Lecture 25:

The Light Field and Computational Cameras

Computer Graphics

CMU 15-462/15-662, Spring 2016

Credit: light-field camera slides courtesy of Ren Ng

Takeaway from today's lecture

The values of pixels in photographs you see on screen are quite different than the values output by the photosensor in a modern digital camera.

Computation is now a fundamental aspect of producing high-quality good pictures.

Output of sensor ("RAW") → **Computation** →

Beautiful image

Camera's lens system focuses light on sensor

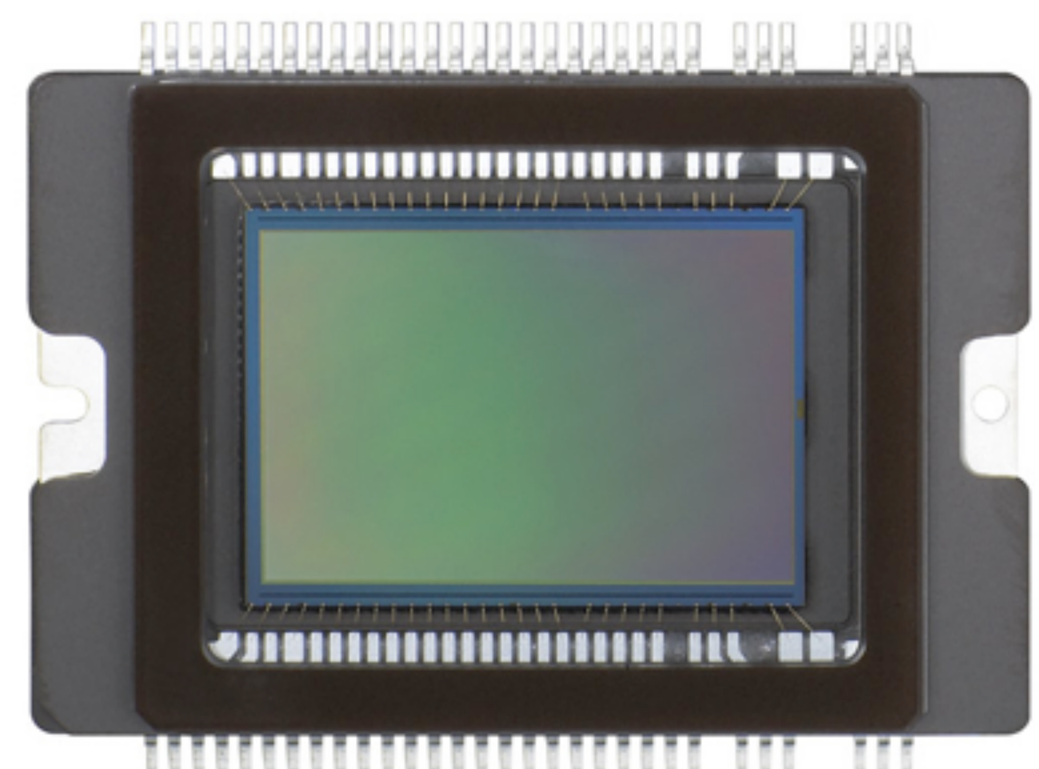

Canon 14 MP CMOS Sensor
(14 bits per pixel)

Image credit: Canon

CMOS sensor pixel

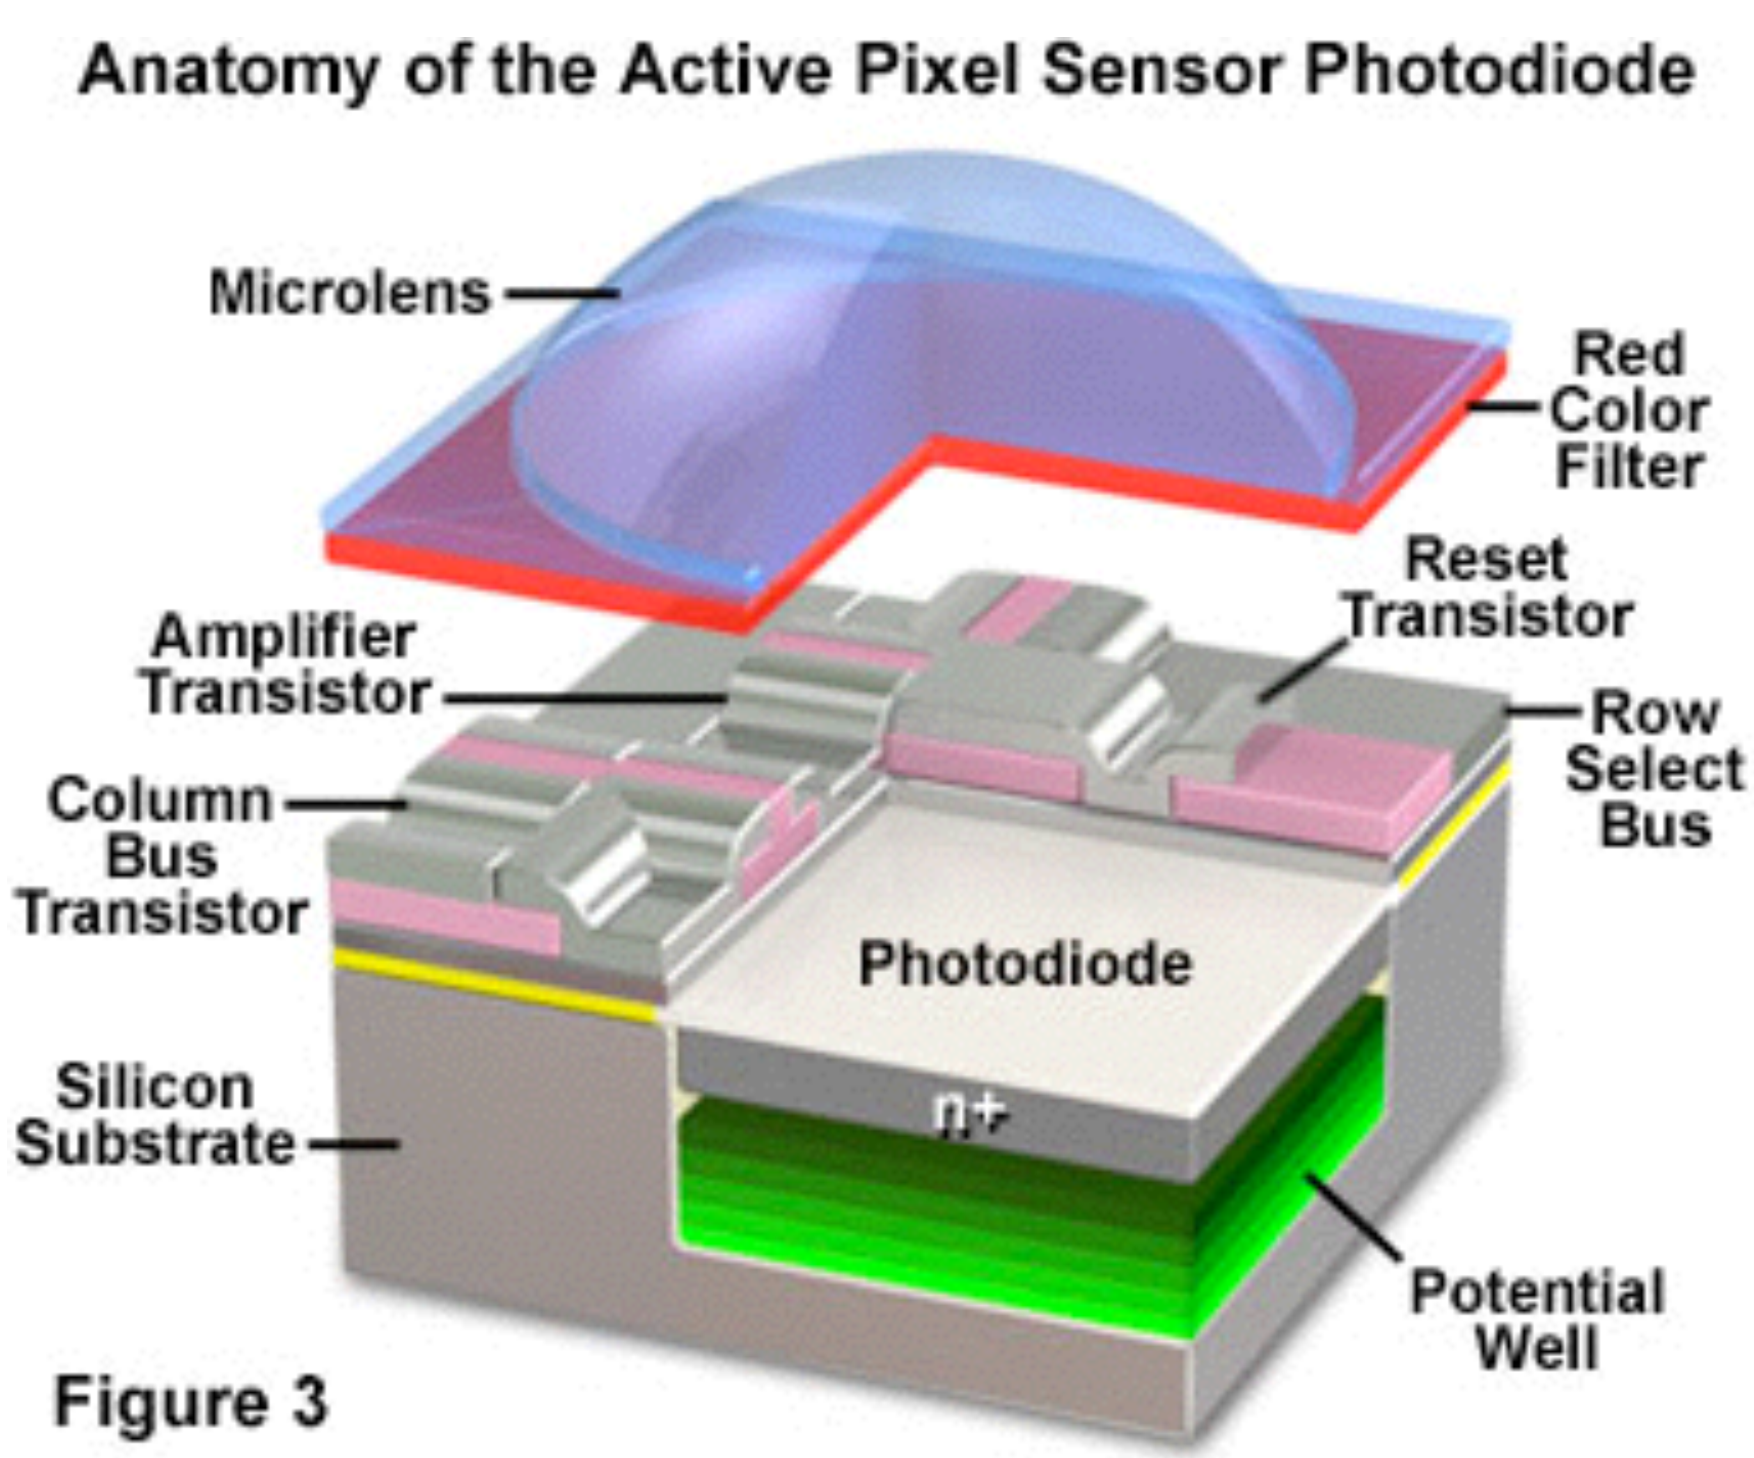

Color filter attenuates light (more on this soon)

Fill factor: fraction of surface area used for light gathering

Microlens (a.k.a. lenslet) steers light toward photo-sensitive region (increases light-gathering capability)

Microlens also serves to prefilter signal. Why?

Quantum efficiency of photodiode in typical digital camera ~ 50%

Illustration credit: Molecular Expressions (http://micro.magnet.fsu.edu/primer/digitalimaging/cmosimagesensors.html)

Reading sensed signal off sensor

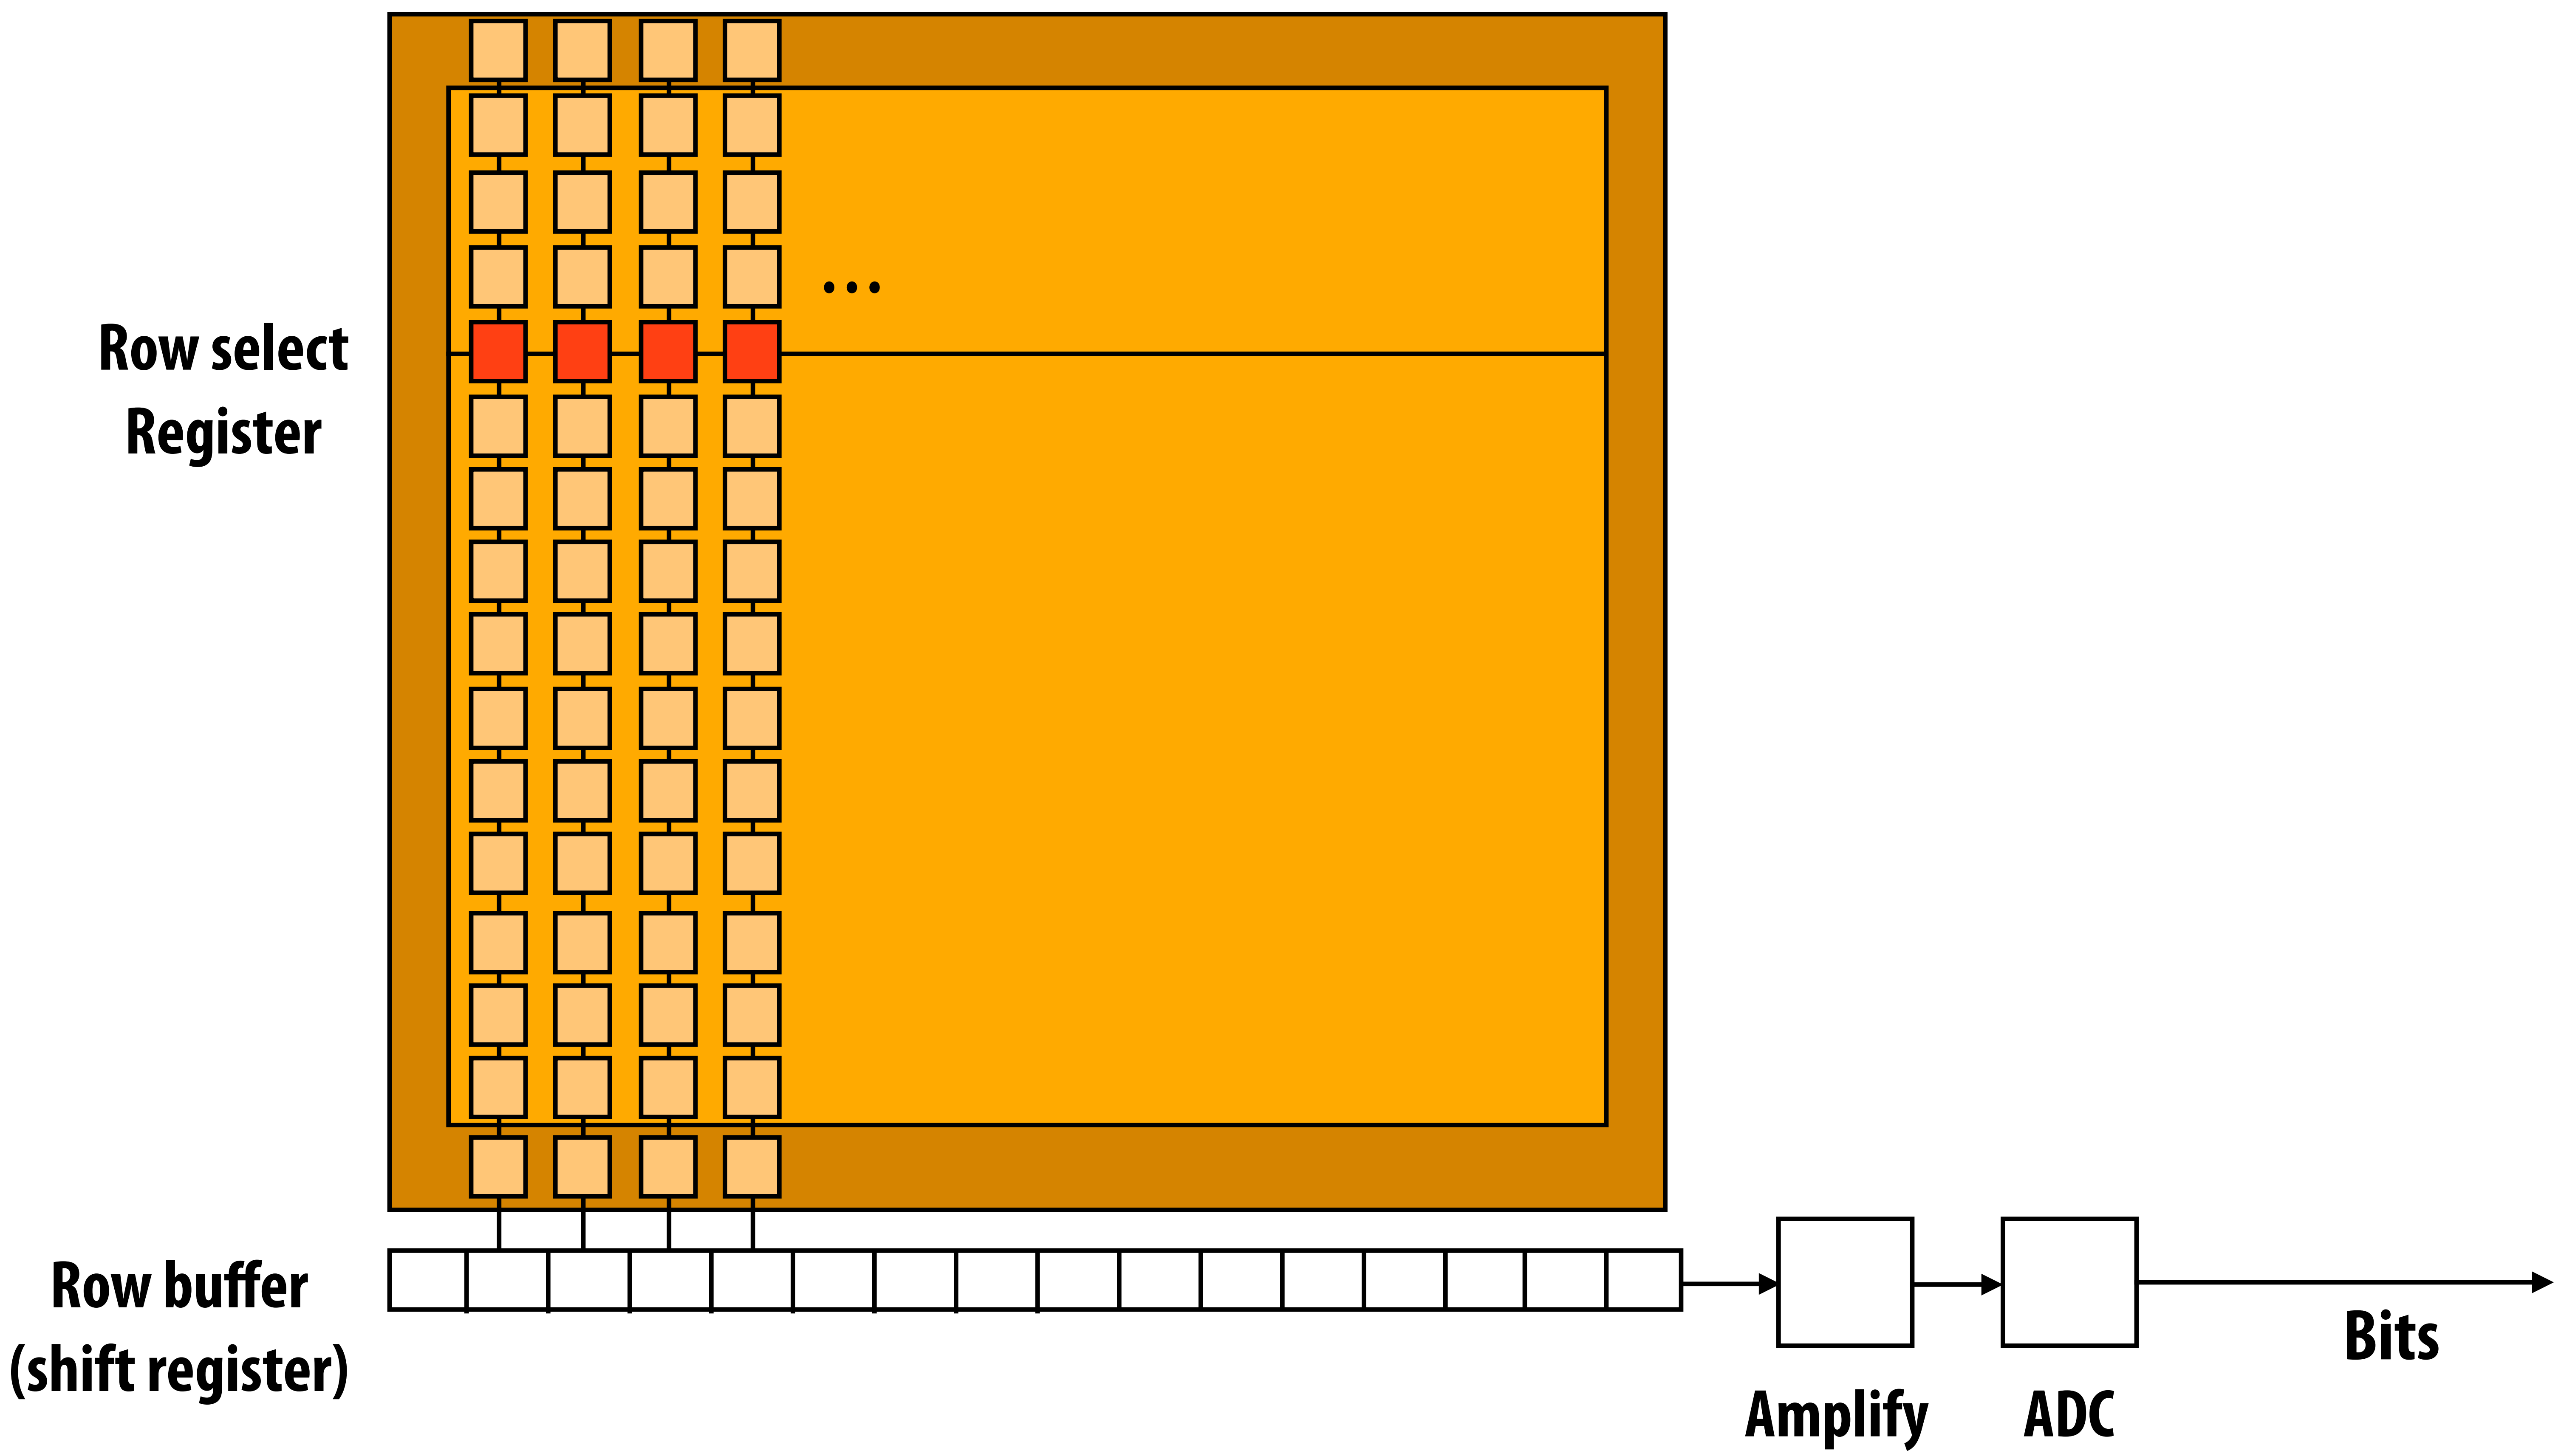

Measurement noise

- **Photon shot noise:**
  - Photon arrival rates feature poisson distribution
  - Standard deviation = sqrt(N)
  - Signal-to-noise ratio: N/sqrt(N)
- **Dark-shot noise**
  - Due to leakage current
- **Non-uniformity of pixel sensitivity**
- **Read noise**
  - e.g., due to amplification prior to ADC

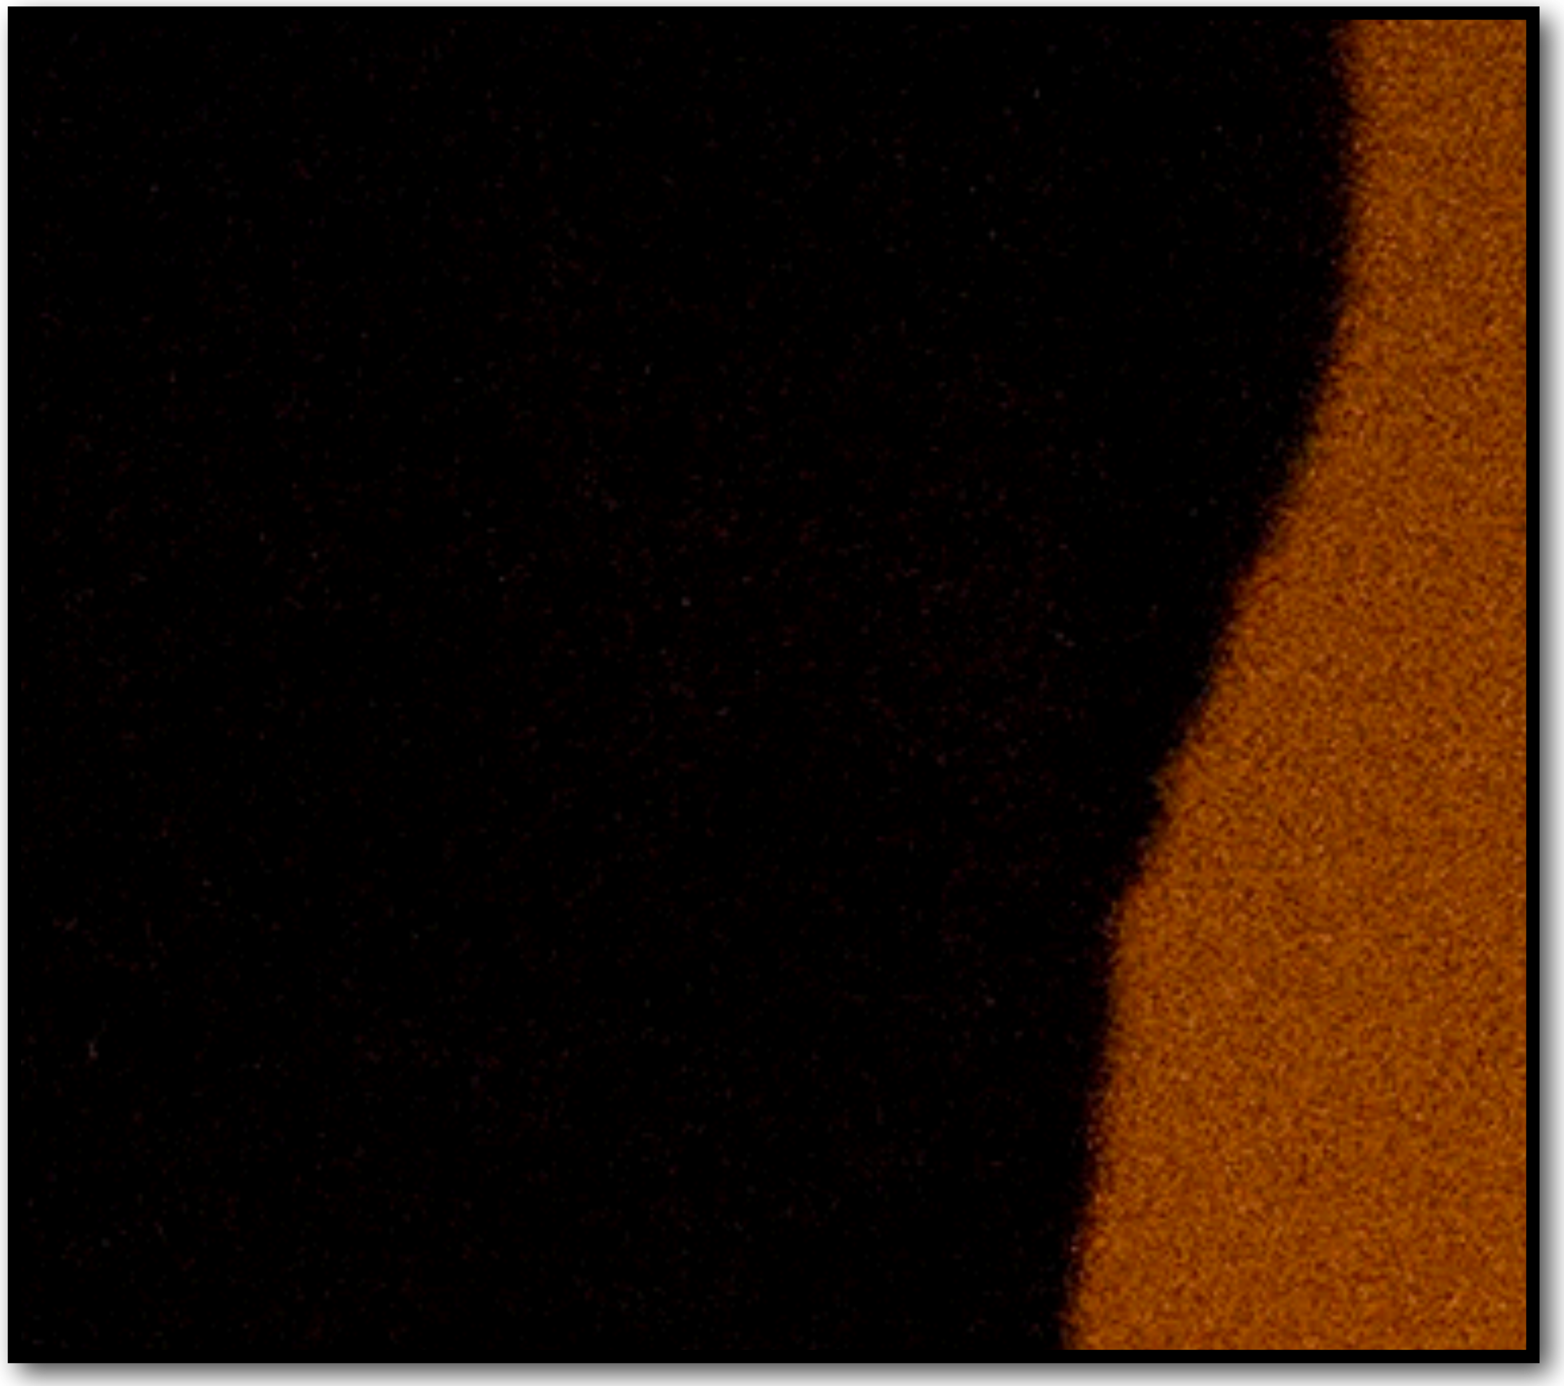

Illustration credit: Molecular Expressions (http://micro.magnet.fsu.edu/primer/digitalimaging/cmosimagesensors.html)

