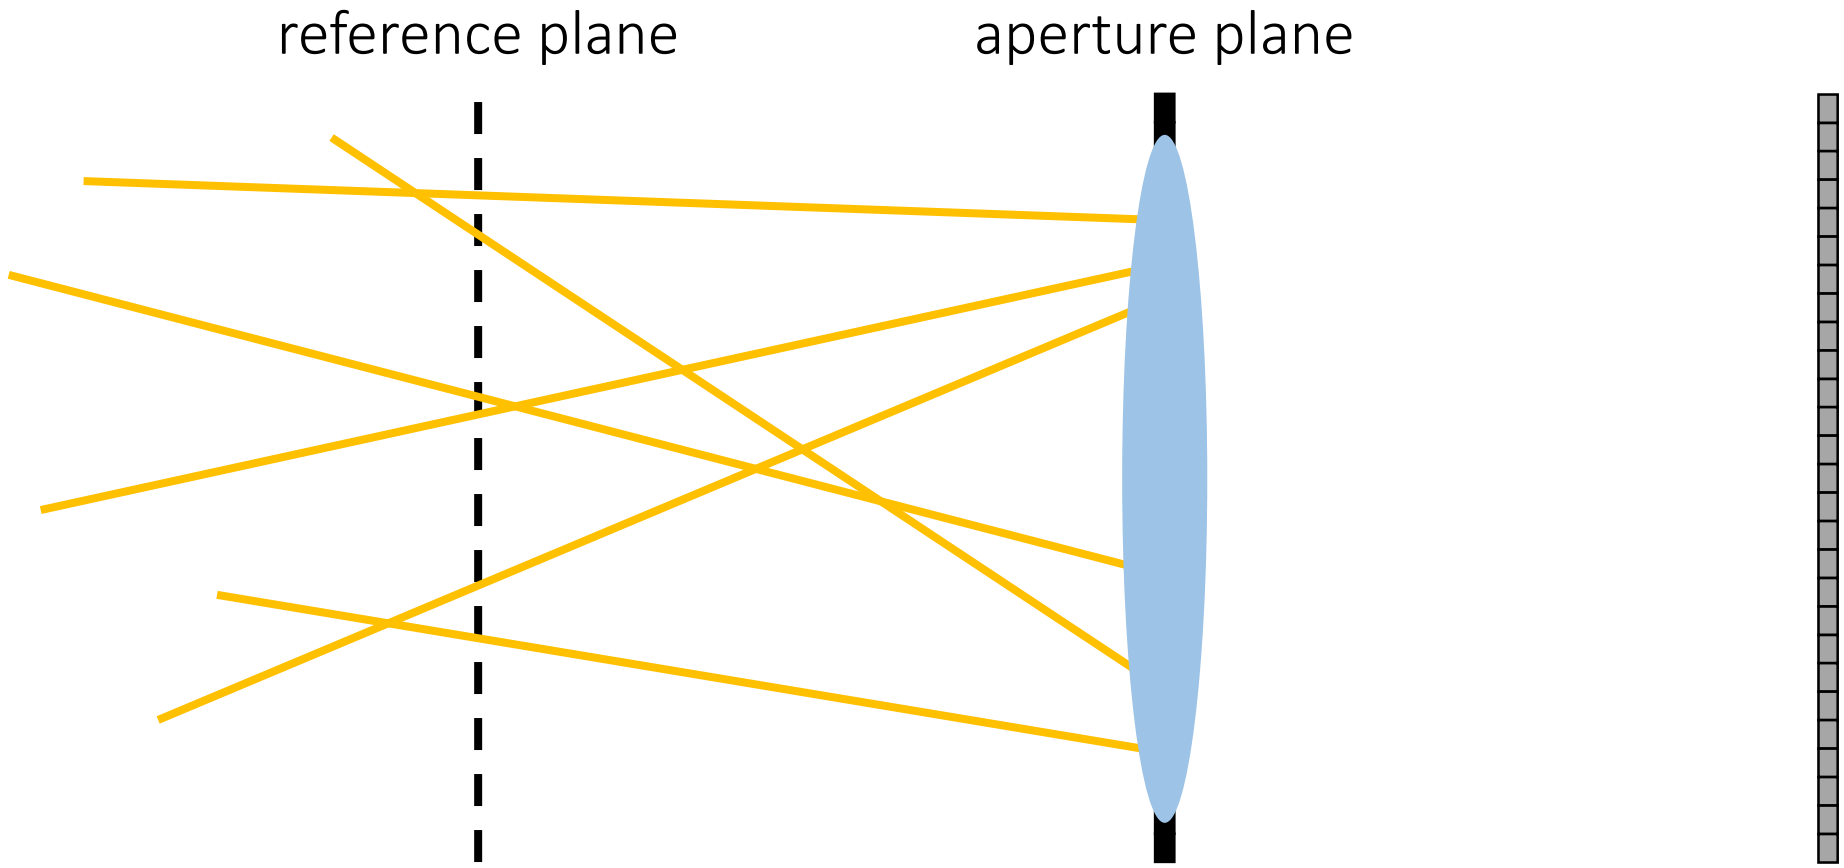

What is the dimension of the lightfield?

Parameterize every ray based on its intersections with two planes.

Lightfield: all rays in a scene

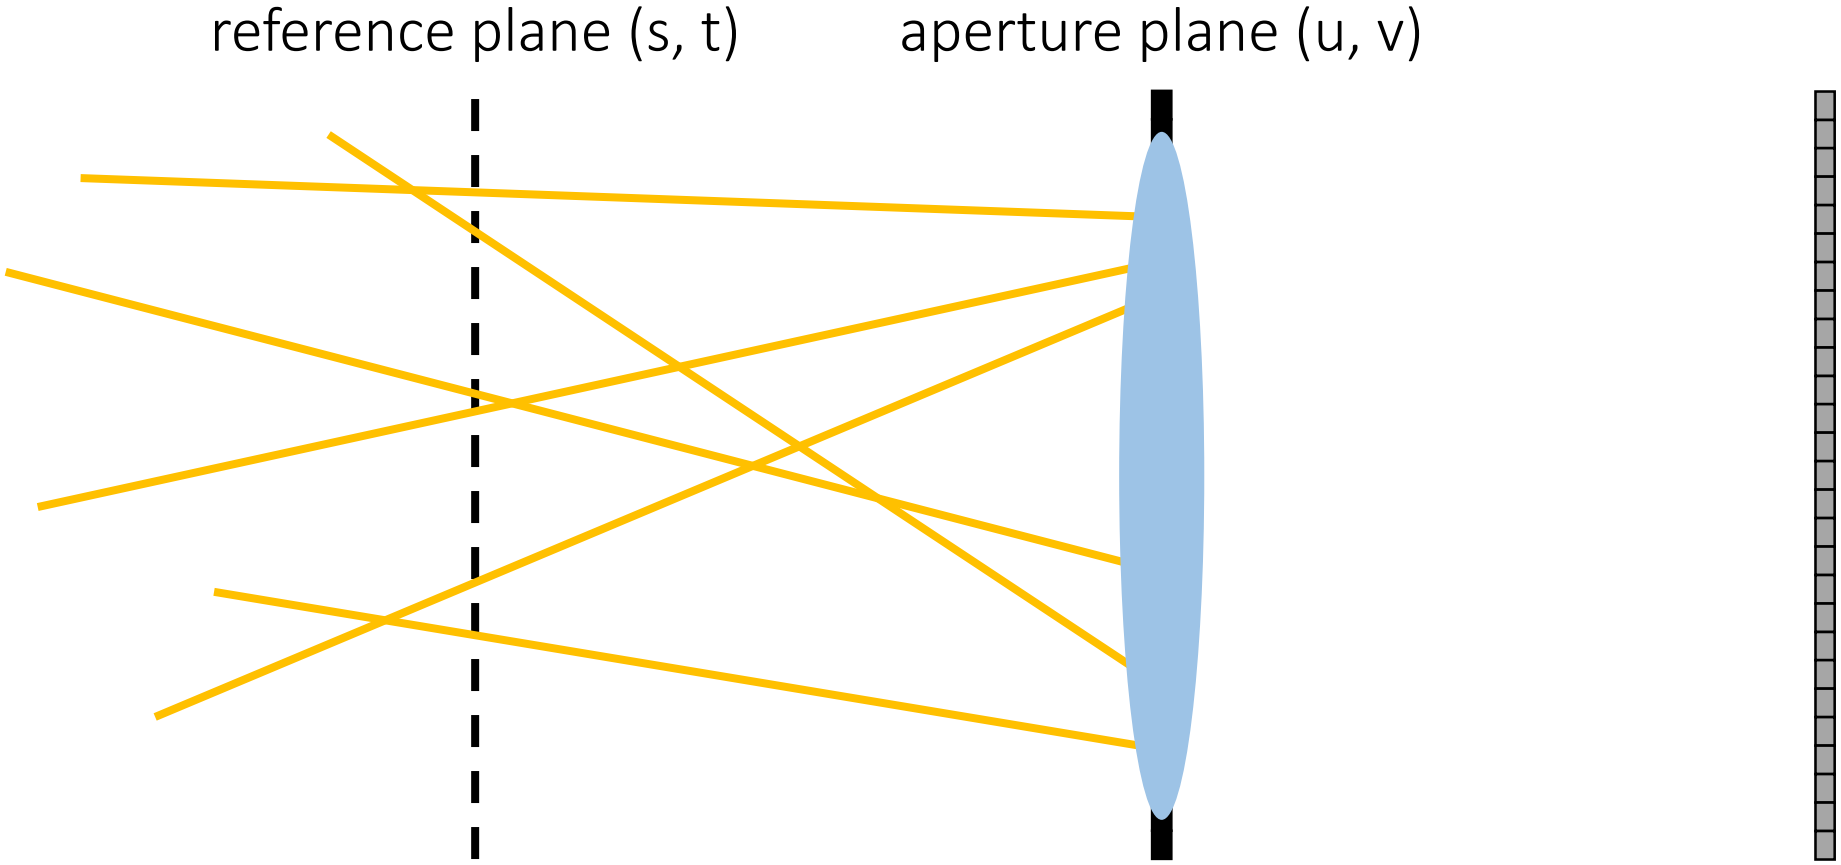

4-dimensional function L(u, v, s, t)

Parameterize every ray based on its intersections with two planes.

Lightfield: all rays in a scene

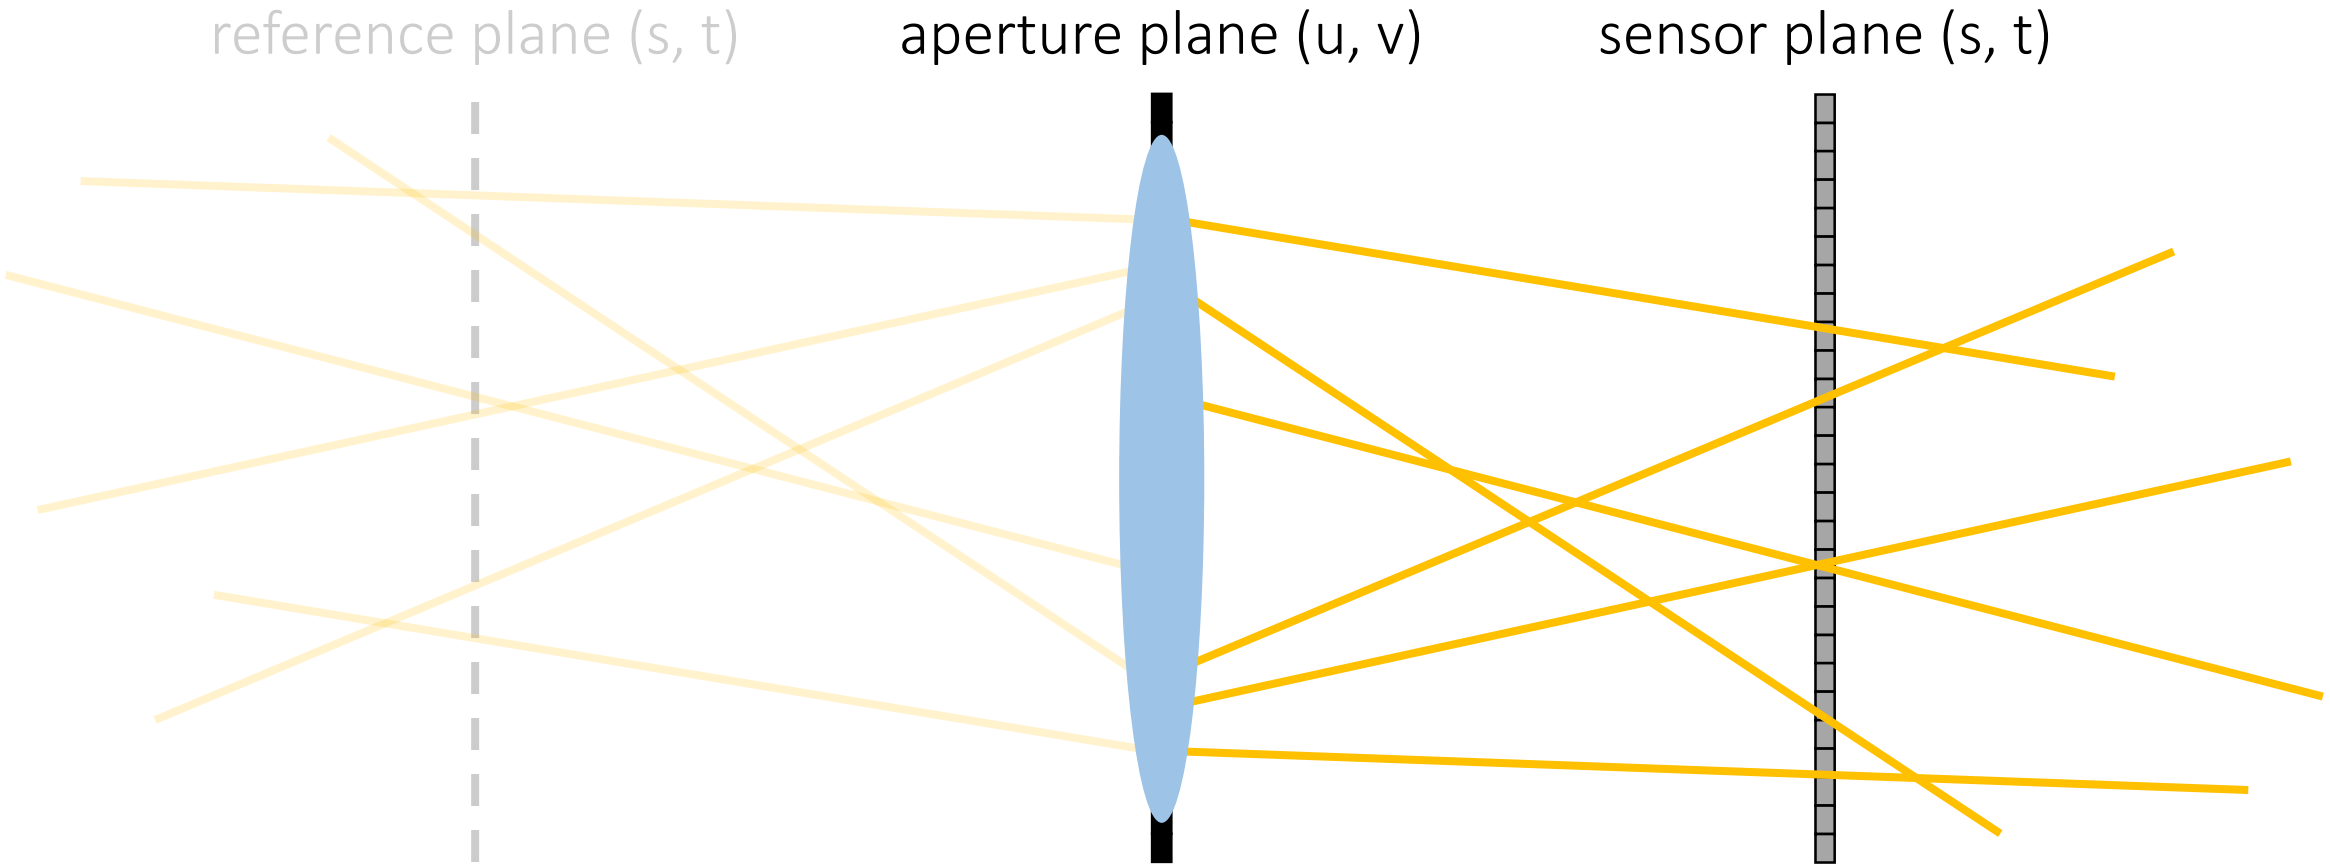

4-dimensional function L(u, v, s, t)
(conjugate of scene-based function)

Parameterize every ray based on its intersections with two planes.

Lightfield slices

What does L(u = u_o, v = v_o, s, t) look like?

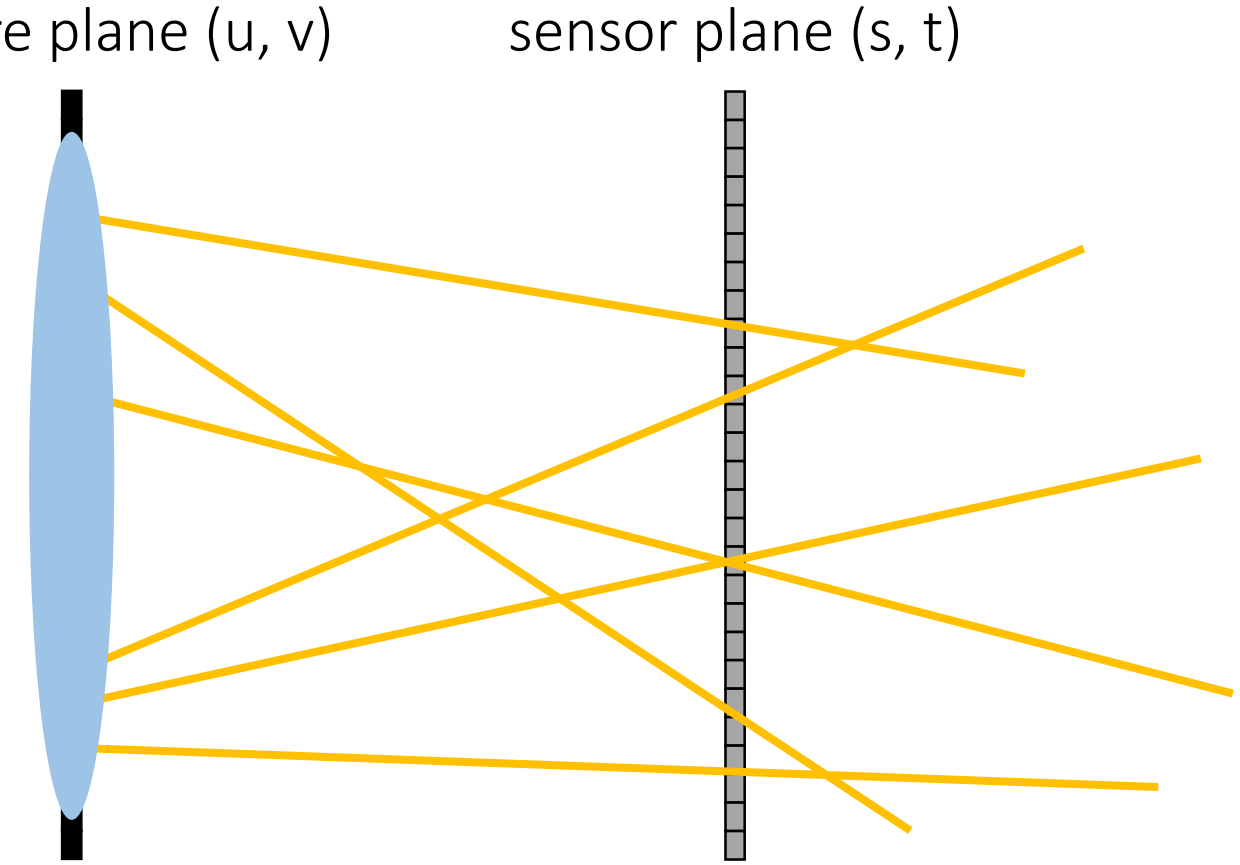

4-dimensional function L(u, v, s, t)
(conjugate of scene-based function)

Lightfield slices

What does L(u = u_o, v = v_o, s, t) look like?

- a pinhole image from a certain viewpoint

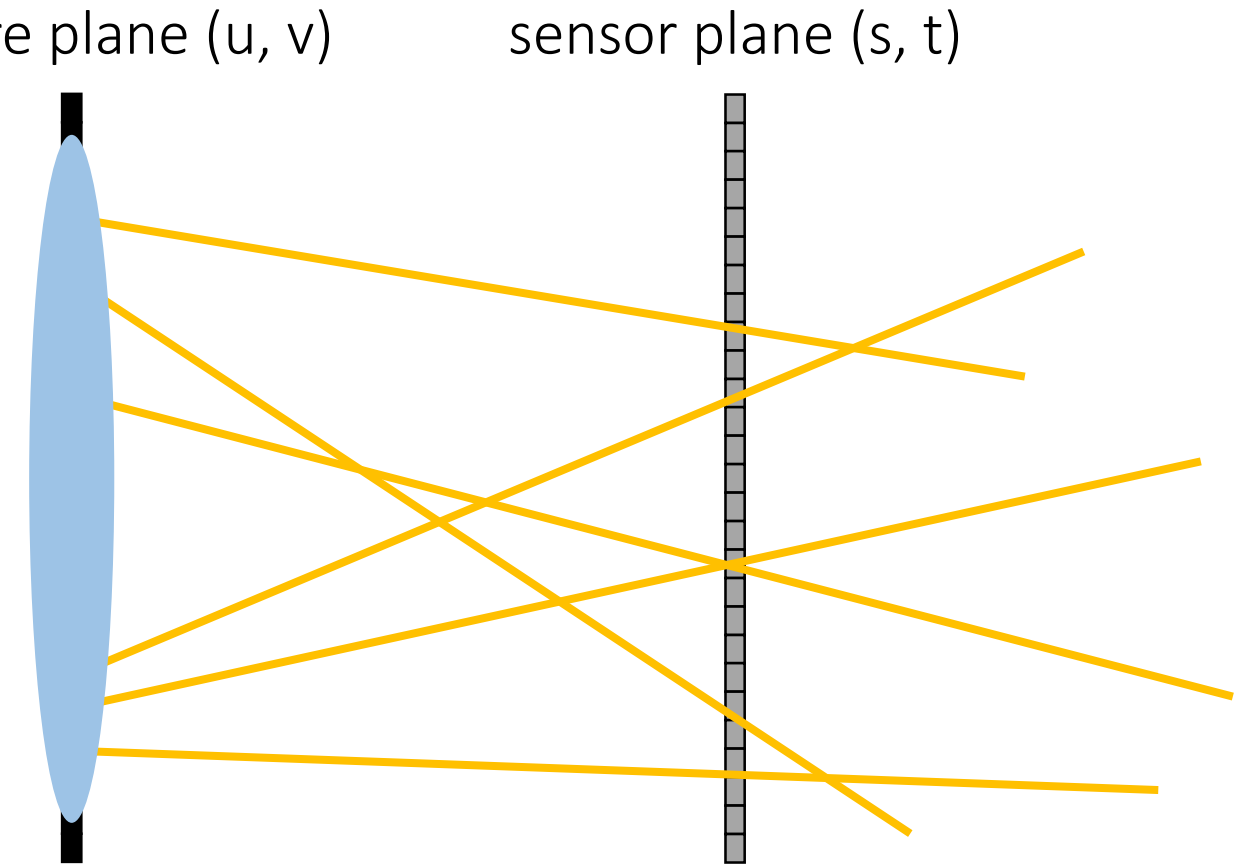

4-dimensional function L(u, v, s, t)
(conjugate of scene-based function)

Lightfield slices

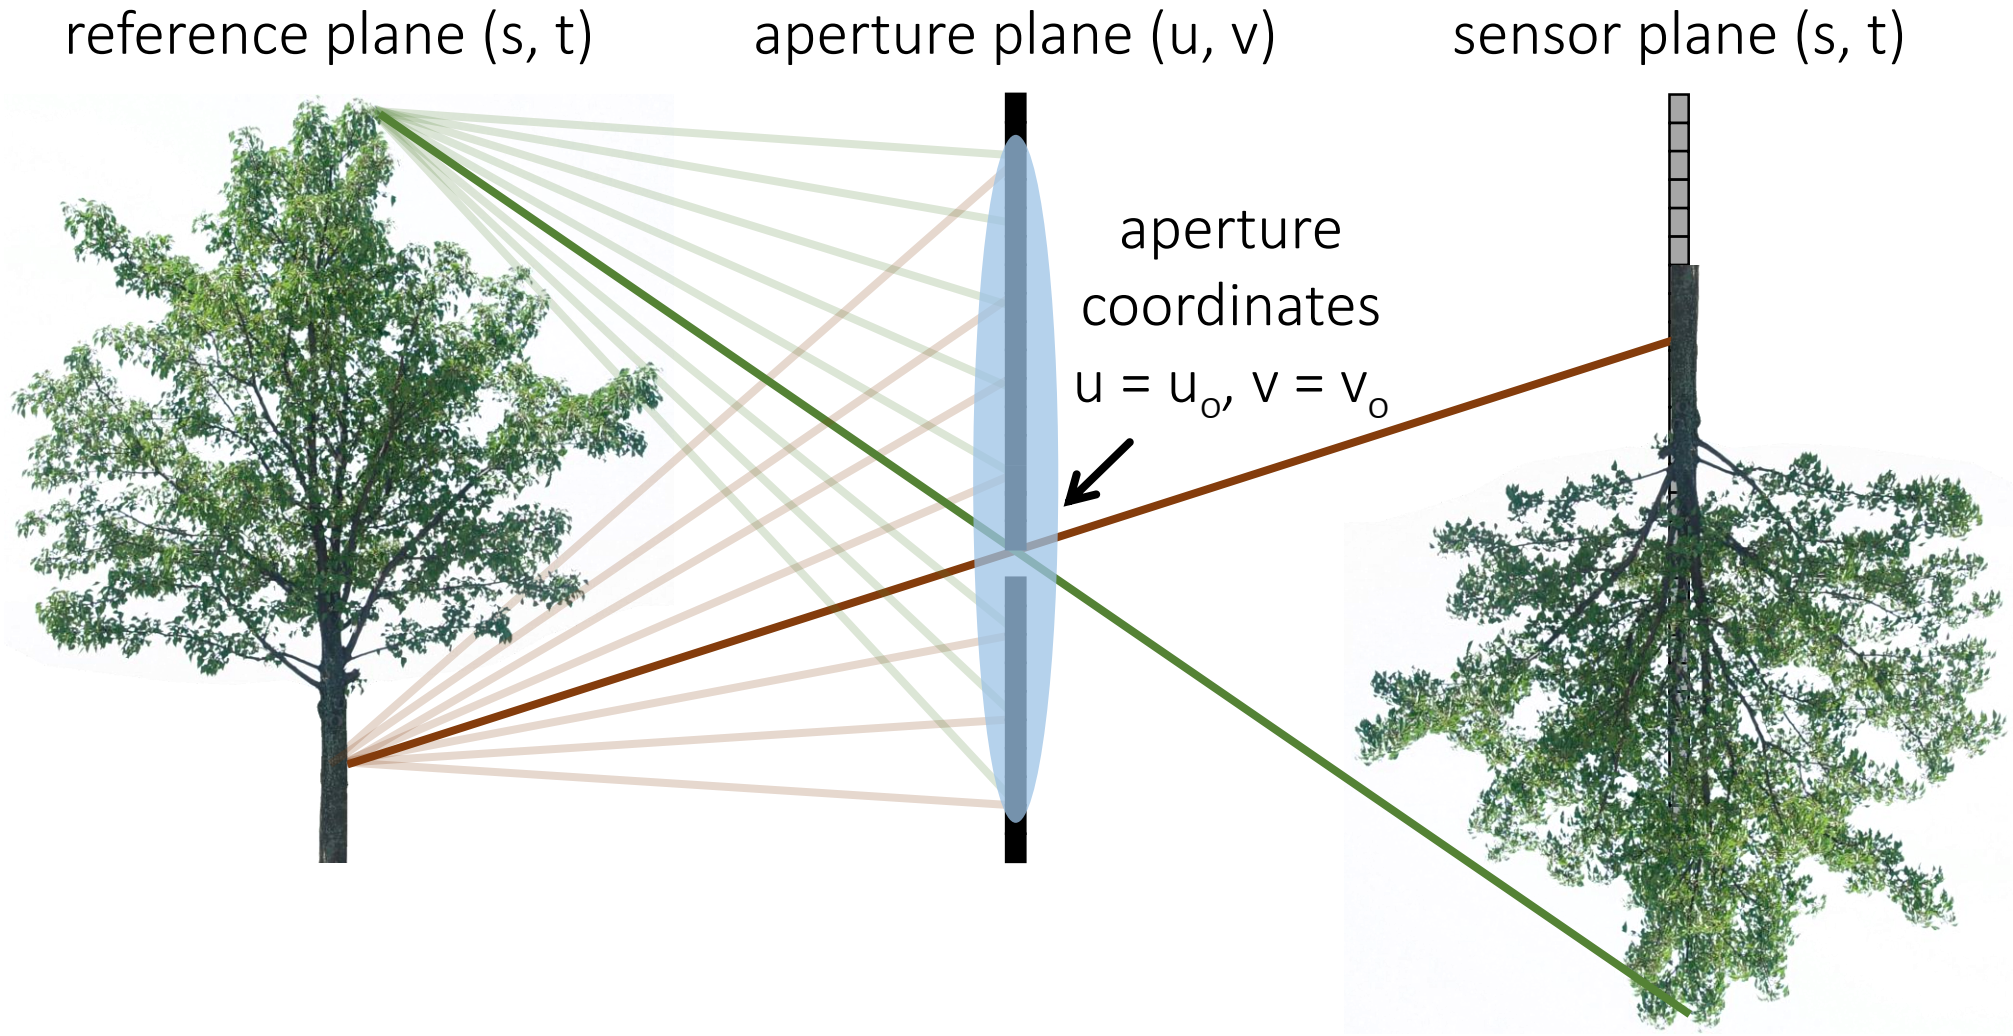

Lightfield slice $L(u = u_o, v = v_o, s, t)$

Lightfield slices

What does L(u = u_o, v = v_o, s, t) look like?

- a pinhole image from a certain viewpoint

What does L(u, v, s = s_o, t = t_o) look like?

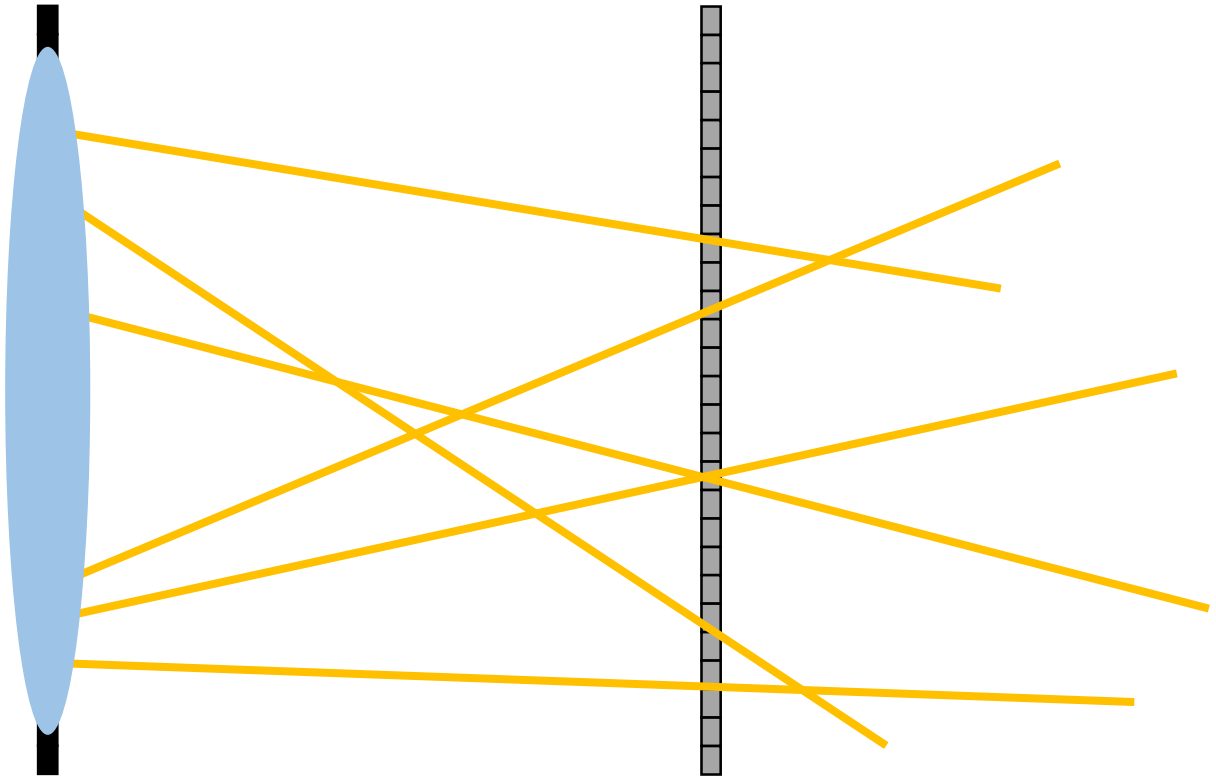

4-dimensional function L(u, v, s, t)
(conjugate of scene-based function)

Lightfield slices

What does L(u = u_o, v = v_o, s, t) look like?

- a pinhole image from a certain viewpoint

What does L(u, v, s = s_o, t = t_o) look like?

- radiance emitted by a certain (in-focus) point at various directions

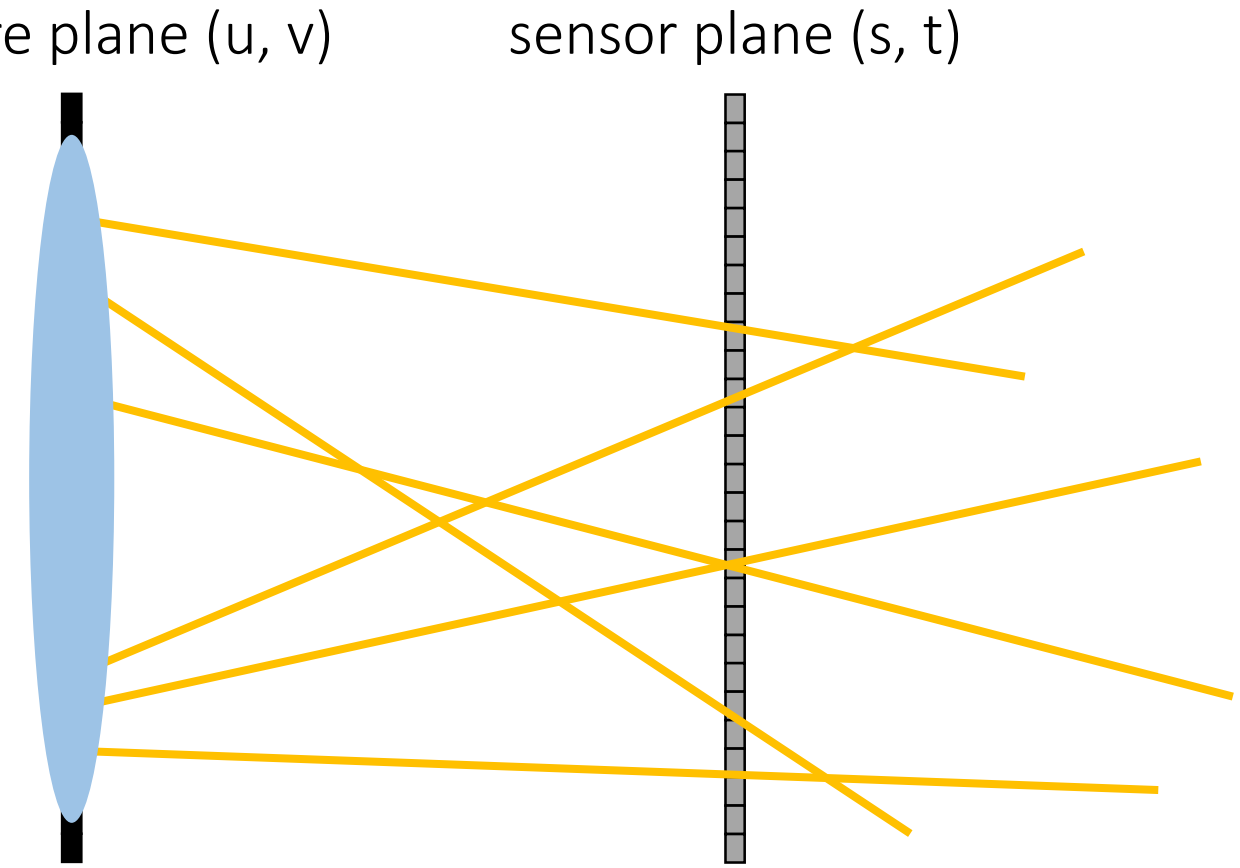

4-dimensional function L(u, v, s, t)
(conjugate of scene-based function)

Lightfield slices

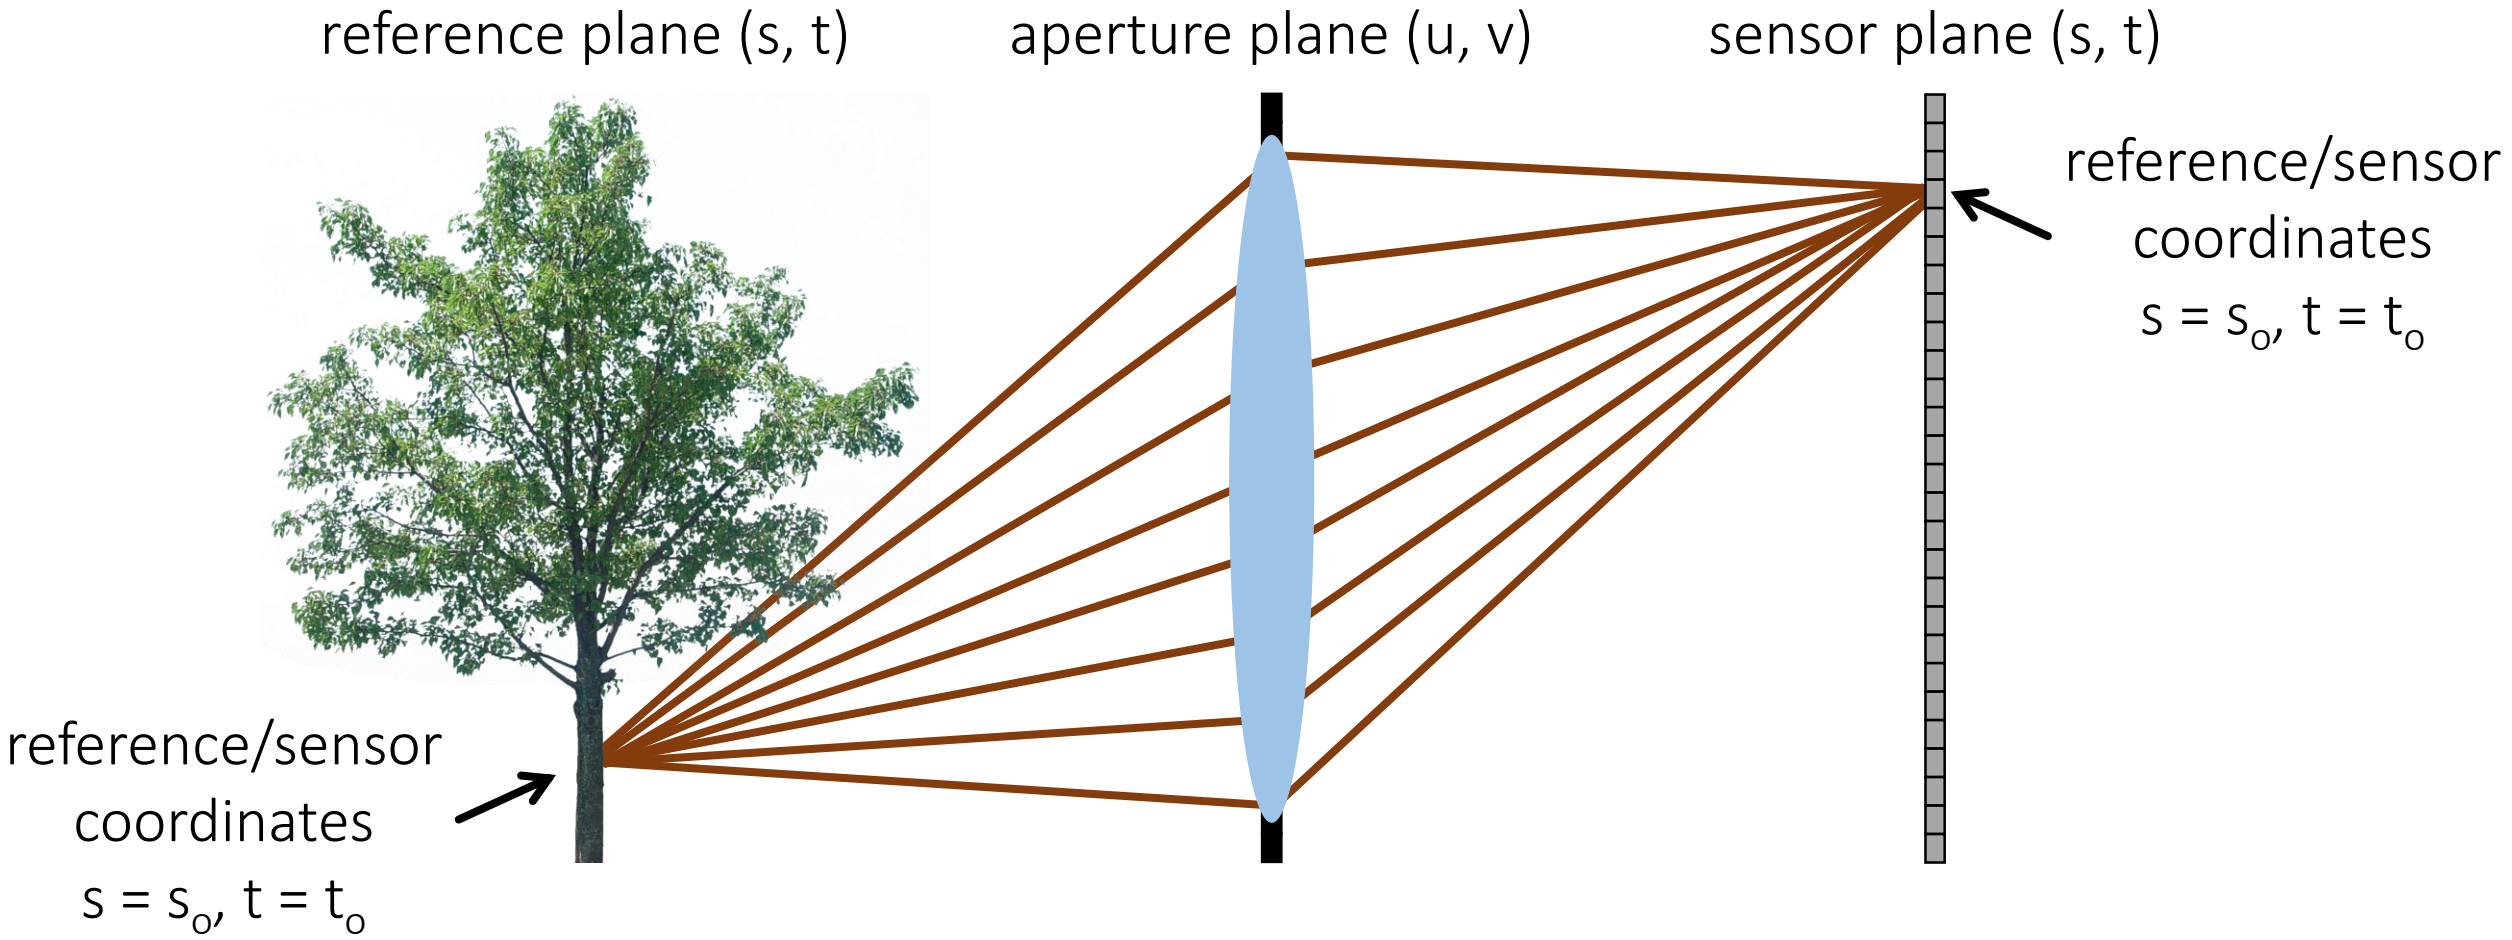

Lightfield slice $L(u, v, s = s_o, t = t_o)$

Lightfield slices

What does L(u = u_o, v = v_o, s, t) look like?

- a pinhole image from a certain viewpoint

What does L(u, v, s = s_o, t = t_o) look like?

- radiance emitted by a certain (in-focus) point in various directions

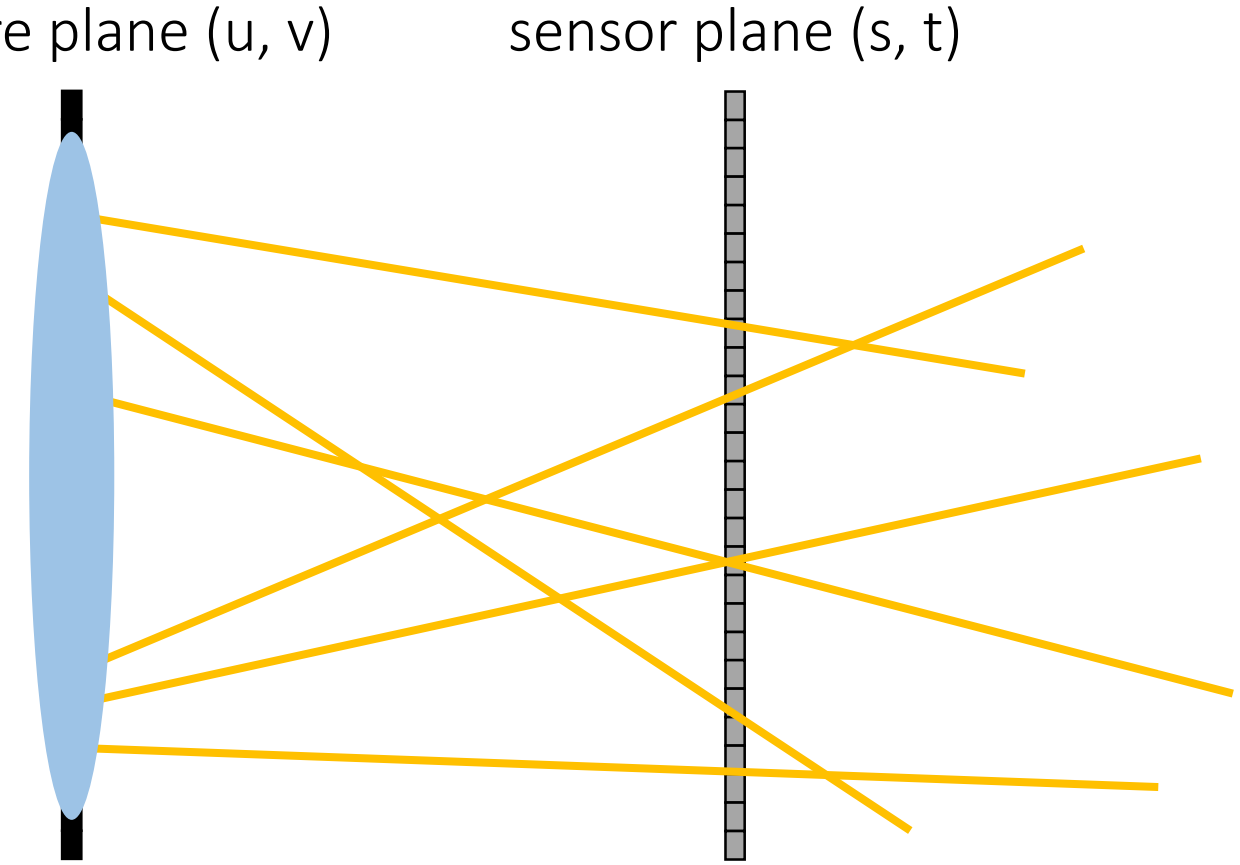

4-dimensional function L(u, v, s, t)
(conjugate of scene-based function)

Lightfield visualization

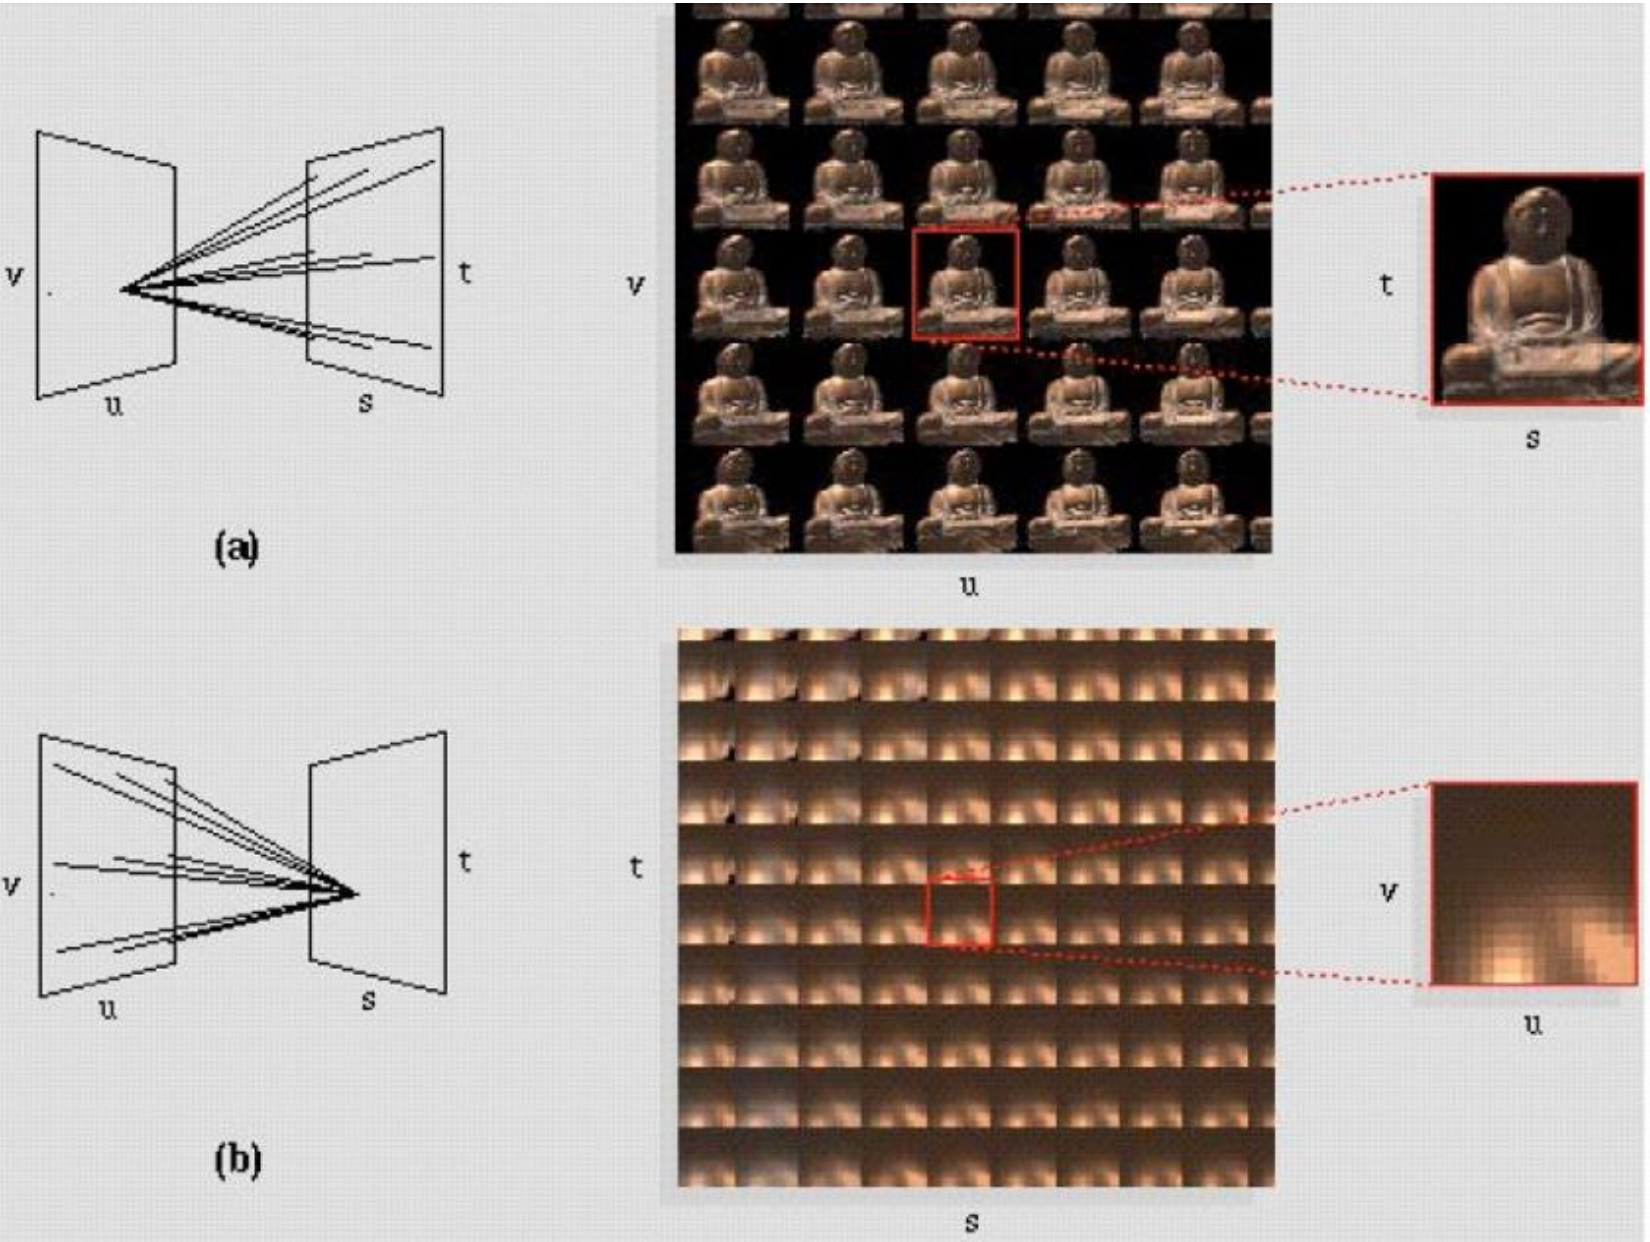

$L(u = u_o, v = v_o, s, t)$ is a pinhole image from a certain viewpoint

$L(u, v, s = s_o, t = t_o)$ is the radiance emitted by a certain (in-focus) point at various directions