Noise

Black image examples: Nikon D7000, High ISO

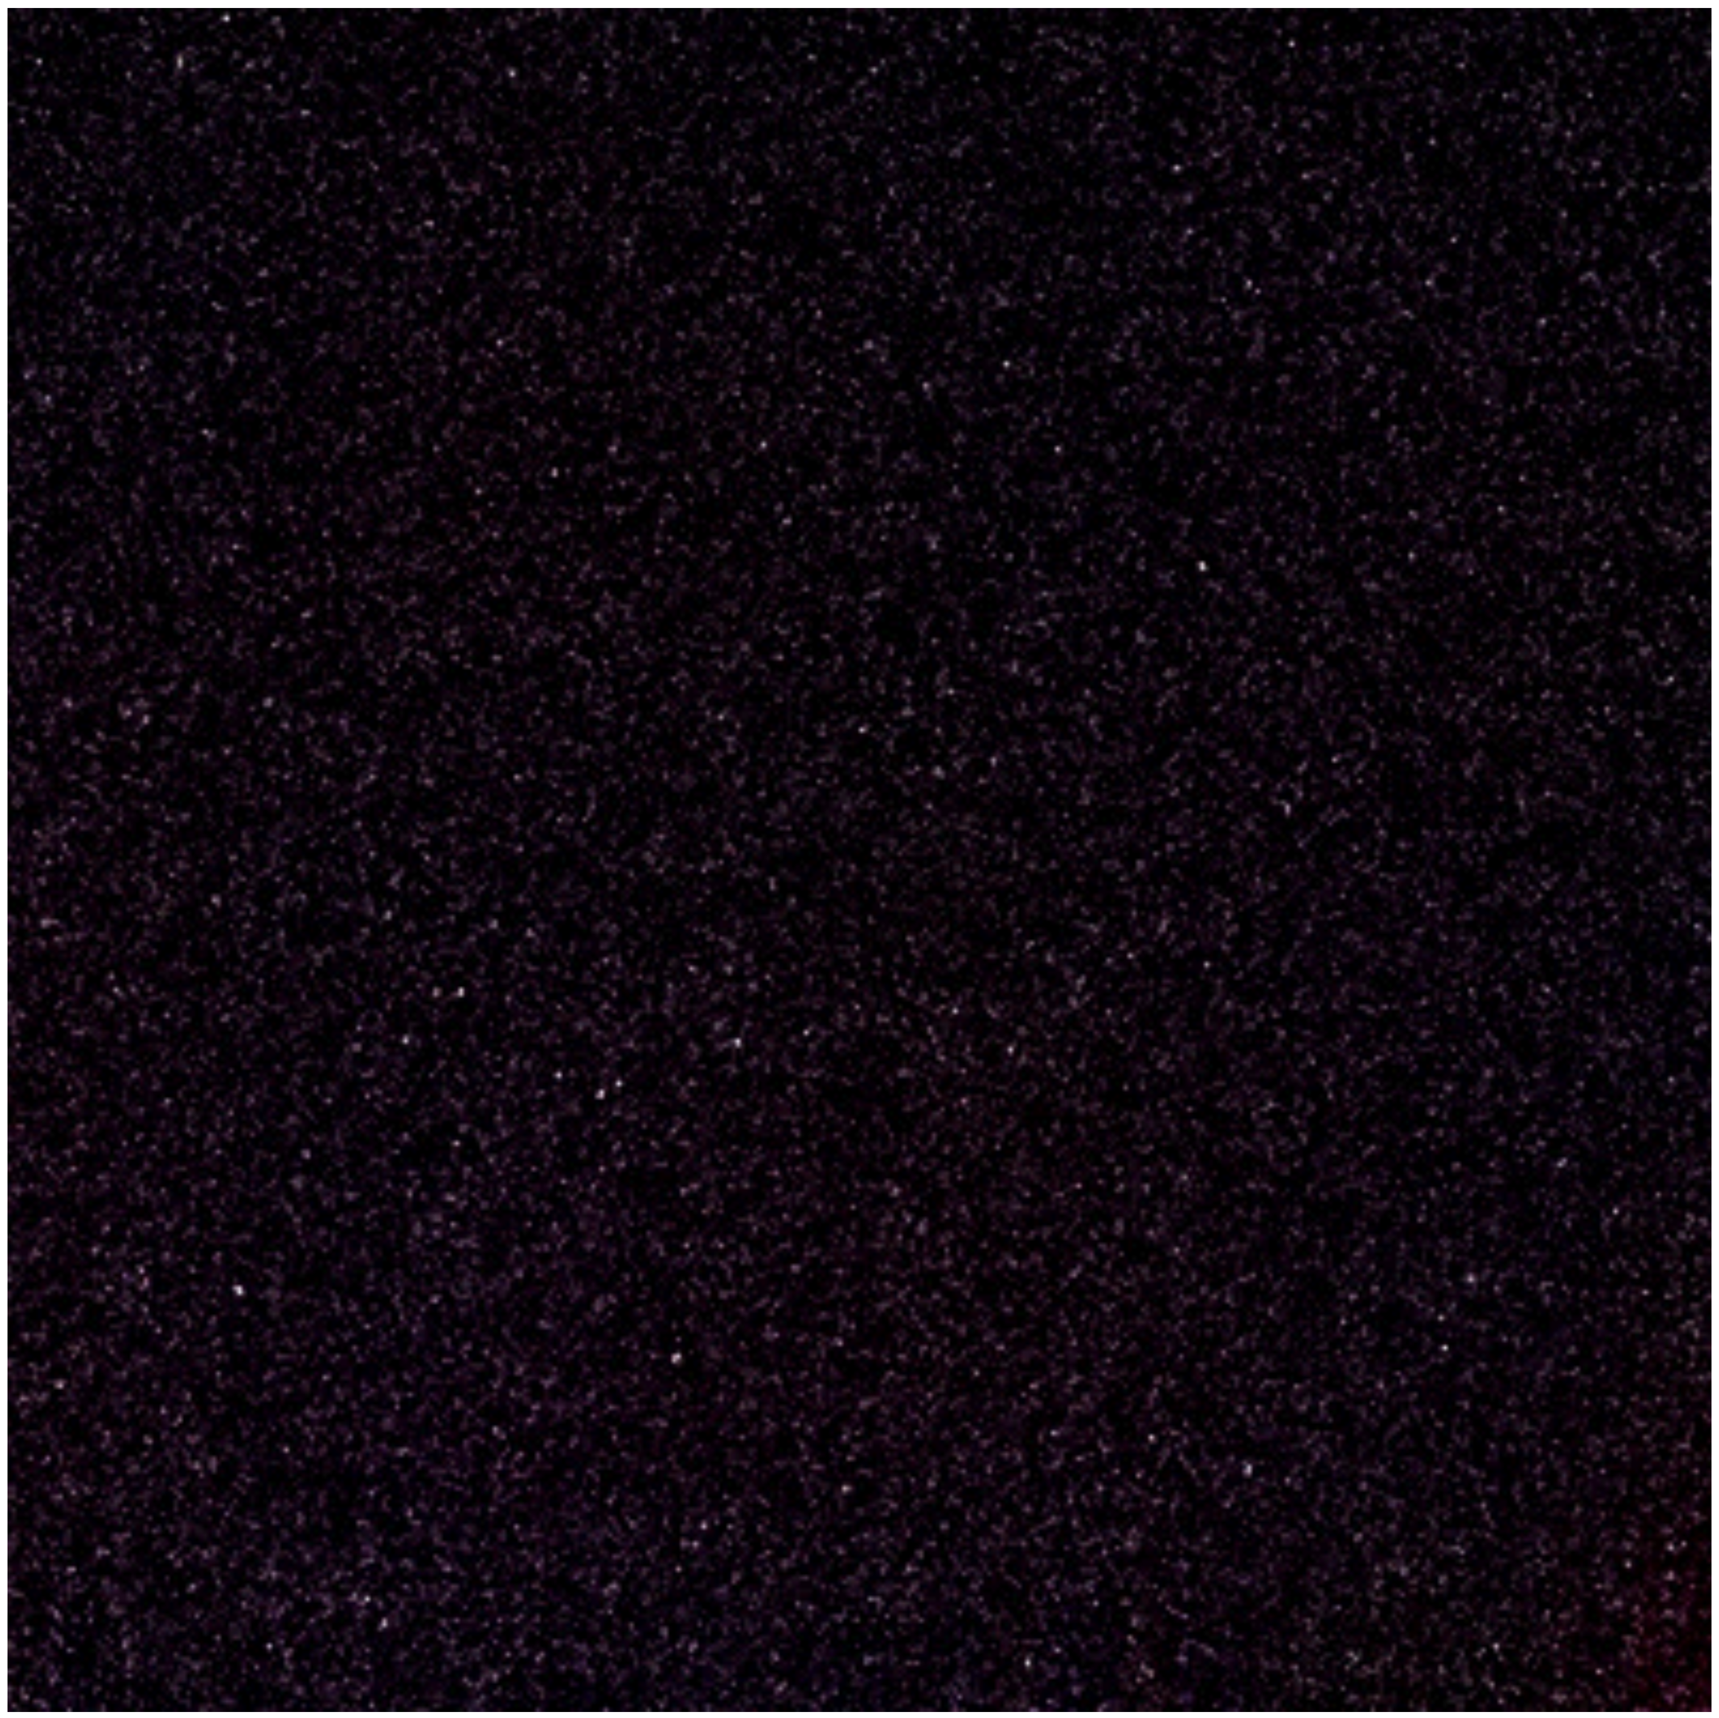

1/60 sec exposure

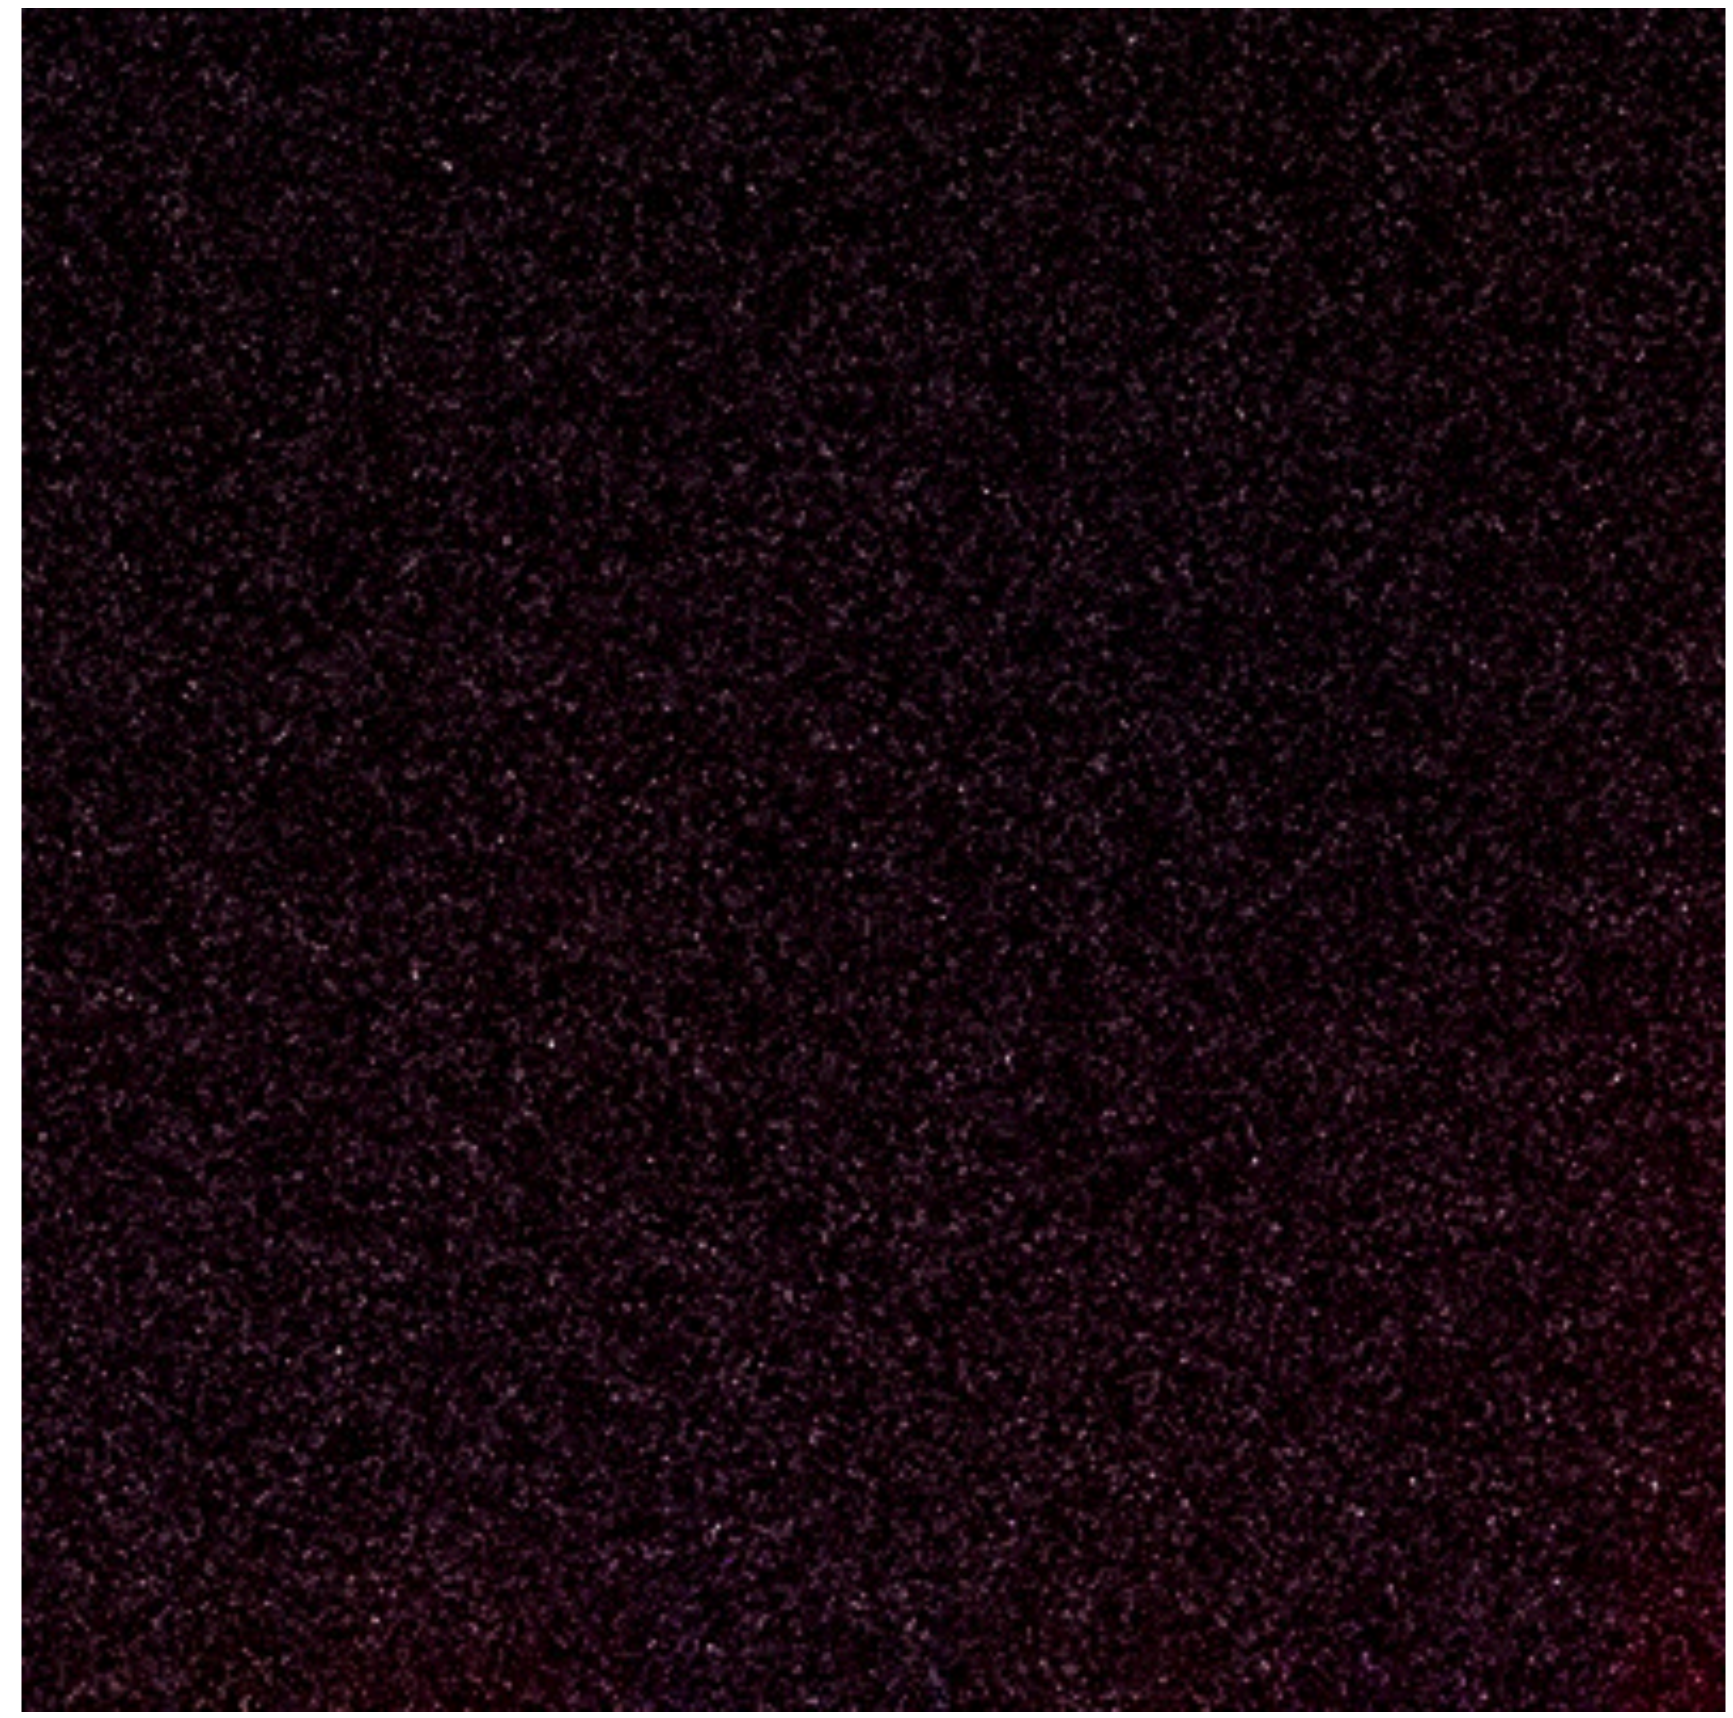

1 sec exposure

Color filter array (Bayer mosaic)

- Color filter array placed over sensor
- Result: each pixel measures red, green, or blue light
- 50% of pixels are green pixels

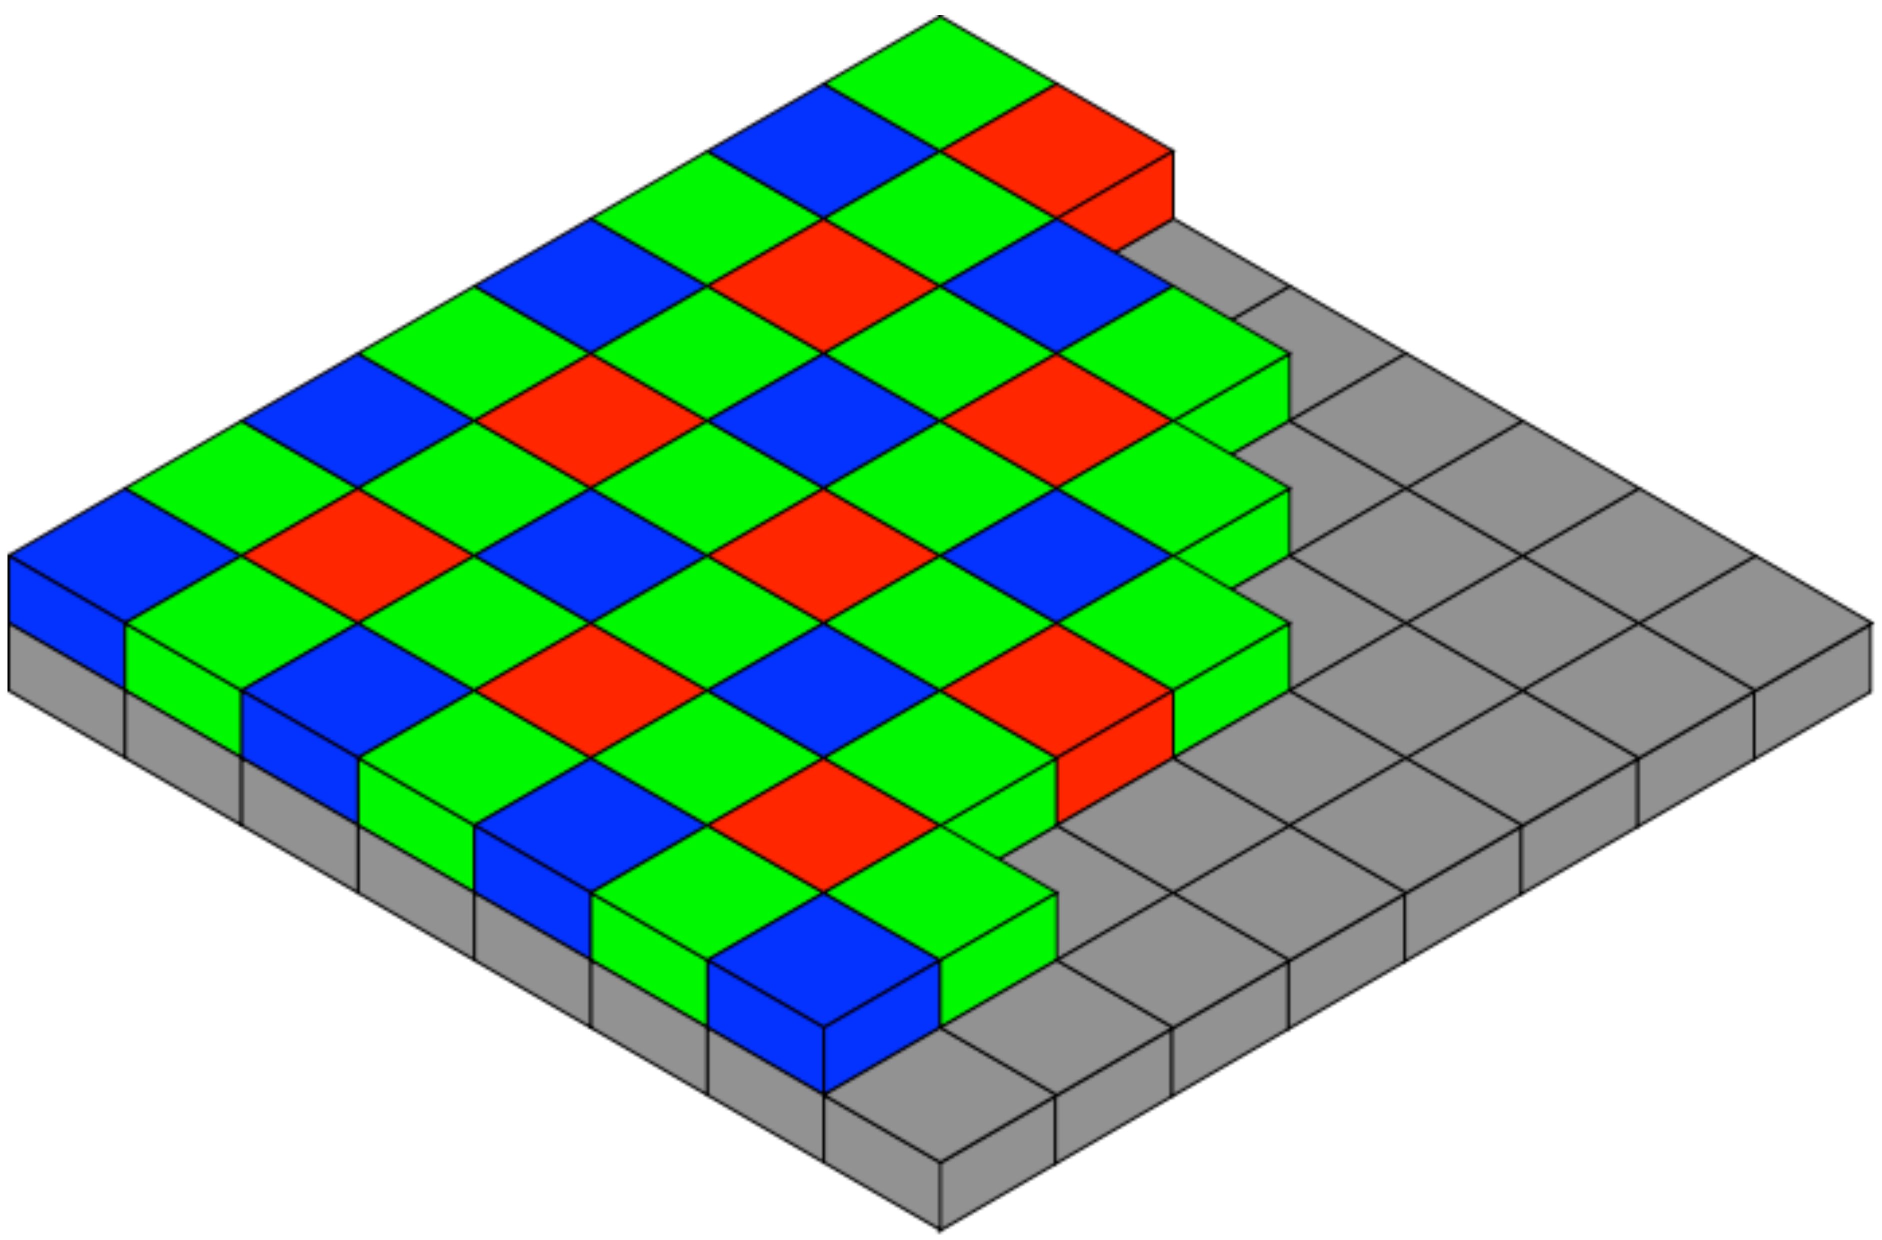

Traditional Bayer mosaic
(other filter patterns exist: e.g., Sony's RGBE)

Pixel response curve: Canon 40D/50D

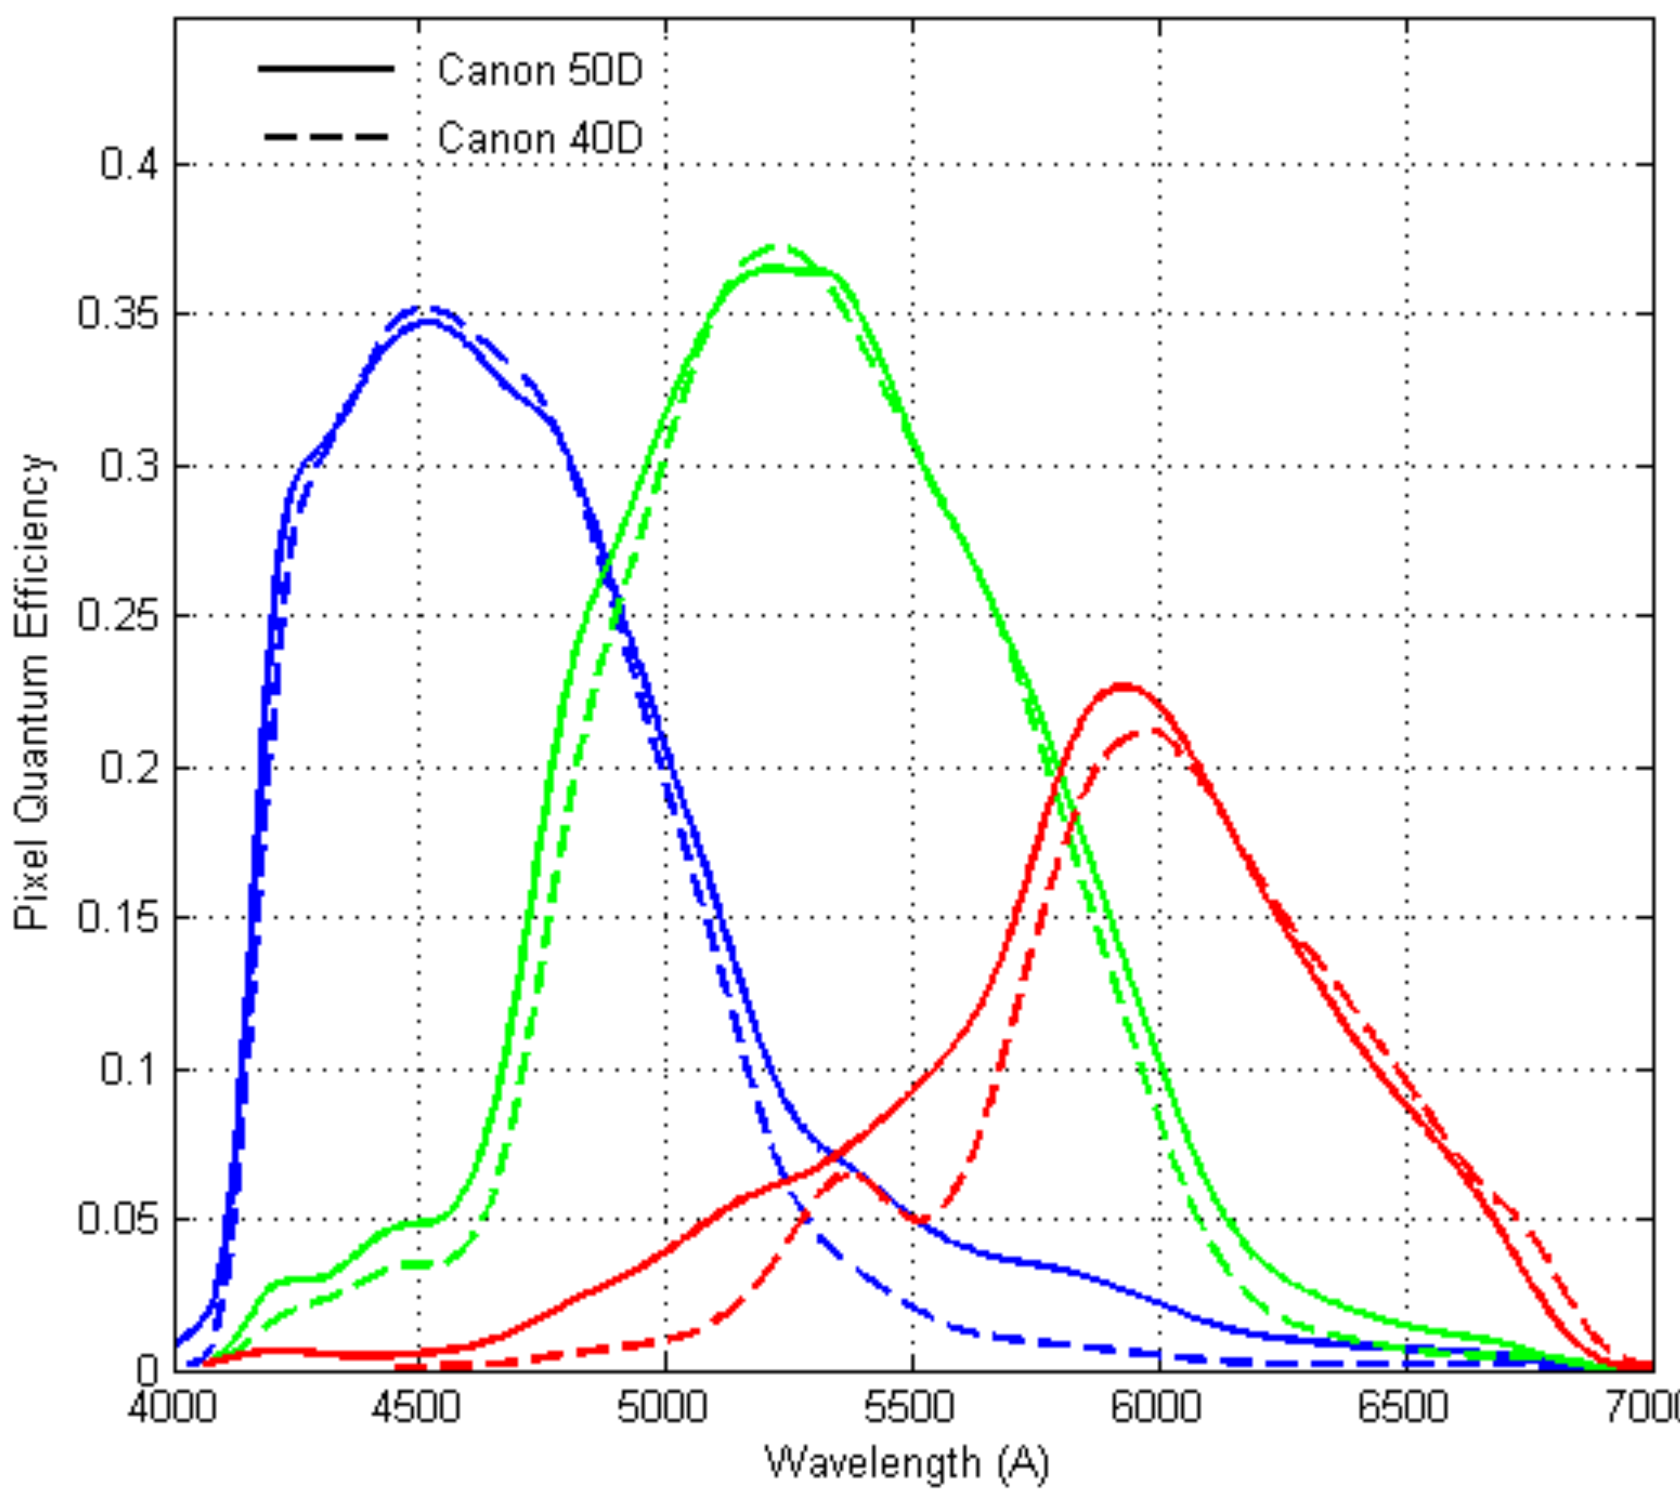

Image credit:
Wikipedia, Christian Buil (http://www.astrosurf.com/~buil/cameras.htm)

Demosiac

- Produce RGB image from mosaiced input image
- Basic algorithm: bilinear interpolation of mosaiced values (need 4 neighbors)
- More advanced algorithms:
  - Bicubic interpolation (wider filter support region… may overblur)
  - Good implementations attempt to find and avoid interpolation across edges

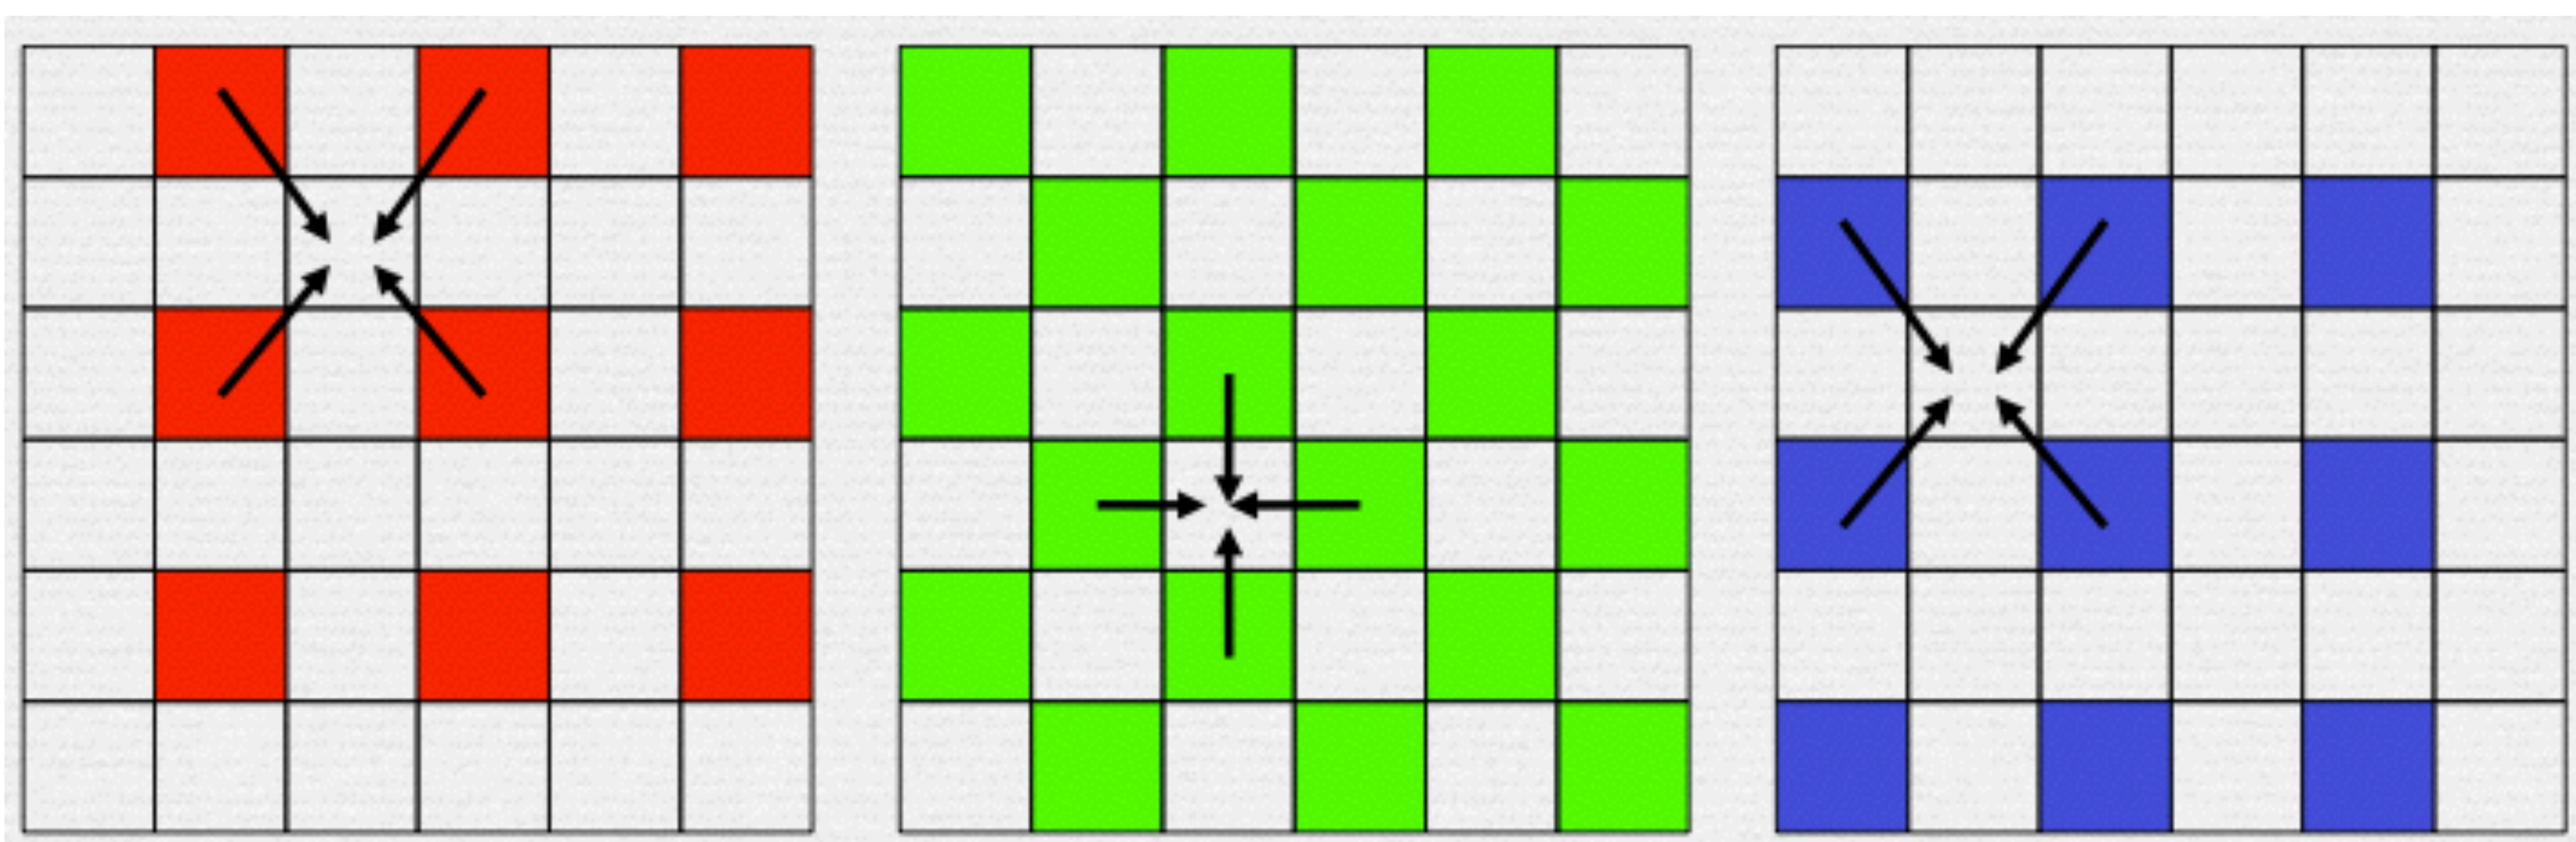

Image credit: Mark Levoy

RAW sensor output (simulated data)