Demo:
http://lightfield.stanford.edu/lfs.html

Lightfield: all rays in a scene

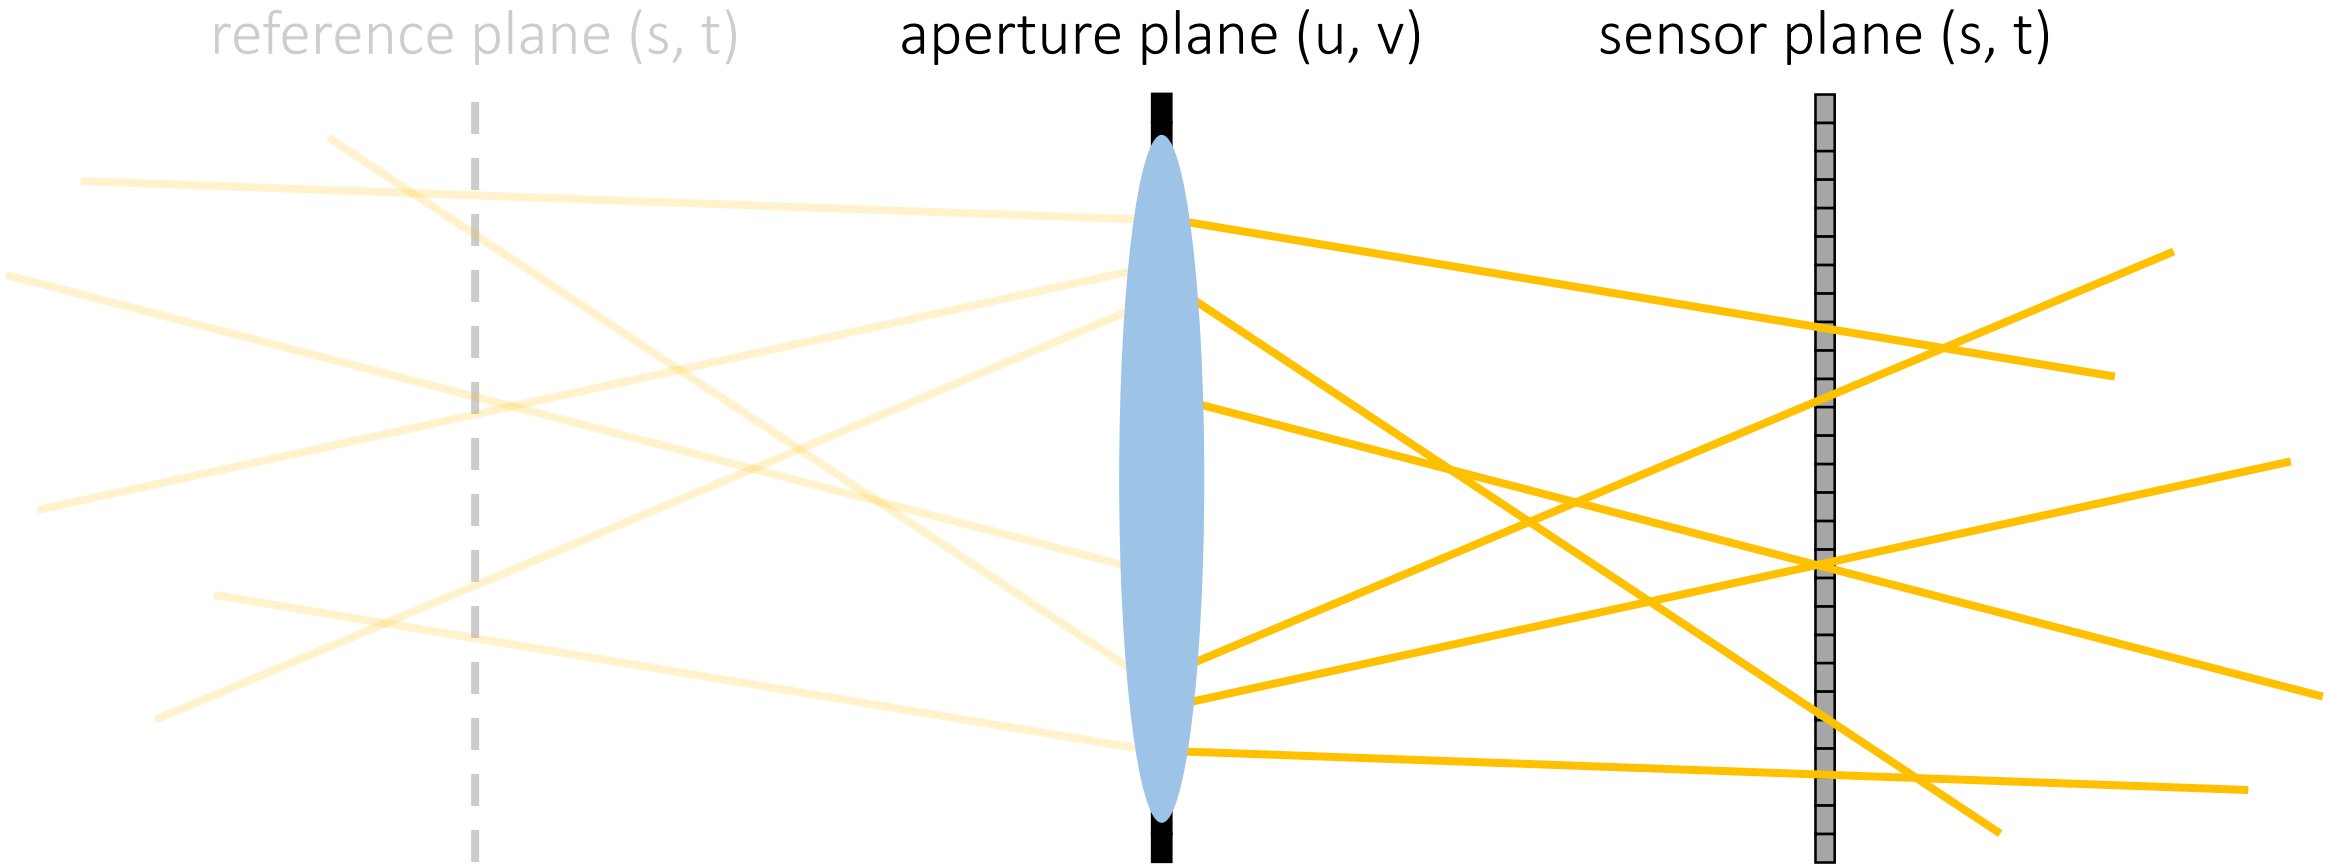

4-dimensional function L(u, v, s, t)
(conjugate of scene-based function)

How can you capture the lightfield of a scene?

Measuring lightfields

How to capture a lightfield?

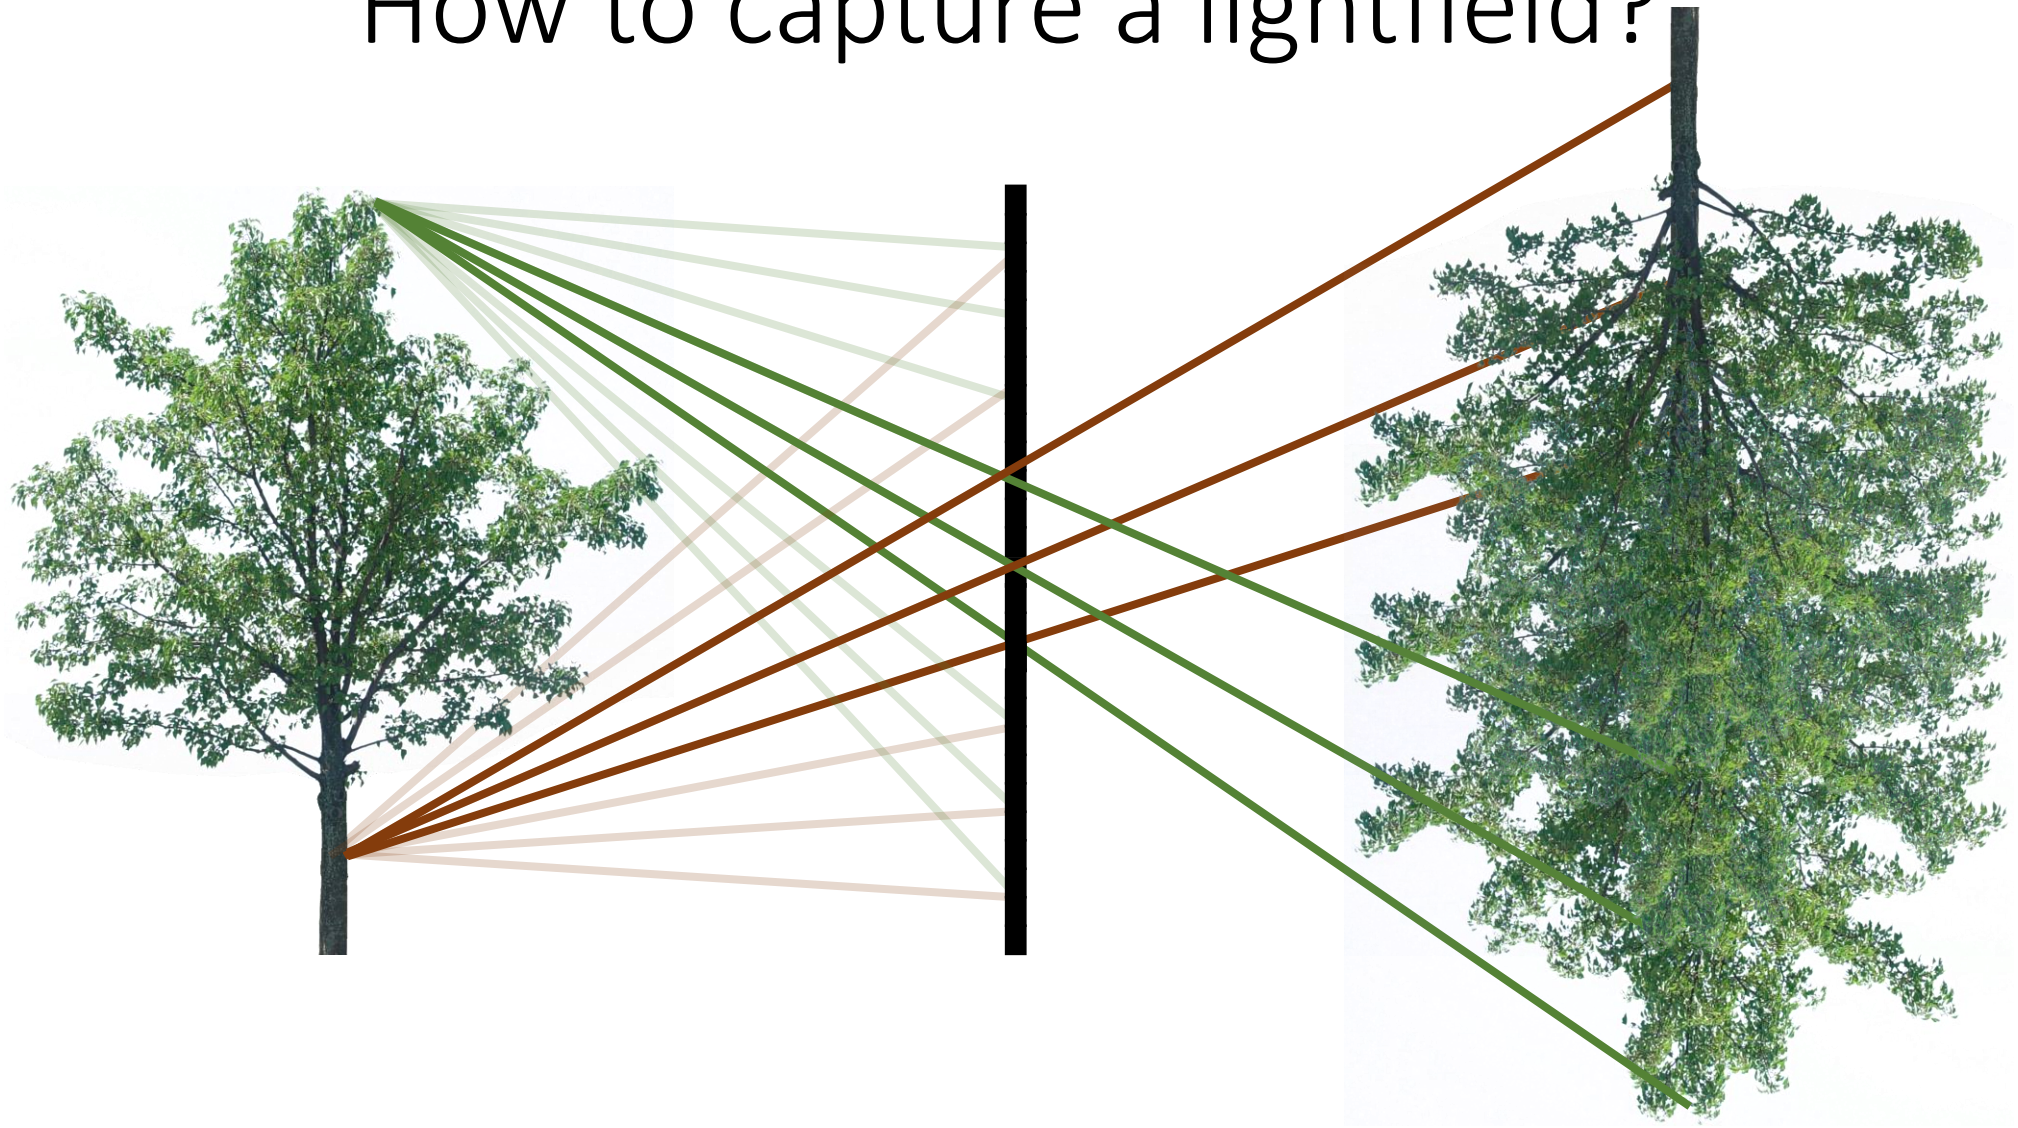

How can you do this?

Option 1: use multiple cameras

Stanford camera array

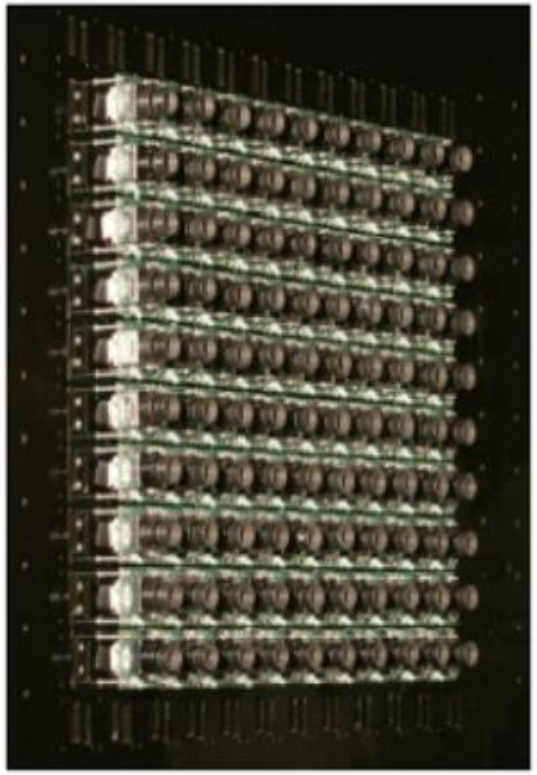

(“synthetic aperture”)

[Willburn et al., SIGGRAPH 2005]

Option 1: use multiple cameras

Stanford camera array

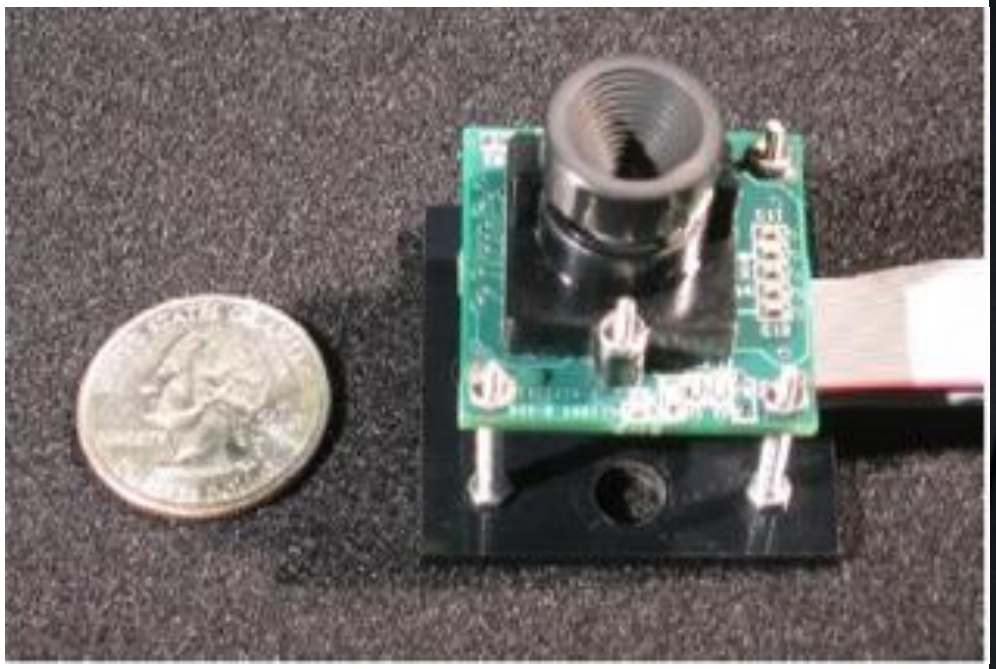

What kind of lens would you use for this?

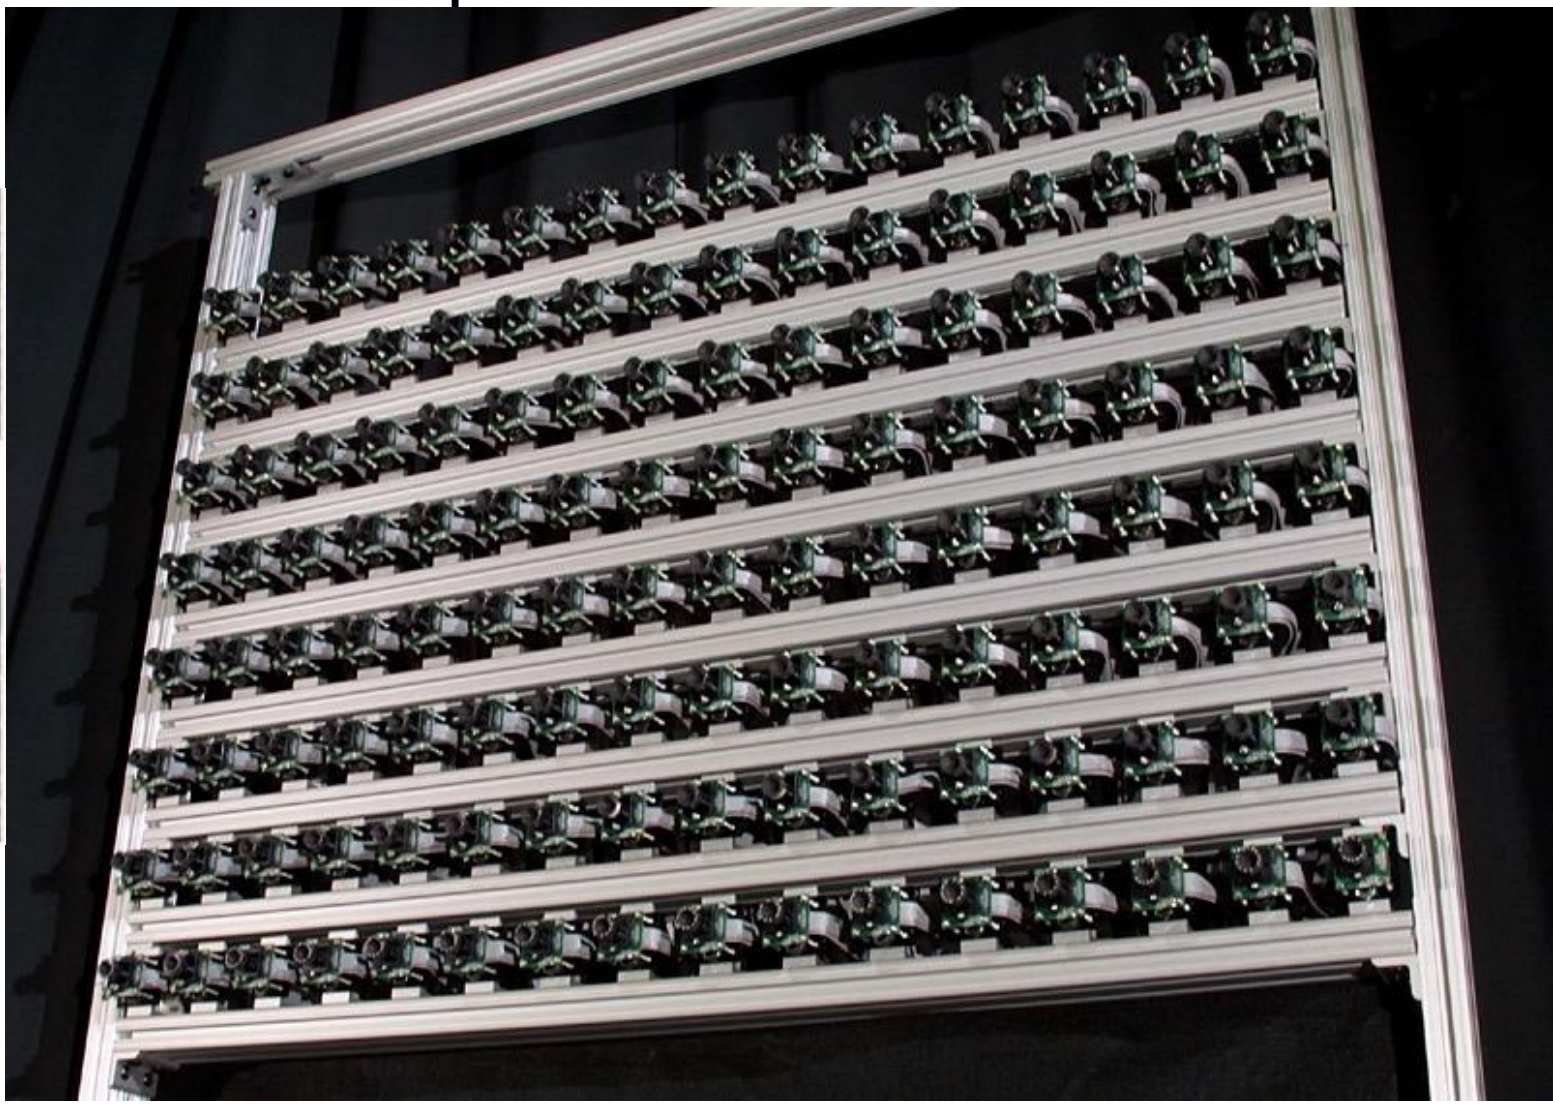

(“synthetic aperture”)

[Willburn et al., SIGGRAPH 2005]

Compound Eye vs Light Field

Image: National Geographic

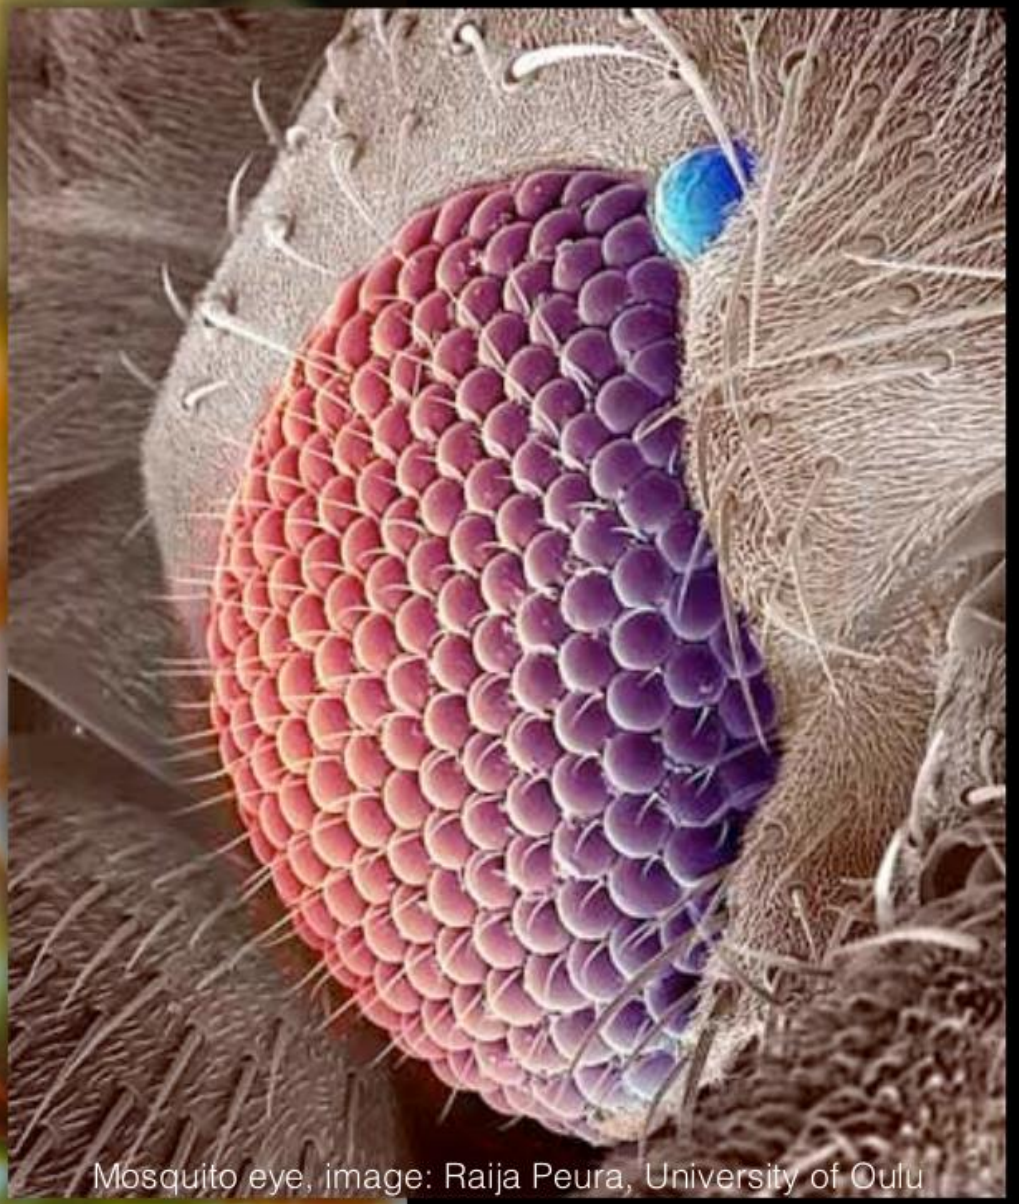

Mosquito eye, image: Raija Peura, University of Oulu

Option 2: take multiple images with one camera

Single camera mounted on LEGO motor. Demo: http://lightfield.stanford.edu/aperture.swf?lightfield=data/self_portrait_lf/preview.zip&zoom=1

How would you move the camera?

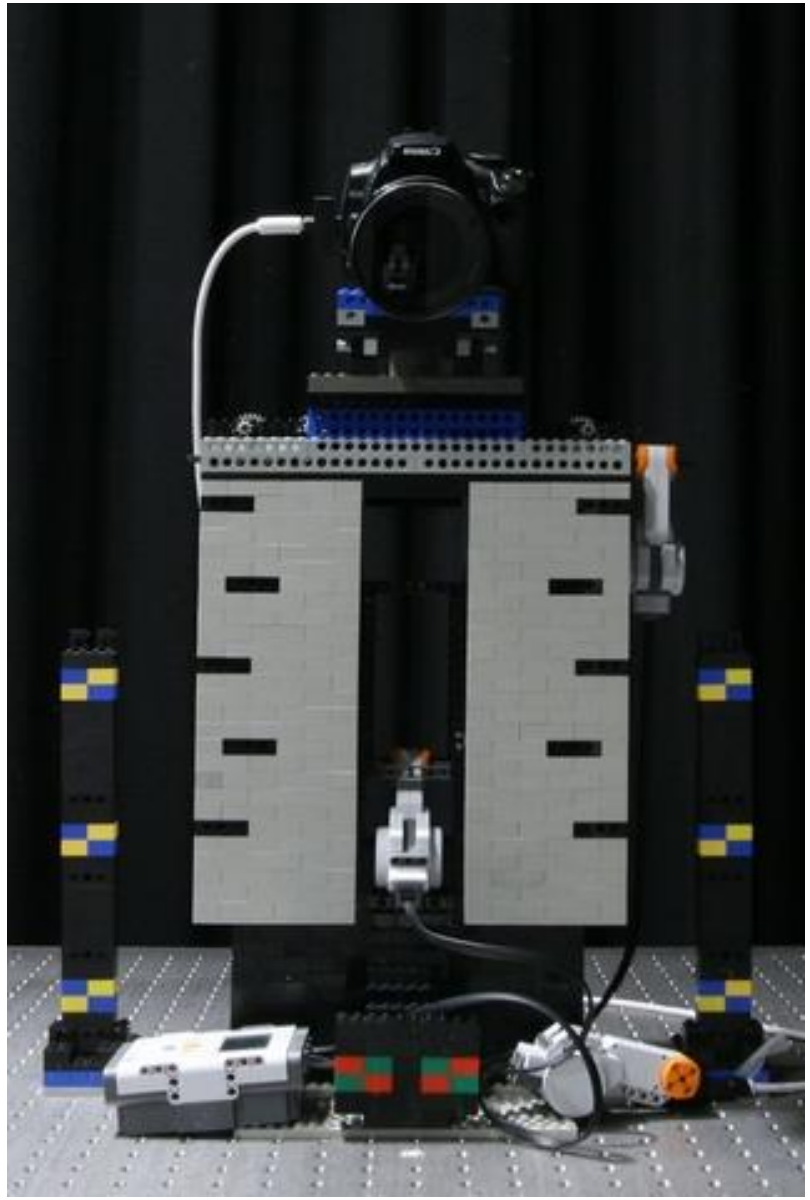

Plenoptic camera

Option 3: use a plenoptic camera

plenoptic = plenus (Latin for “full”) + optic (Greek for “seeing”, in this case)

Making a plenoptic camera

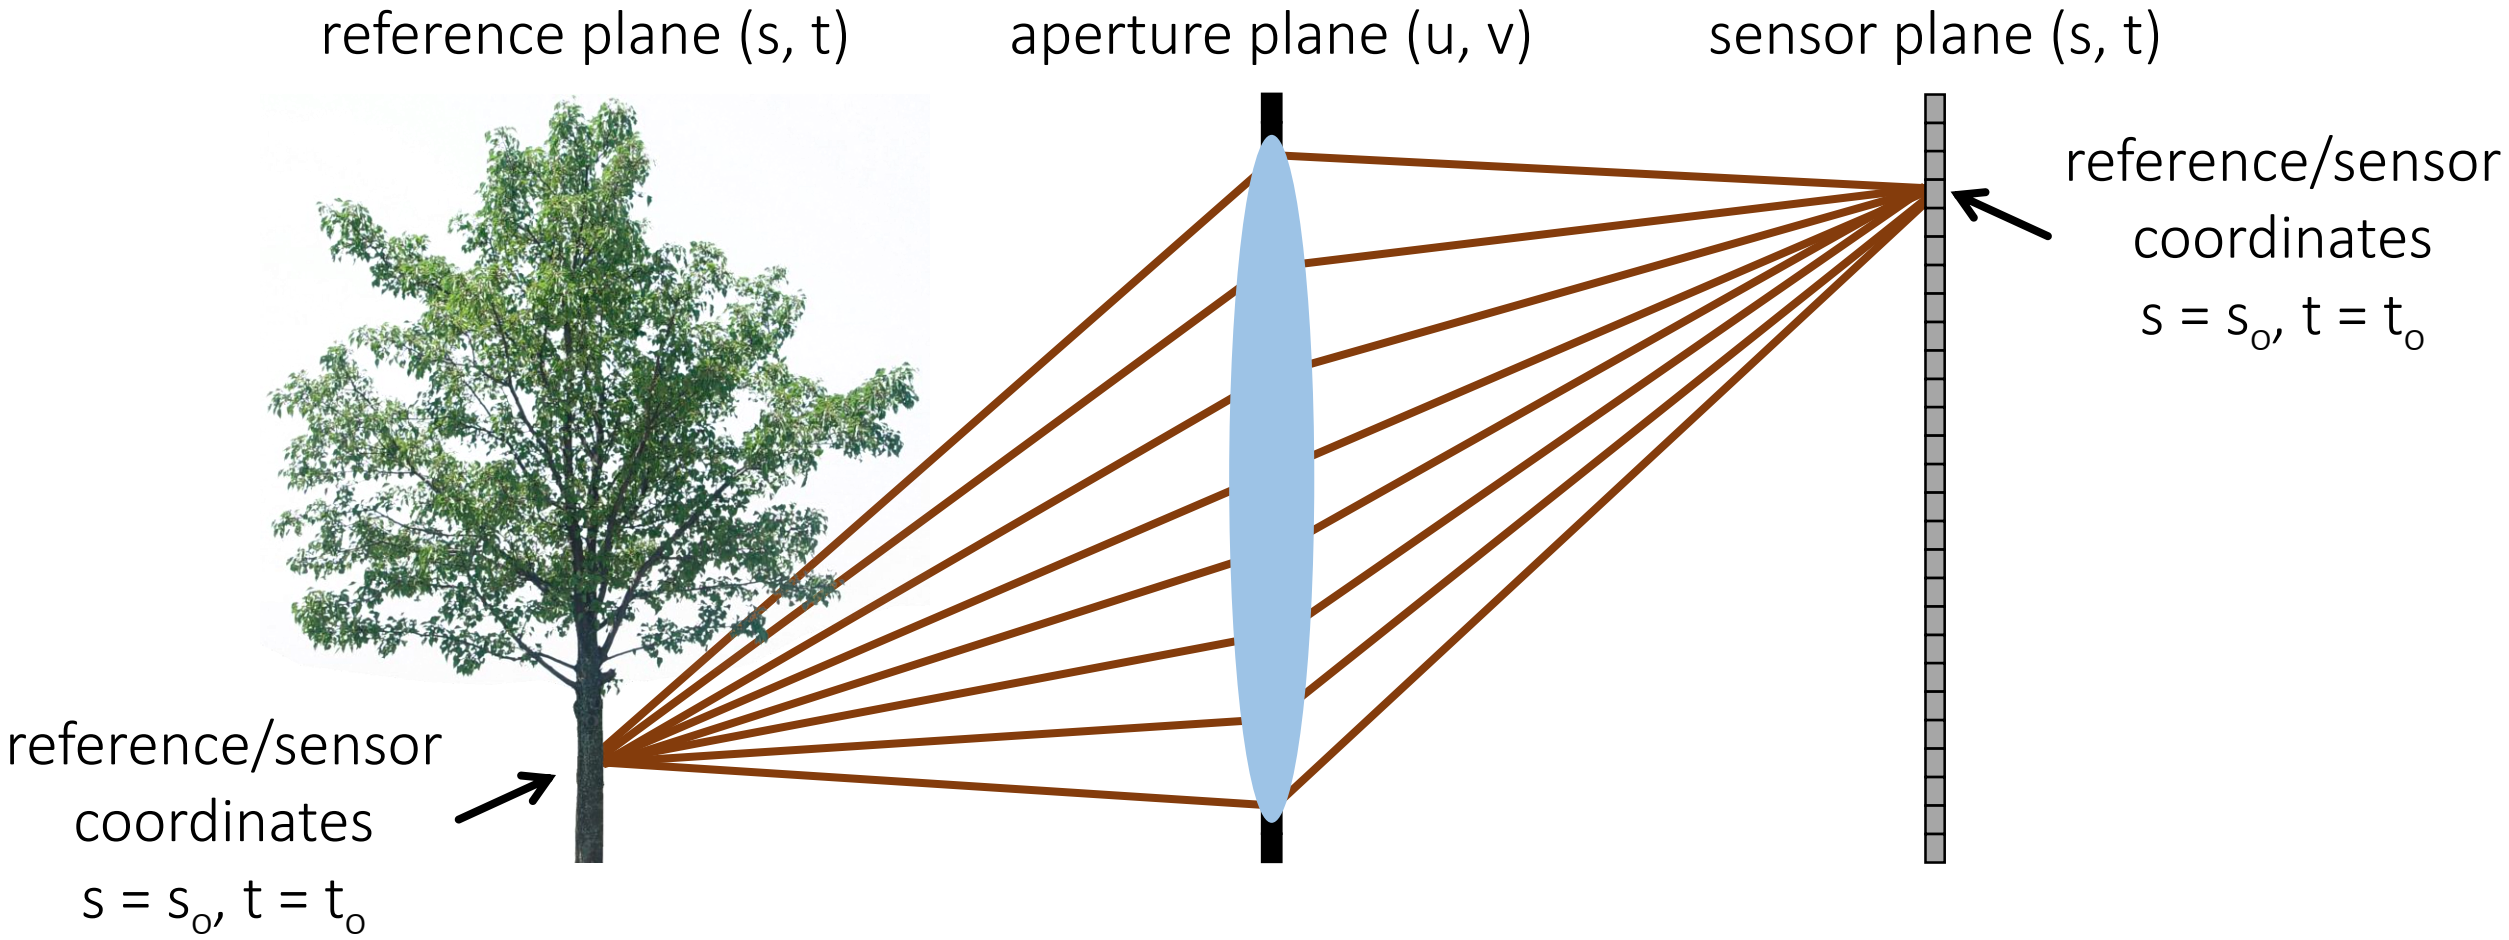

Lightfield slice $L(u, v, s = s_o, t = t_o)$

Making a plenoptic camera

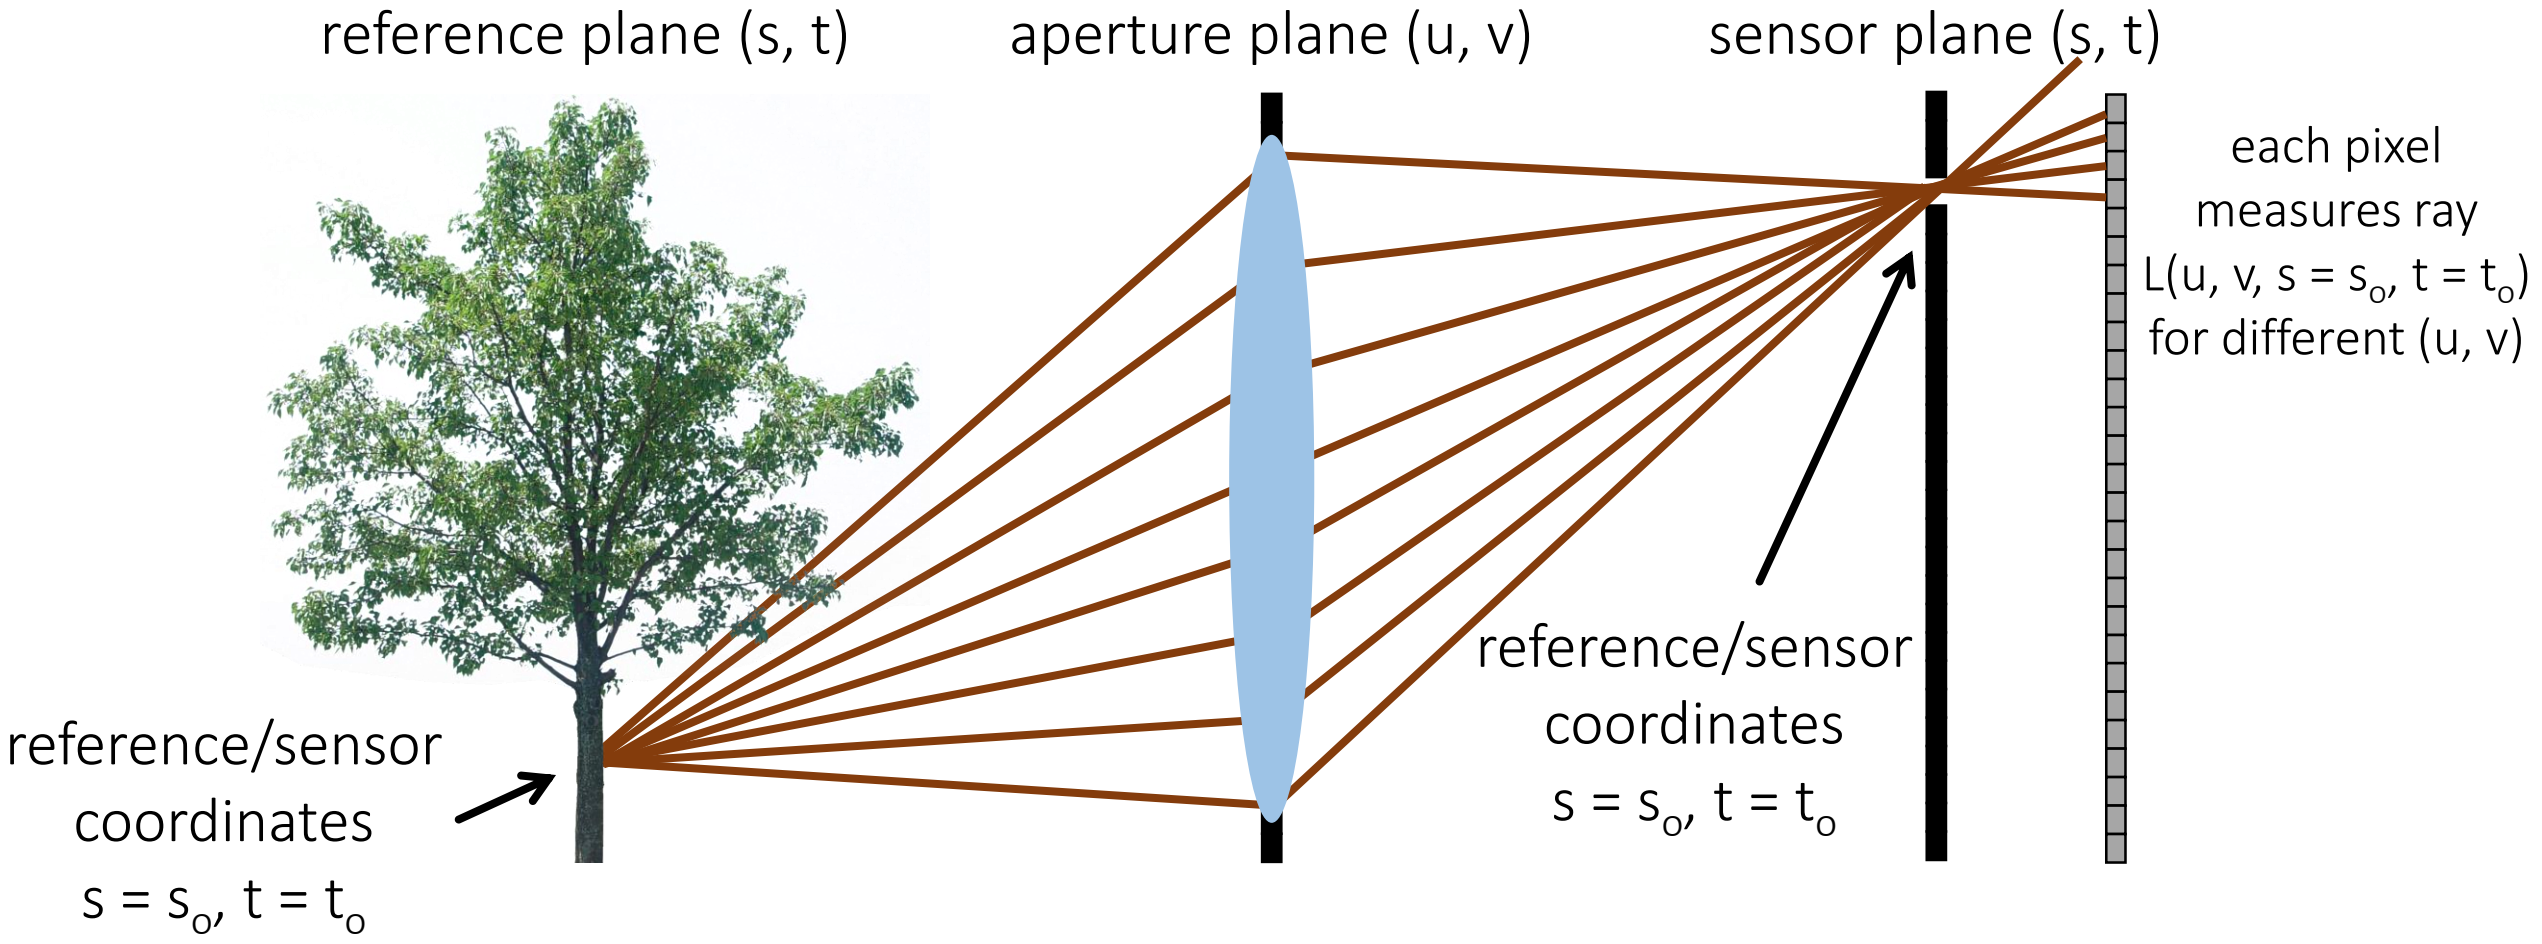

Lightfield slice $L(u, v, s = s_o, t = t_o)$

Making a plenoptic camera

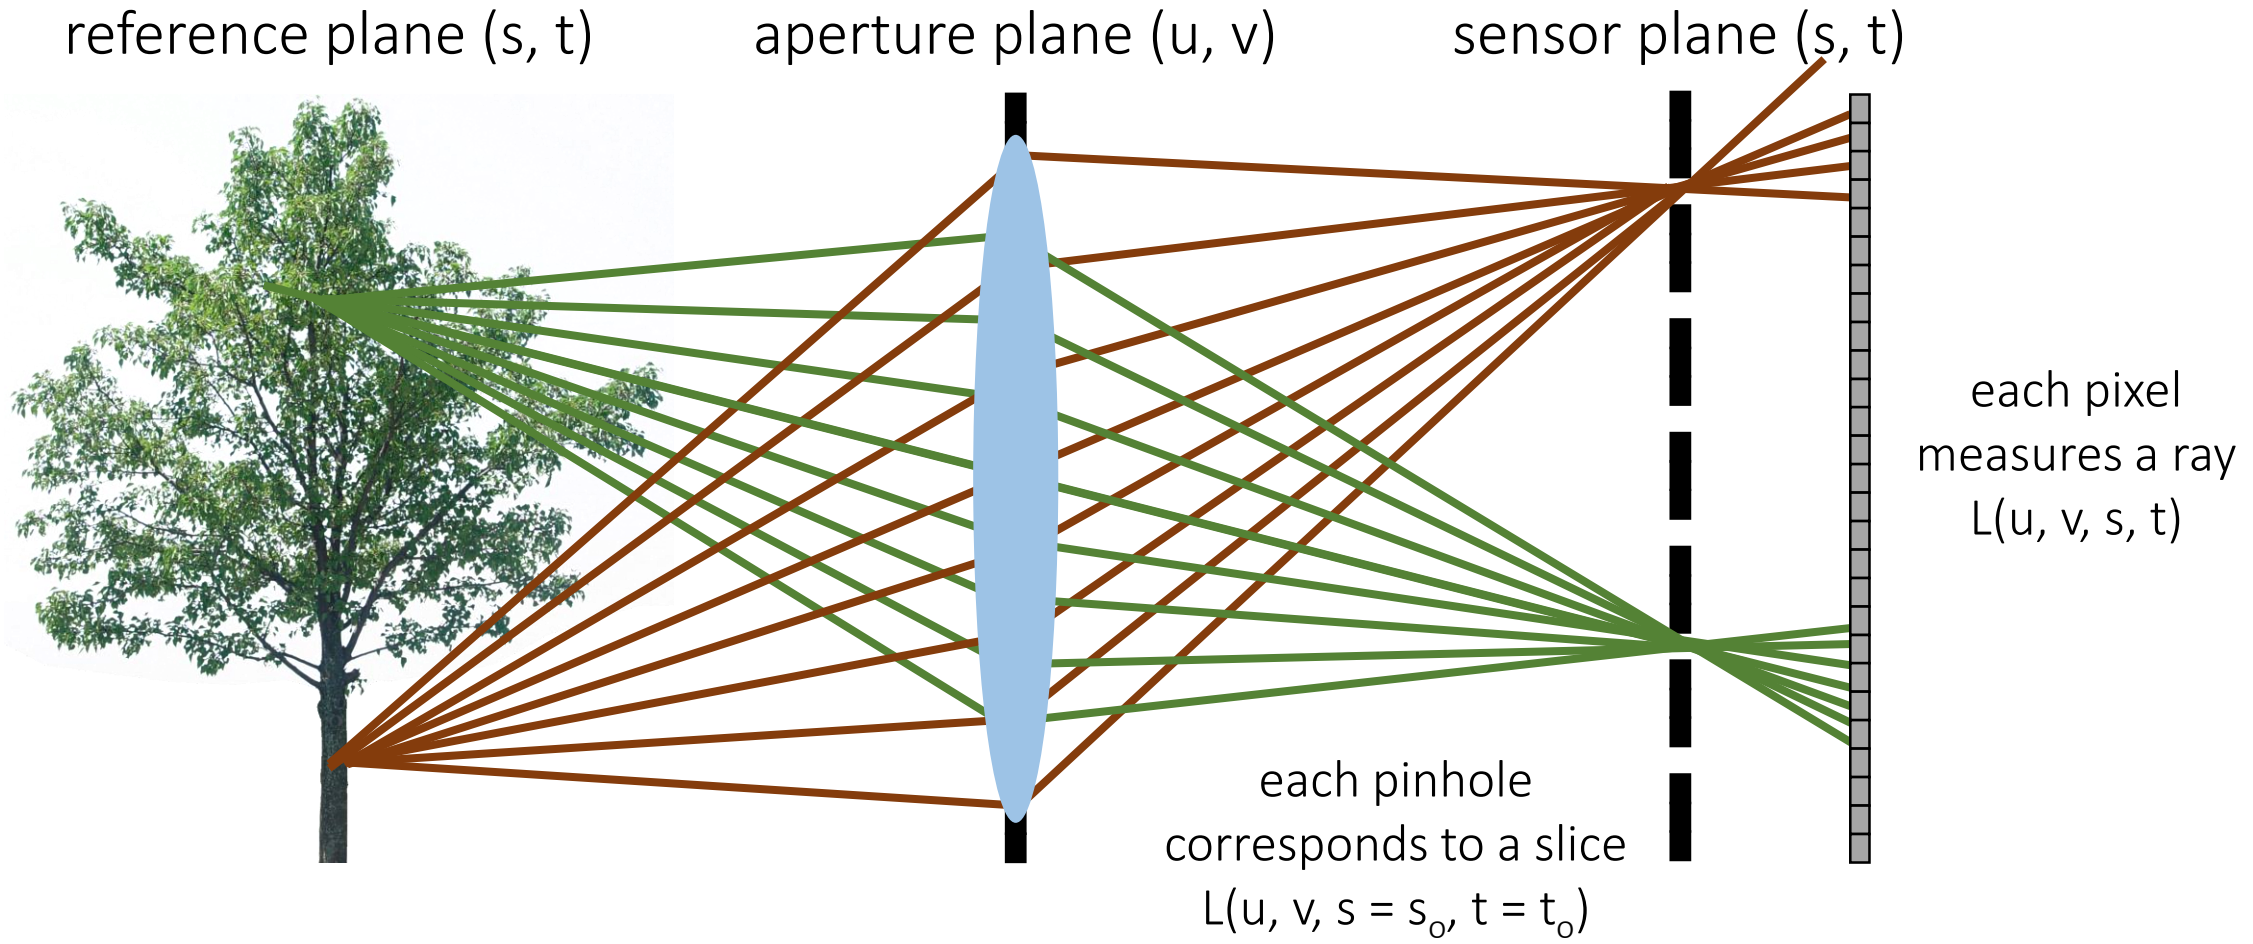

Lightfield L(u, v, s, t)

How can we make this more light efficient?