Light Hitting Sensor

RAW output of sensor

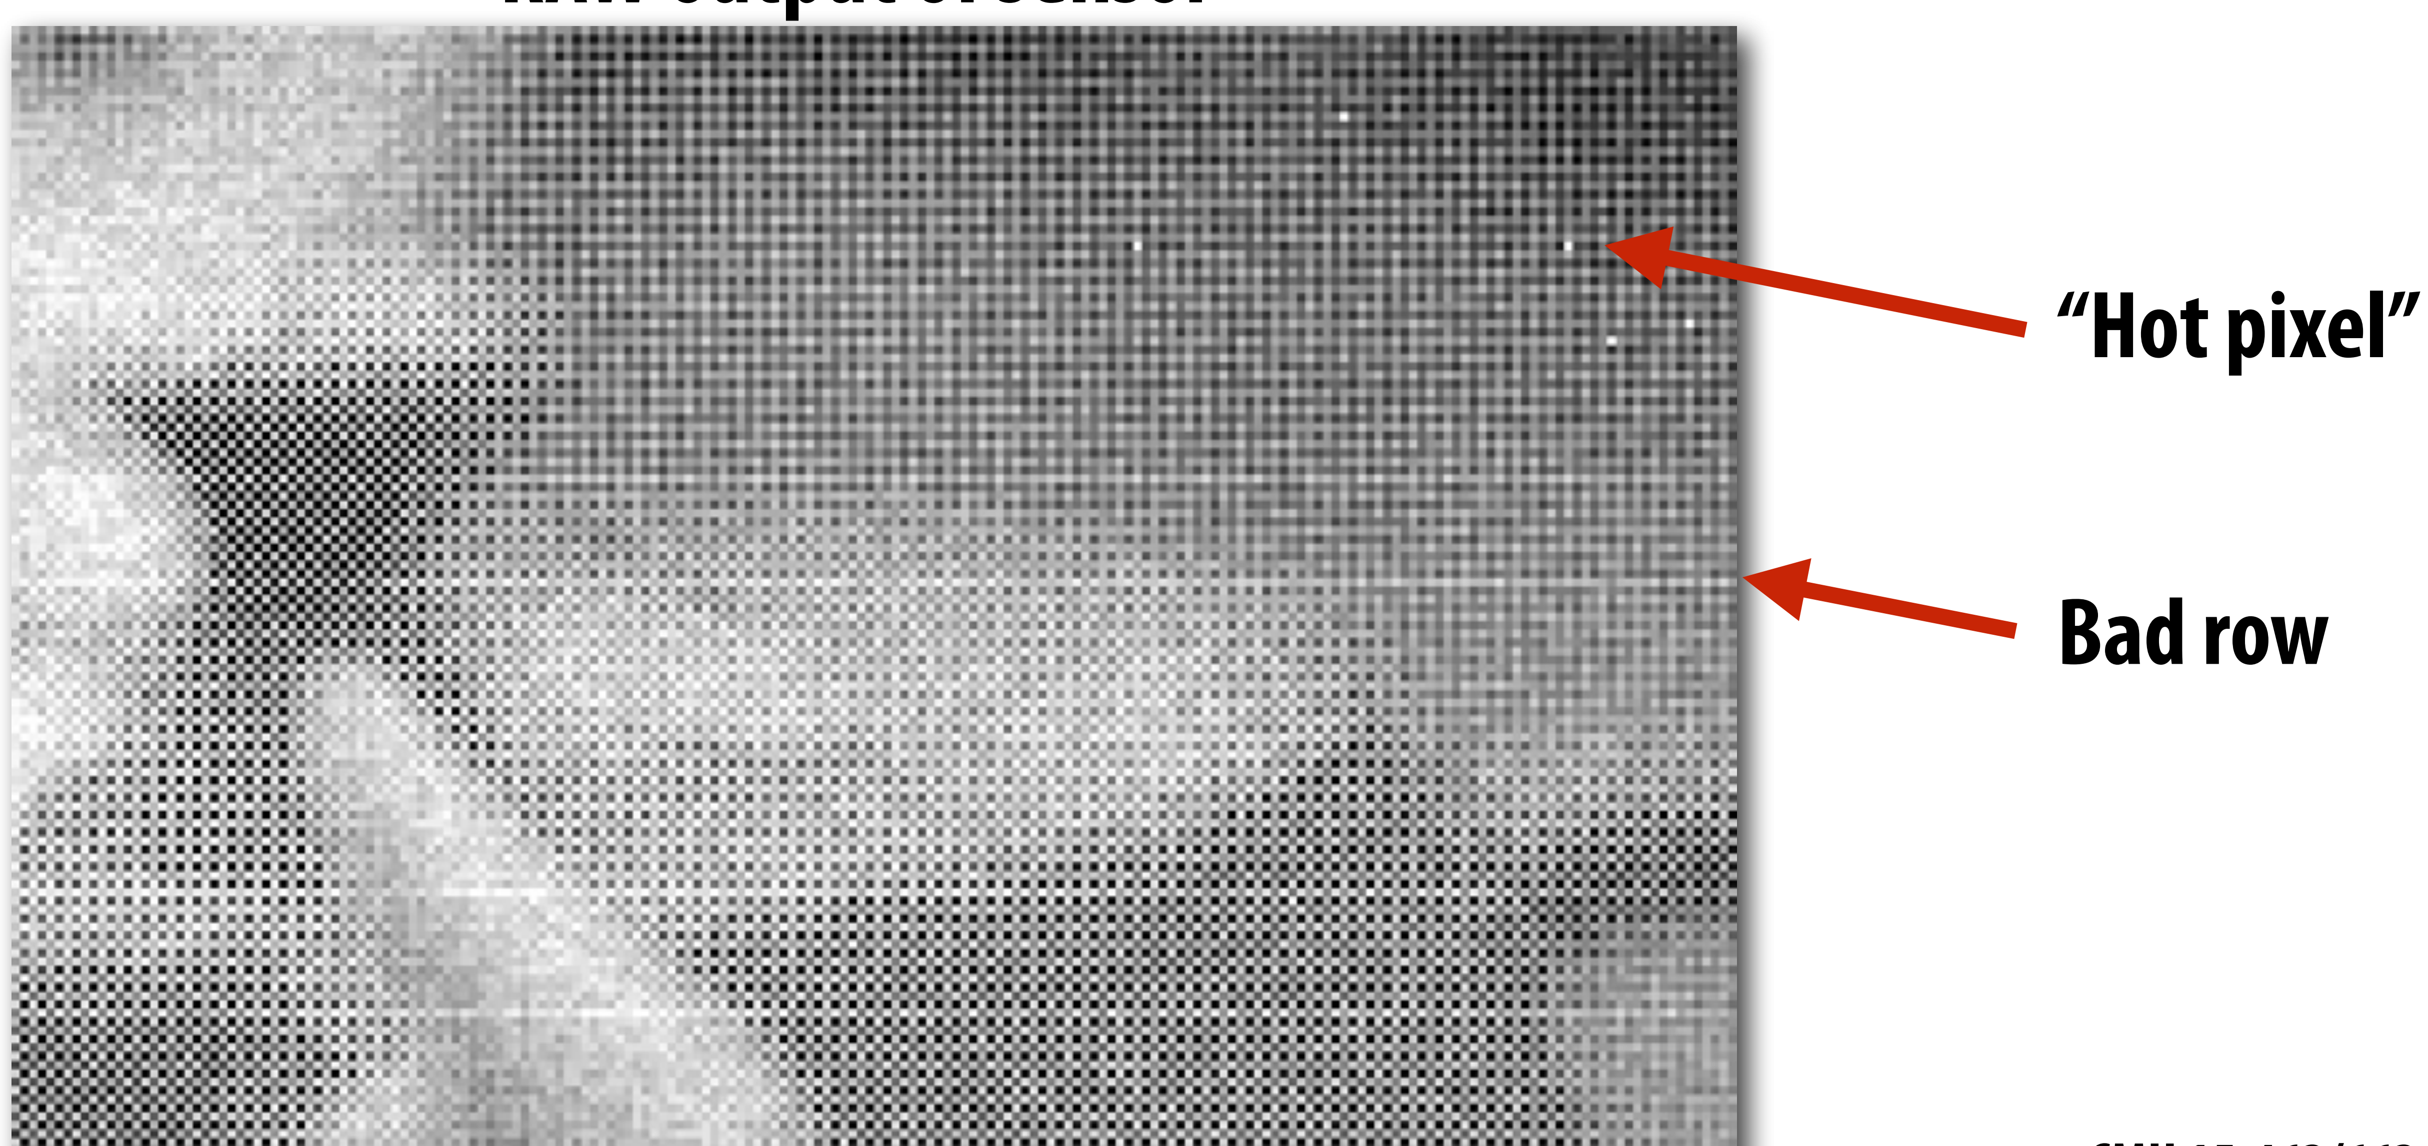

Other corrections

- Dead pixels
- Vignetting (optical and pixel)
- Lens distortion
- …

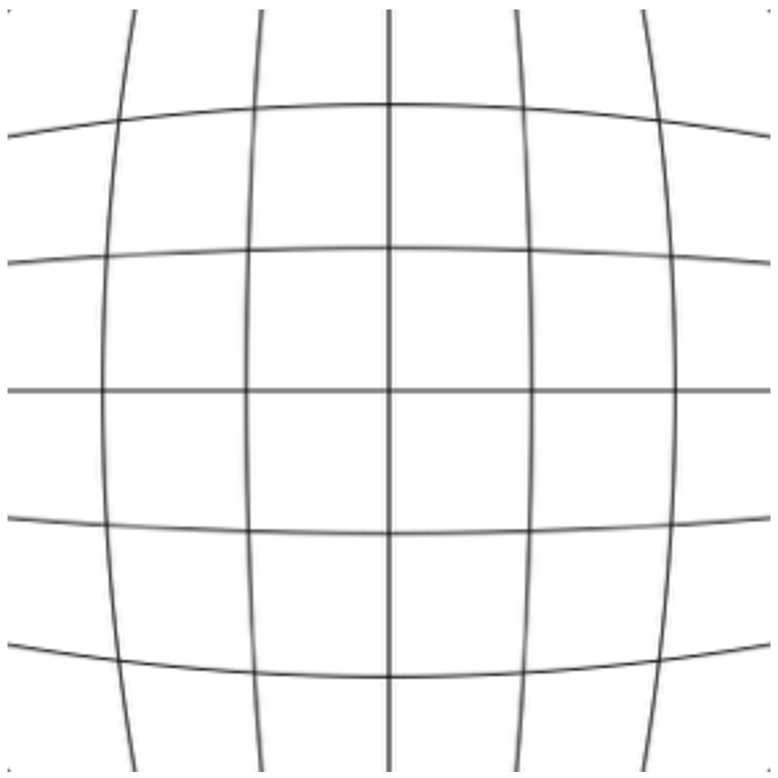

Barrel distortion

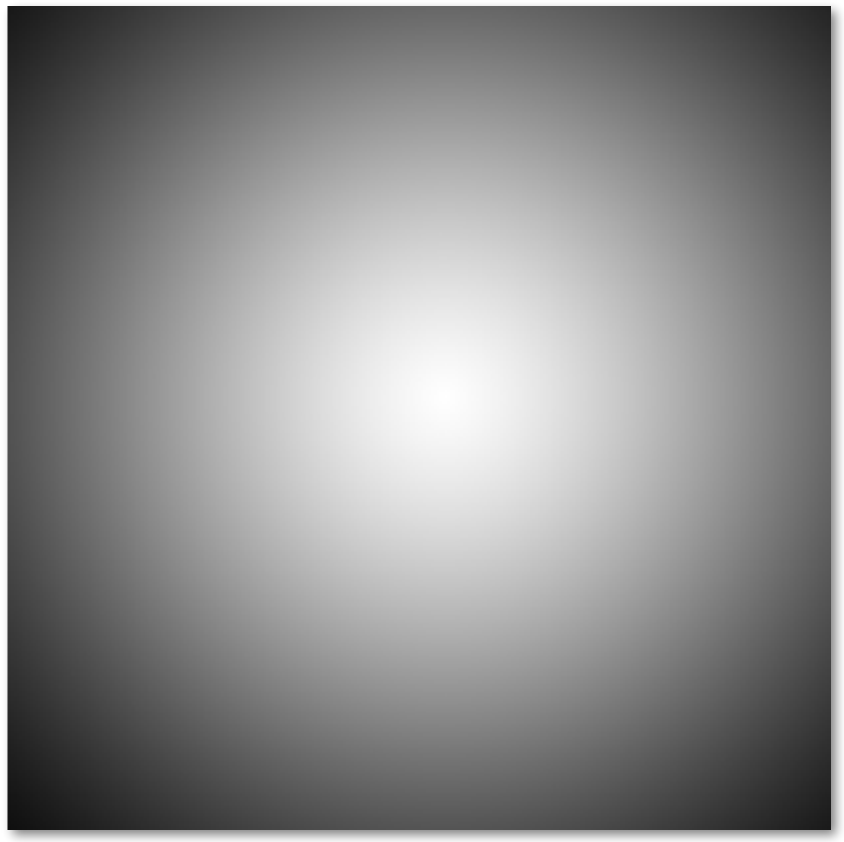

Vignetting
(less light at corner pixels)

Takeaway: a complex pipeline of image processing operations is required to convert the output of a sensor to an image that you'd actually want to see.

Light L16 camera

16 synchronized cameras: allows for varying exposure, zoom, focal depth, etc..

[Image credit: http://www.theverge.com/2015/10/7/9473793/light-l16-16-lens-camera-specs-price]

The light field

[Levoy and Hanrahan 96]

[Gortler et al., 96]

Light-field parameterization

Light field is a 4D function (represents light in free space: no occlusion)

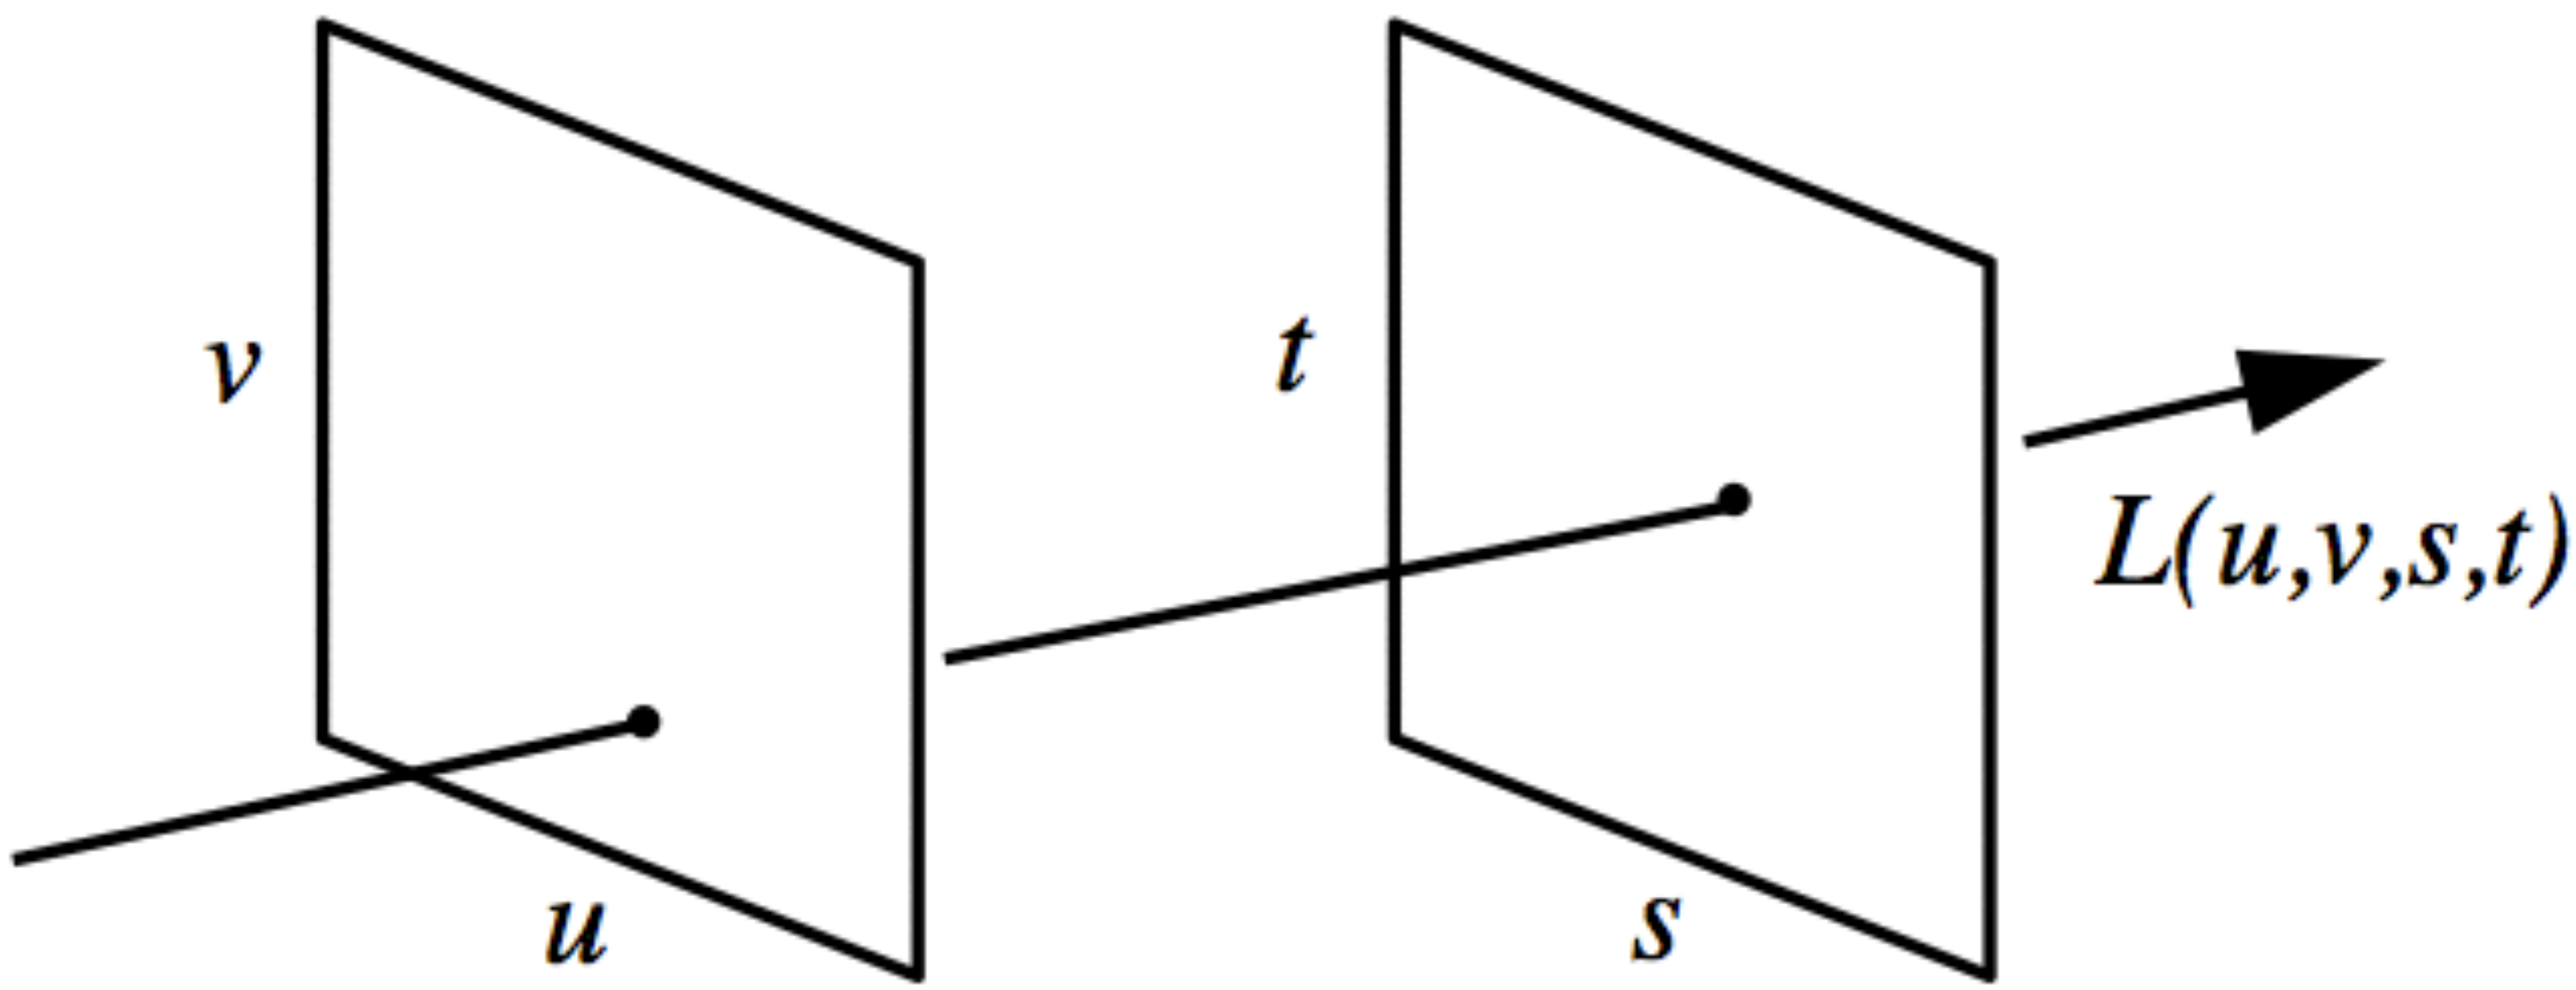

[Image credit: Levoy and Hanrahan 96]

Efficient two-plane parameterization

Line described by connecting point on (u,v) plane with point on (s,t) plane

If one of the planes placed at infinity: point + direction representation

Levoy/Hanrahan refer to representation as a "light slab": beam of light entering one quadrilateral and exiting another

Sampling the light field

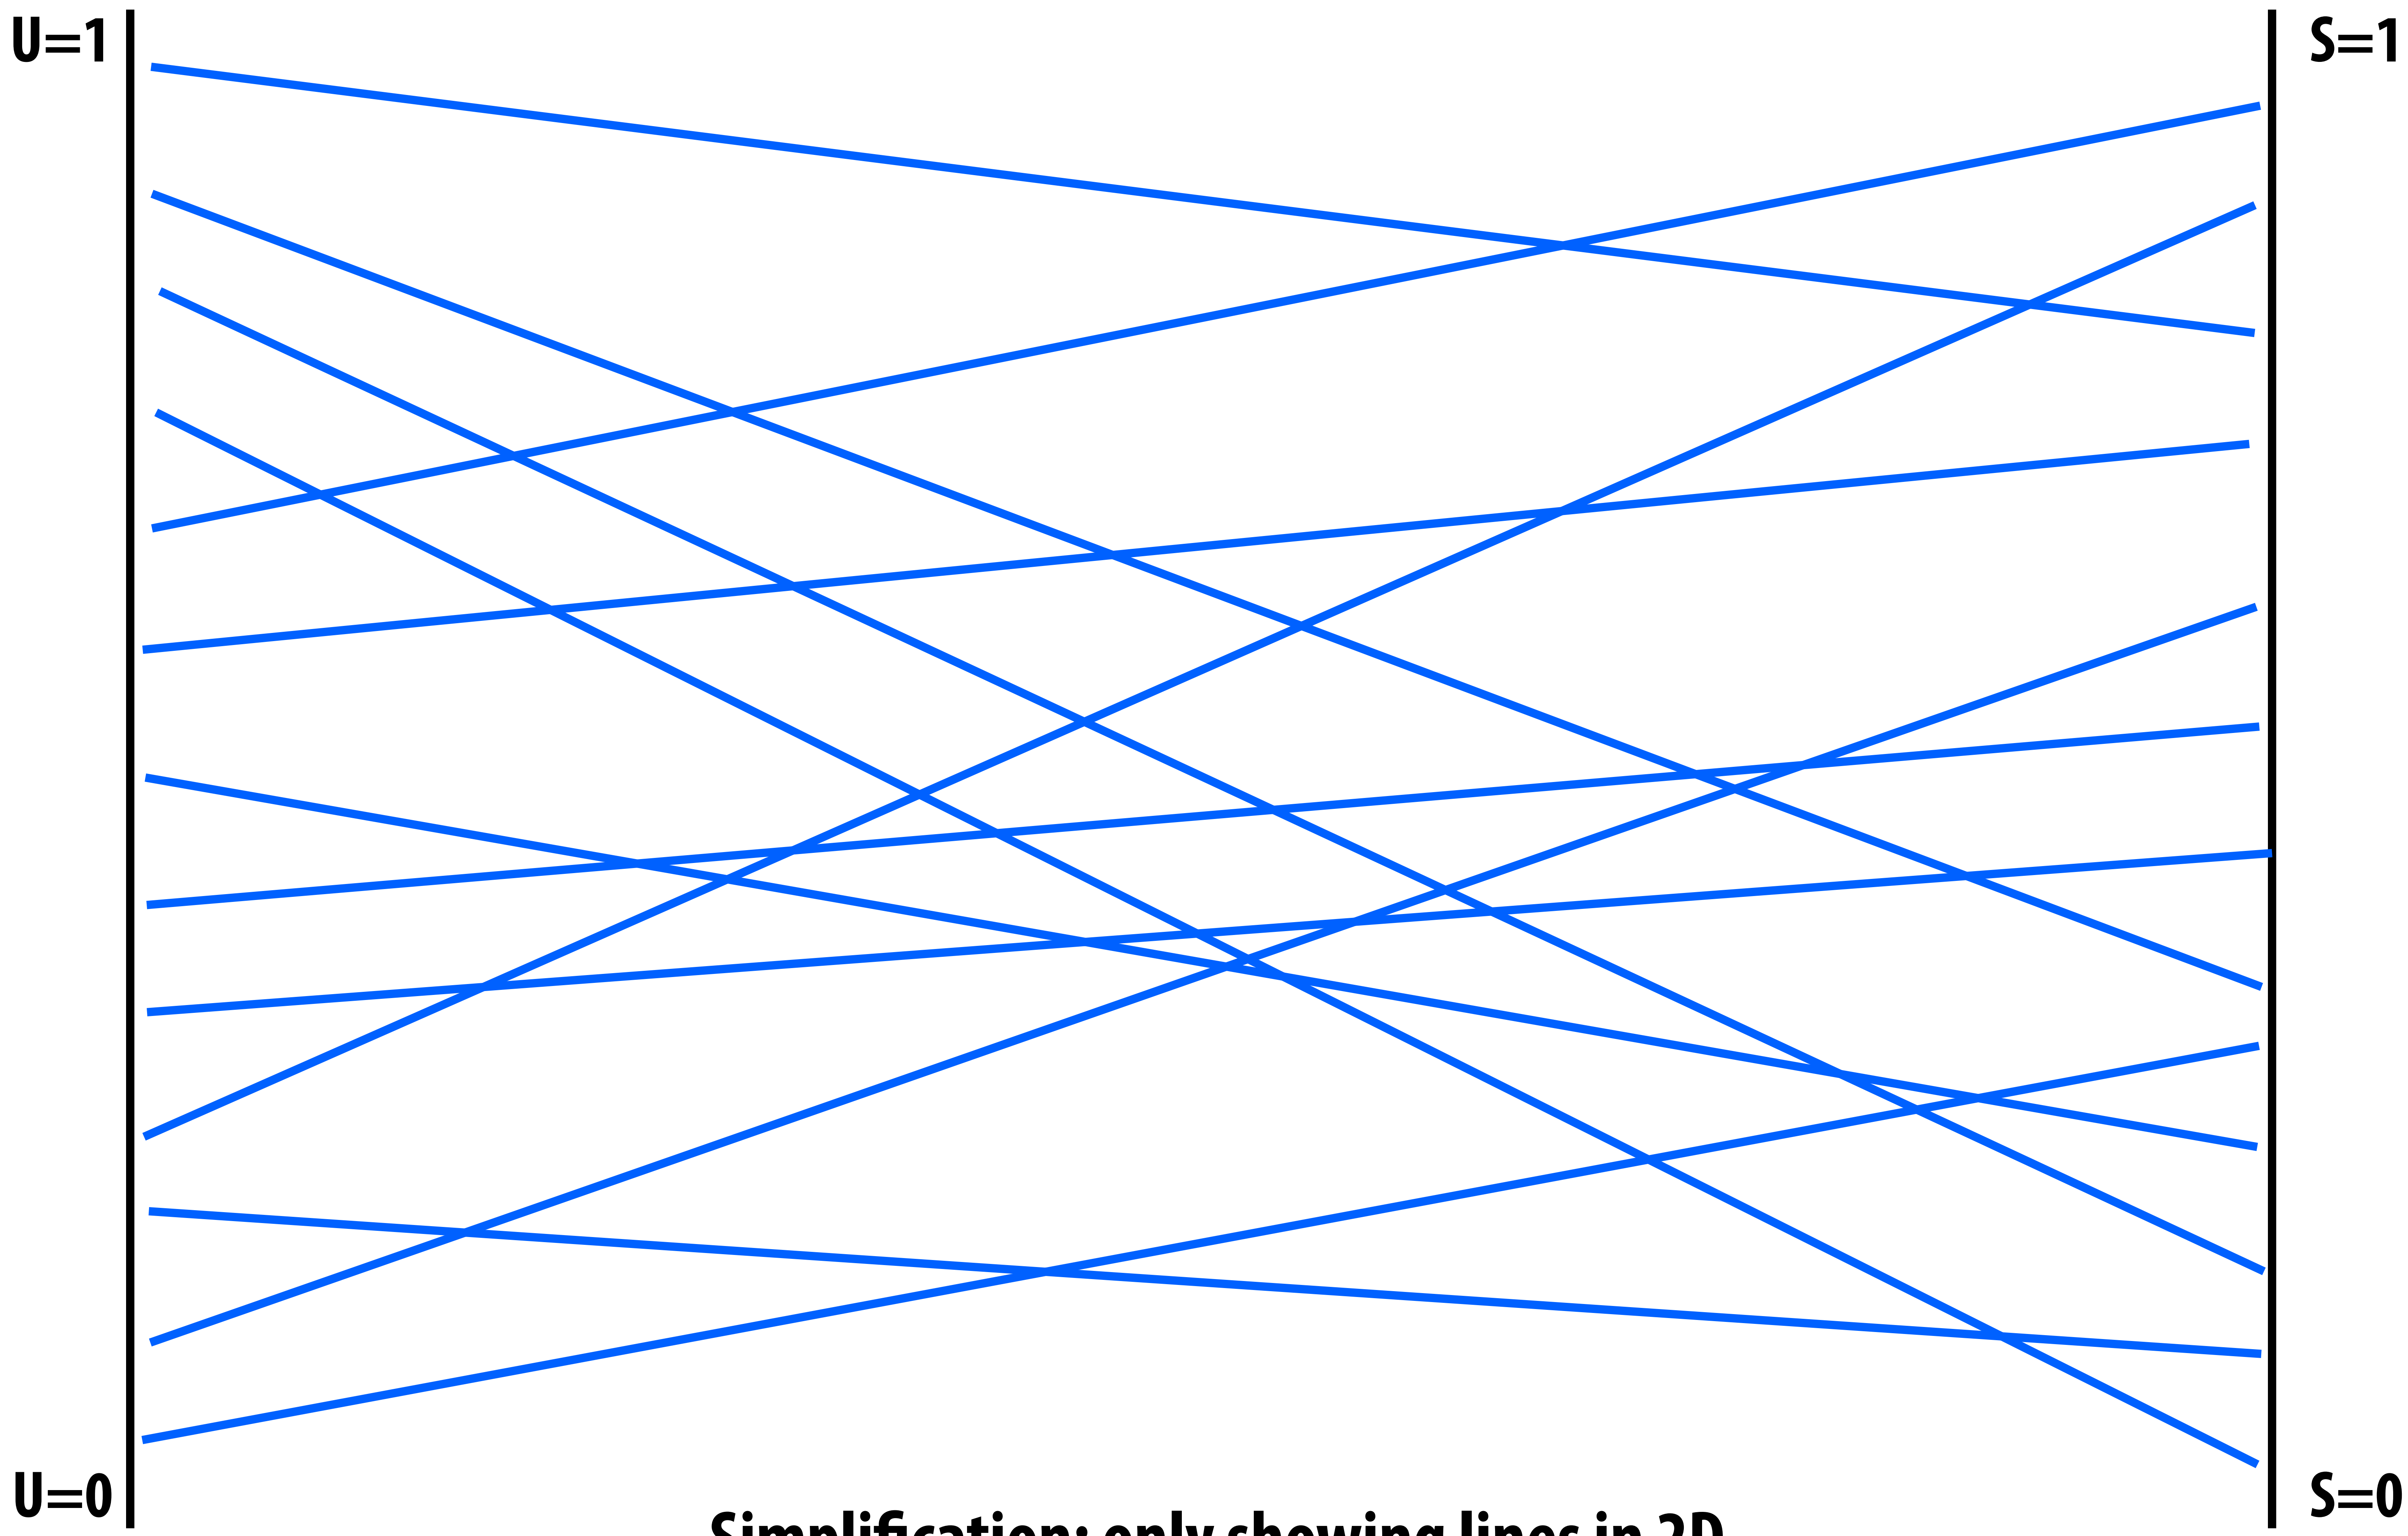

Simplification: only showing lines in 2D
(full light field is 4D function)

Sampling the light field by taking pictures

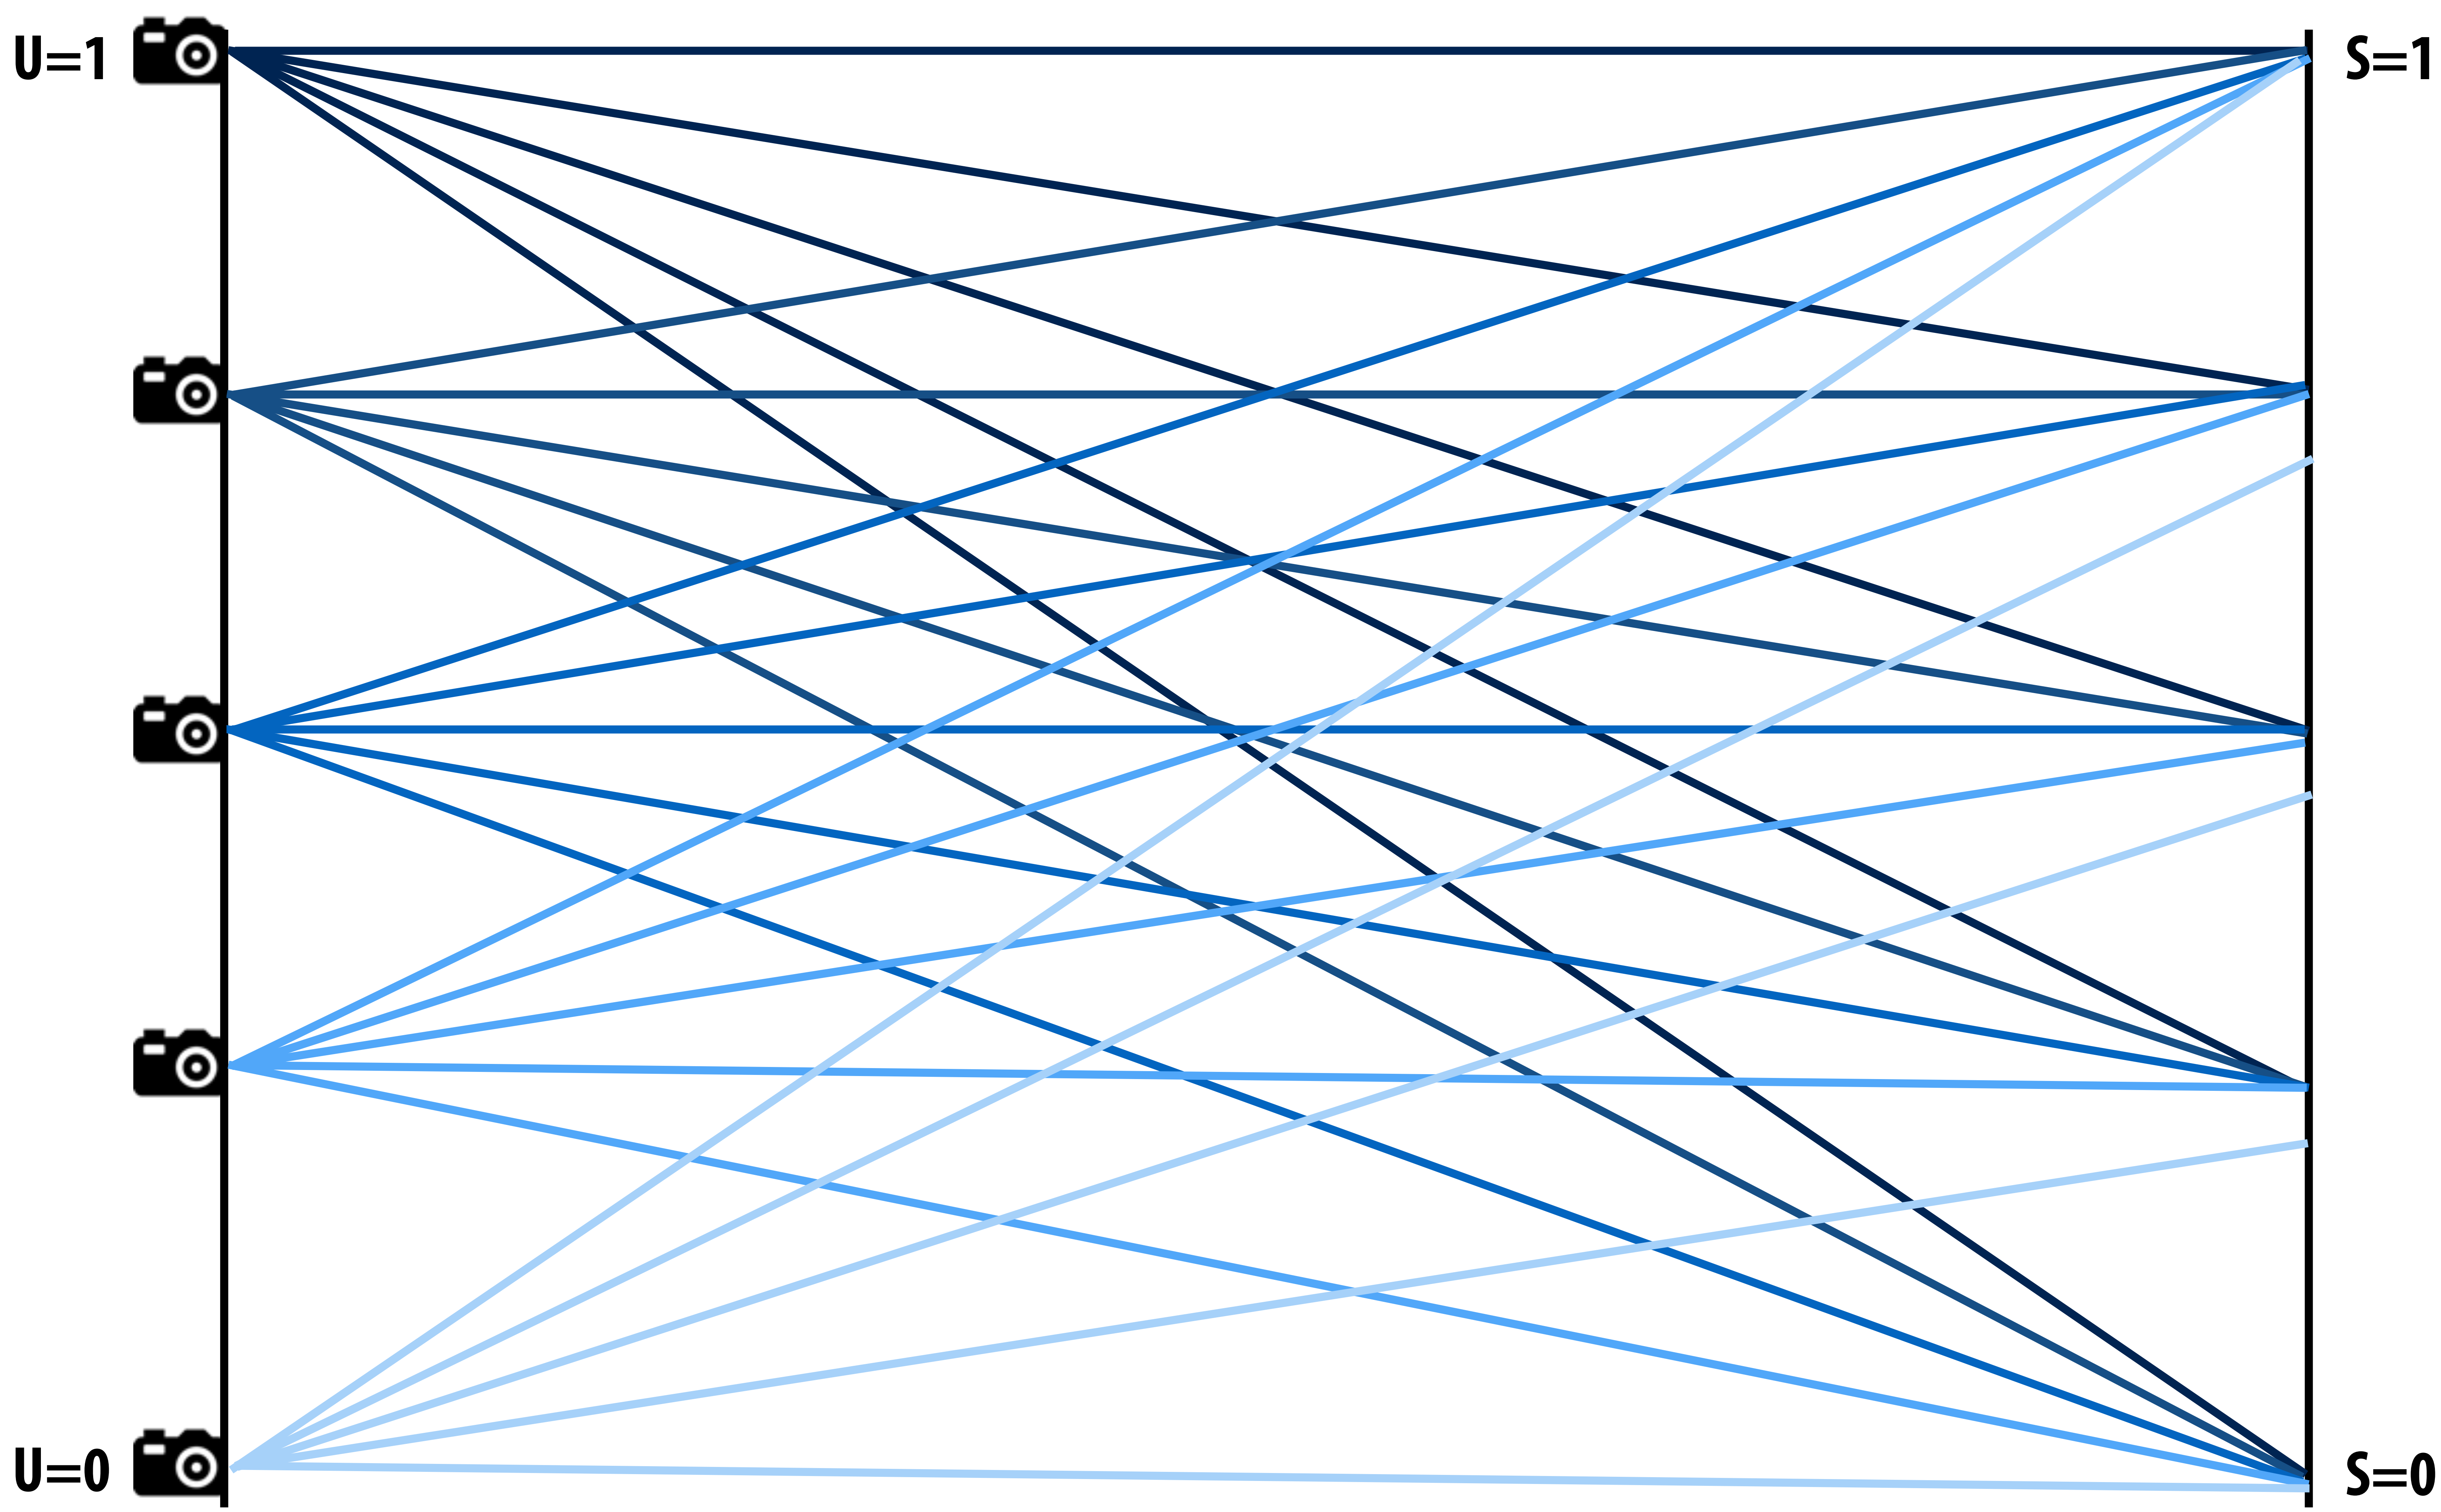

[Credit: Camera icon by Venkatesh Aiyulu from The Noun Project]

Stanford Camera Array

Wilburn et al. 2005

640 x 480 tightly synchronized, repositionable cameras

Custom processing board per camera

Tethered to PCs for additional processing/storage

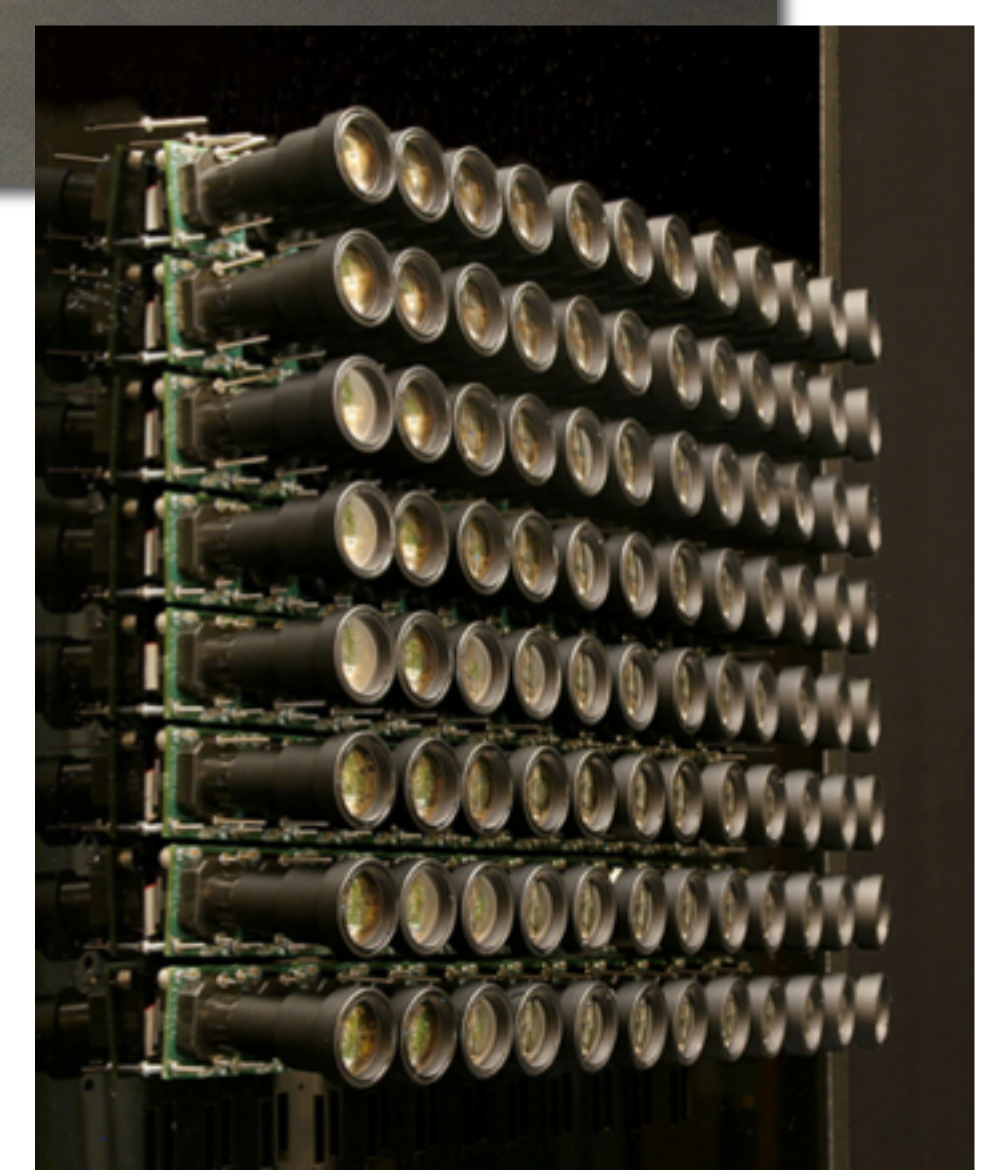

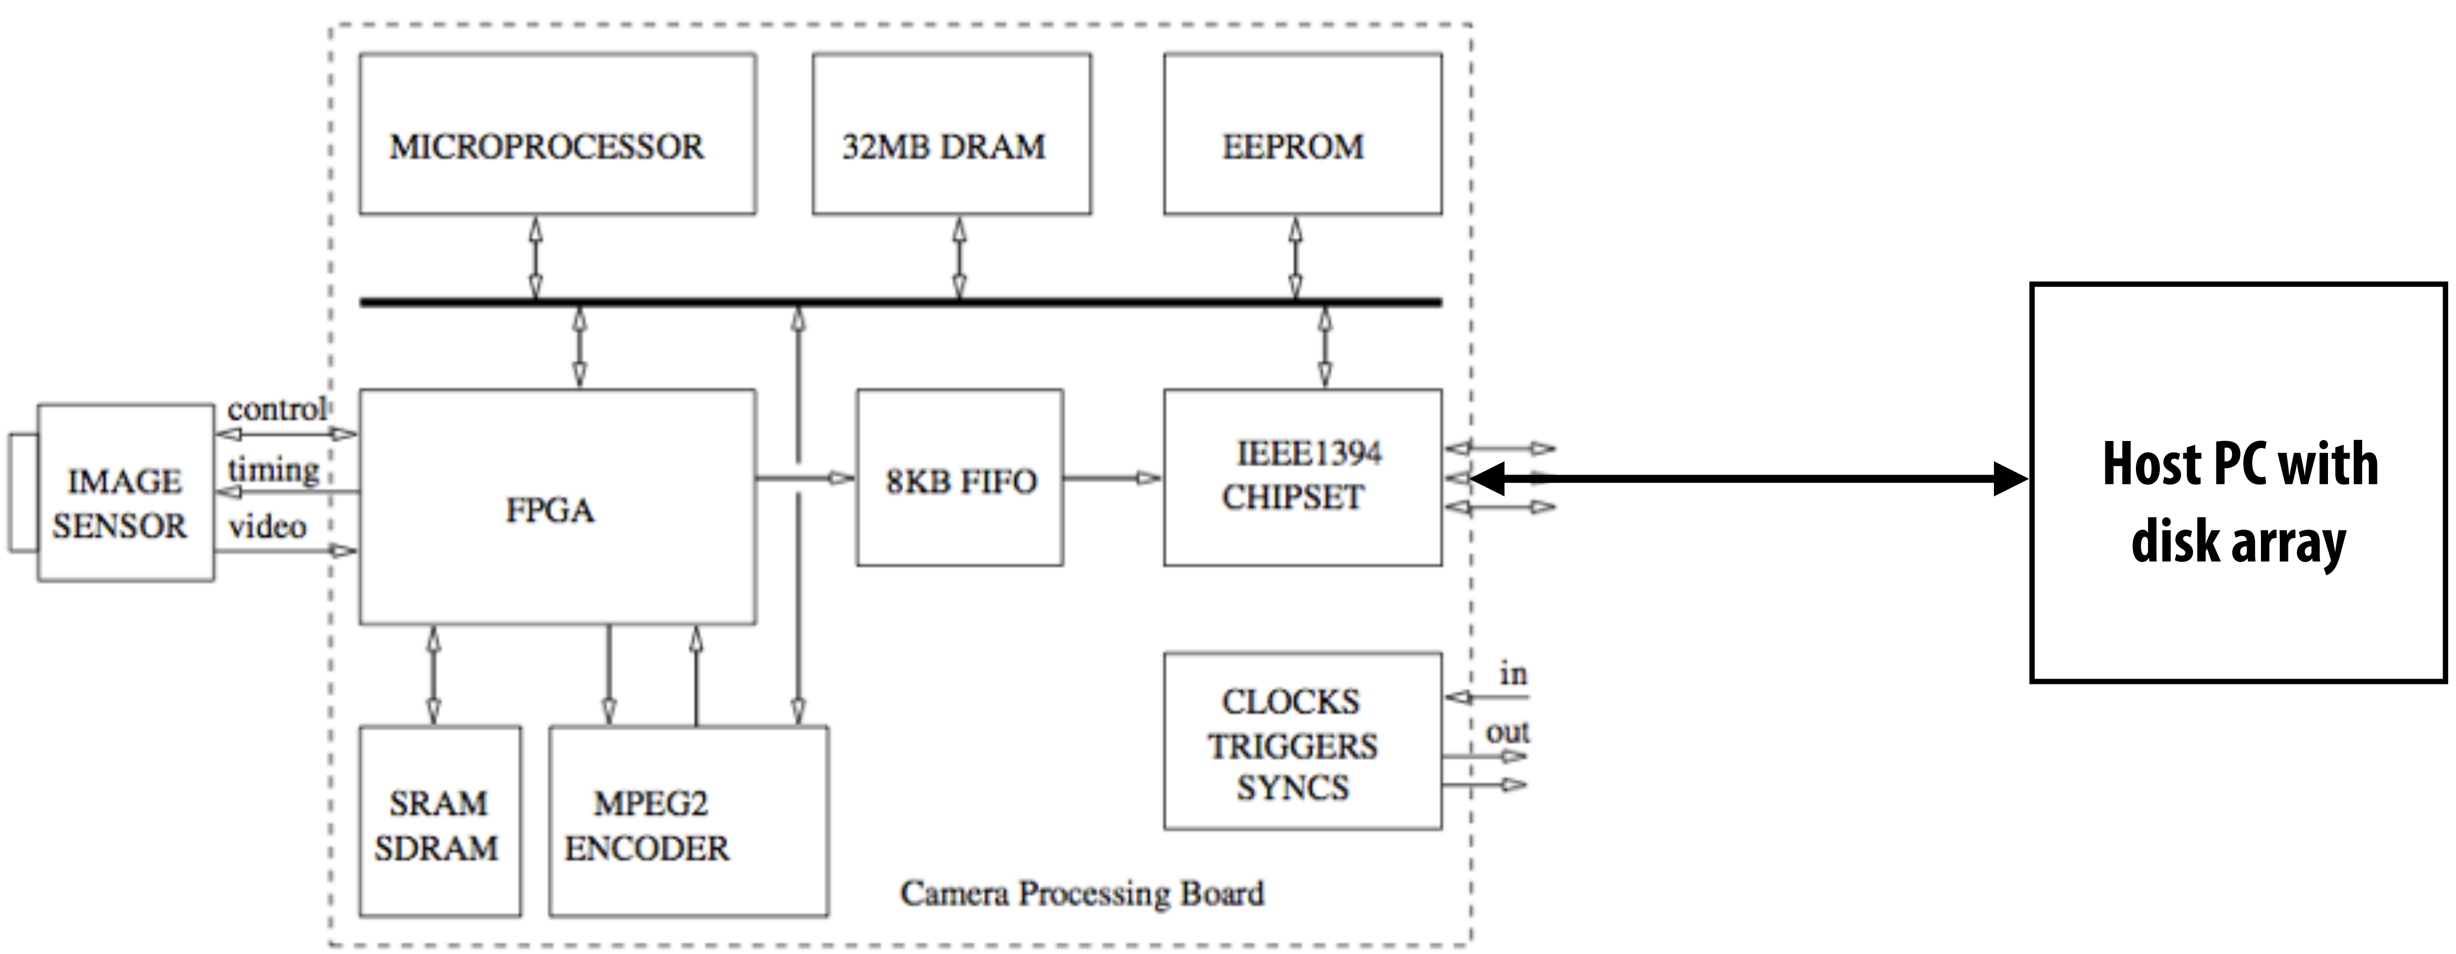

Light field storage layouts

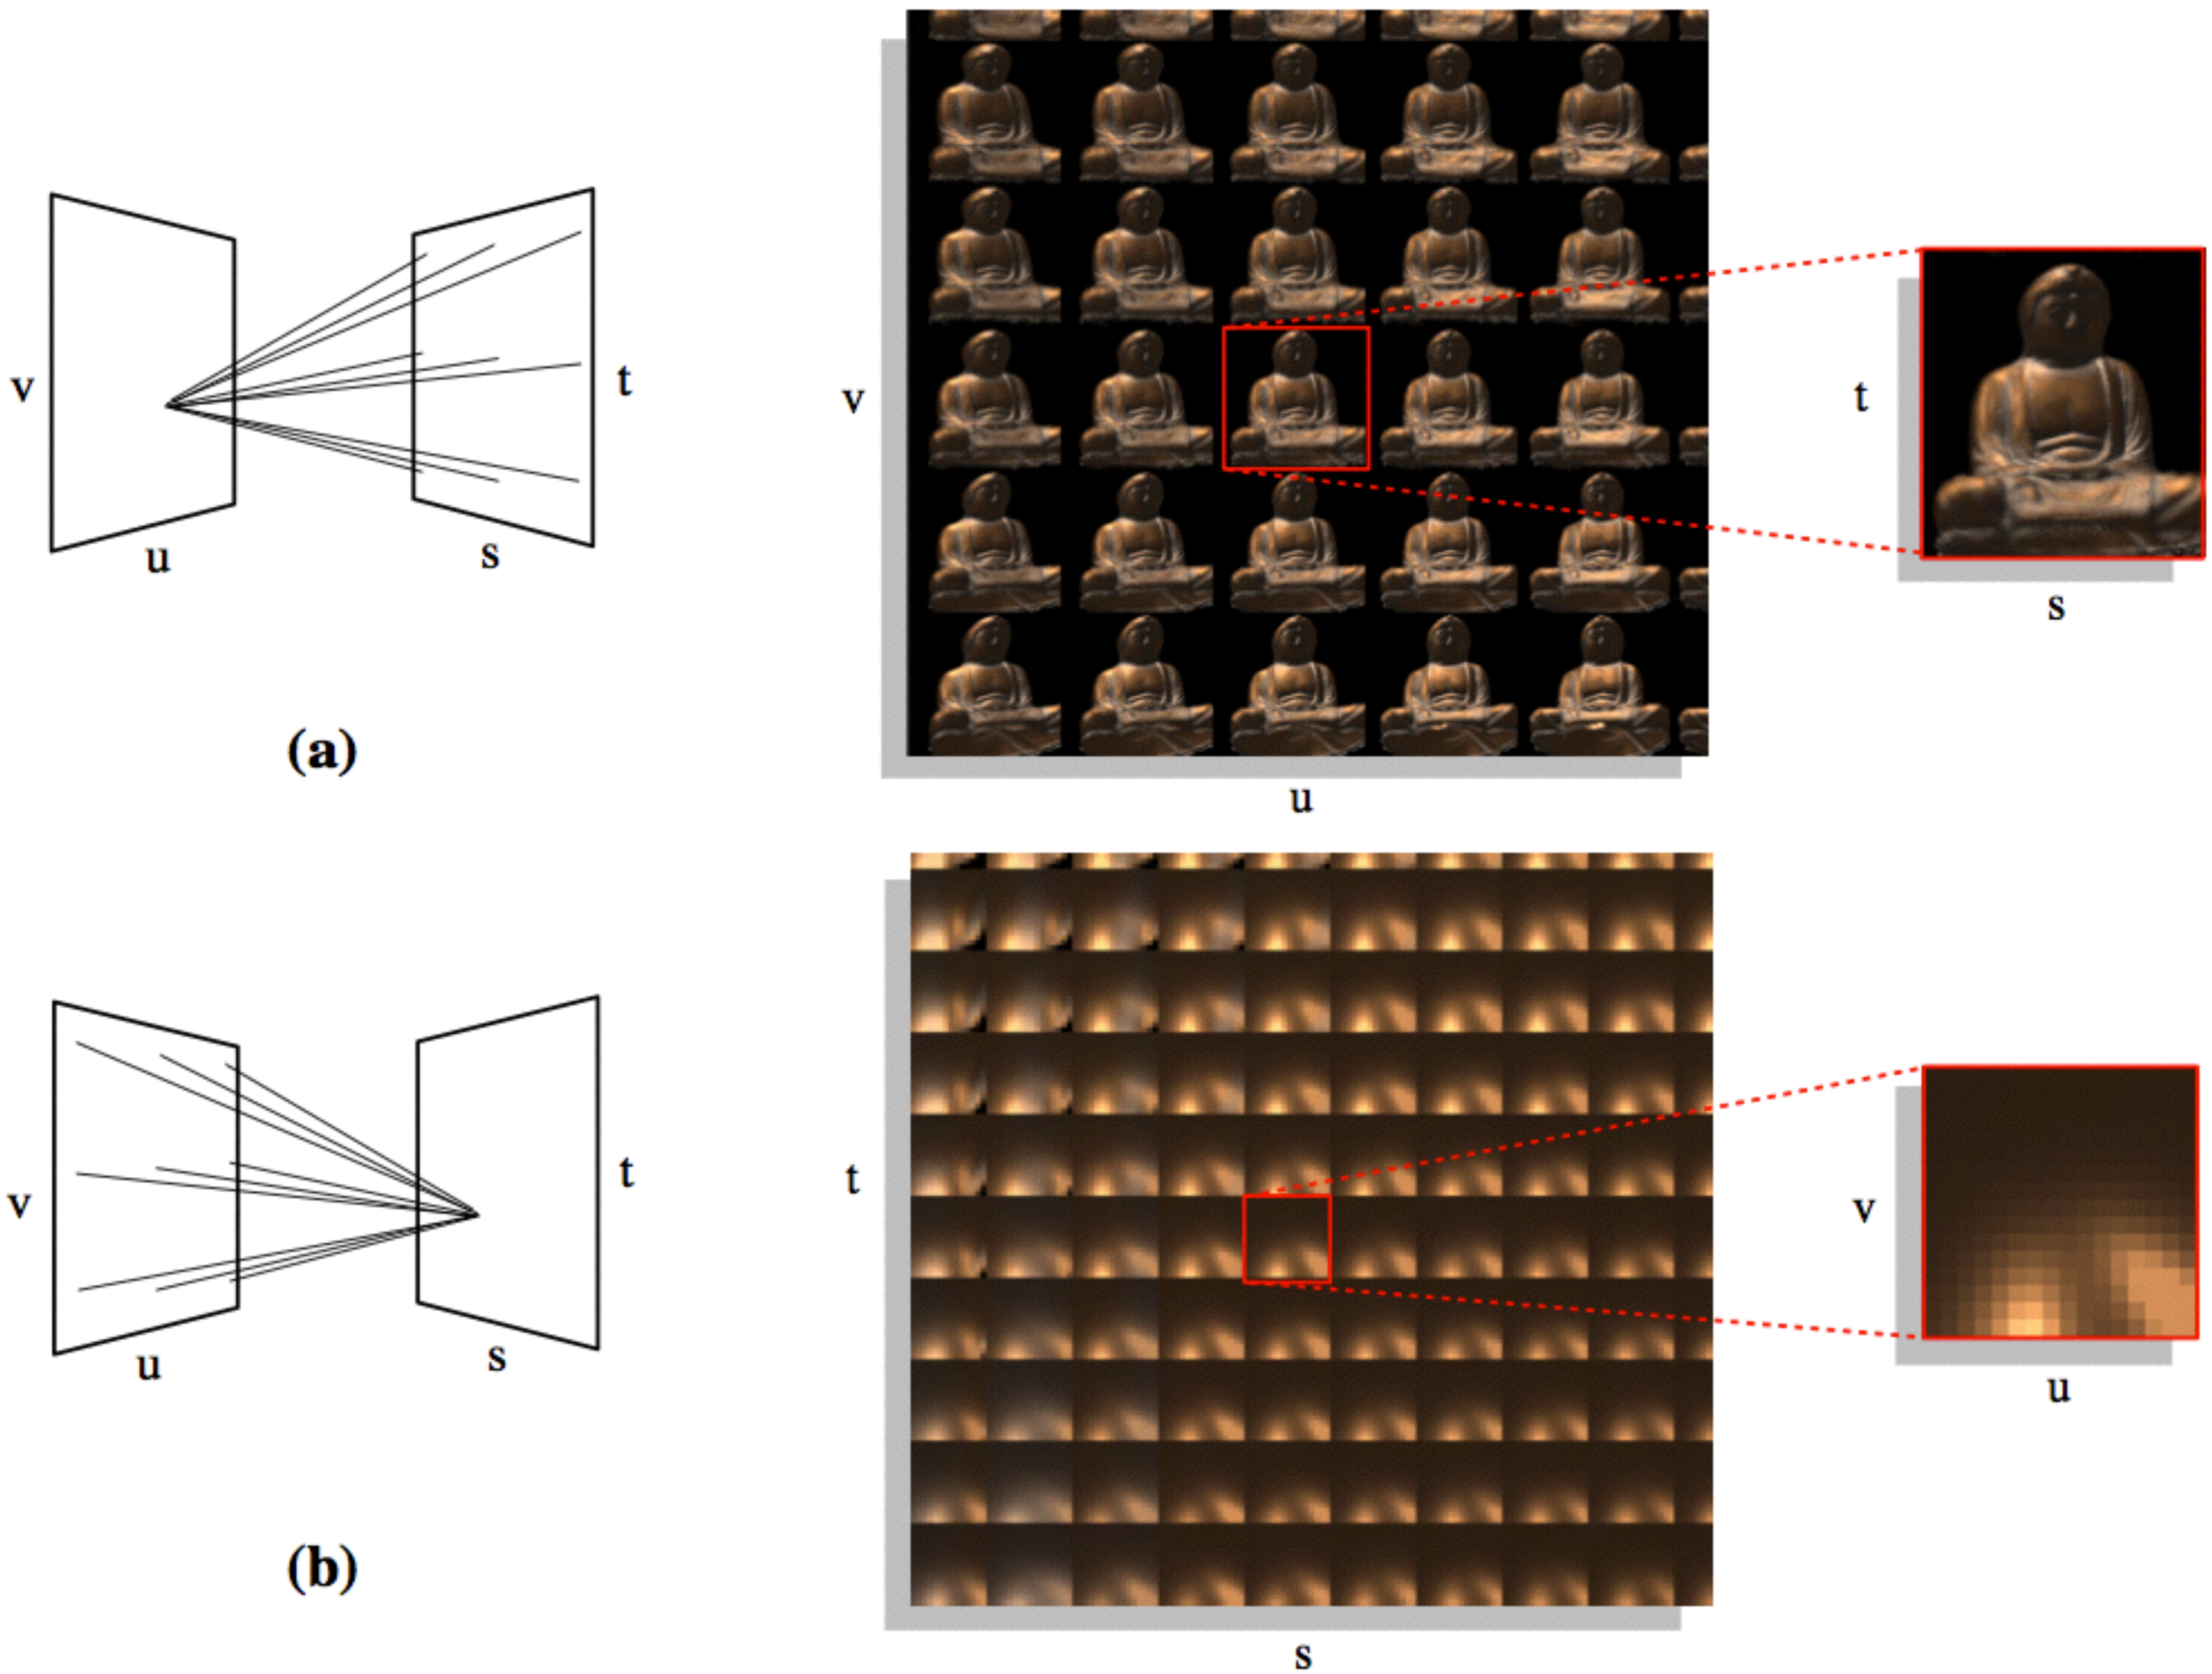

[Image credit: Levoy and Hanrahan 96]

Line-space representation

Each line in Cartesian space* is represented by a point in line space

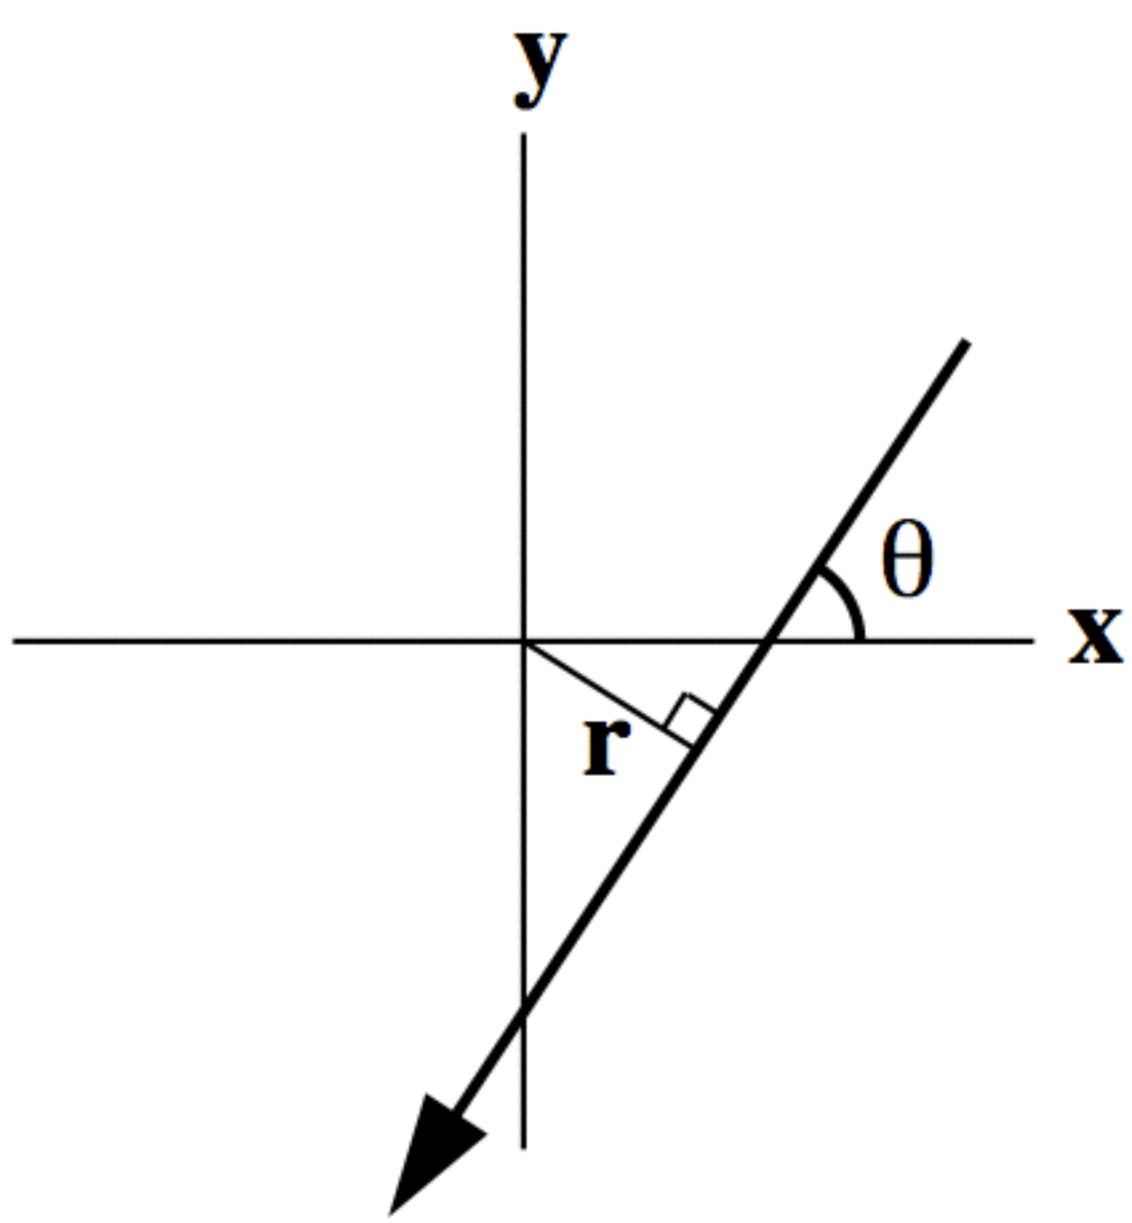

Cartesian space

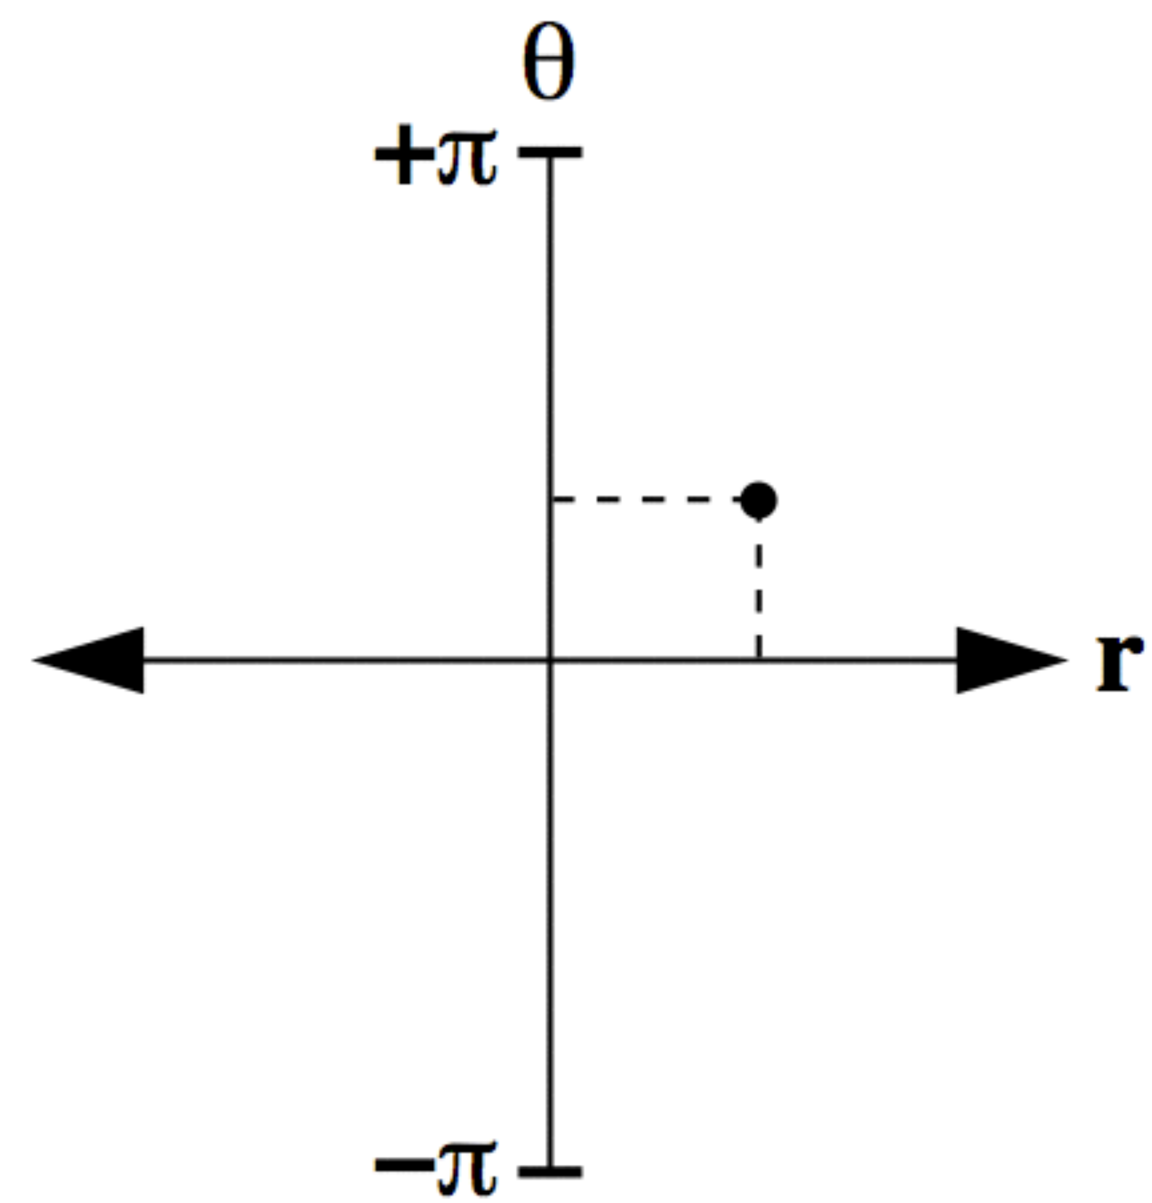

Line space

*** Shown here in 2D, generalizes to 3D Cartesian lines**

[Image credit: Levoy and Hanrahan 96]