Making a plenoptic camera

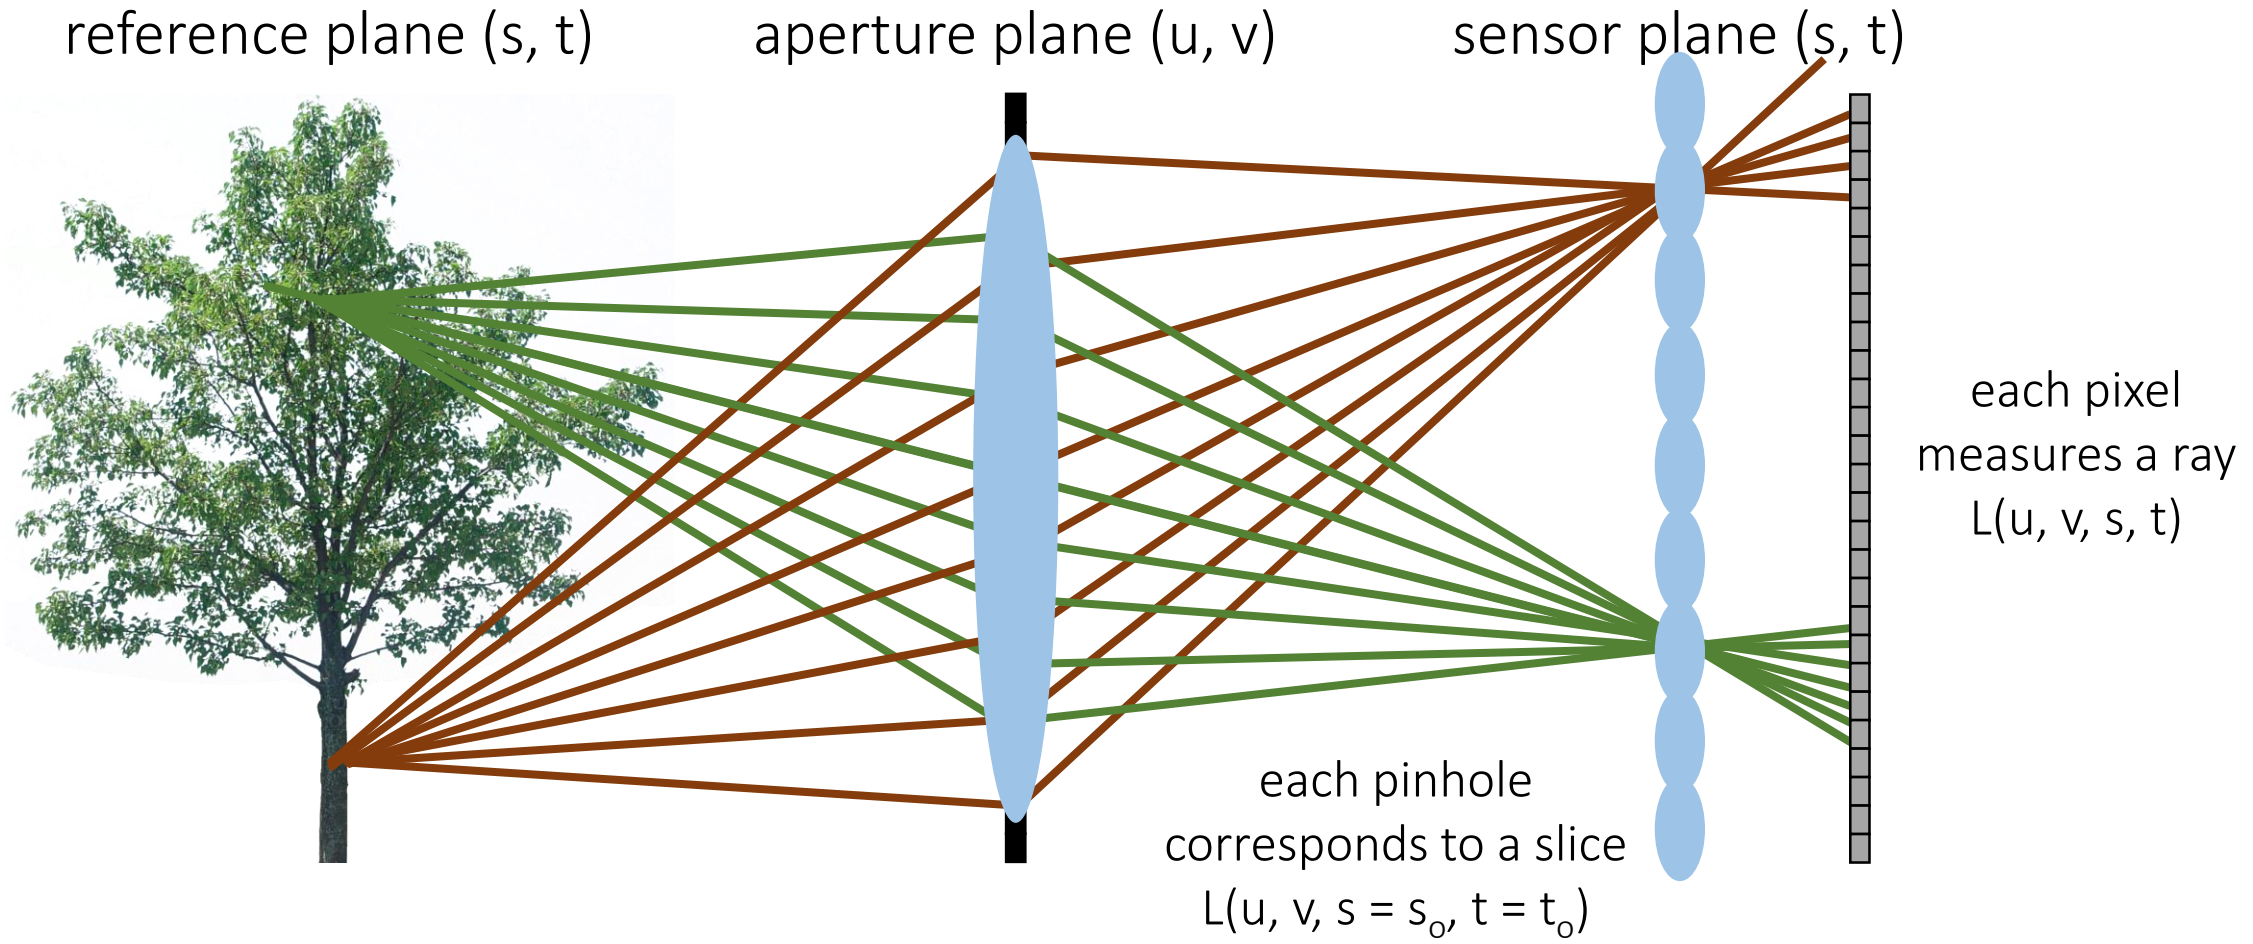

Lightfield L(u, v, s, t)

How can we make this more light efficient?

- replace pinholes with lenslets

Prototype plenoptic camera

- predecessor of Lytro
- resolution: 292x292px, 14x14 light field views

Kodak 16-megapixel sensor

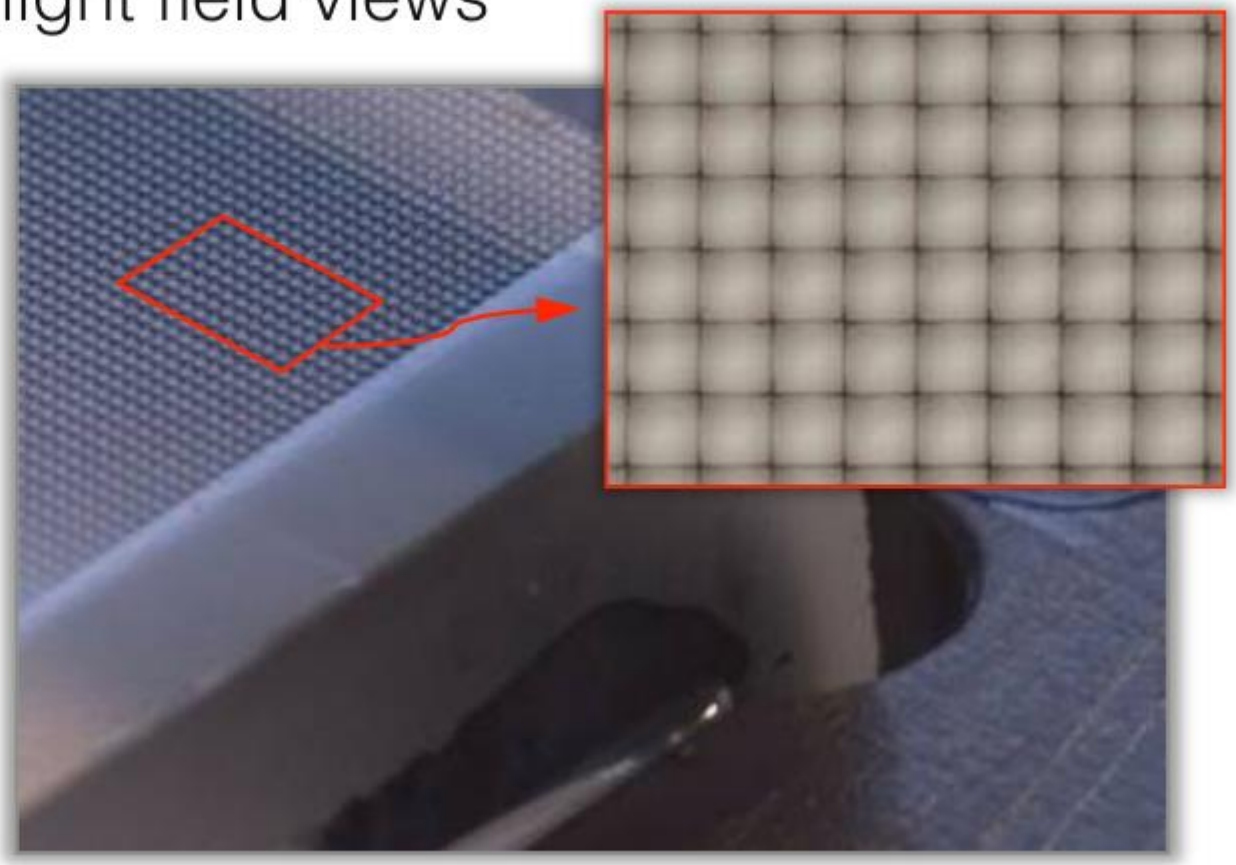

125μ square-sided microlenses

[Ng et al., Stanford Technical Report 2005]

Commercial plenoptic camera

Lens
The Lytro Light Field Camera starts with an 8X optical zoom, f/2 aperture lens. The aperture is constant across the zoom range allowing for unheard of light capture.

Light Field Engine 1.0
The Light Field Engine replaces the supercomputer from the lab and processes the light ray data captured by the sensor.

The Light Field Engine travels with every living picture as it is shared, letting you refocus pictures right on the camera, on your desktop and online.

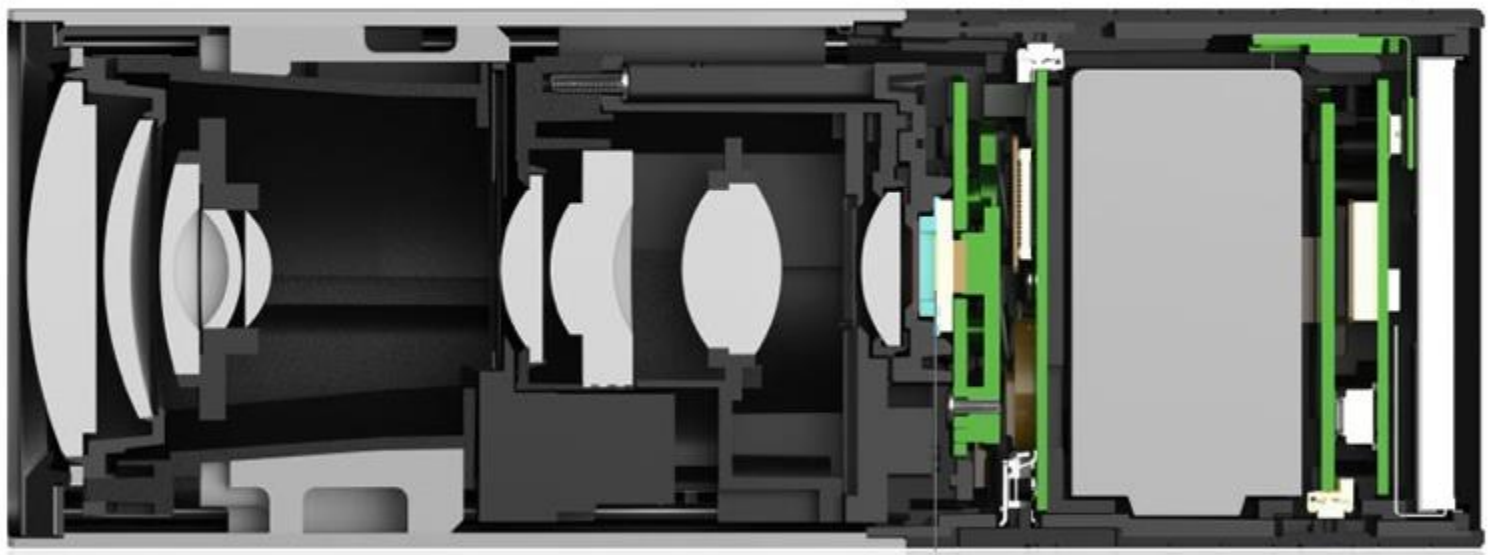

Light Field Sensor
From a roomful of cameras to a micro-lens array specially adhered to a standard sensor, the Lytro's Light Field Sensor captures 11 million light rays.

Commercial plenoptic camera

Lens
The Lytro Light Field Camera starts with an 8X optical zoom, f/2 aperture lens. The aperture is constant across the zoom range allowing for unheard of light capture.

Light Field Engine 1.0
The Light Field Engine replaces the supercomputer from the lab and processes the light ray data captured by the sensor.

The Light Field Engine travels with every living picture as it is shared, letting you refocus pictures right on the camera, on your desktop and online.

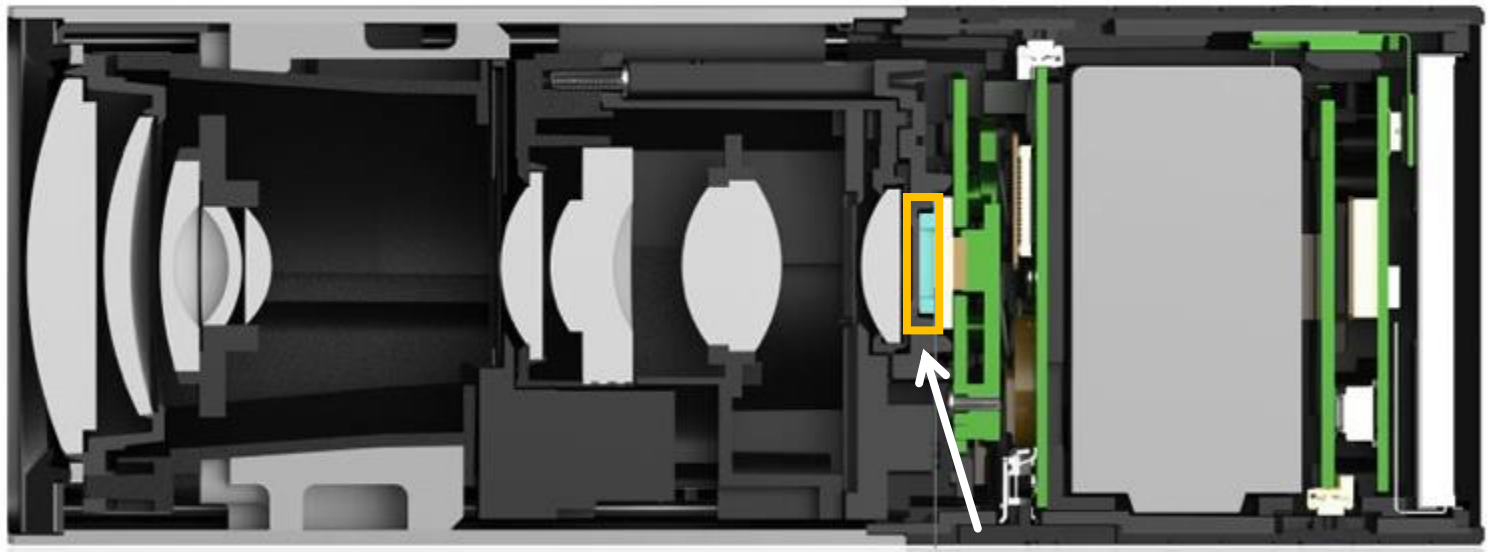

lenslet array

Light Field Sensor
From a roomful of cameras to a micro-lens array specially adhered to a standard sensor, the Lytro's Light Field Sensor captures 11 million light rays.

Commercial plenoptic camera

newer version with
higher resolution

Making a plenoptic camera

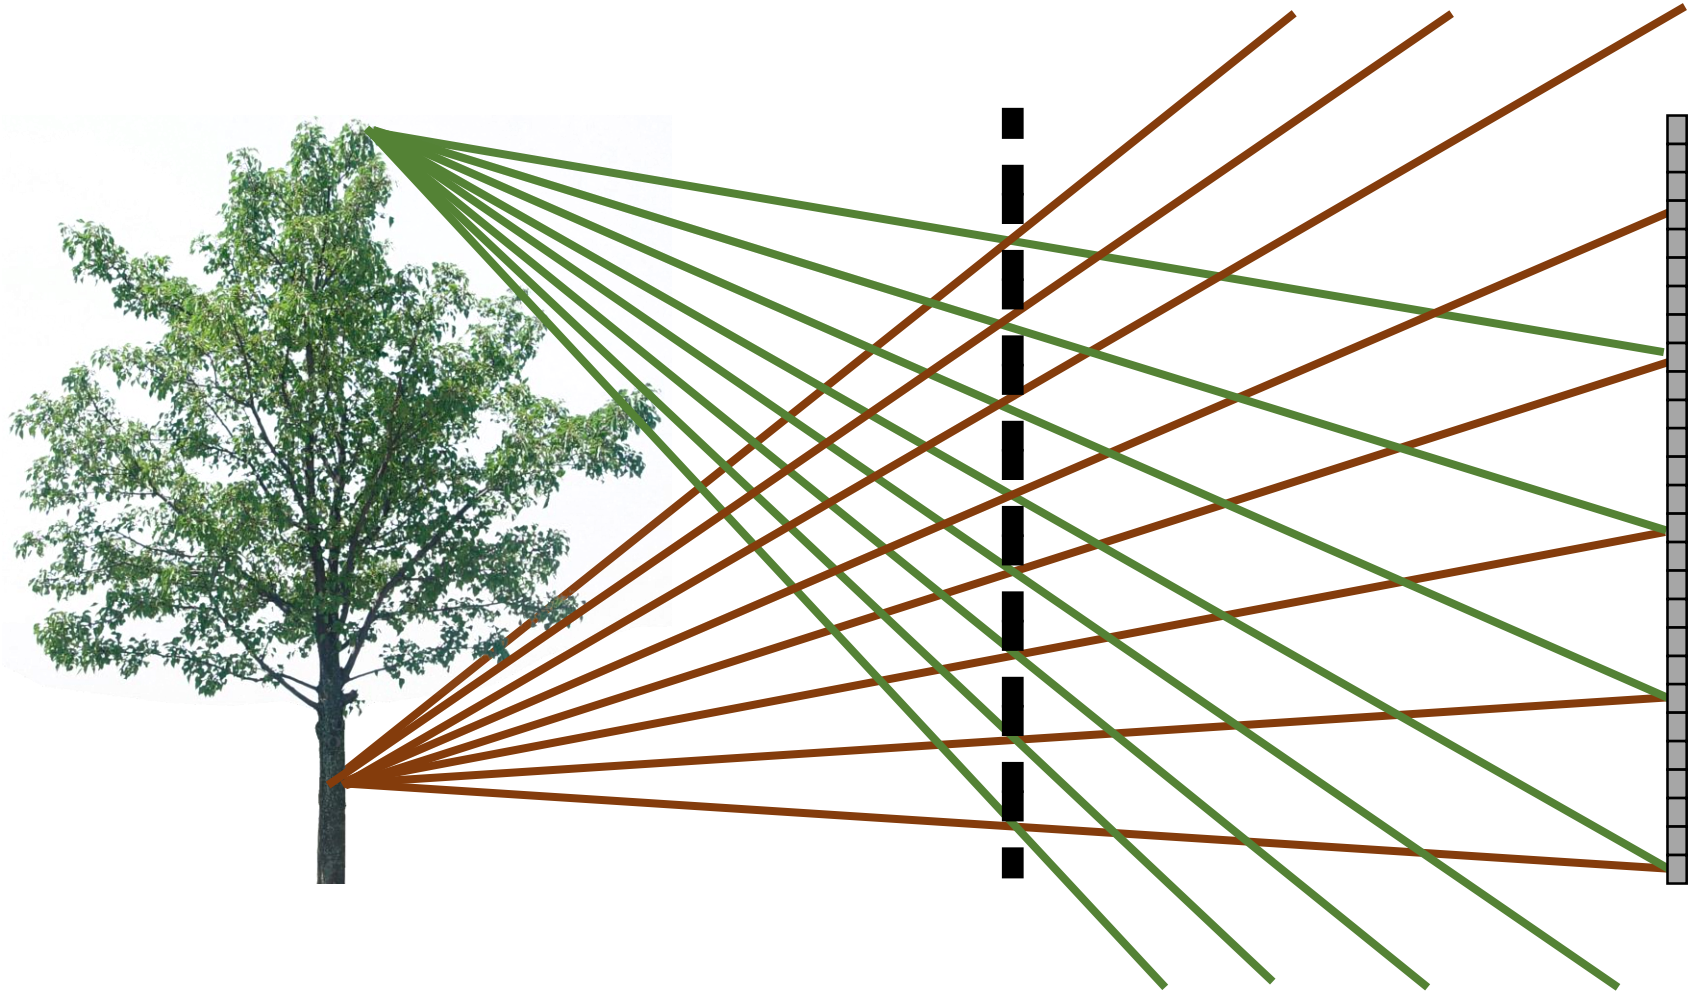

Is it possible to make a plenoptic camera using an array of pinholes?

Making a plenoptic camera

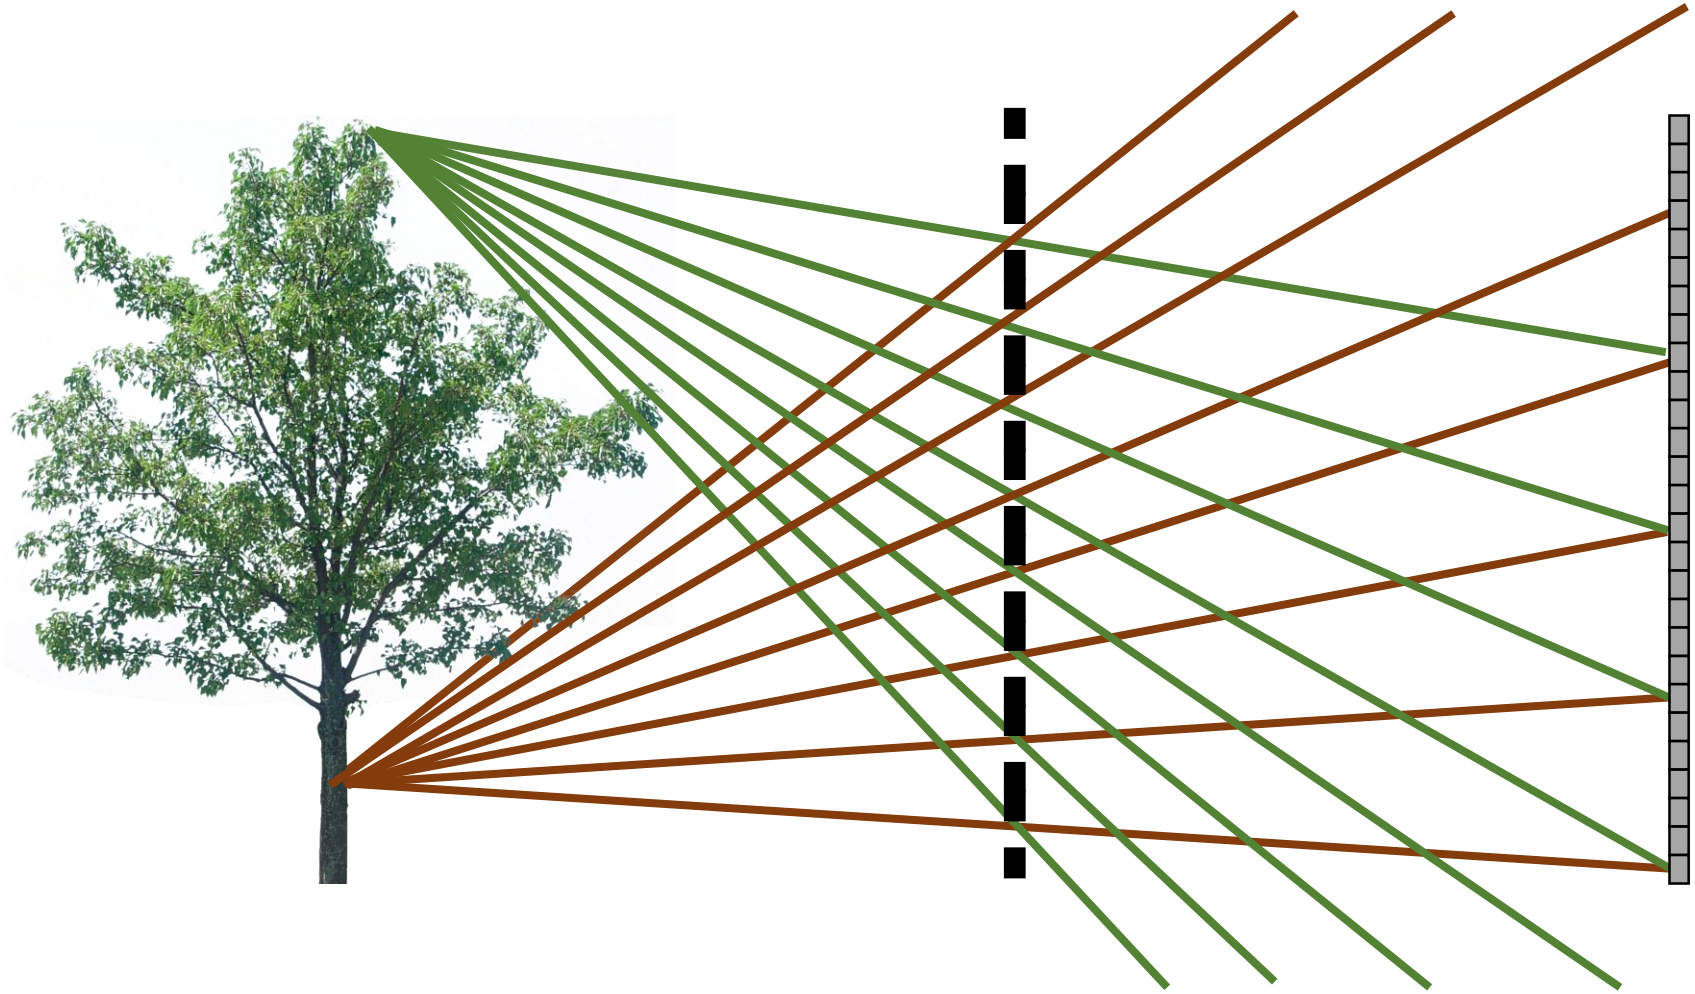

Need to be able to separate different rays falling on same pixel.