Light field inside a pinhole camera

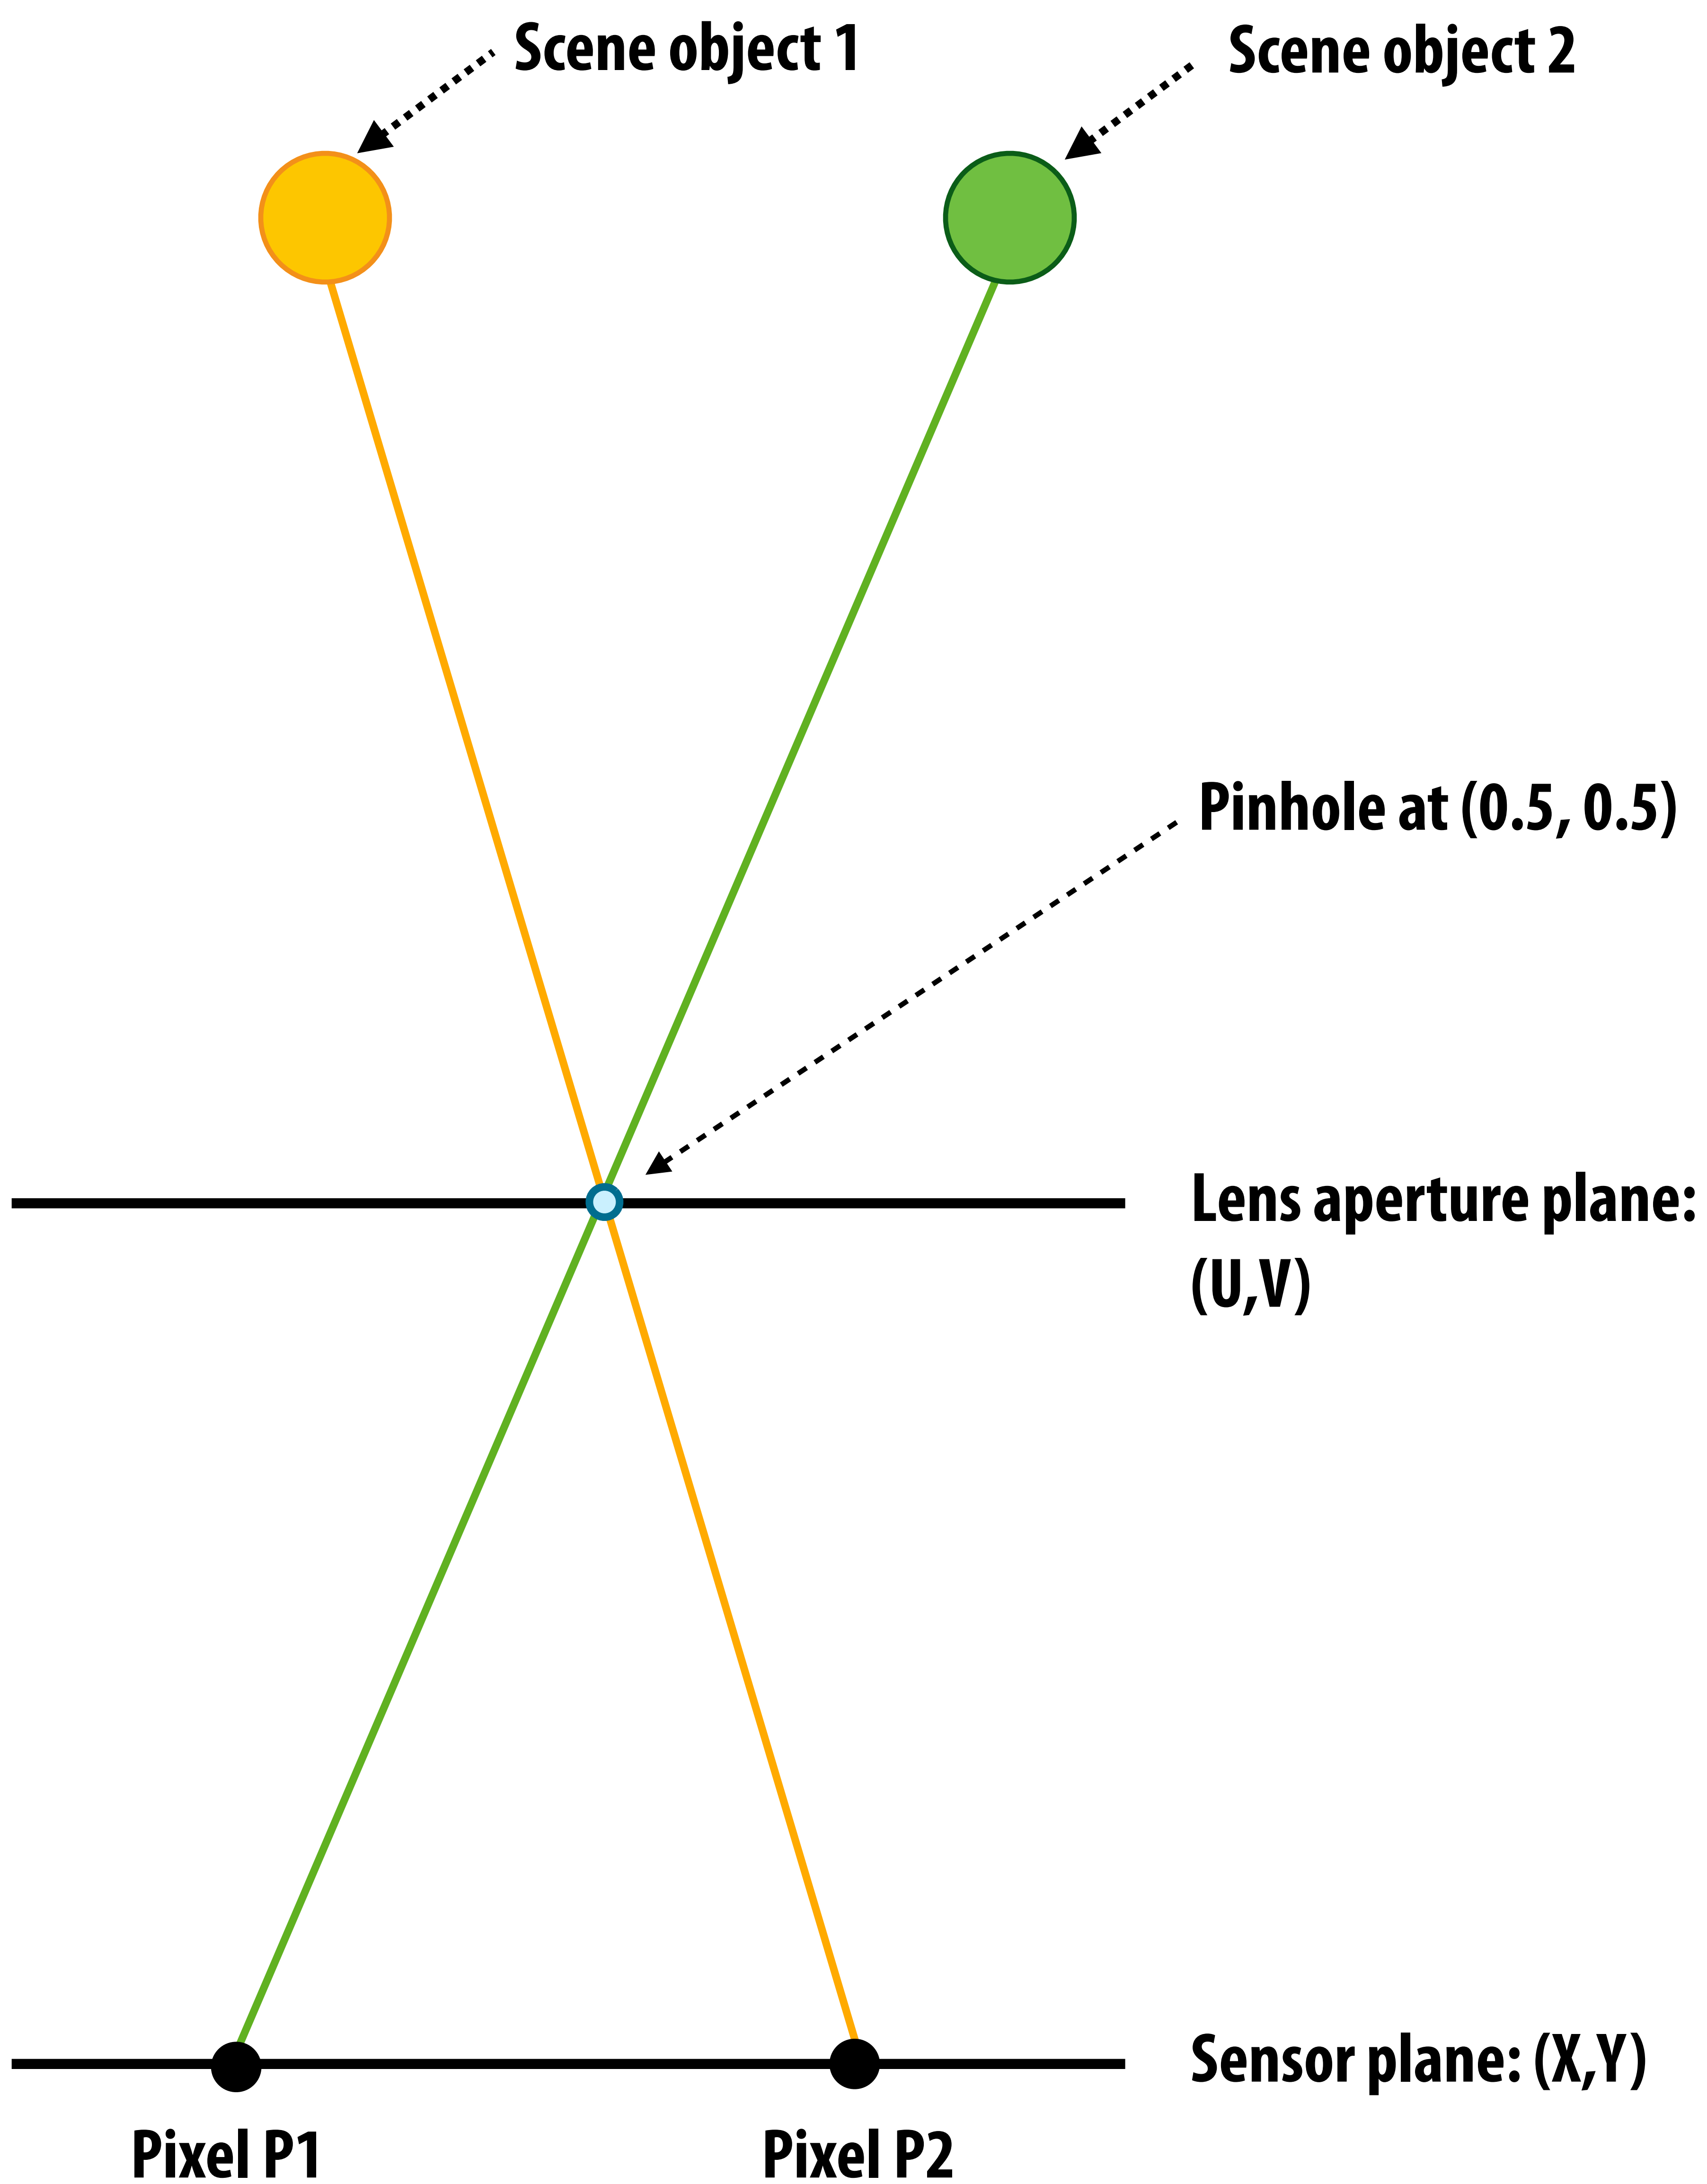

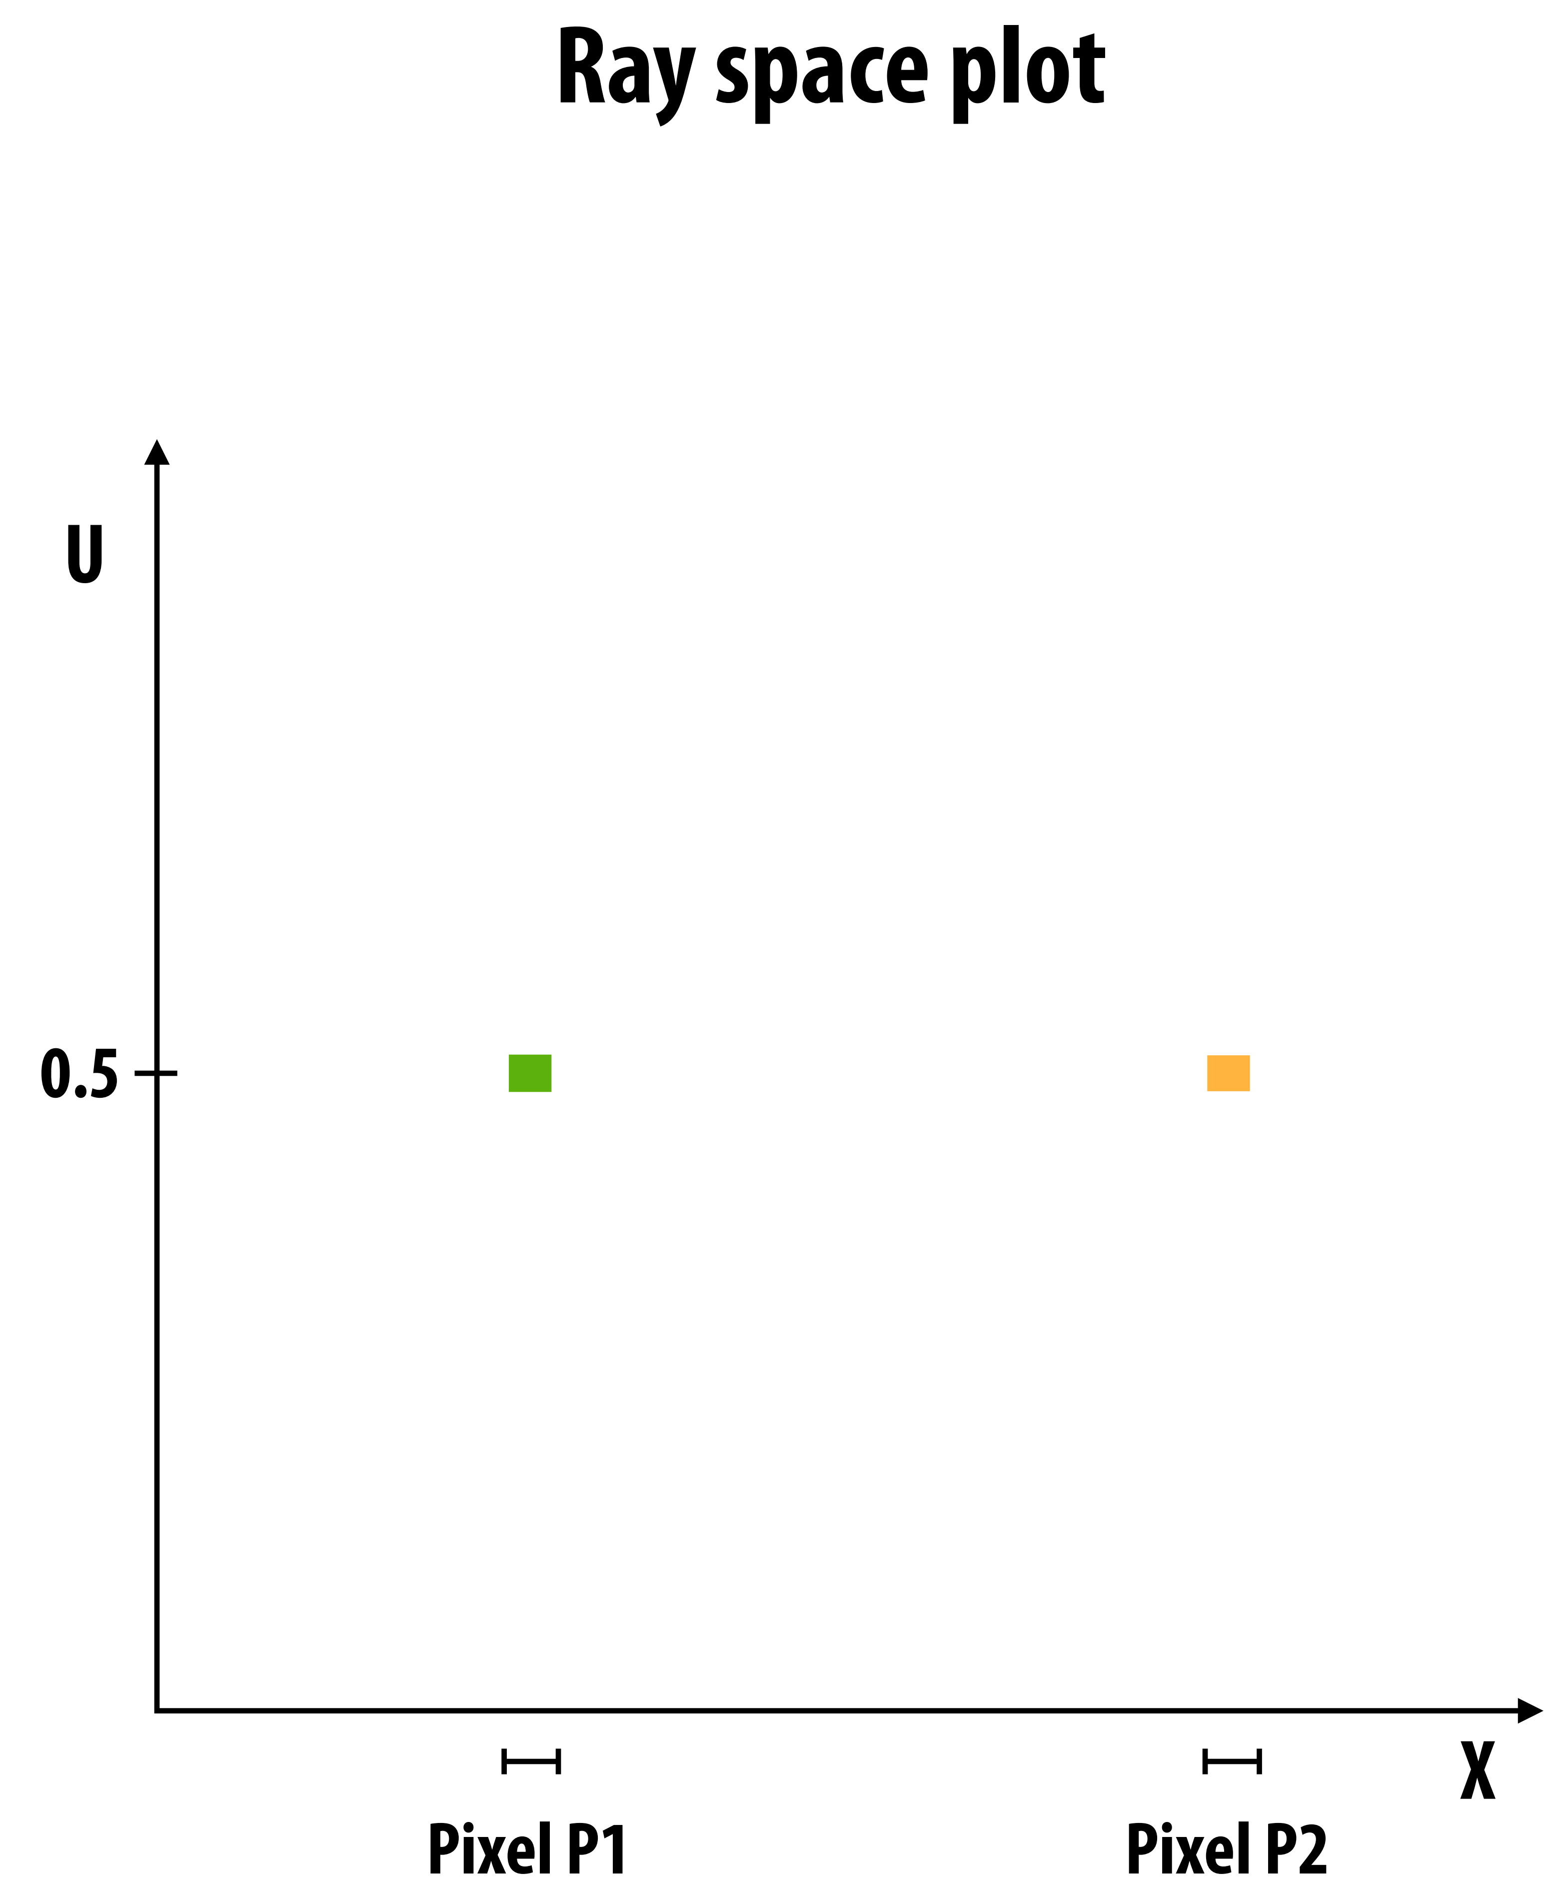

Review: camera with finite aperture

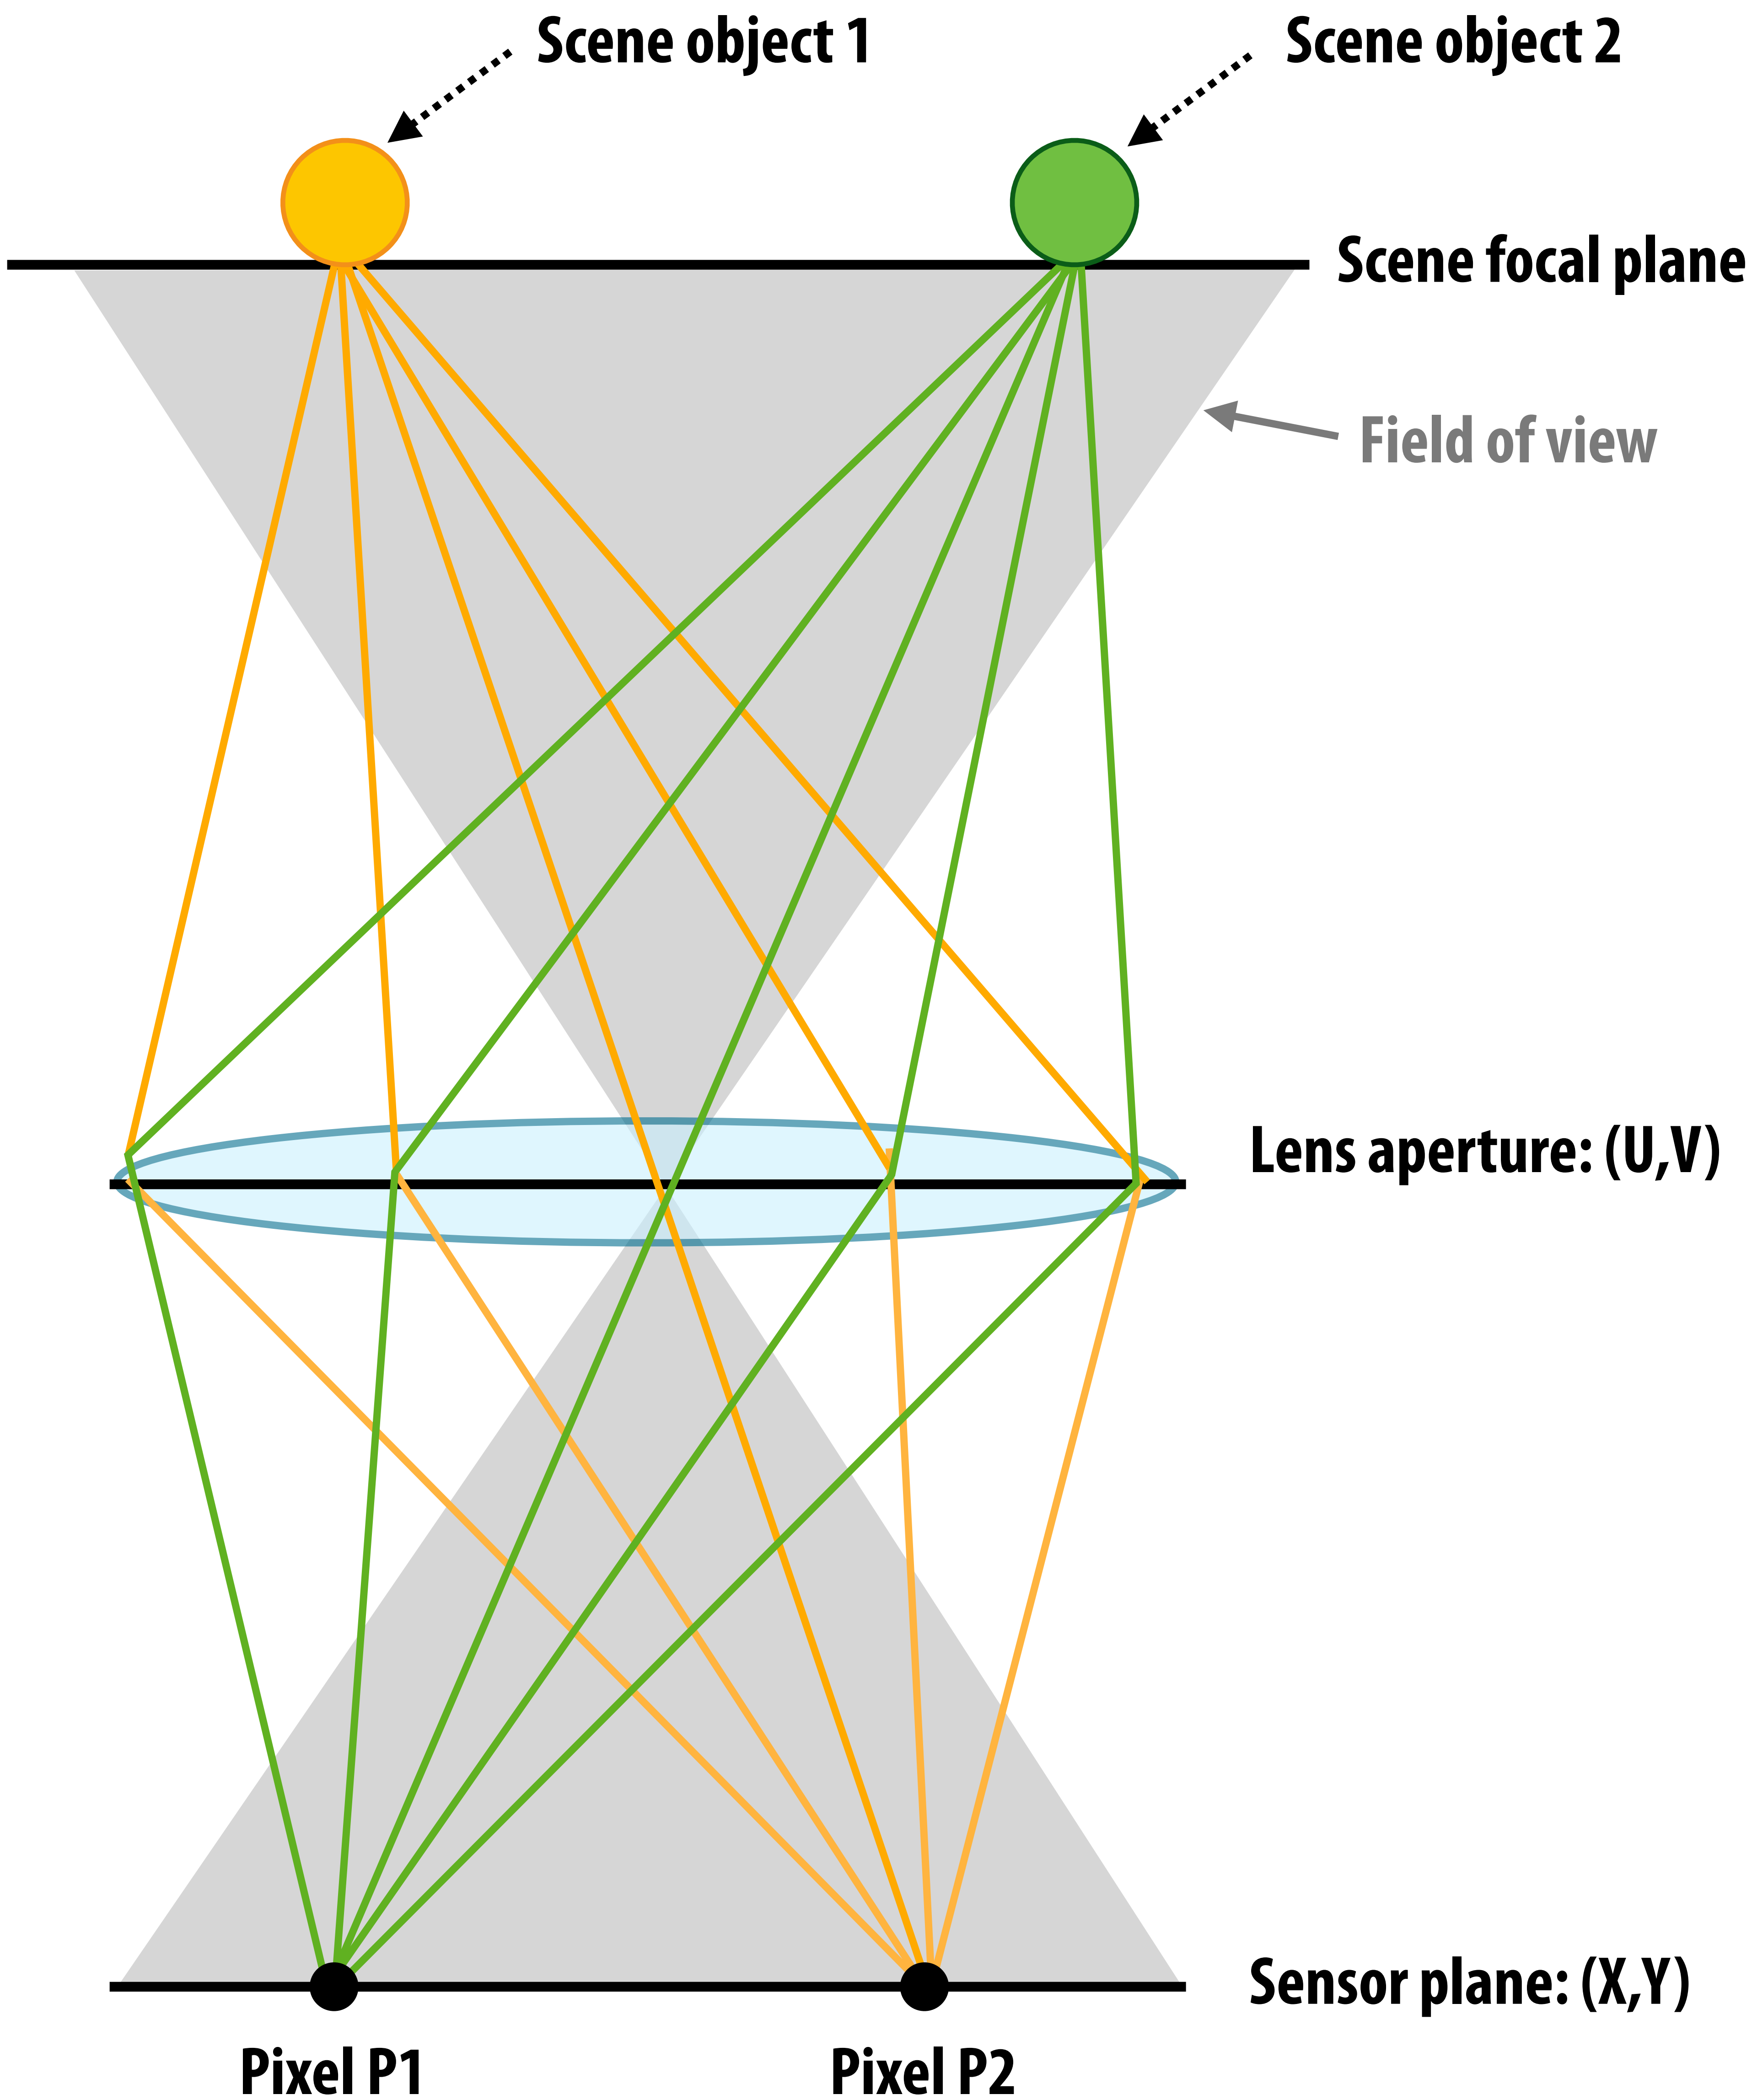

Light field inside a camera

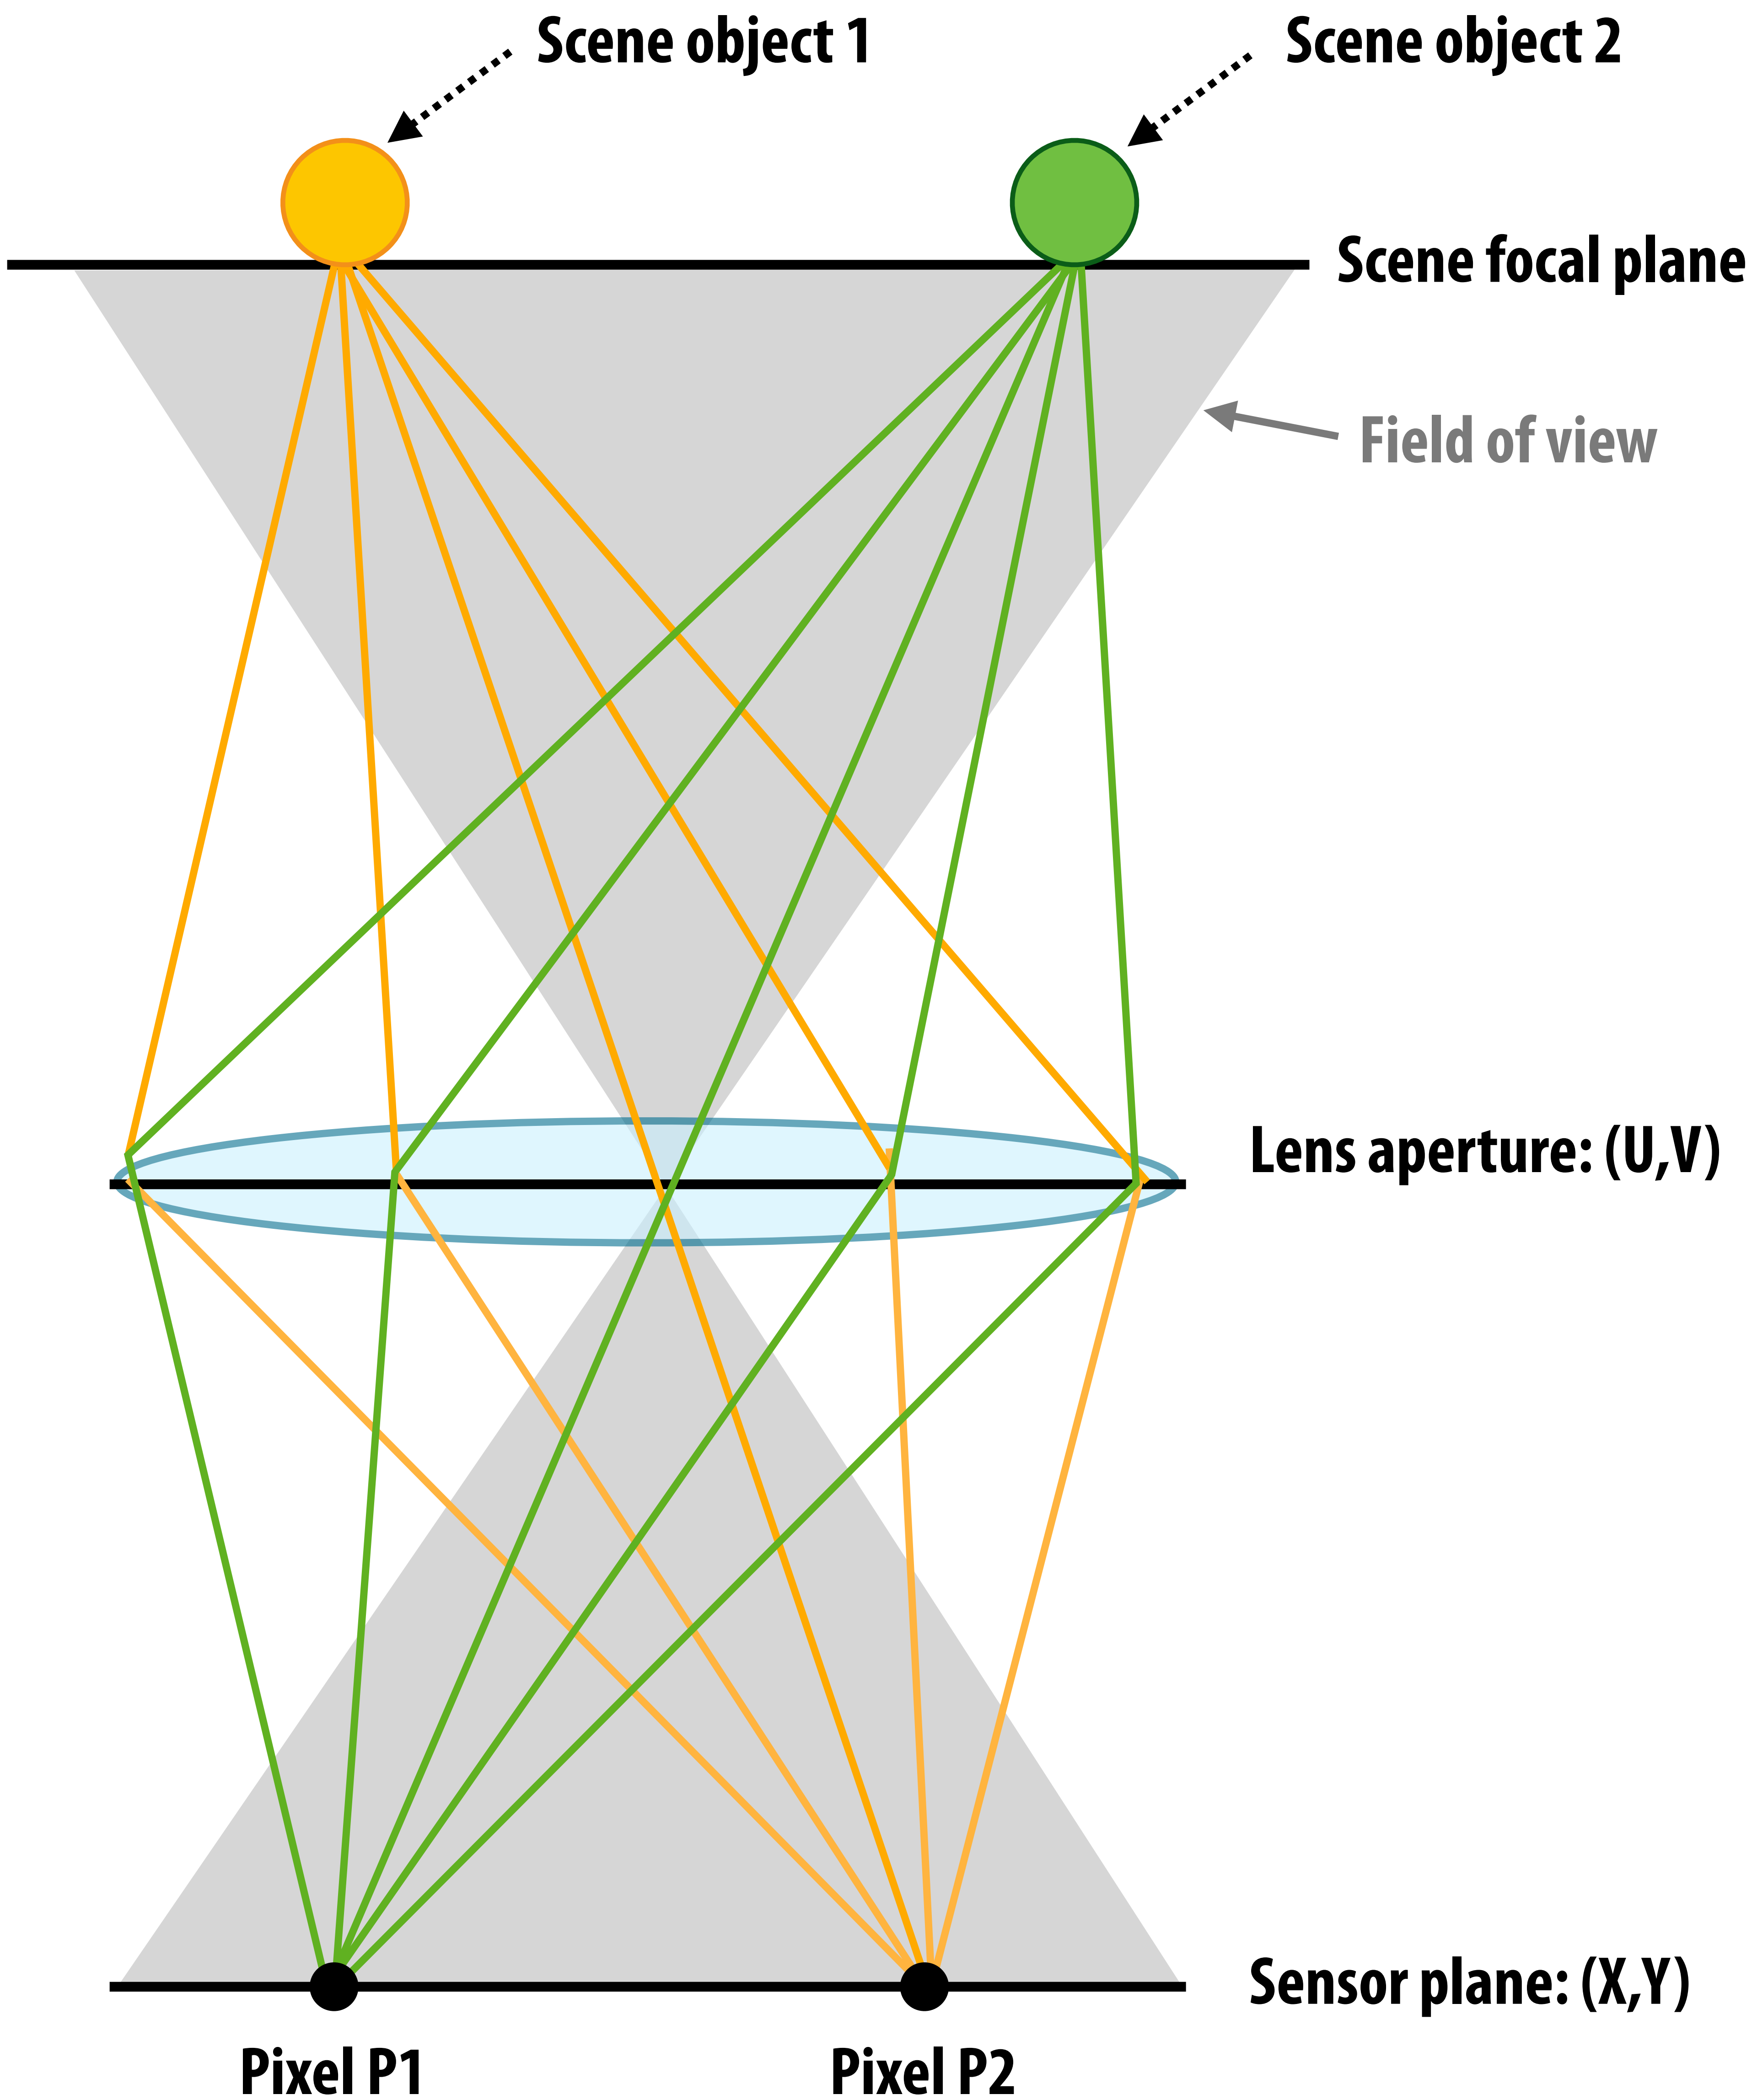

Ray space plot
(only showing X-U 2D projection)

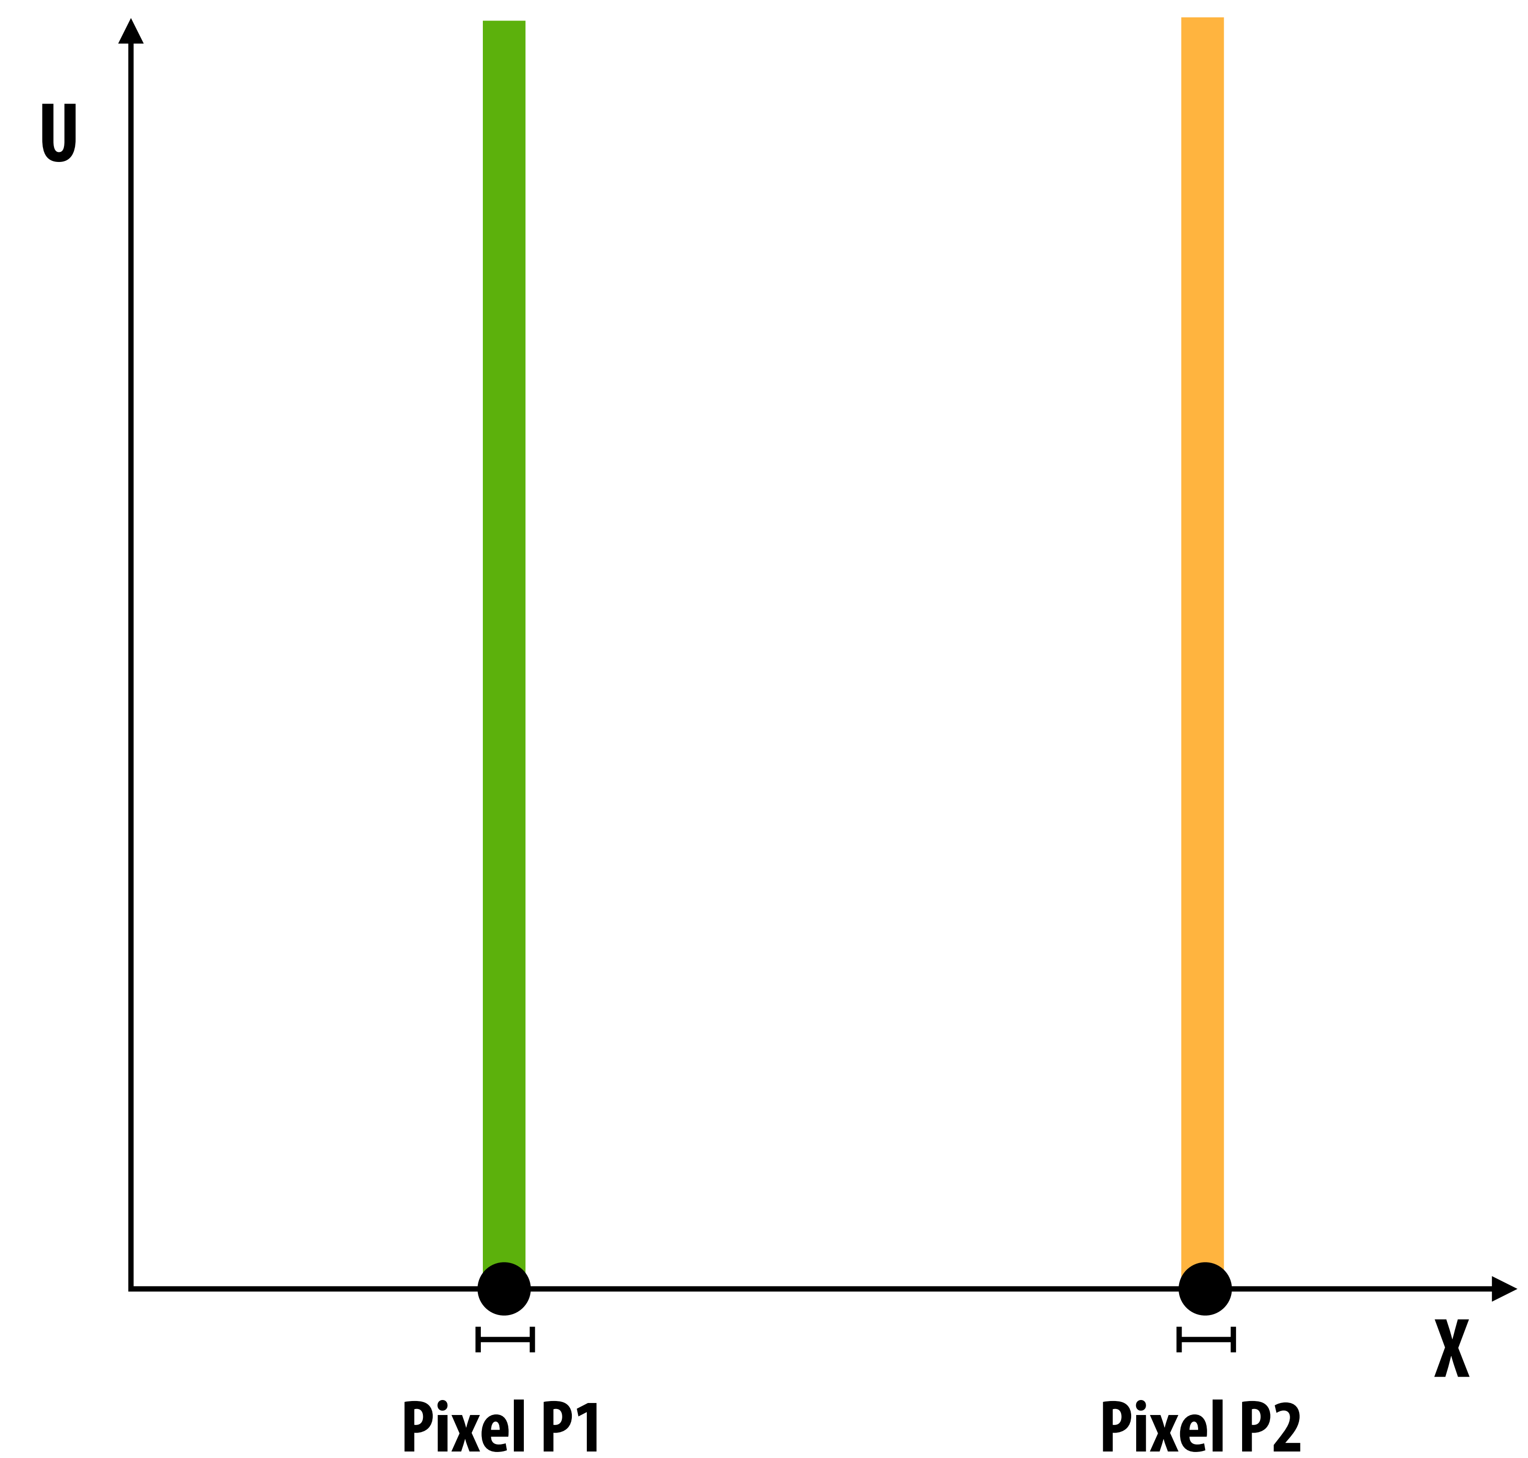

Sensor pixels measure integral of energy from all rays of light passing through points on the aperture and a pixel-sized area of the sensor.