Is it possible to make a plenoptic camera using an array of pinholes?

- Yes, under certain conditions on the images being captured.

[Georgiev, ECCV 2008]

Images from lightfields

A plenoptic “image”

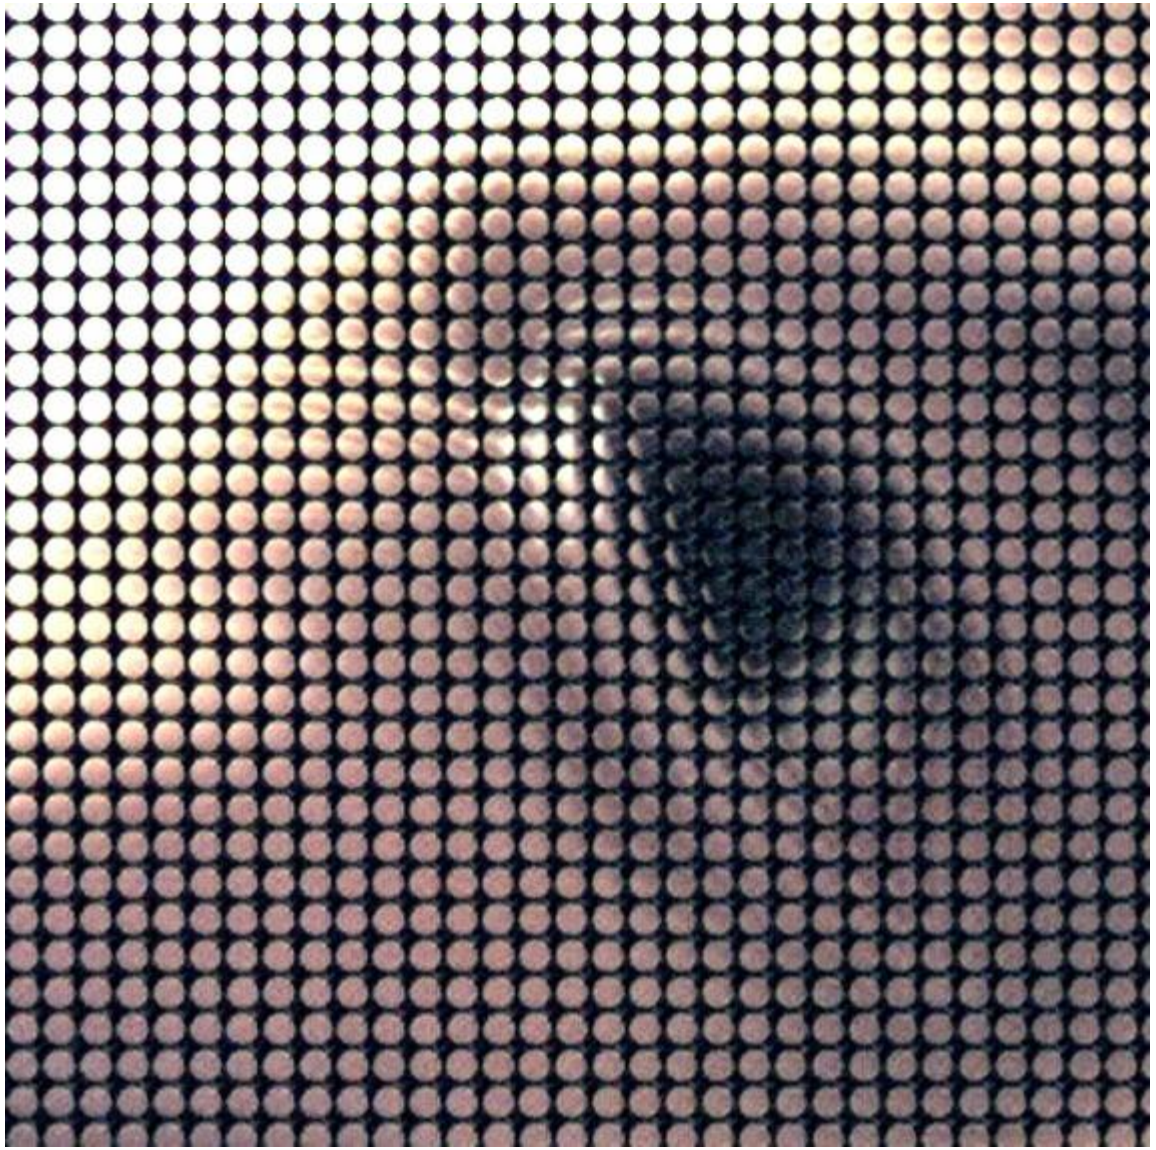

What are these circles?

A plenoptic camera

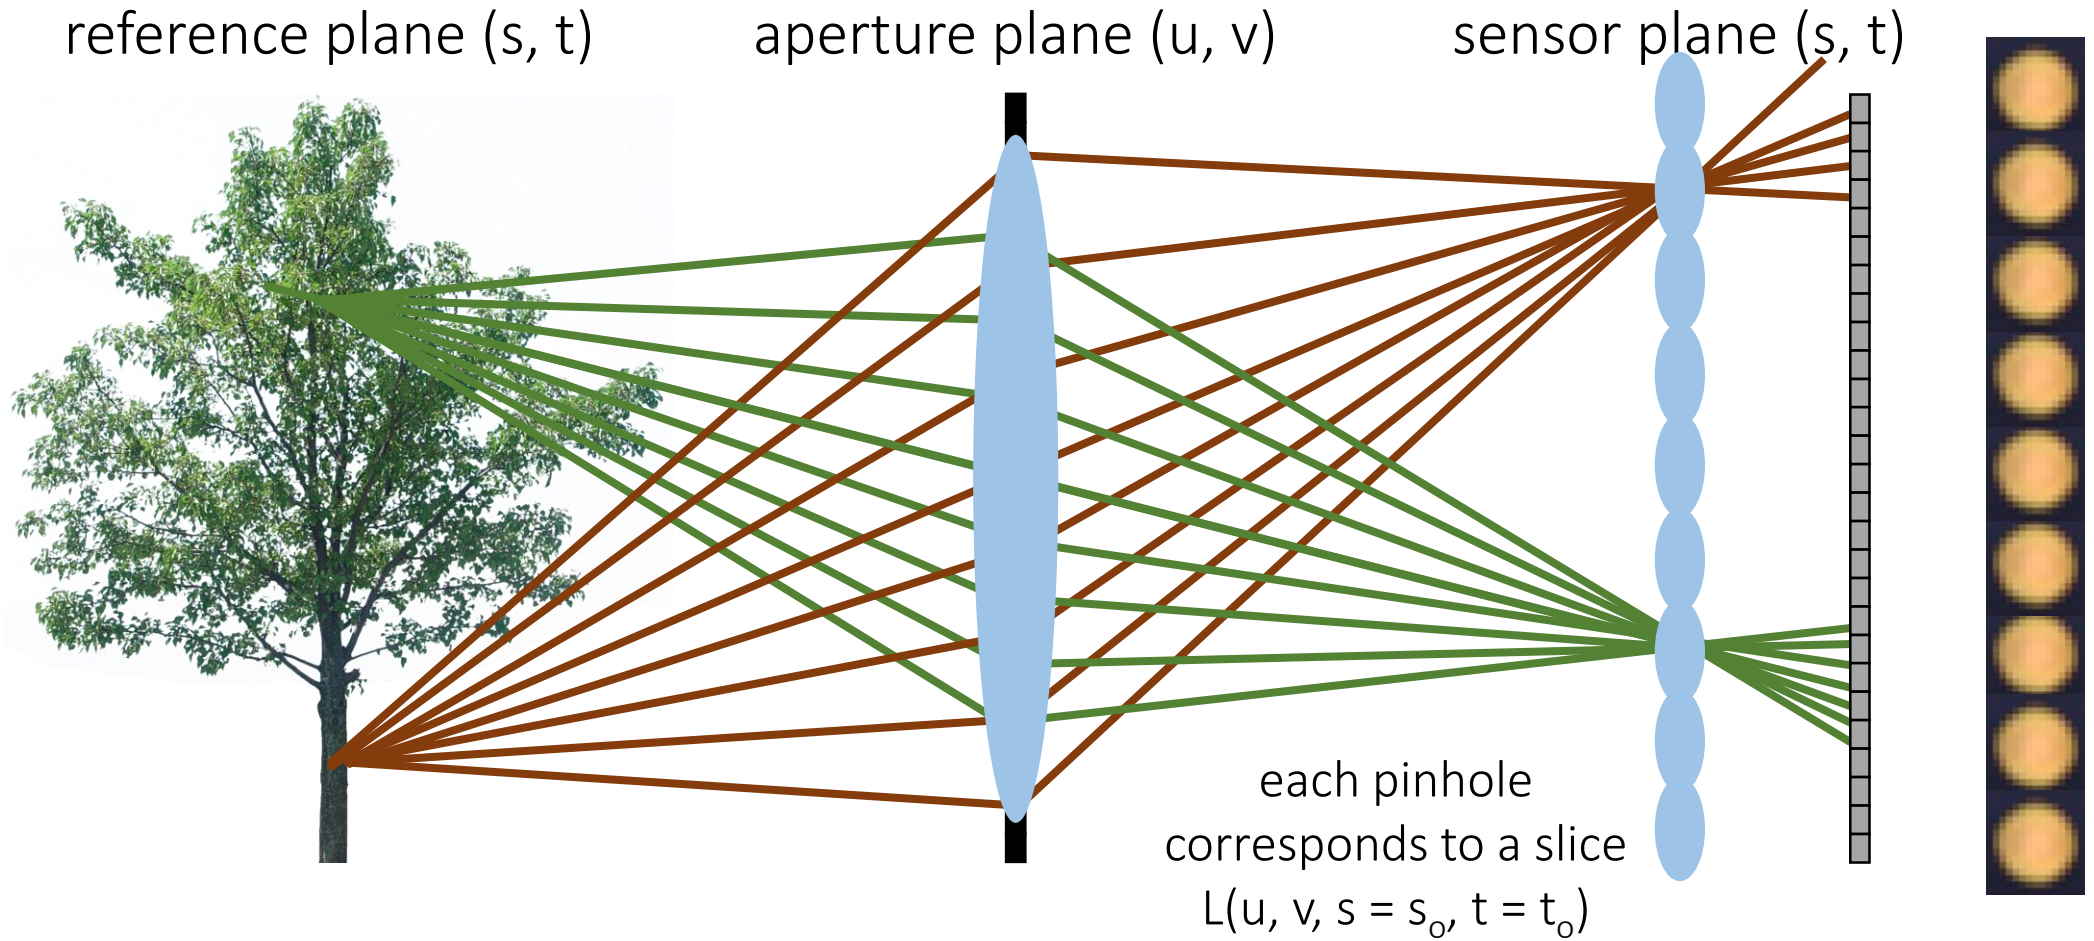

Lightfield L(u, v, s, t)

How do I…

Simulate different viewpoints?

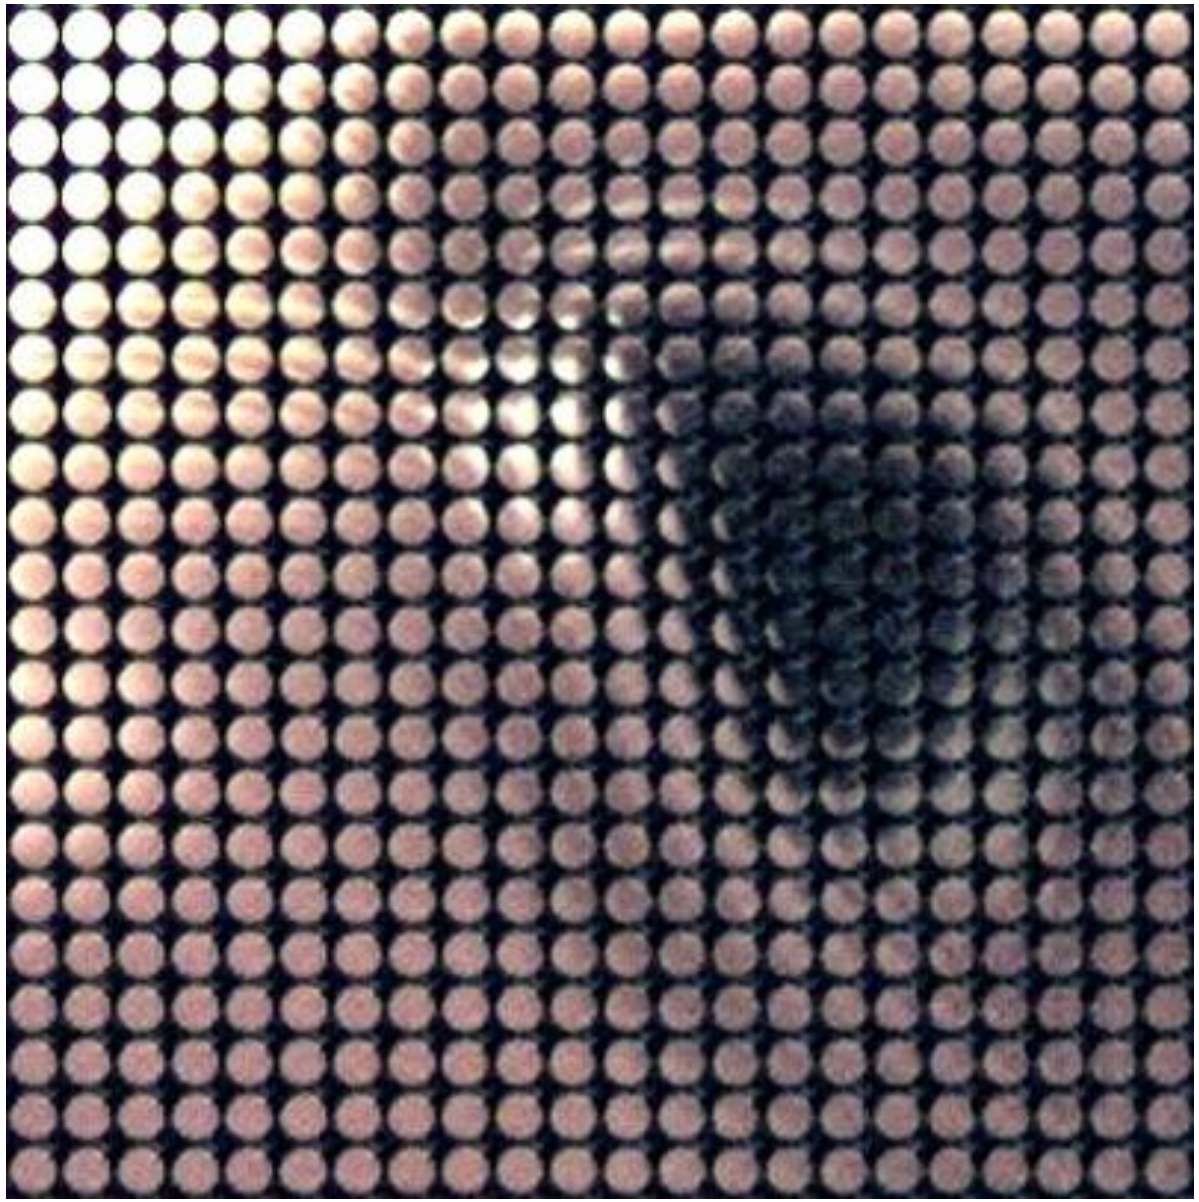

How do I…

Simulate different viewpoints?

- Pick same pixel within each aperture view

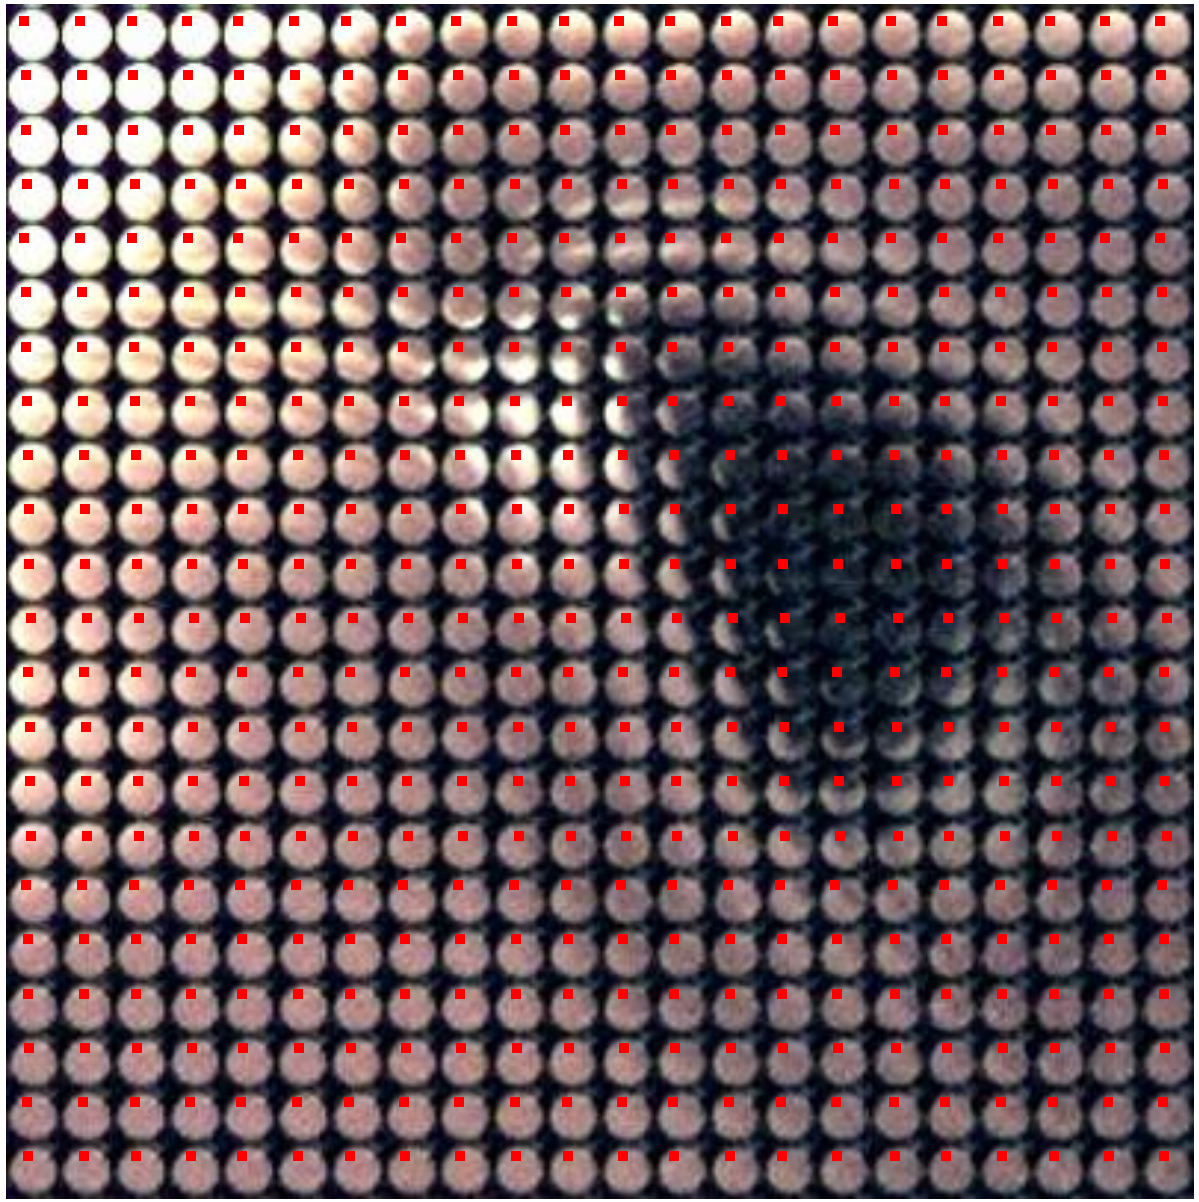

How do I…

Simulate different viewpoints?

- Pick same pixel within each aperture view

Simulate different aperture sizes?

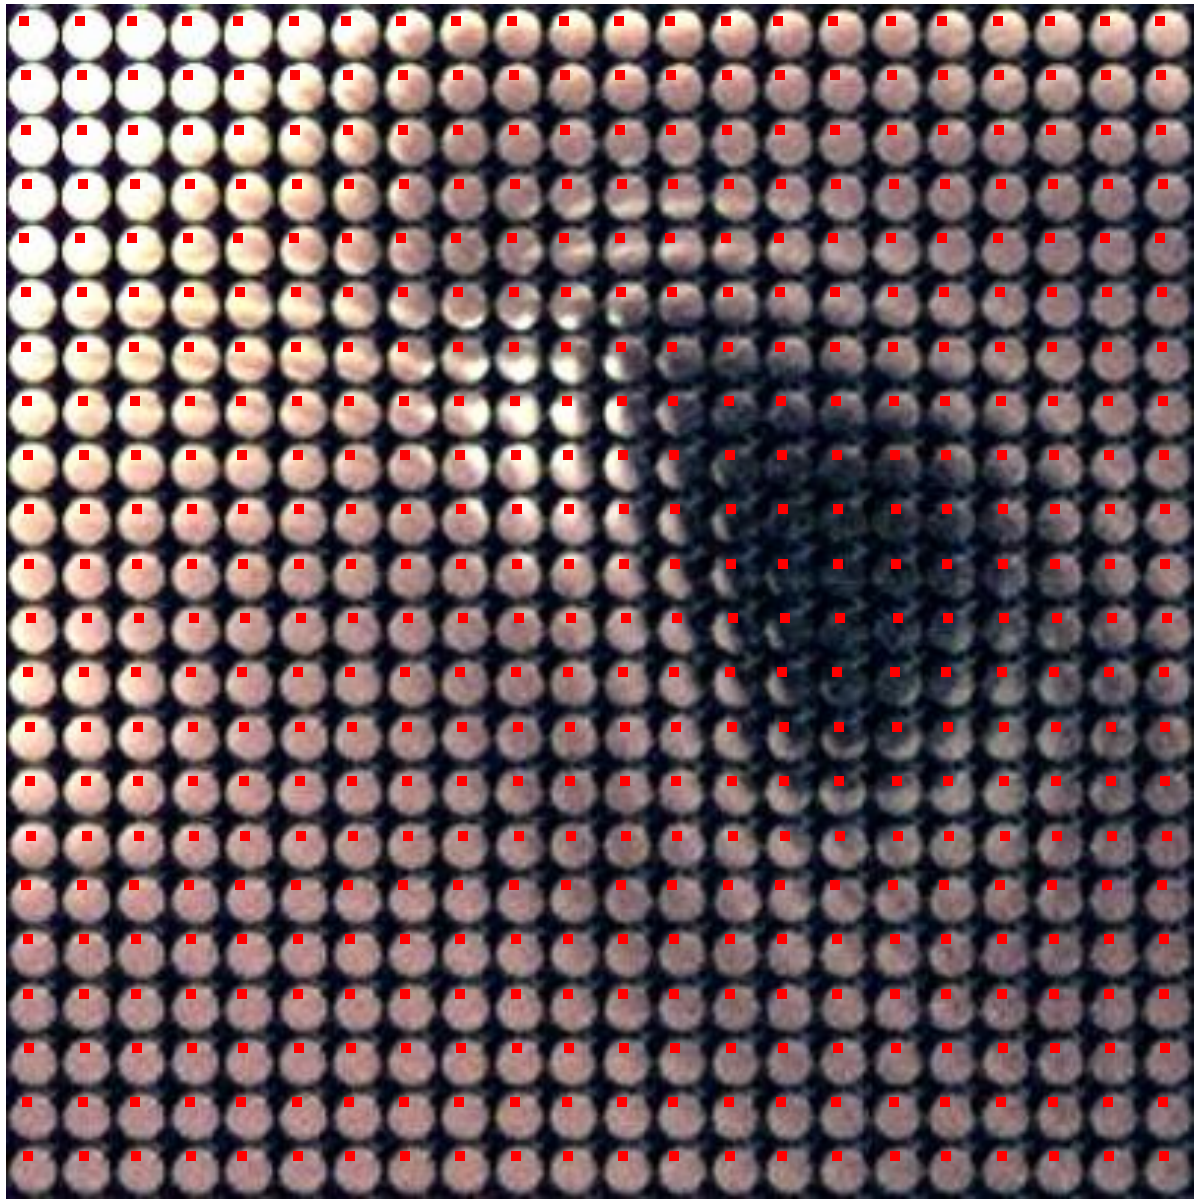

How do I…

Simulate different viewpoints?

- Pick same pixel within each aperture view

Simulate different aperture sizes?

- Sum more than one pixels within each aperture view

How do I…

Simulate different viewpoints?

- Pick same pixel within each aperture view

Simulate different aperture sizes?

- Sum more than one pixels within each aperture view

Simulate different focus depths?

How do I…

Simulate different viewpoints?

- Pick same pixel within each aperture view

Simulate different aperture sizes?

- Sum more than one pixels within each aperture view

Simulate different focus depths?

- Shift the areas you integrate relative to each other

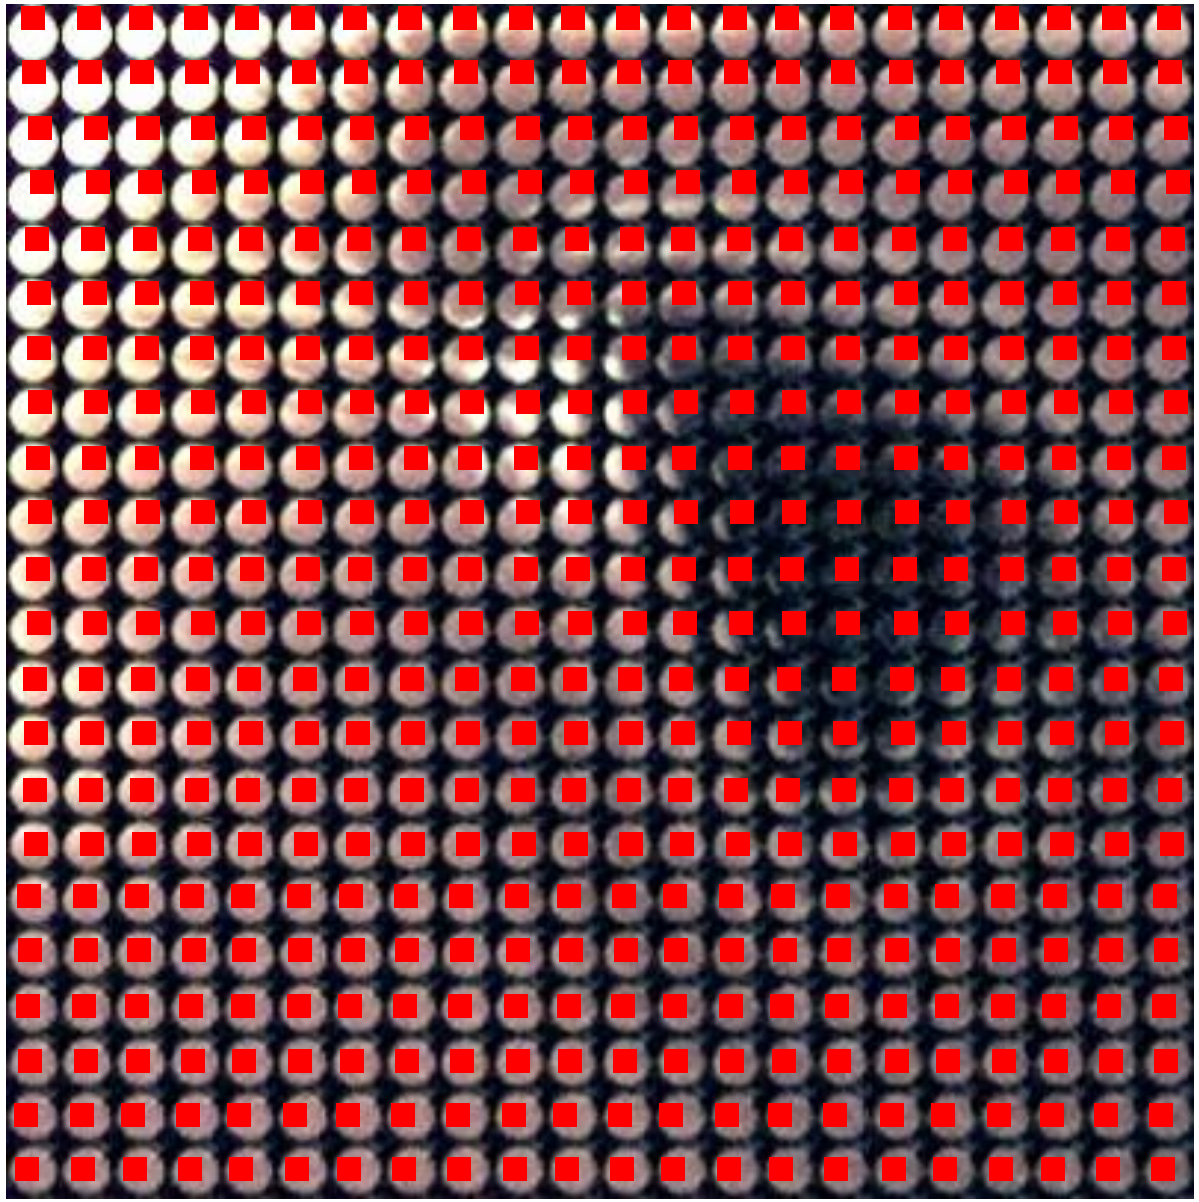

Refocus

- select and average correct rays

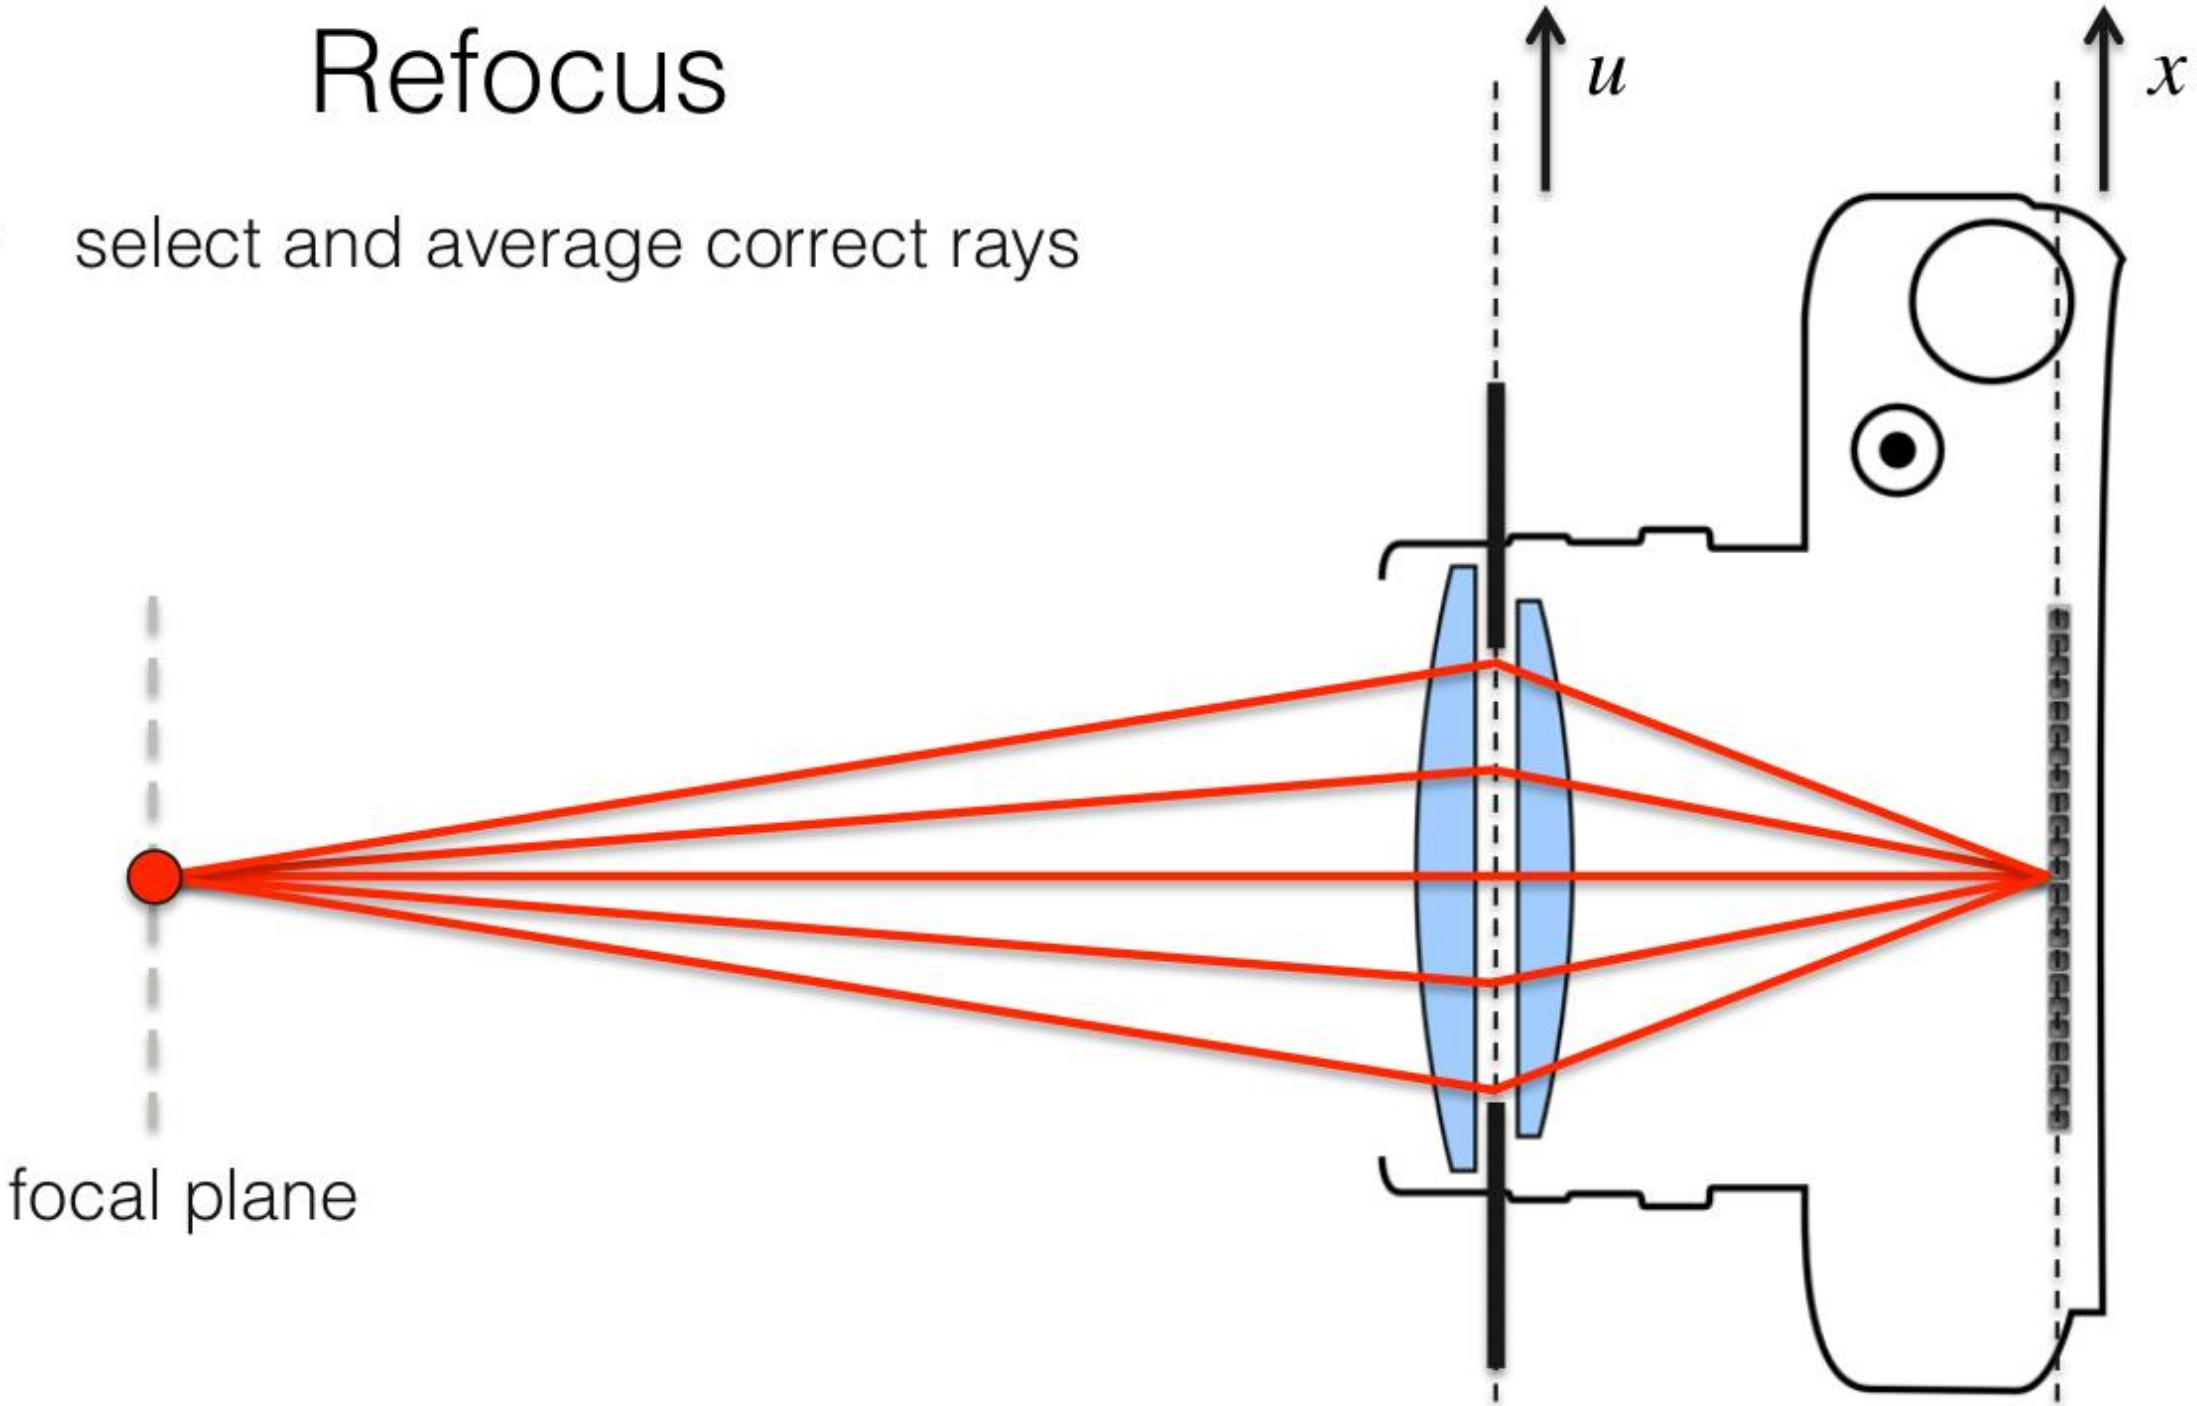

Refocus

- select and average correct rays

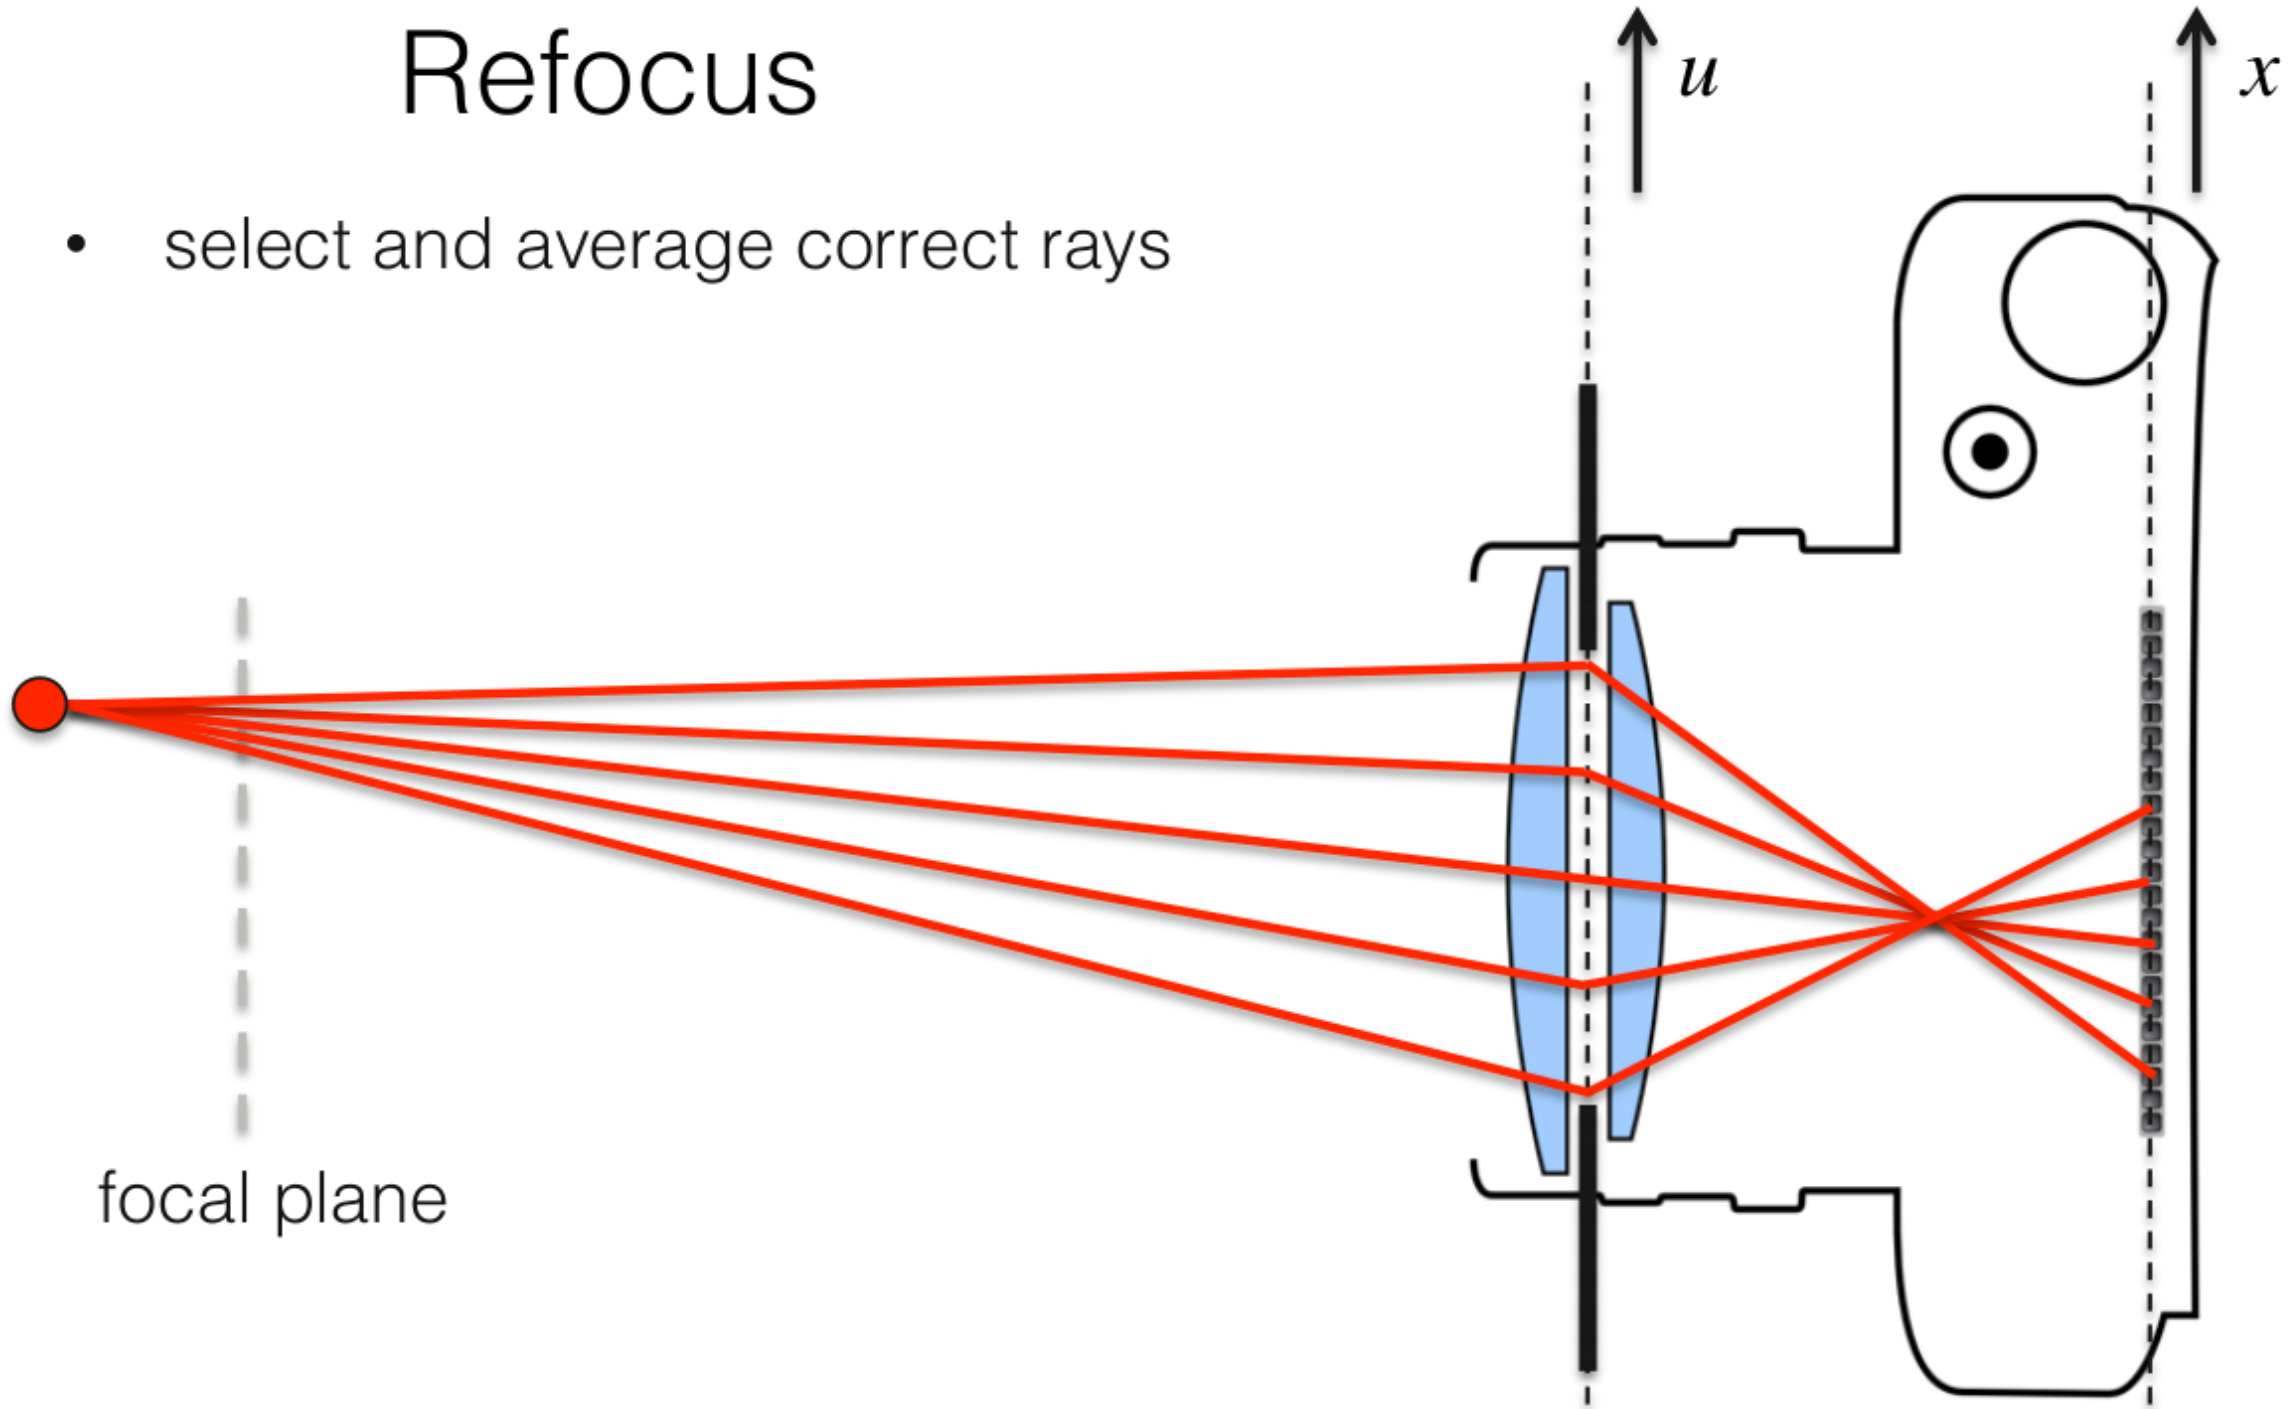

Refocus

- select and average correct rays

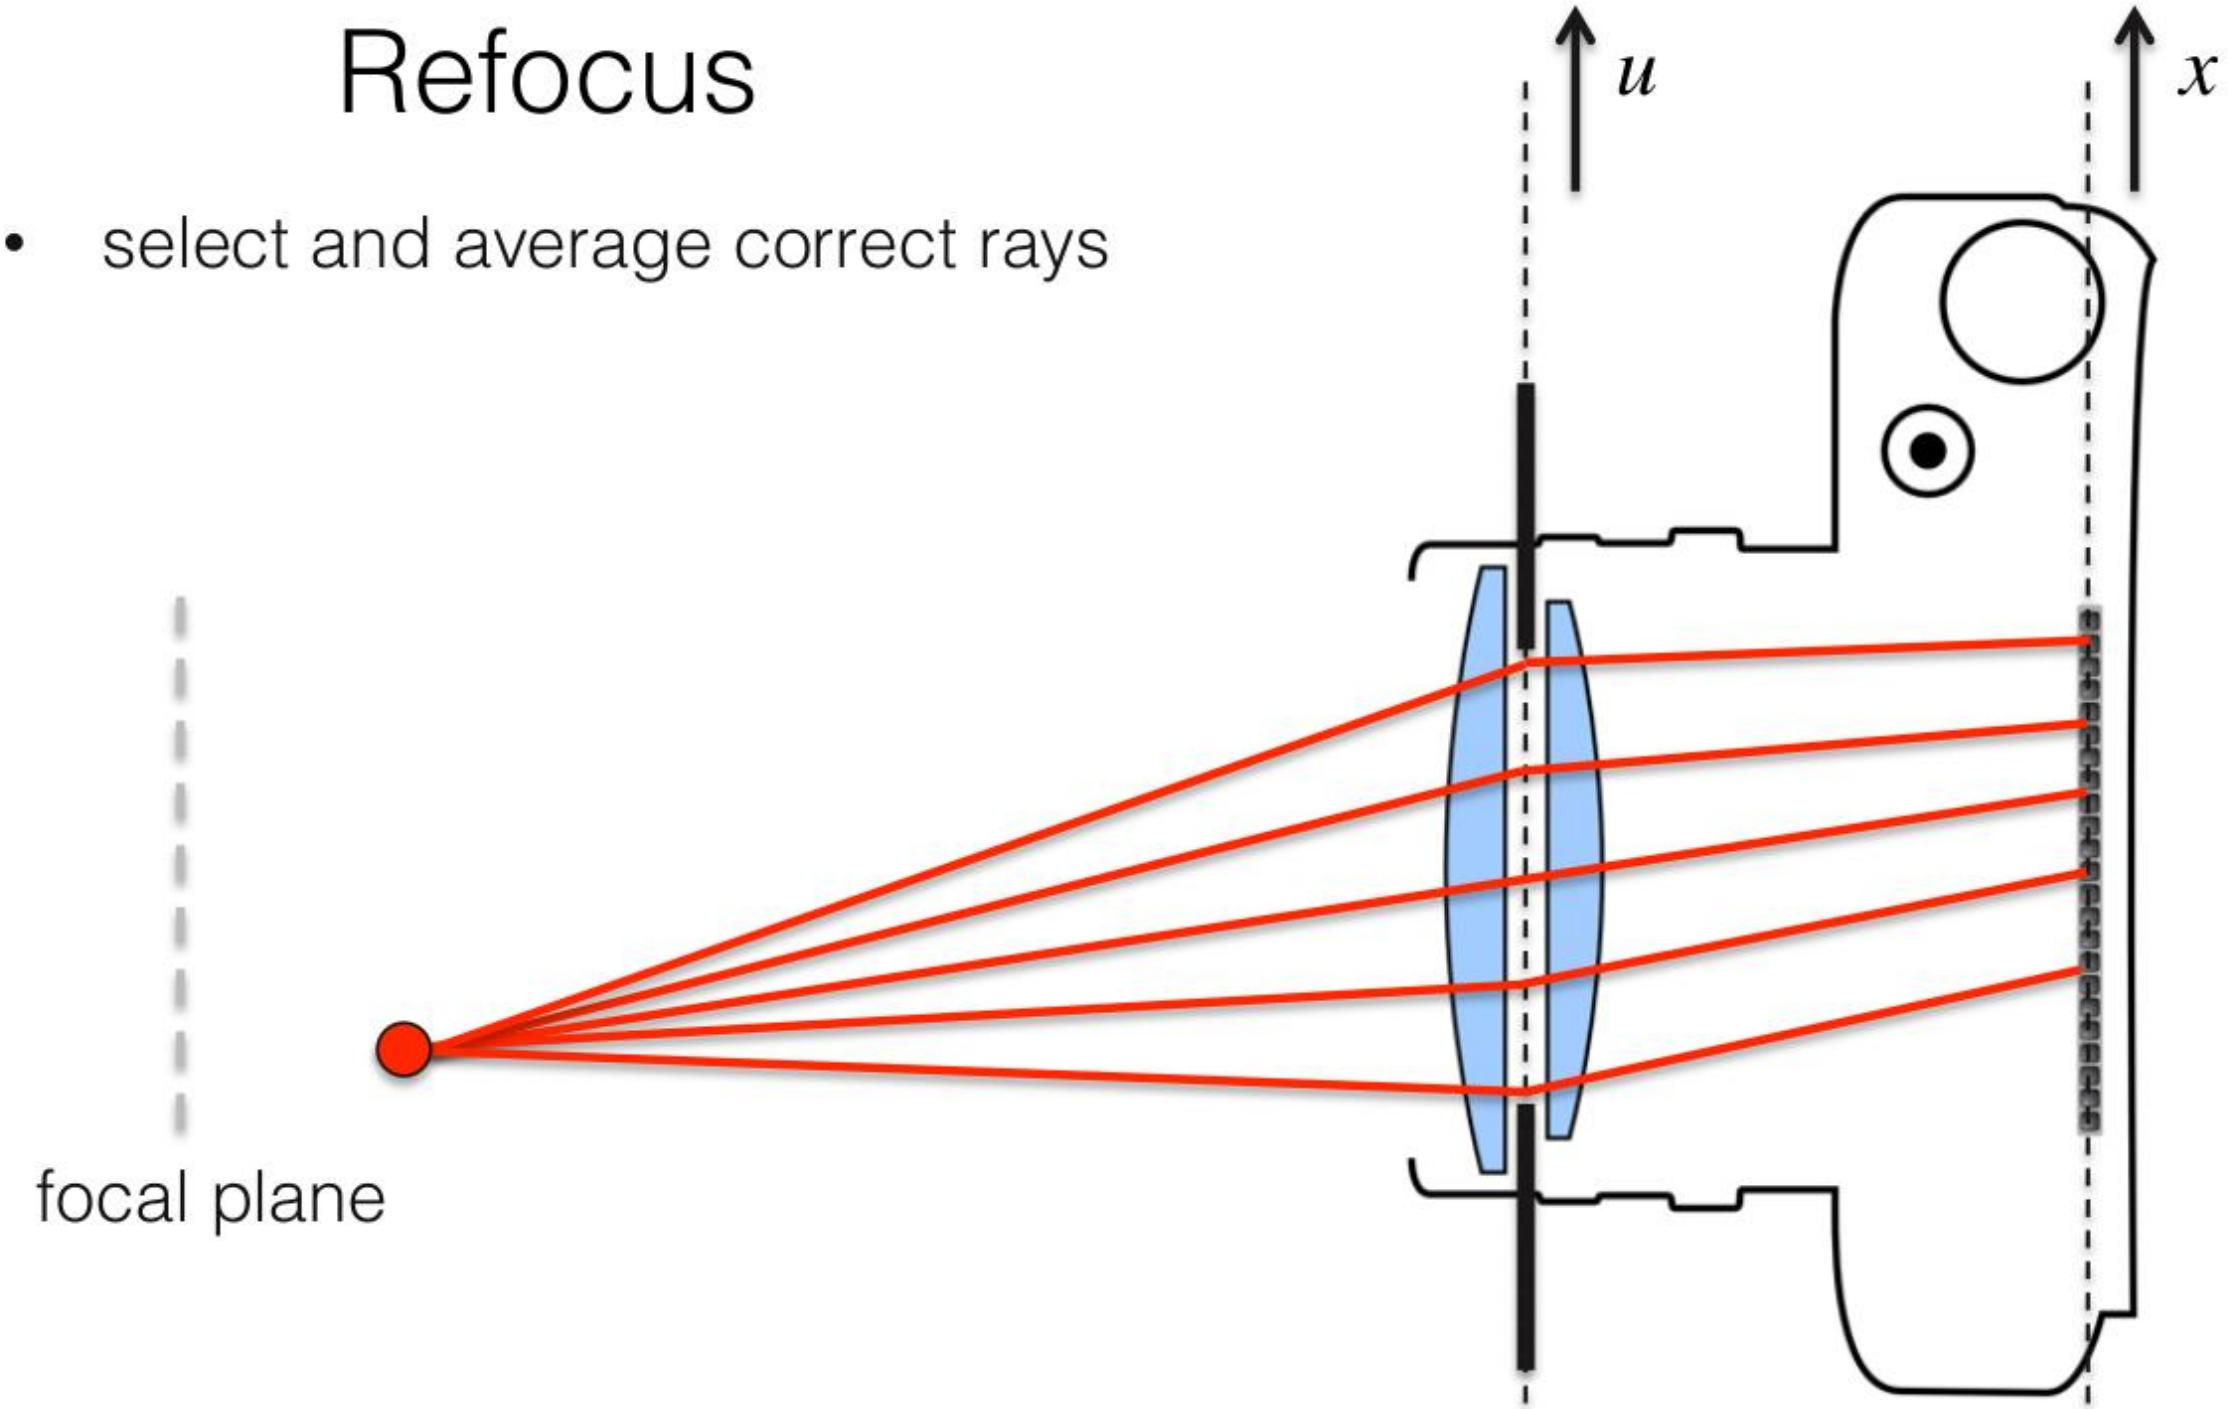

Understanding Refocus

- consider light field inside camera
- synthesize image on sensor $i_{d=0}(x) = \int_{\Omega} l(x,\upsilon) d\upsilon$

$$i_d(x) = \int_{\Omega} l(x + d\upsilon, \upsilon) d\upsilon$$

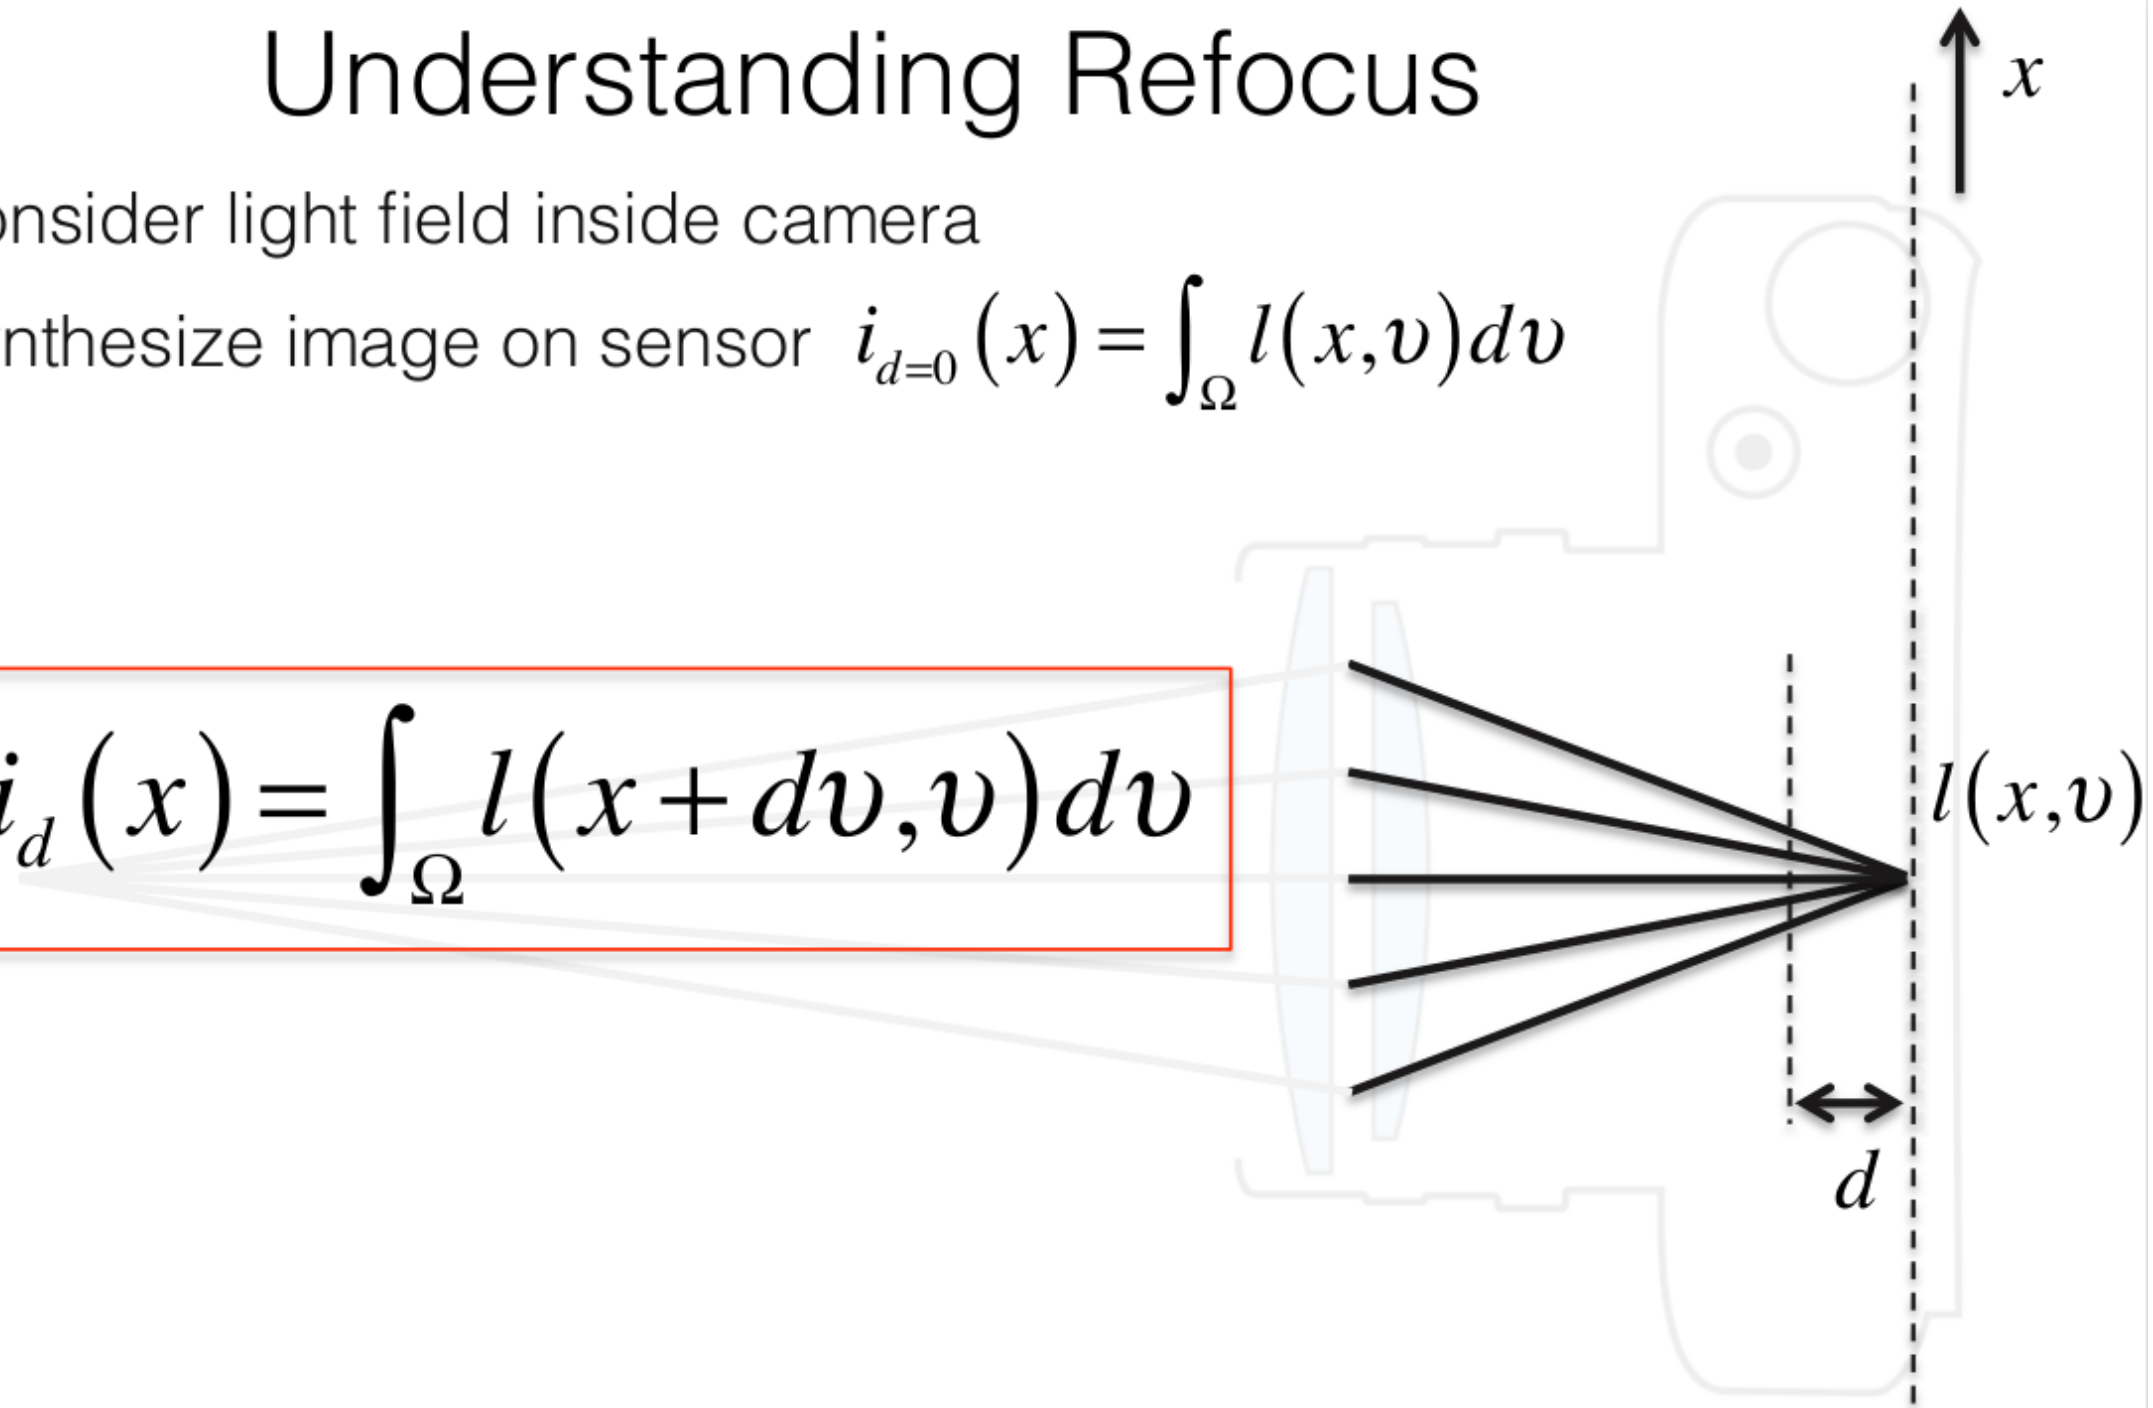

Understanding Refocus – Fourier Slicing

- Fourier slice theorem: projection in primal is slicing in Fourier space

$$i_d(x) = \int_\Omega l(x + dv, v)\,dv \quad \longleftrightarrow \quad \hat{i}_d(f_x) = \hat{l}(f_x, f_v - df_x)$$

primal domain

Fourier domain

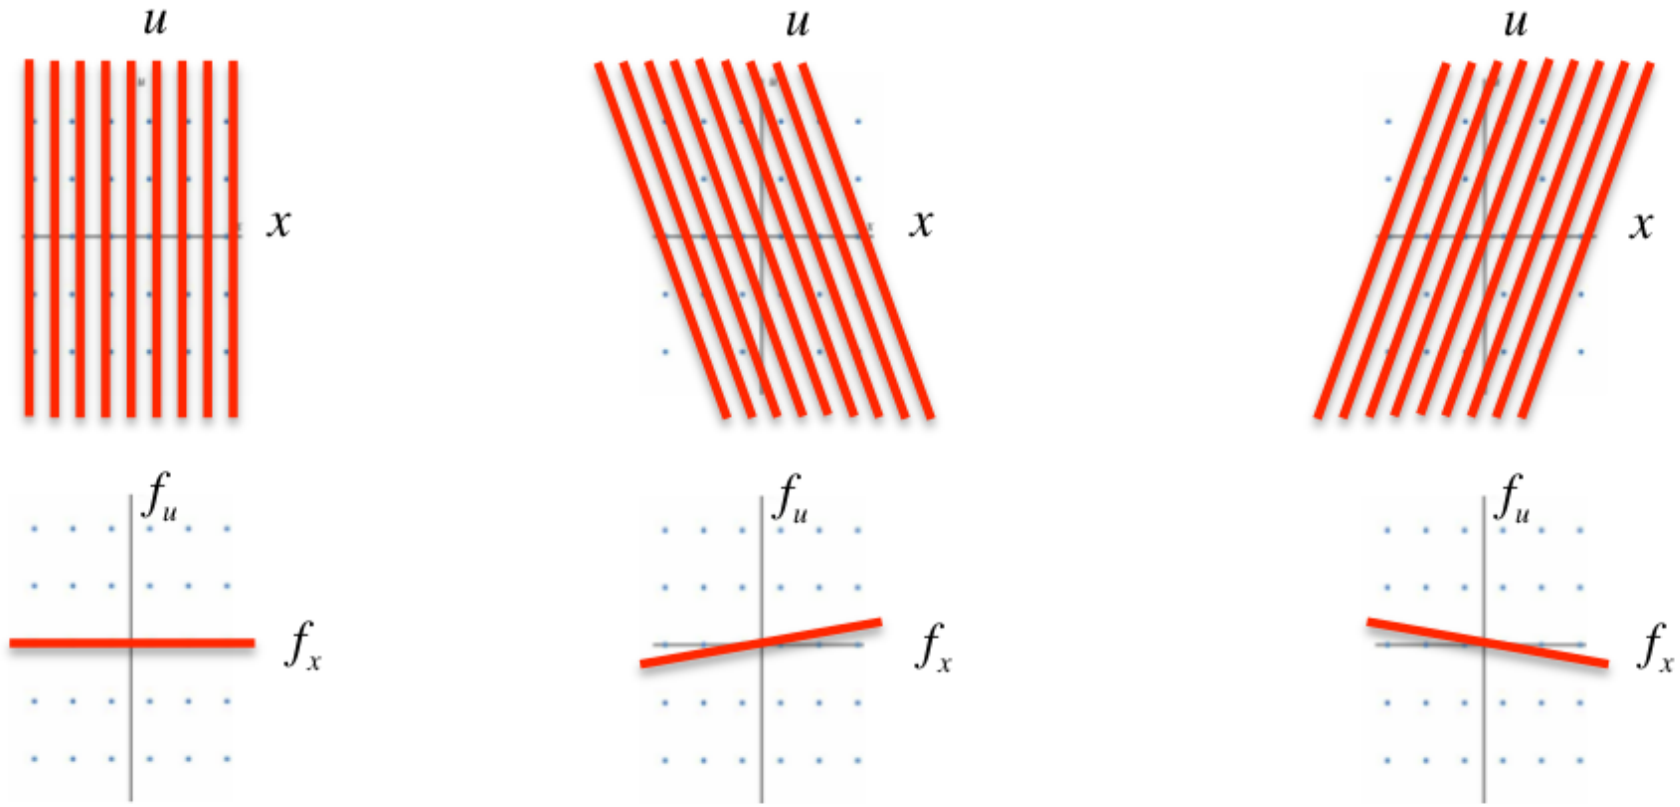

[Ng, SIGGRAPH 2006]

Refocusing example

Refocusing example

Refocusing example

Refocusing example

Many more examples with demo: http://lightfields.stanford.edu/

Three ways to measure a lightfield

1) Use a plenoptic camera

2) Use a camera array

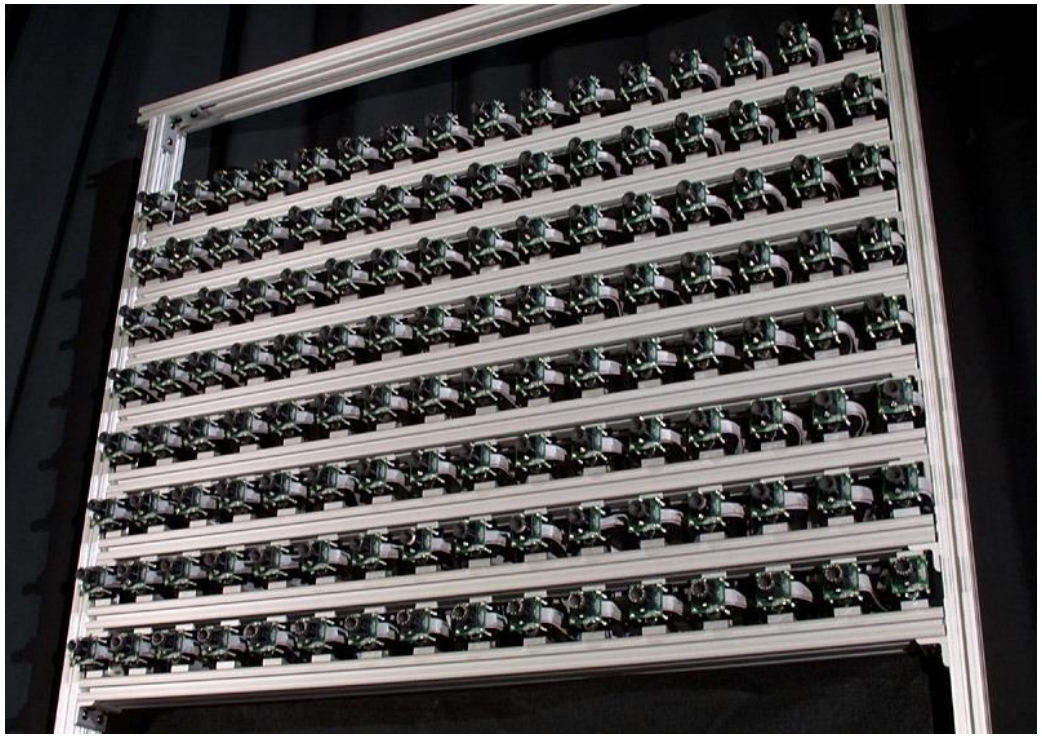

3) Use one camera multiple times

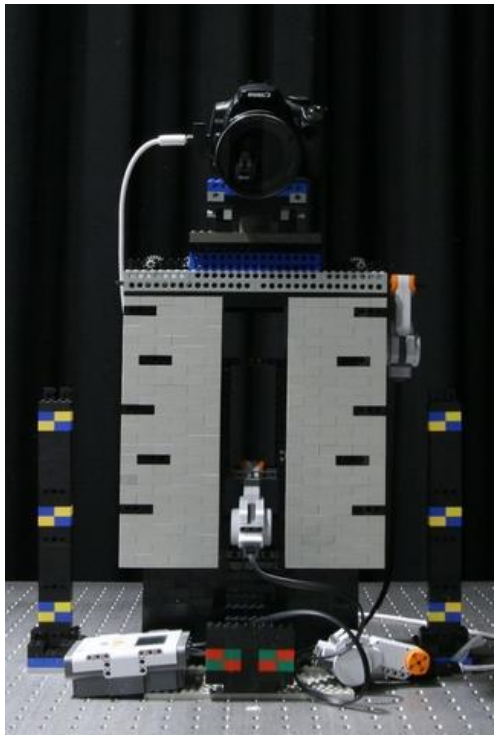

What are the pros and cons of each?

References

Basic reading:

- Szeliski textbook, Section 12.1.3, 13.3.

Additional reading:

- Nayar et al., “Real-Time Focus Range Sensor,” PAMI 1996.