Decrease aperture size

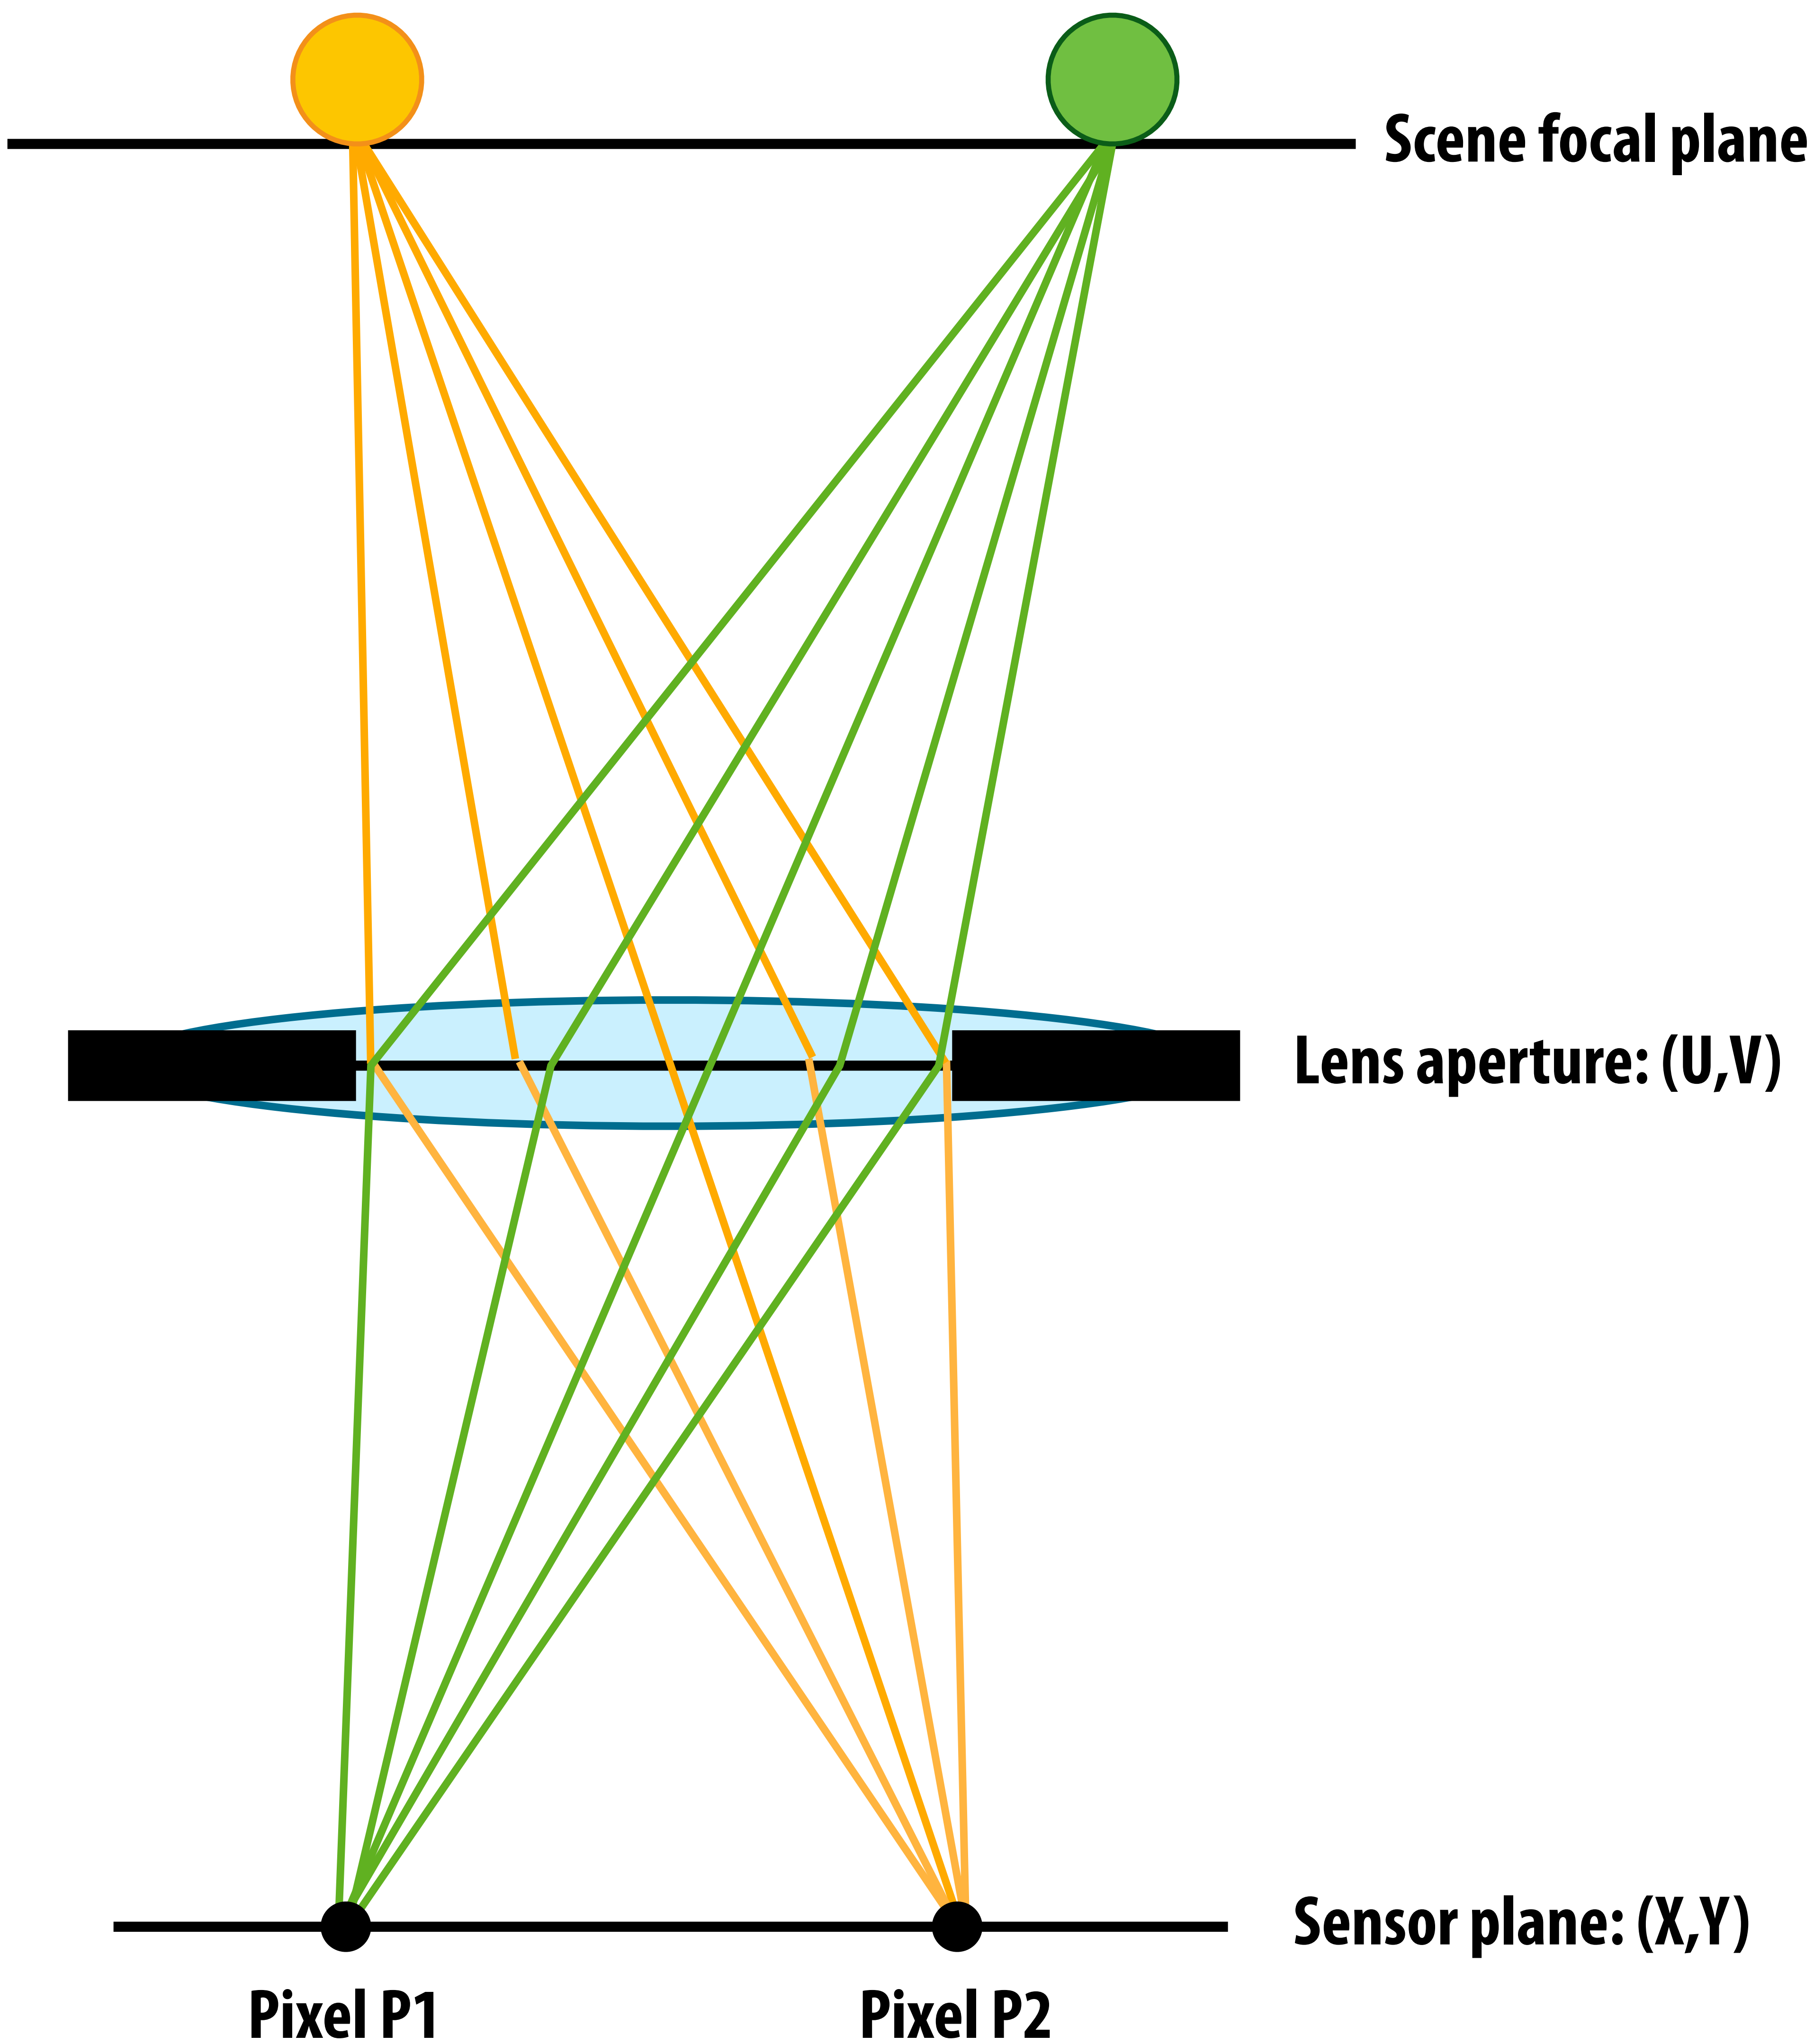

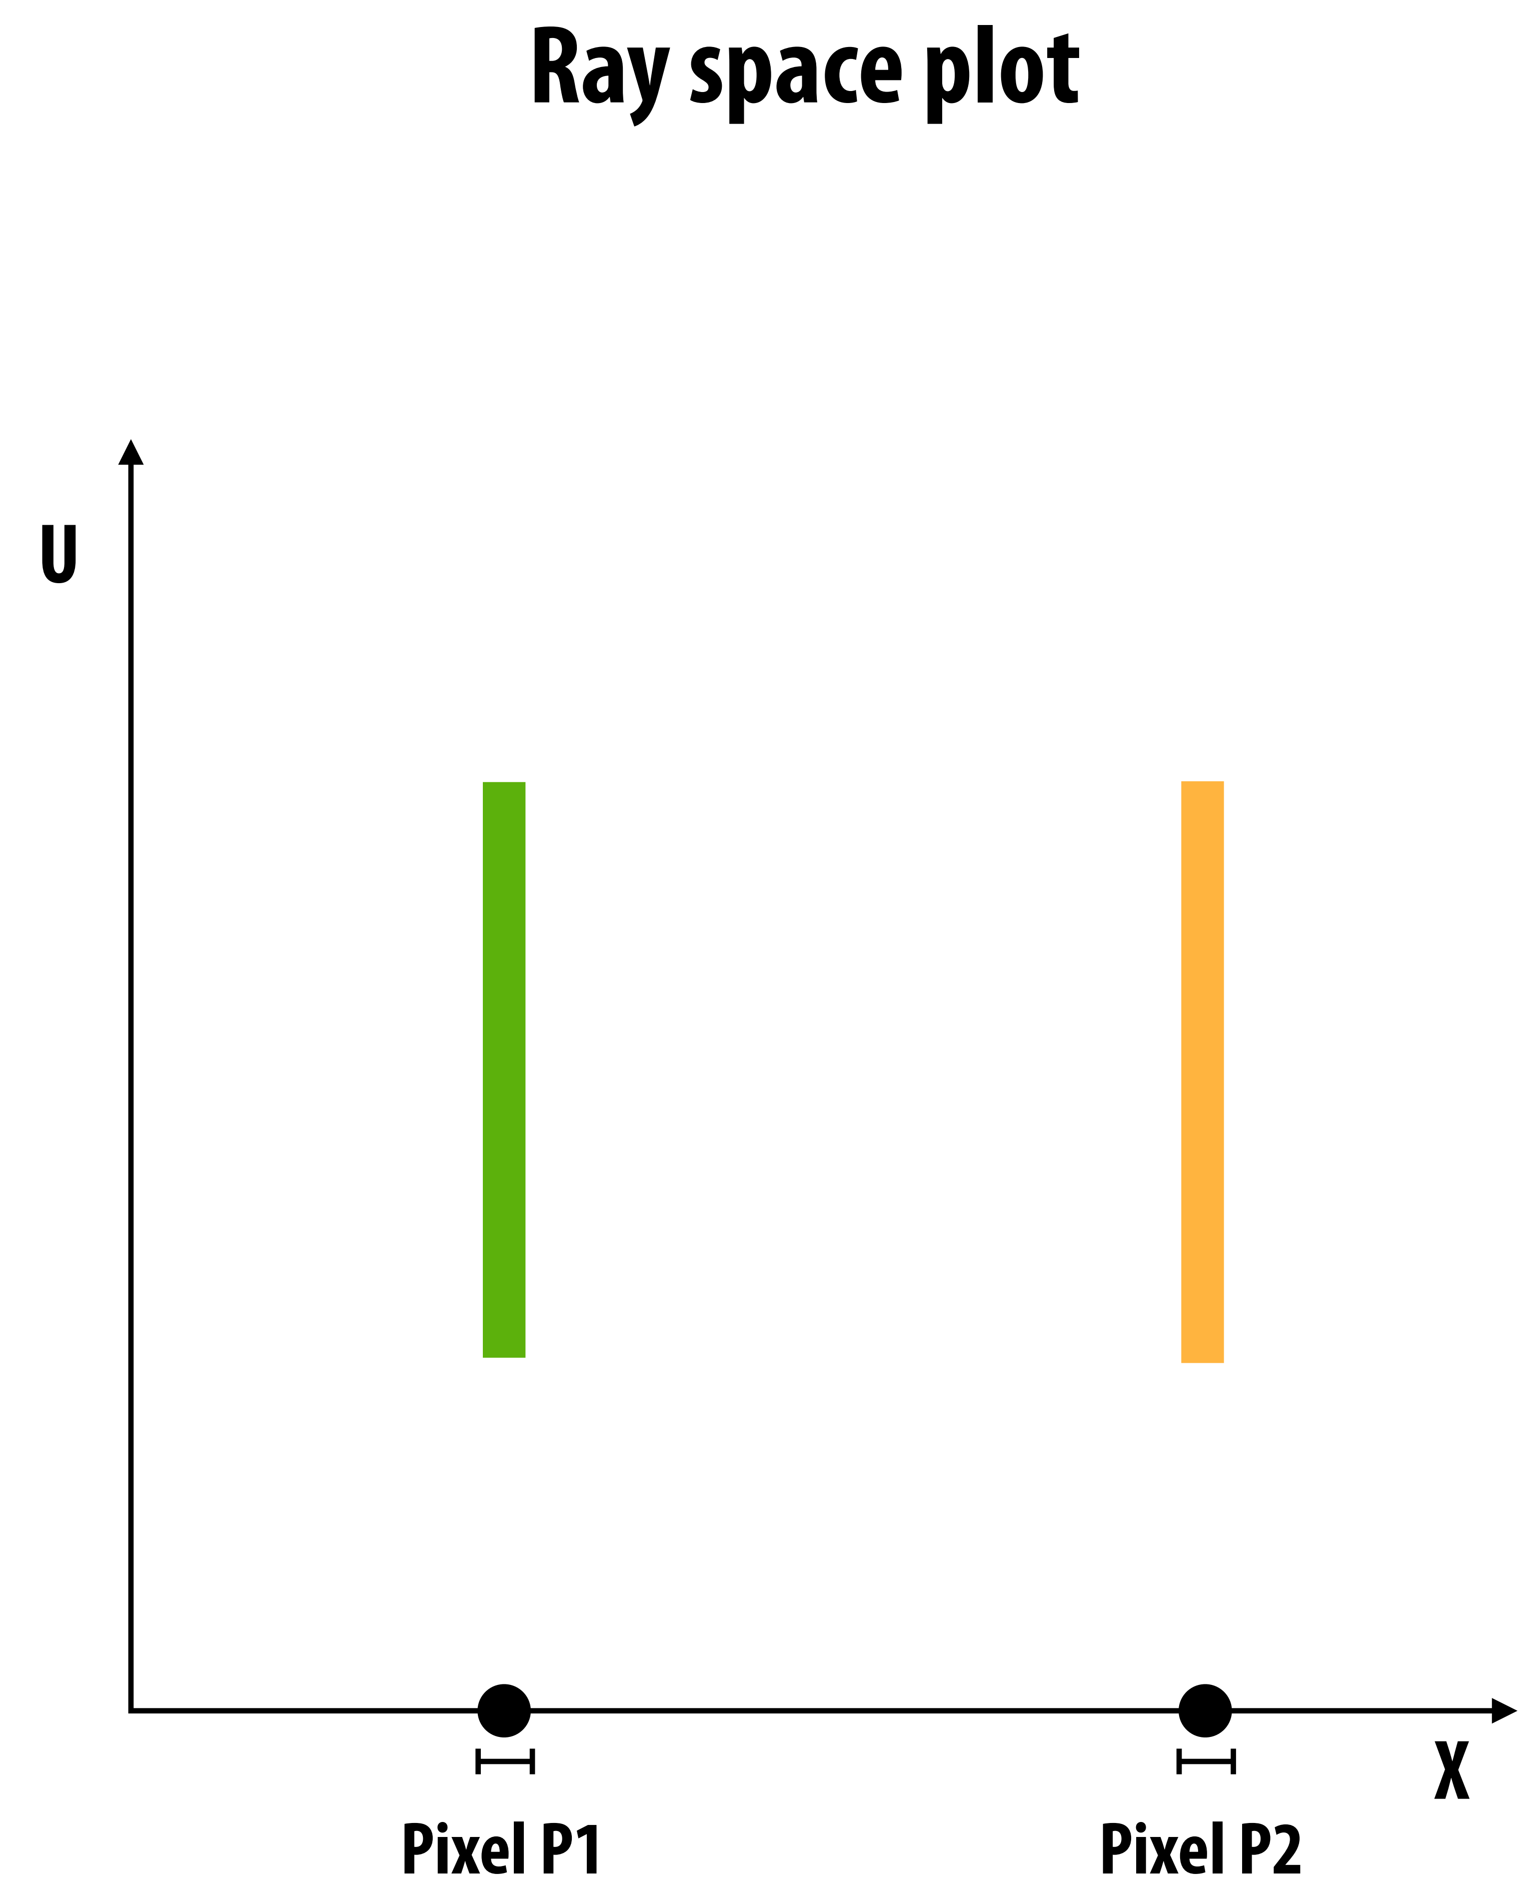

Defocus

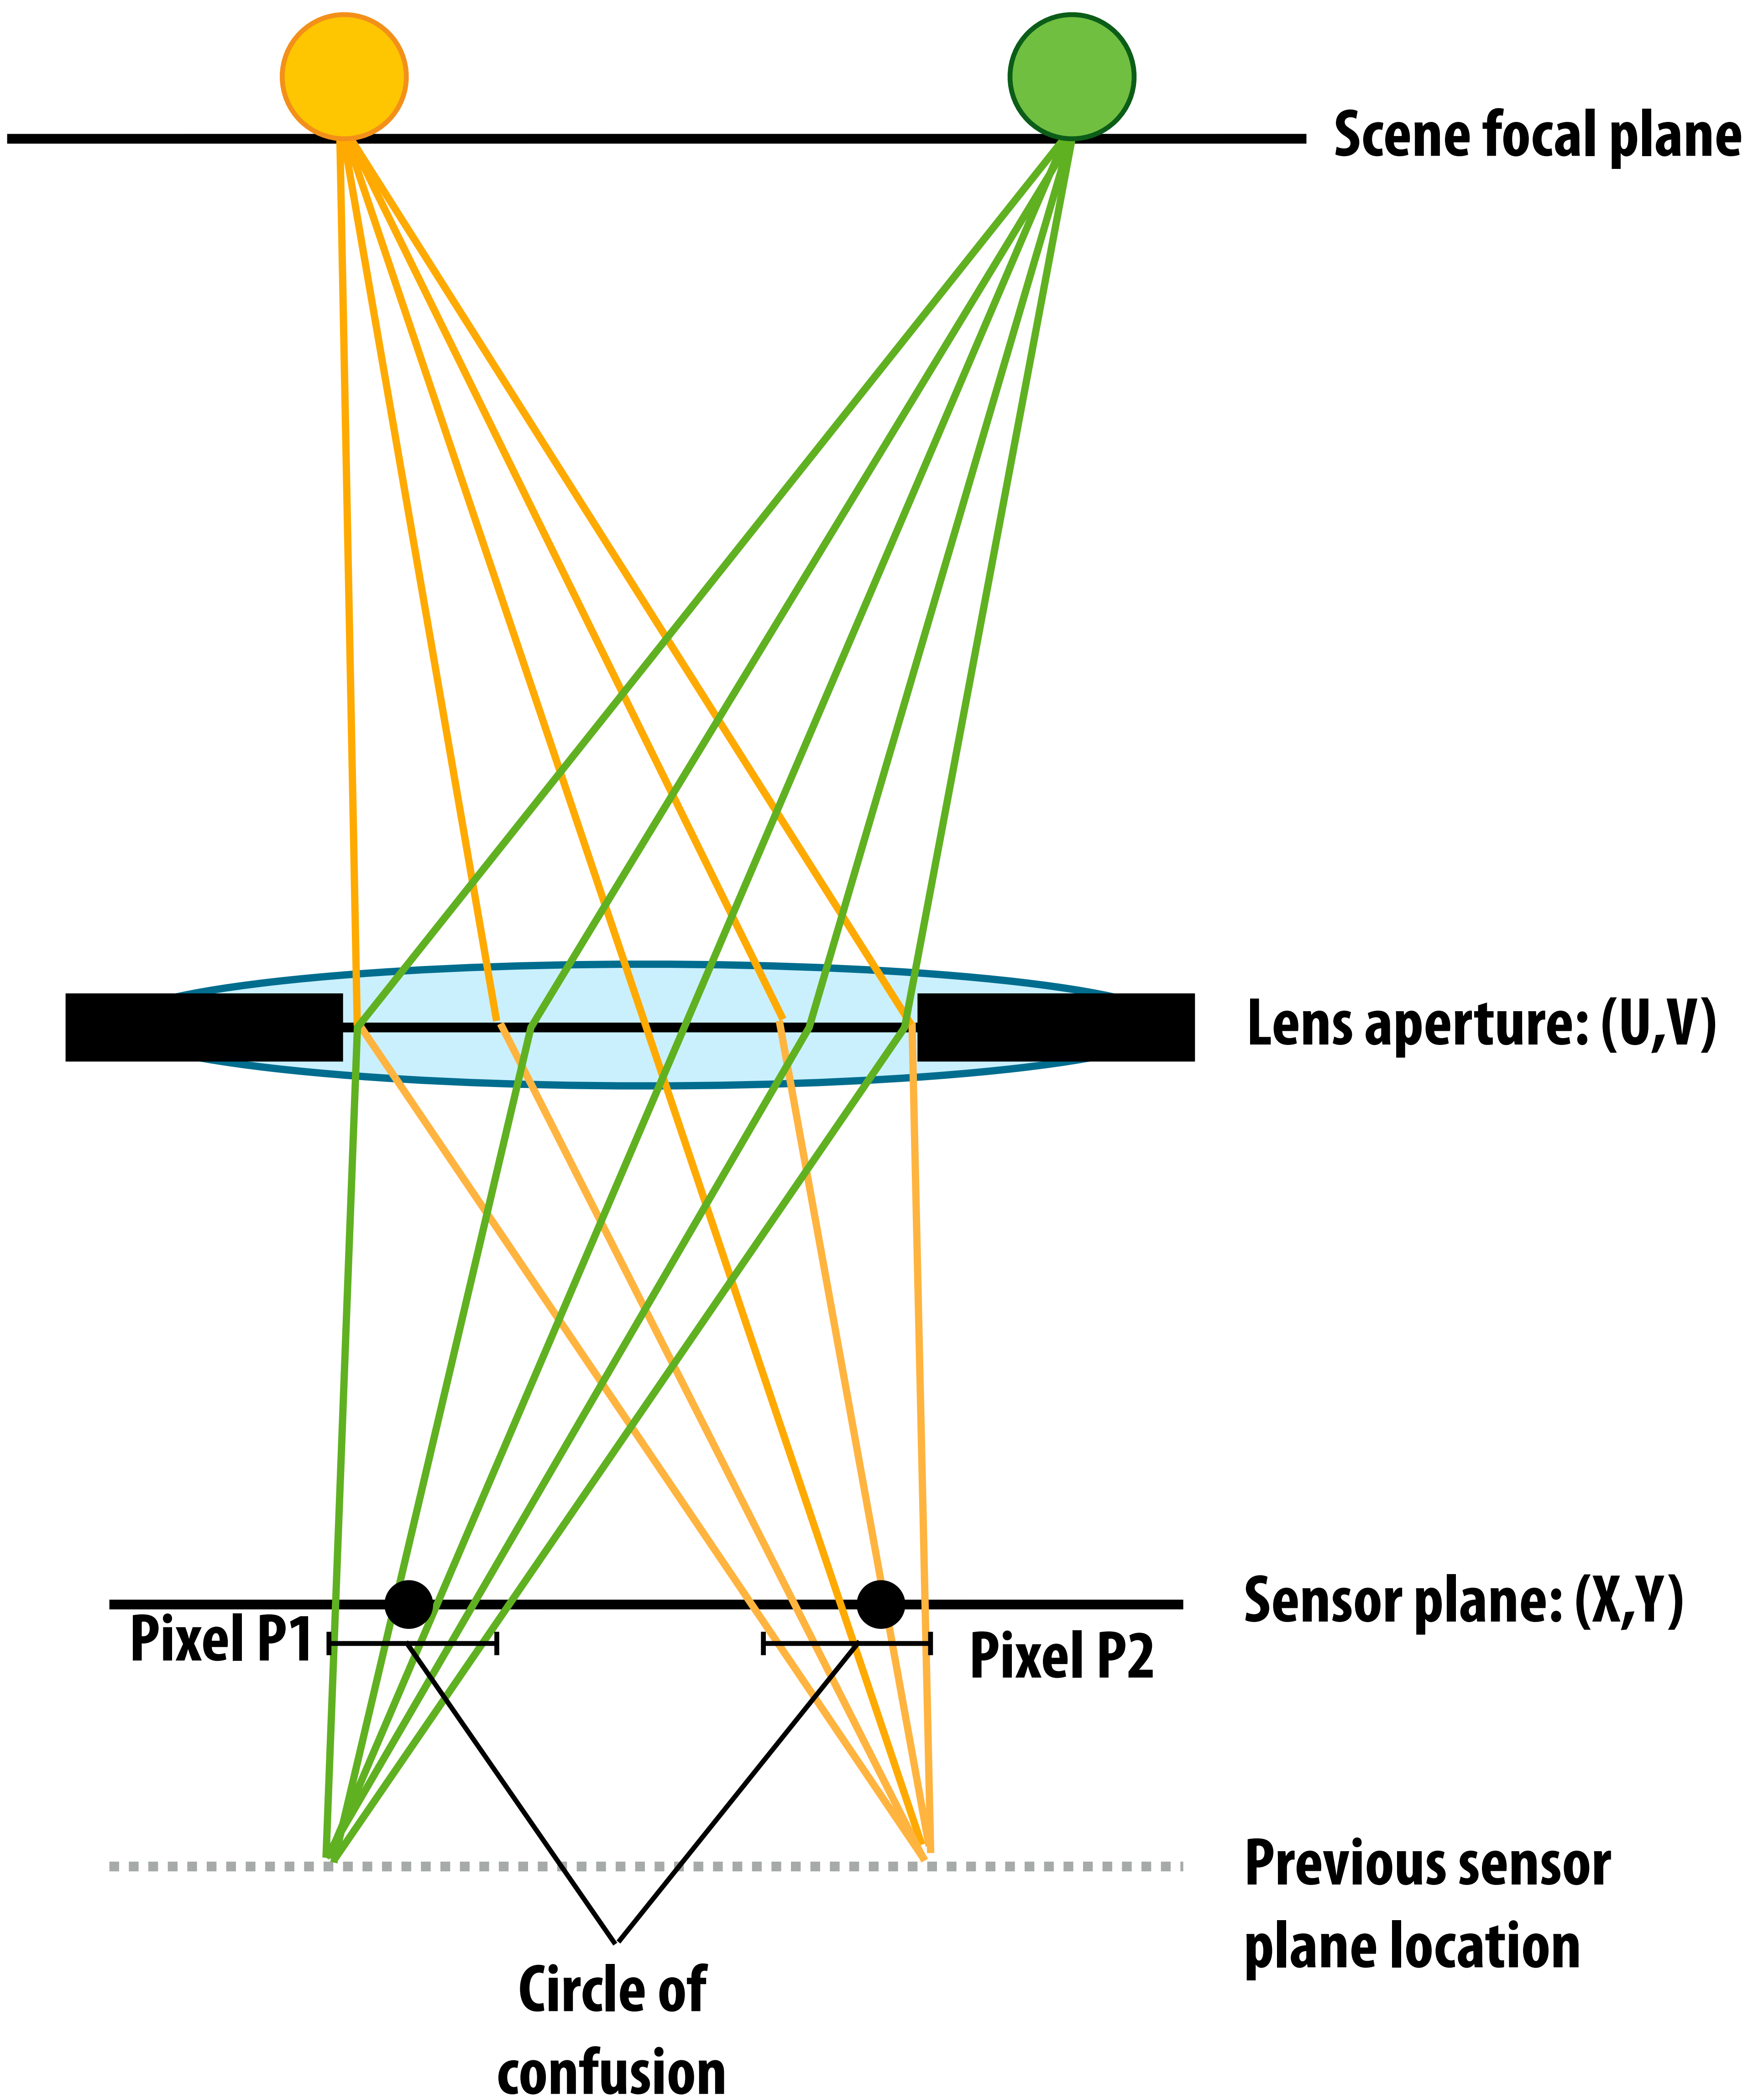

Ray space plot

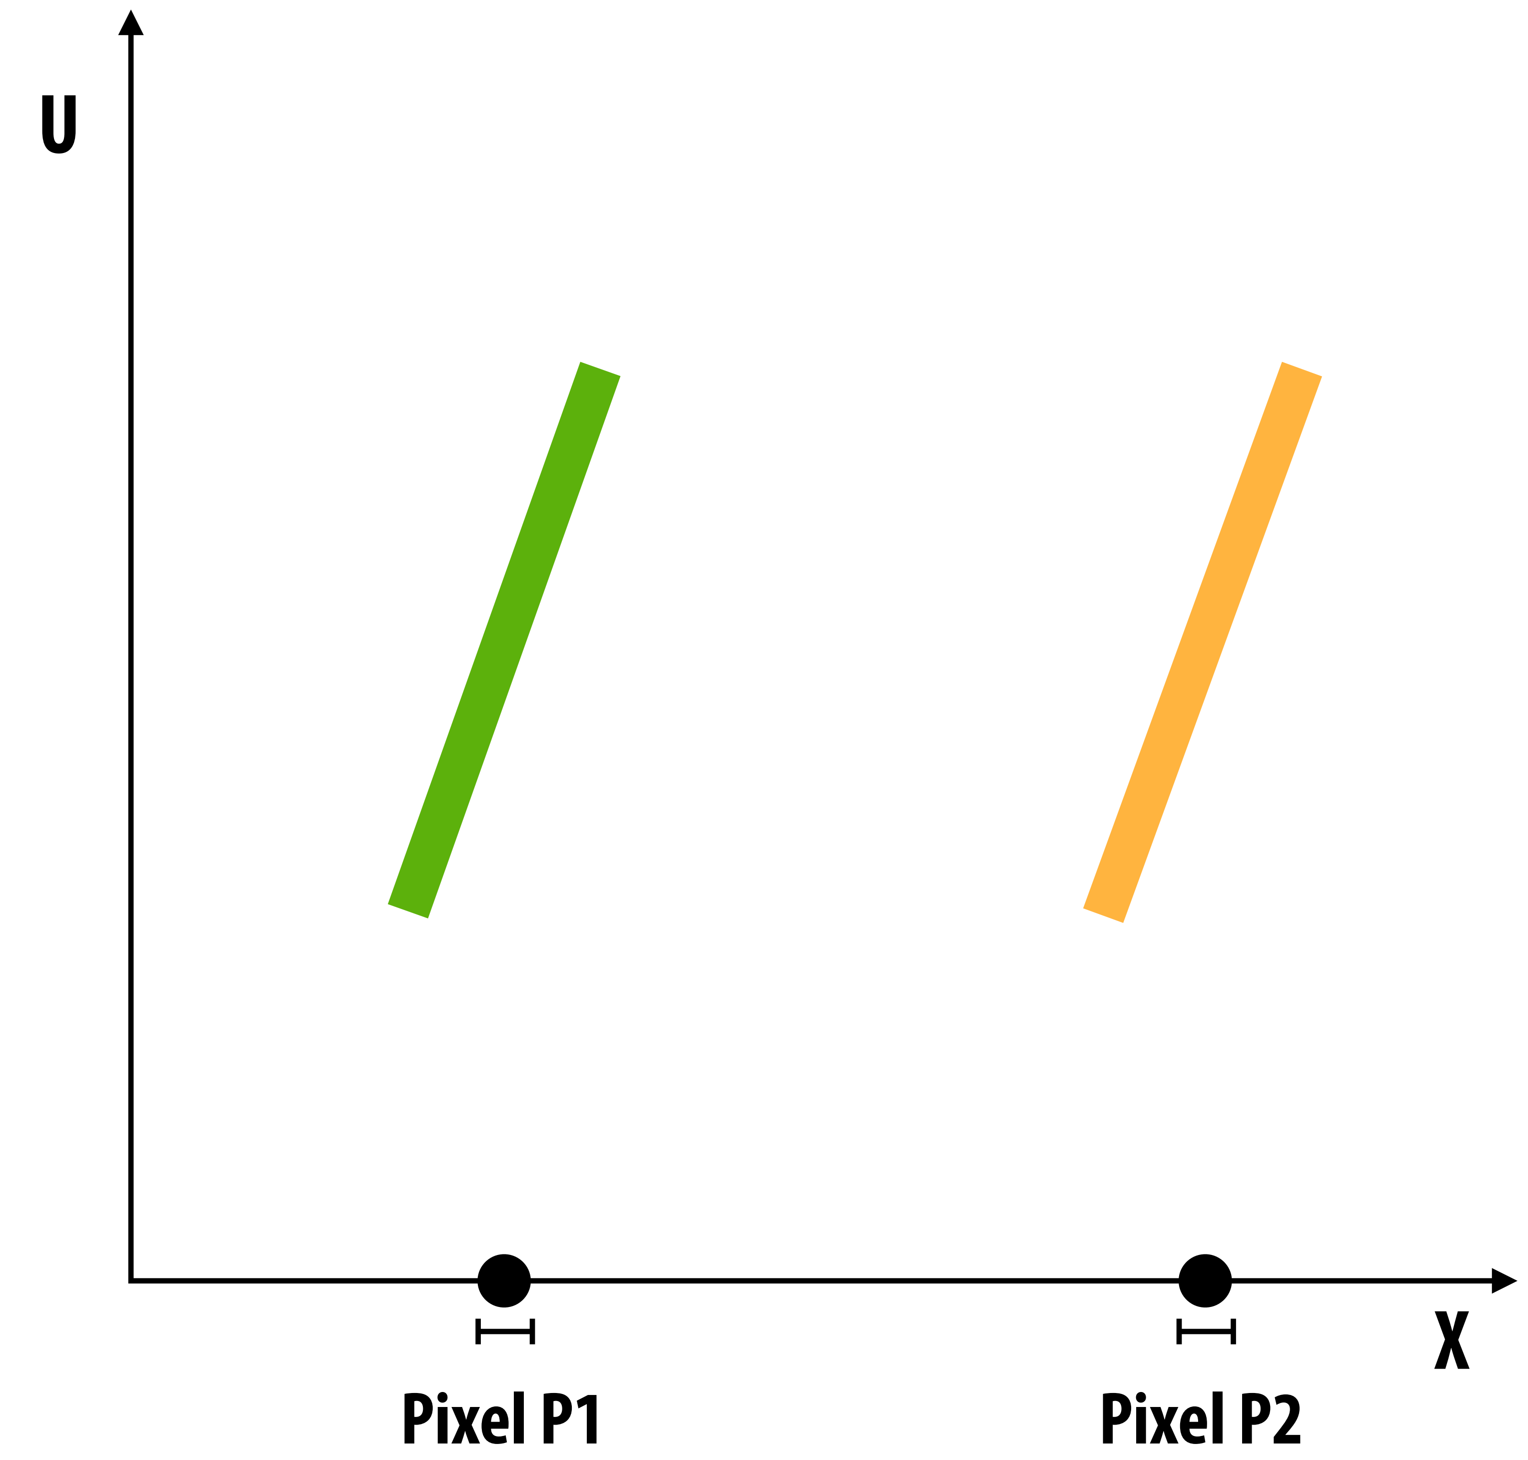

Defocus

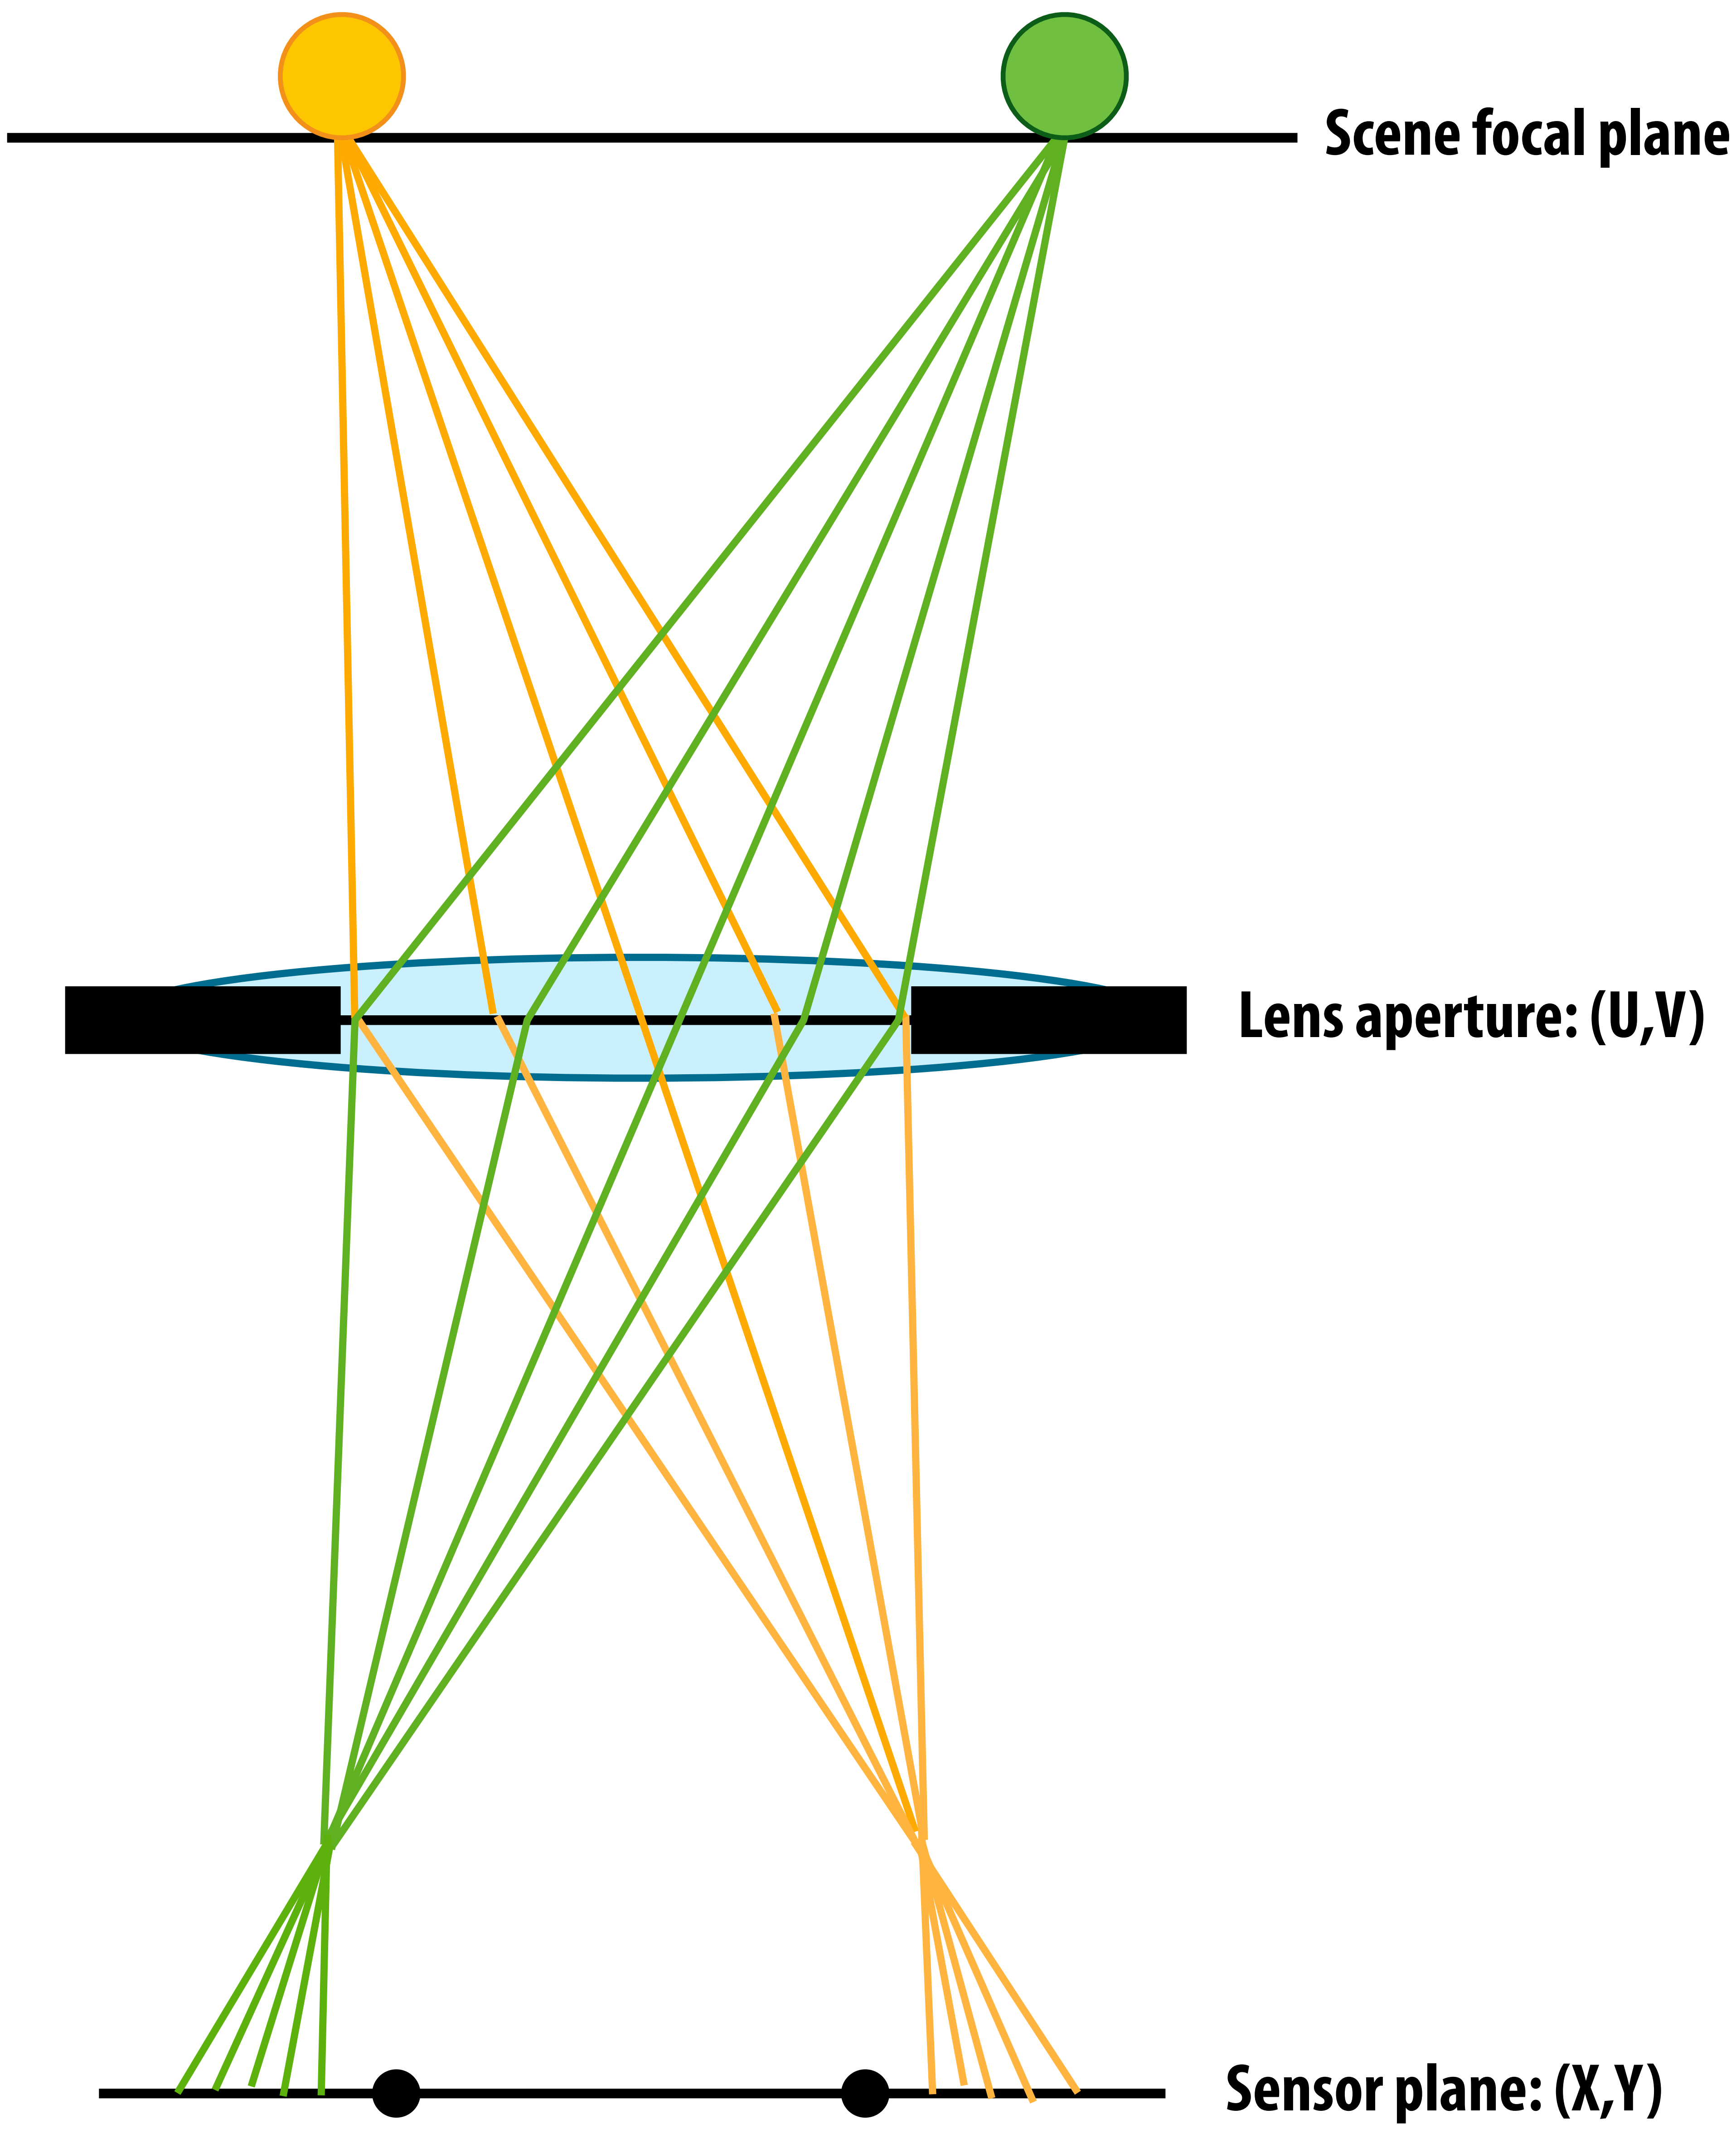

Scene focal plane
Lens aperture: (U,V)
Sensor plane: (X,Y)