  one of the (relatively) early papers on depth from defocus.
- Suwajanakorn et al., “Depth from Focus with Your Mobile Phone,” CVPR 2015.
  implementing depth from defocus on a mobile phone using the autofocus focal stack.
- Hasinoff and Kutulakos, “Confocal Stereo,” IJCV 2009.
  the paper on high resolution depth from a focus and aperture stack.
- Kuthirummal et al., “Flexible Depth of Field Photography,” PAMI 2010.
  continuously change focus within one exposure, without stopping to capture a stack.
- Levoy and Hanrahan, “Light Field Rendering,” SIGGRAPH 1996.
- Gortler et al., “The Lumigraph,” SIGGRAPH 1996.
  the two papers introducing the light field.
- Adelson and Wang, “Single Lens Stereo with a Plenoptic Camera,” PAMI 1992.
  the paper (re)-introducing the plenoptic camera to computer vision and graphics.
- Ng et al., “Light field photography with a hand-held plenoptic camera,” Stanford TR 2005.
  the paper (re)-(re)-introducing the plenoptic camera, and the precursor to Lytro.
- Ng, “Fourier Slice Photography,” SIGGRAPH 2005.
  the paper on frequency-space analysis of refocusing and lightfield measurements.
- Wilburn et al., “High Performance Imaging Using Large Camera Arrays,” SIGGRAPH 2005.
  the camera array paper.
- Georgiev et al., “Unified Frequency Domain Analysis of Lightfield Cameras,” ECCV 2008.
  make a lightfield camera from a pinhole array (and many other interesting stuff about lightfield cameras).