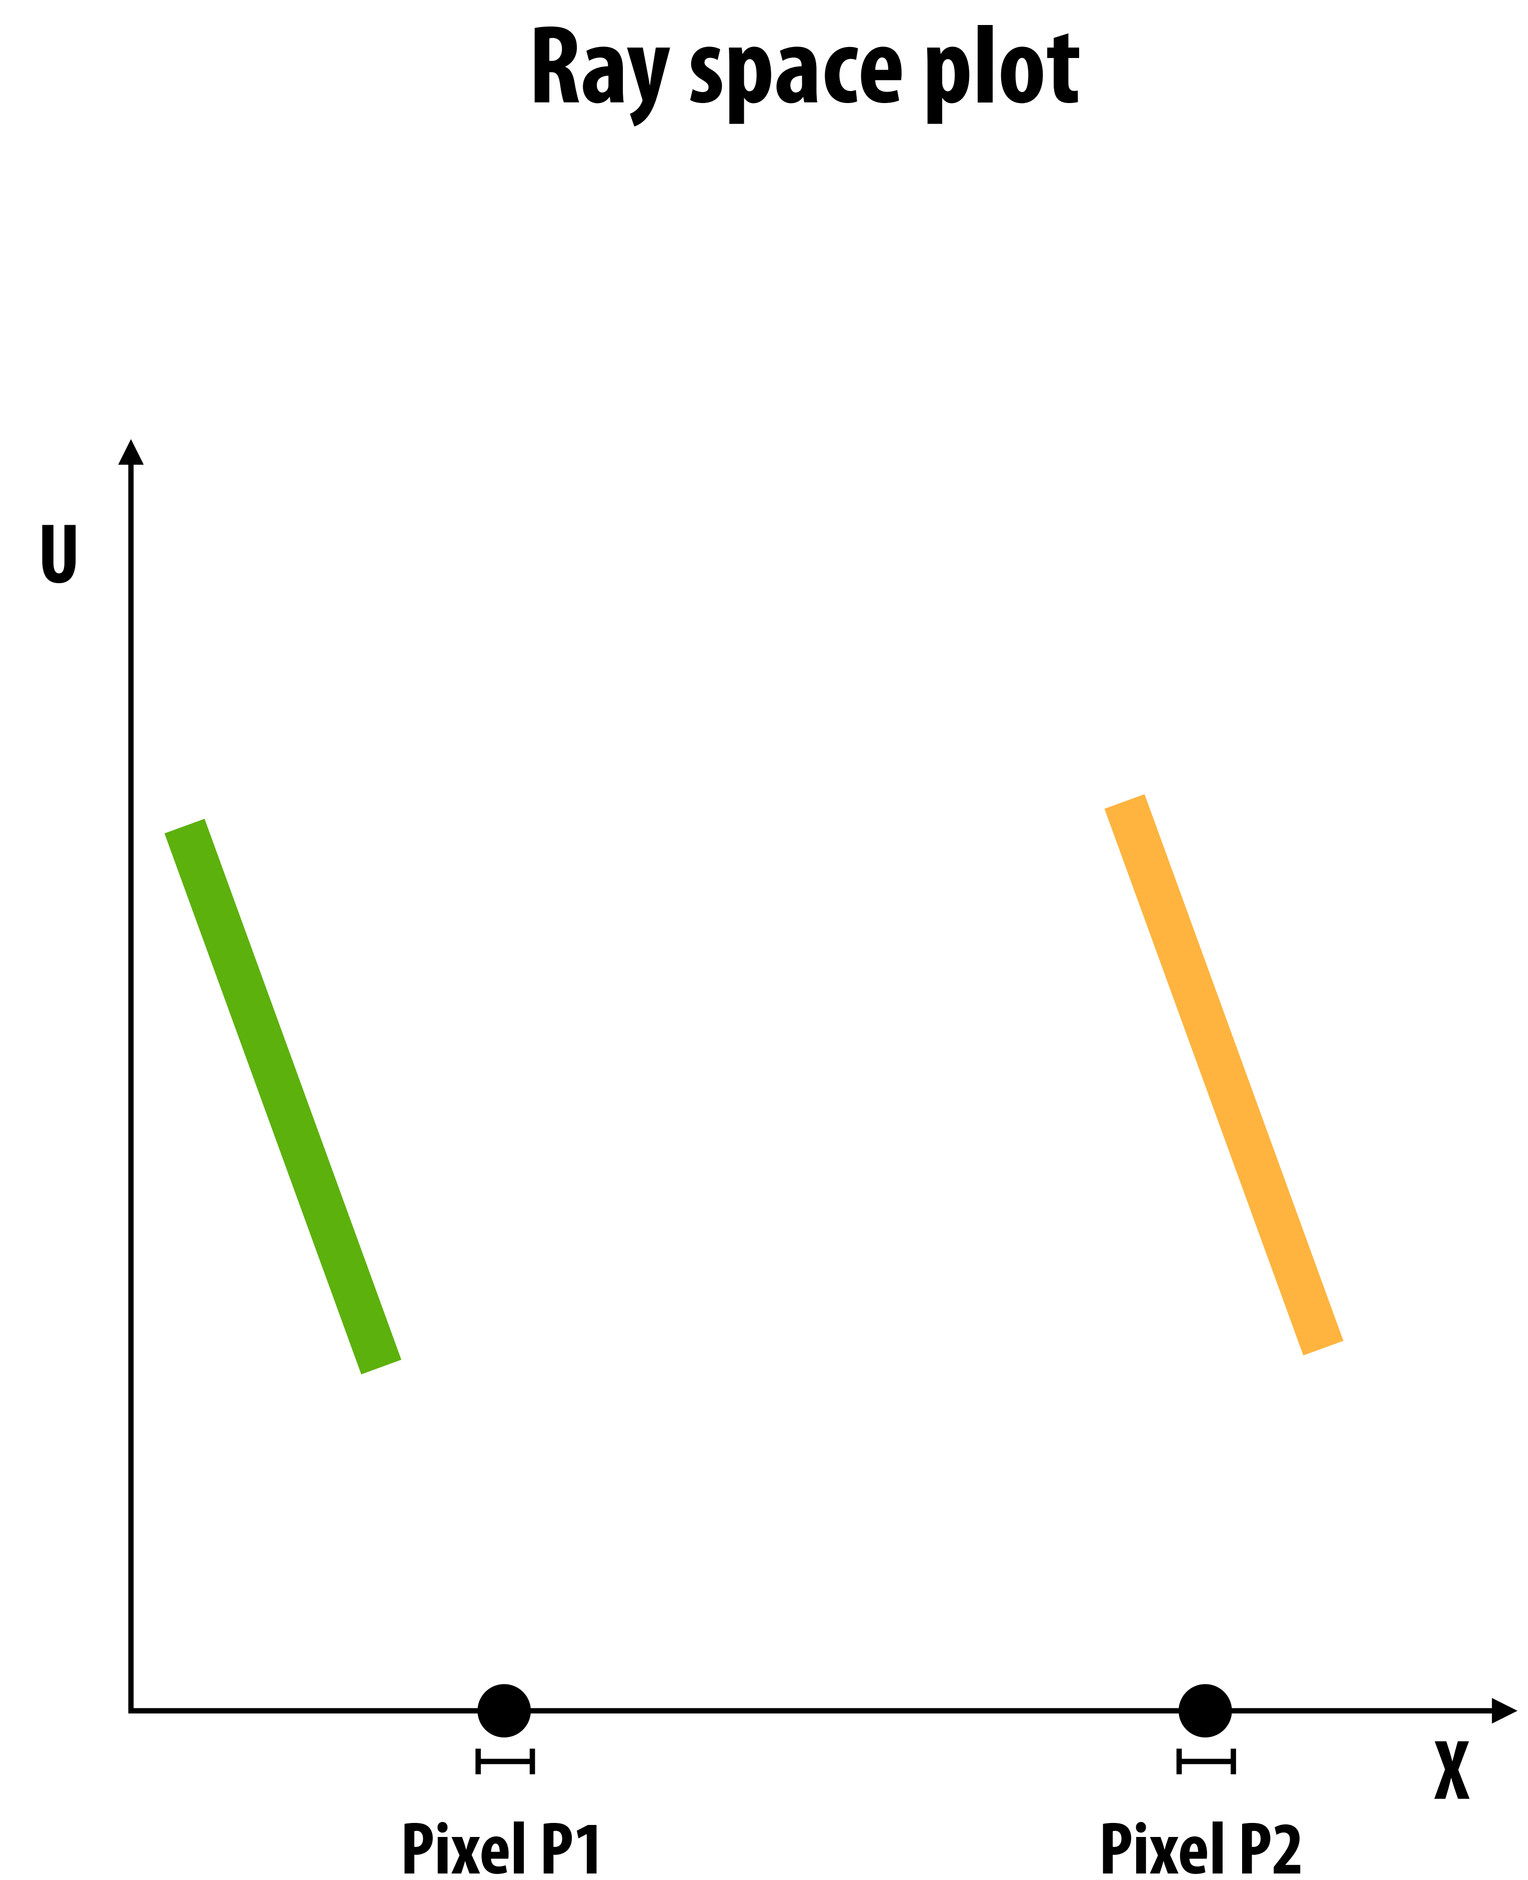

Ray space plot
U
X
Pixel P1
Pixel P2

How might we measure the light field inside a camera?

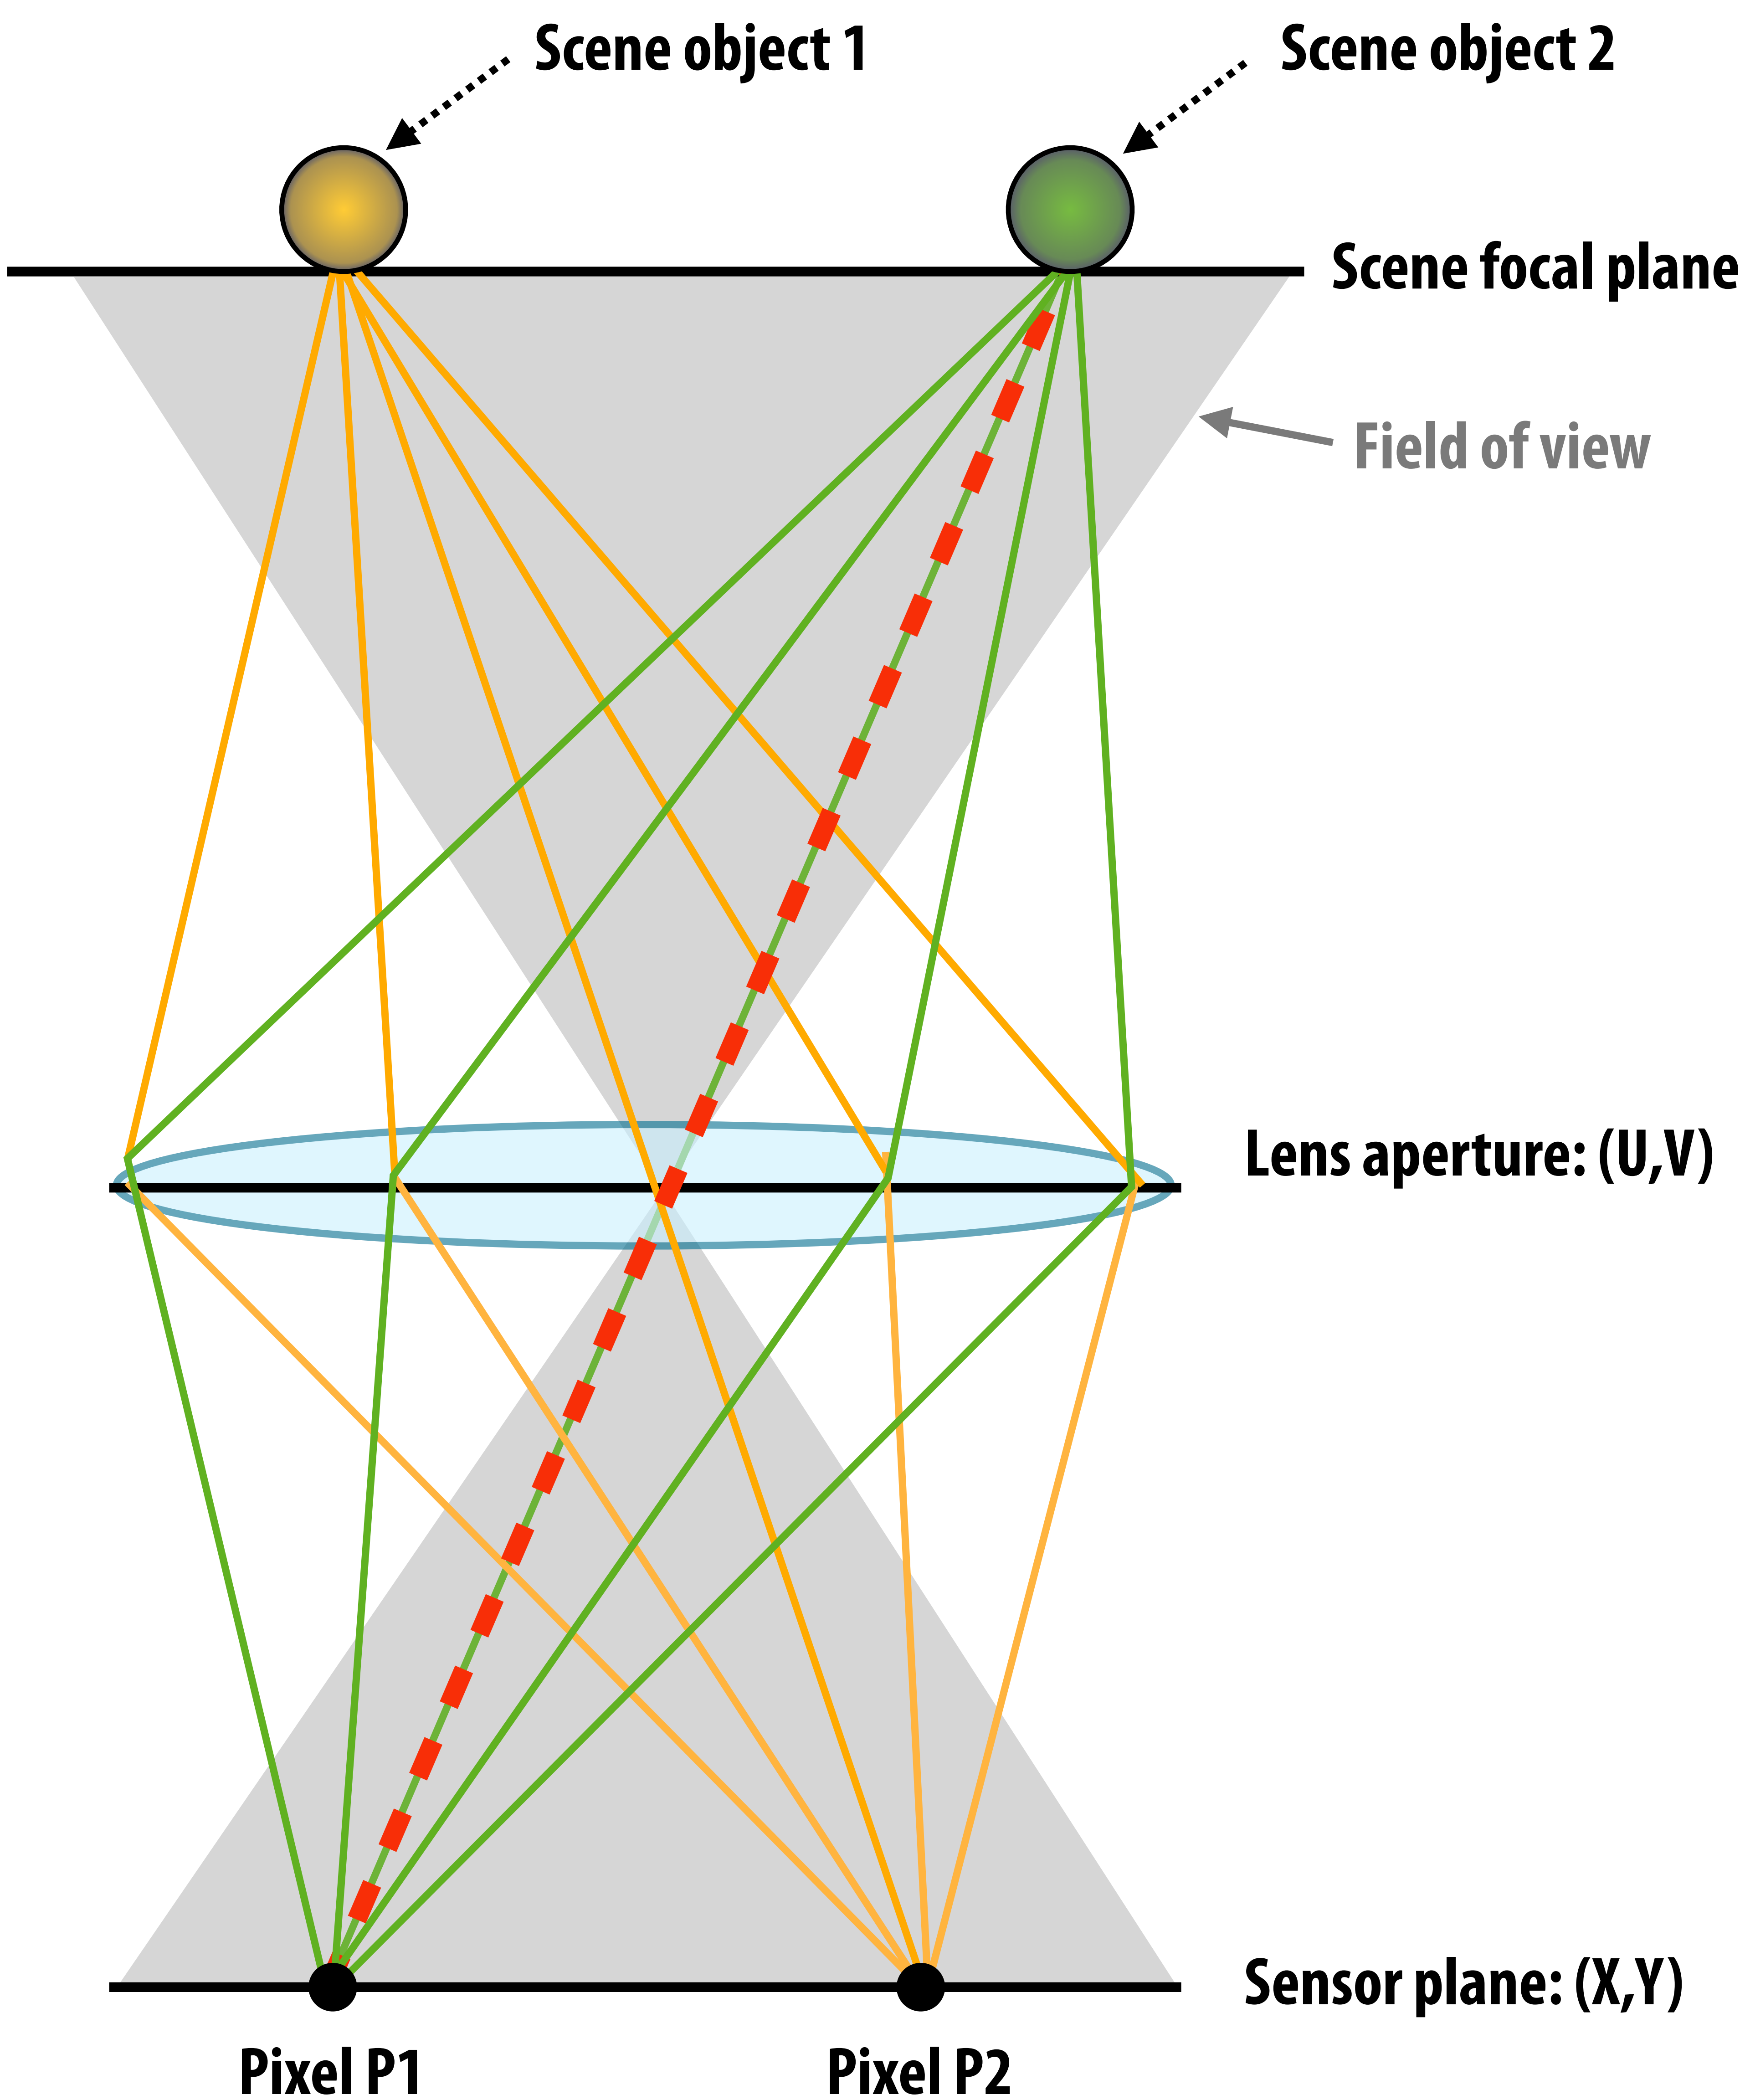

Ray space plot
(only showing X-U 2D projection)

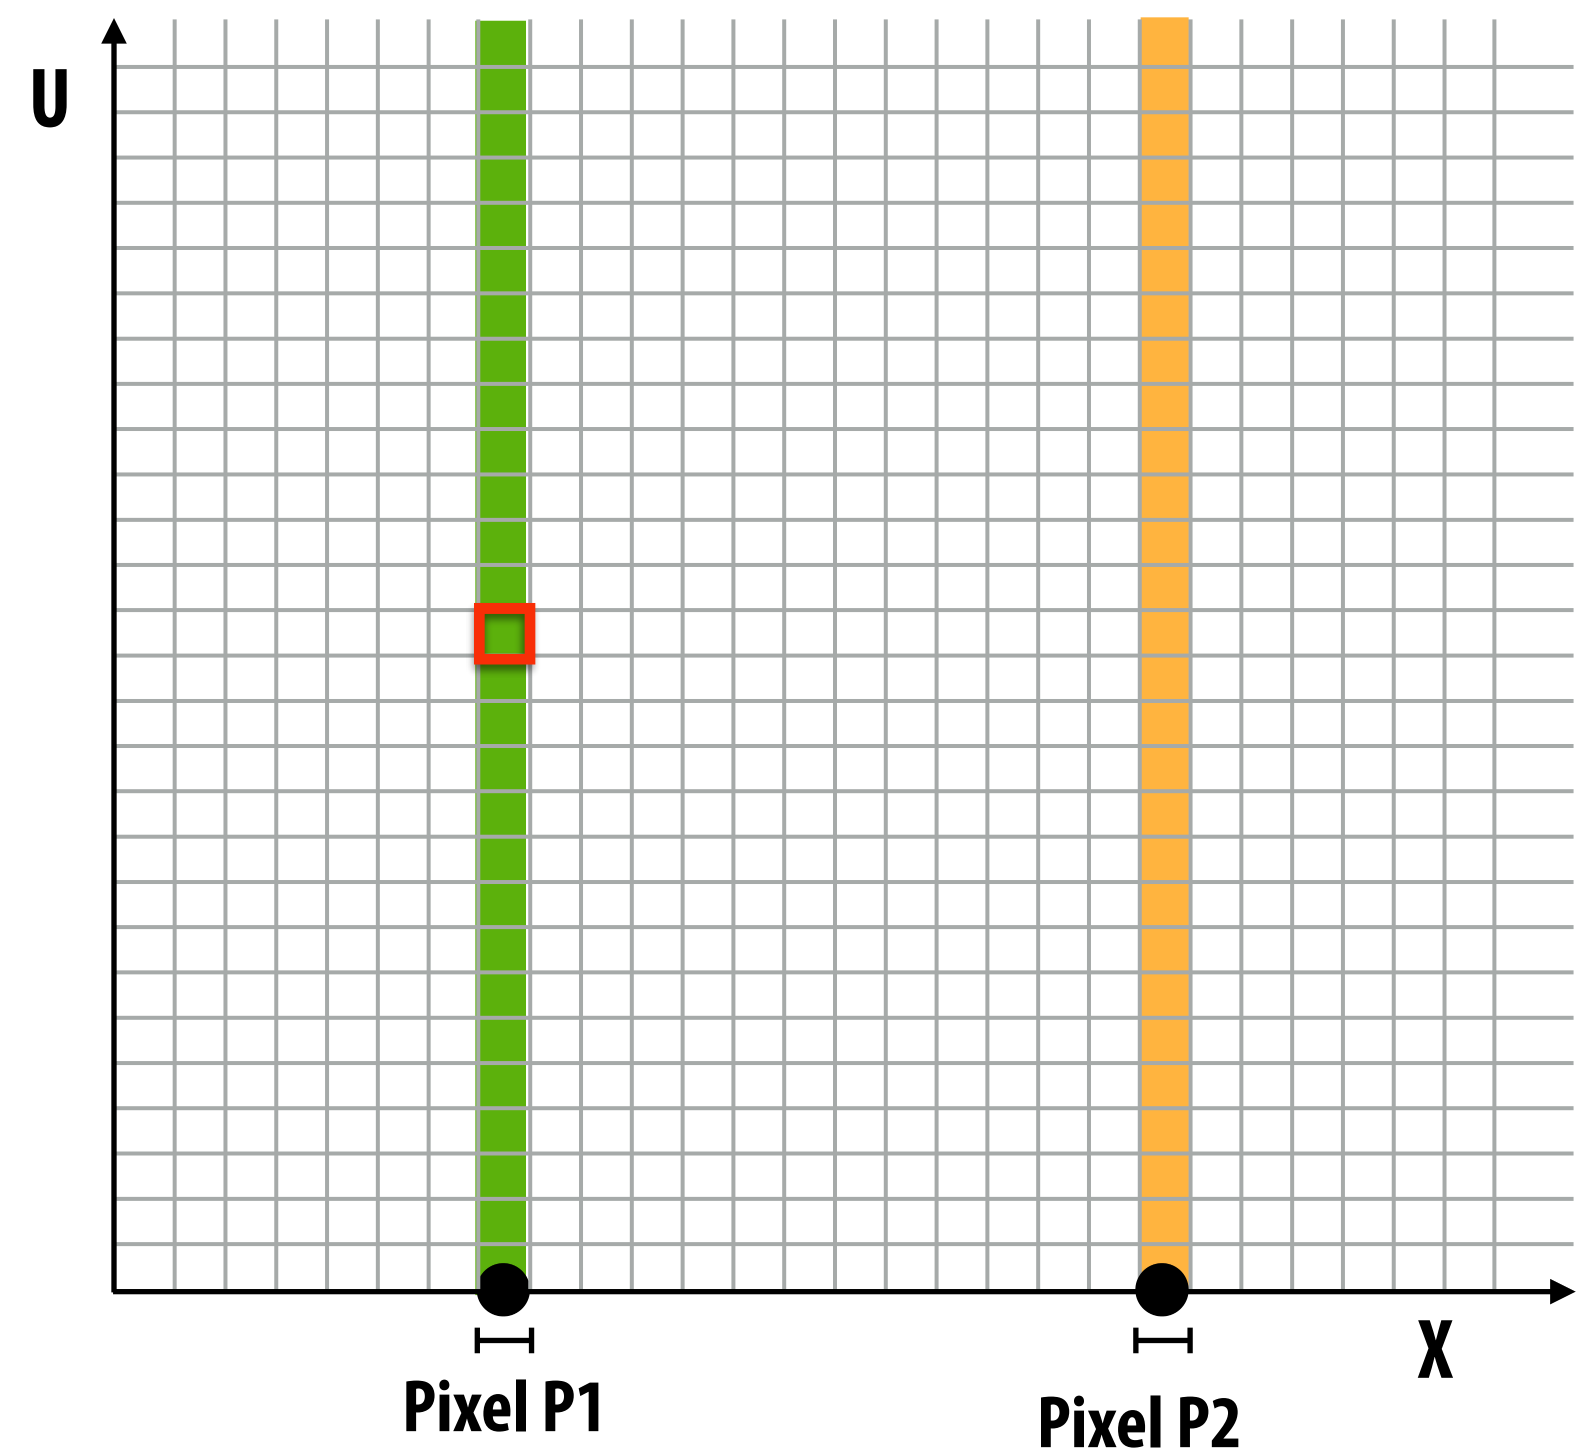

Intuition: handheld light field camera

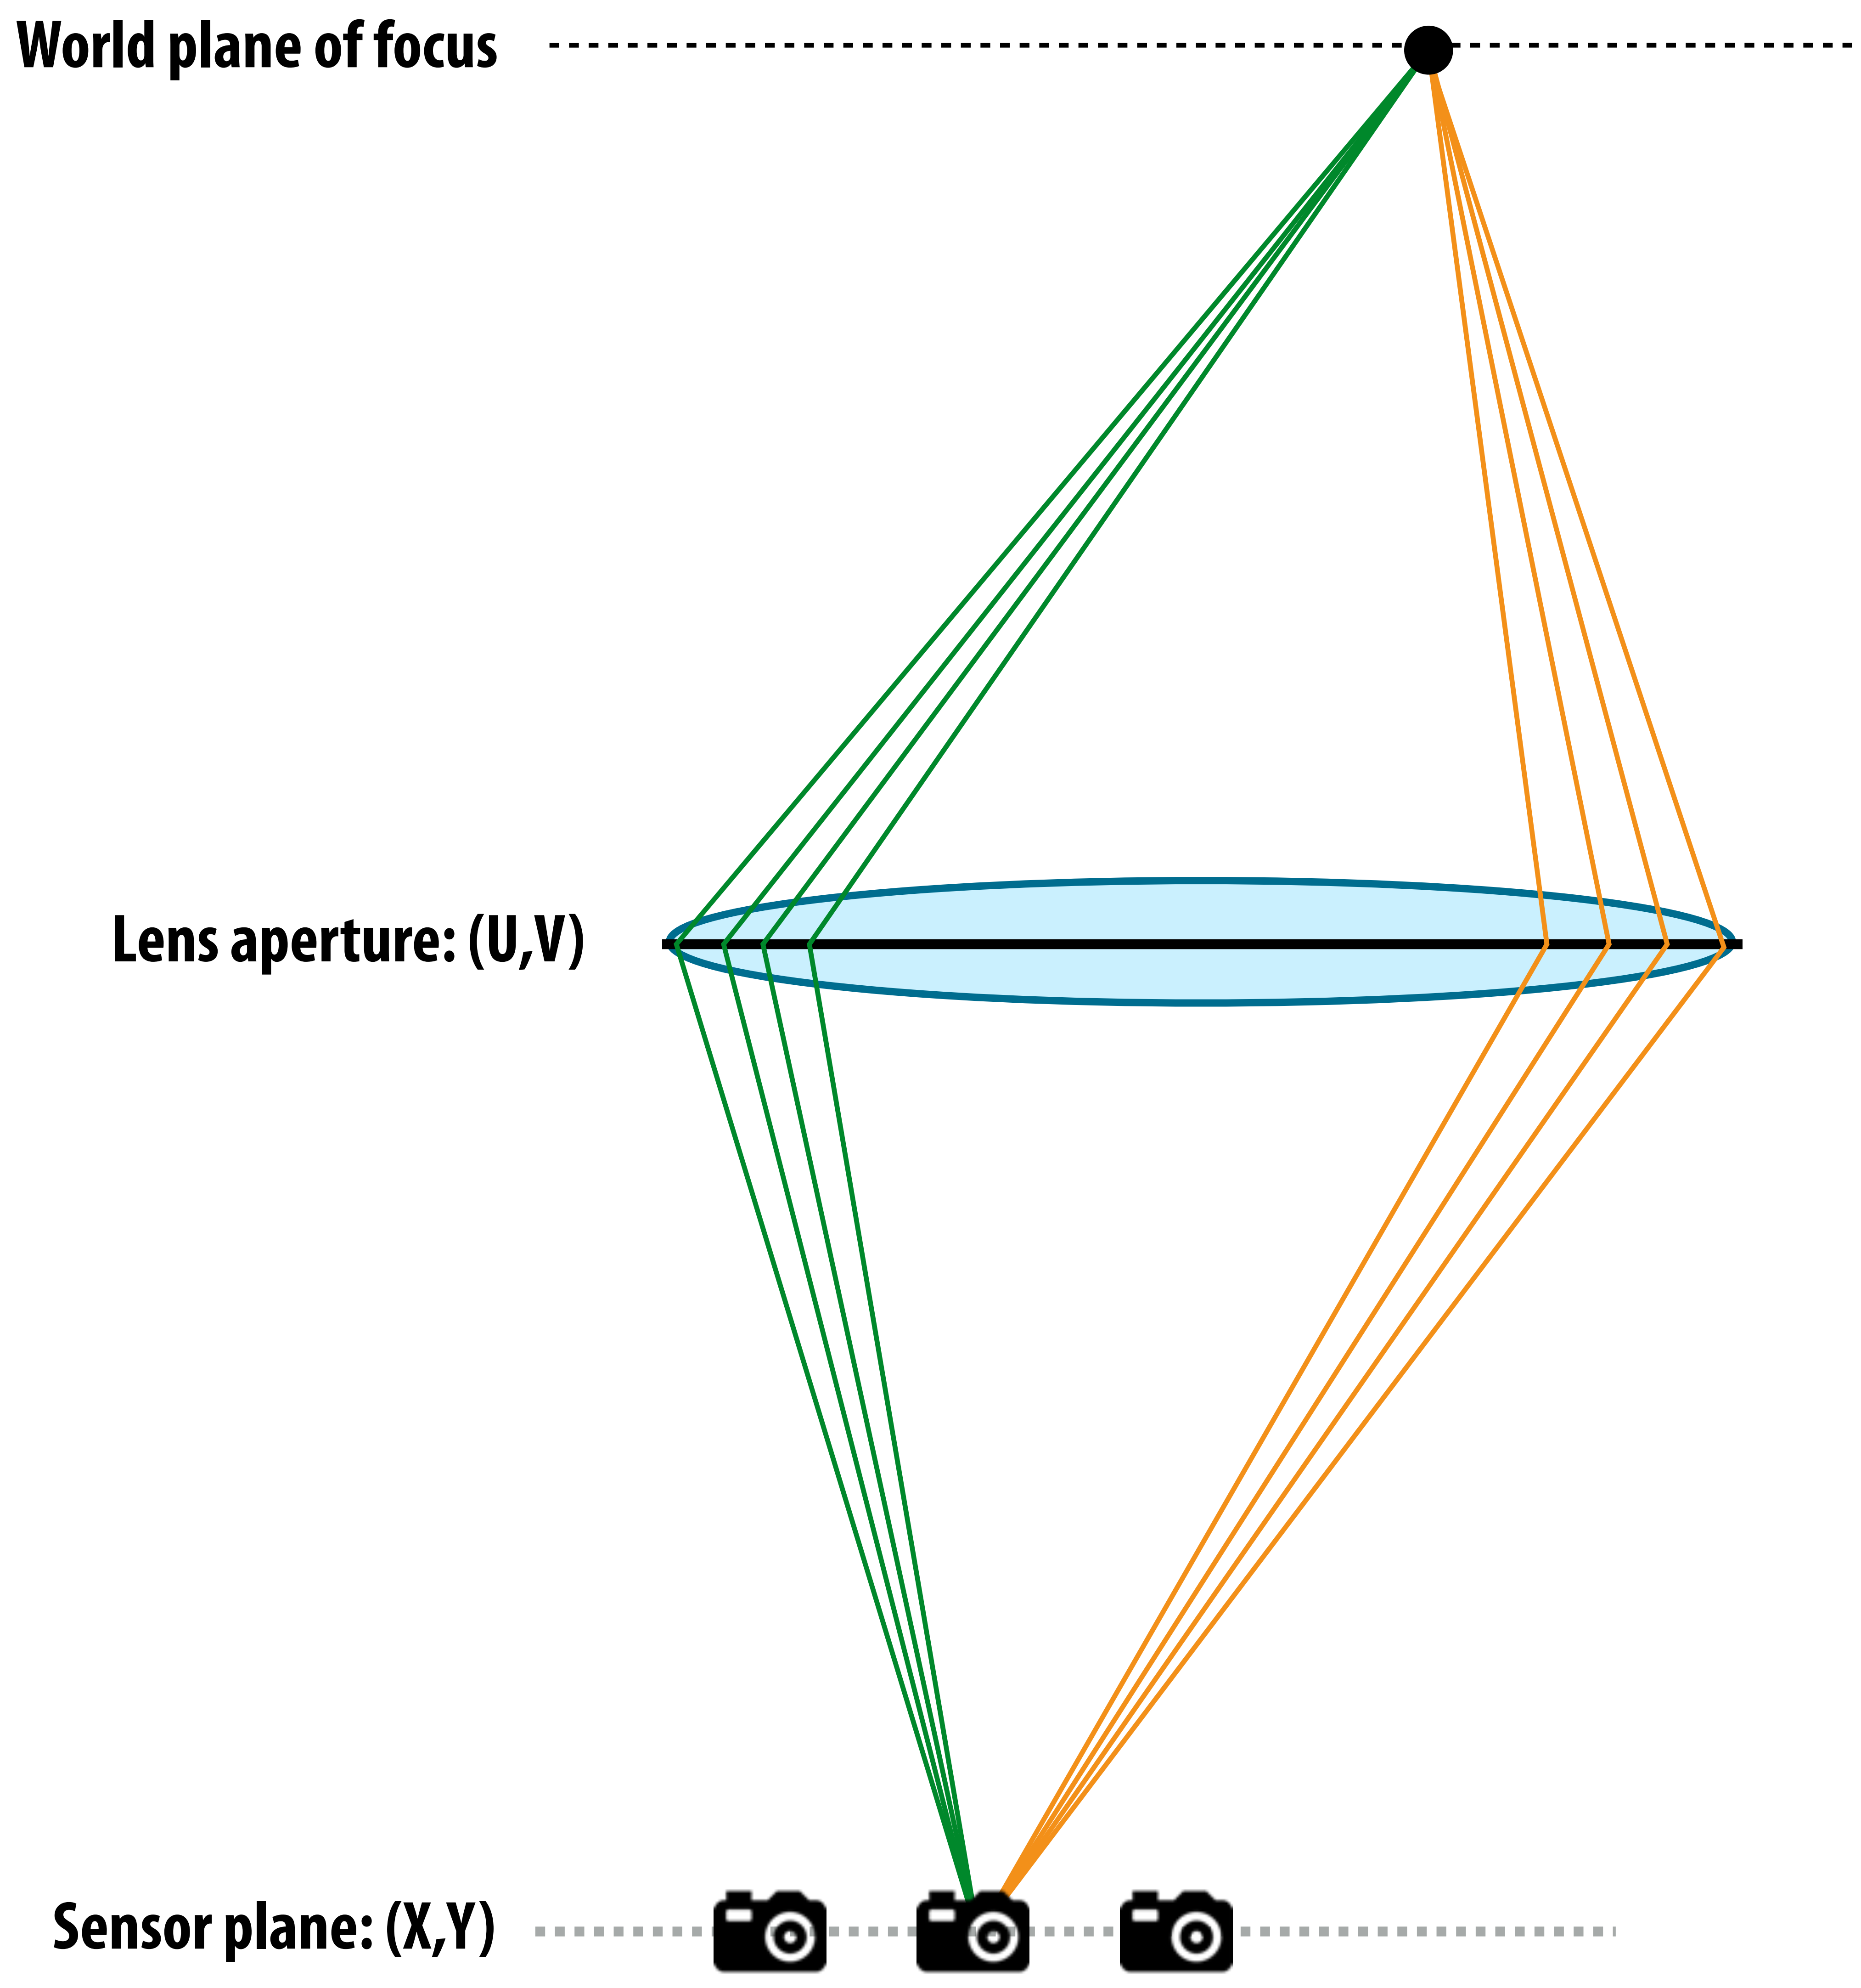

Intuition: build an optical system where each region of the sensor "takes" a picture of the aperture of the main lens

Handheld light field camera

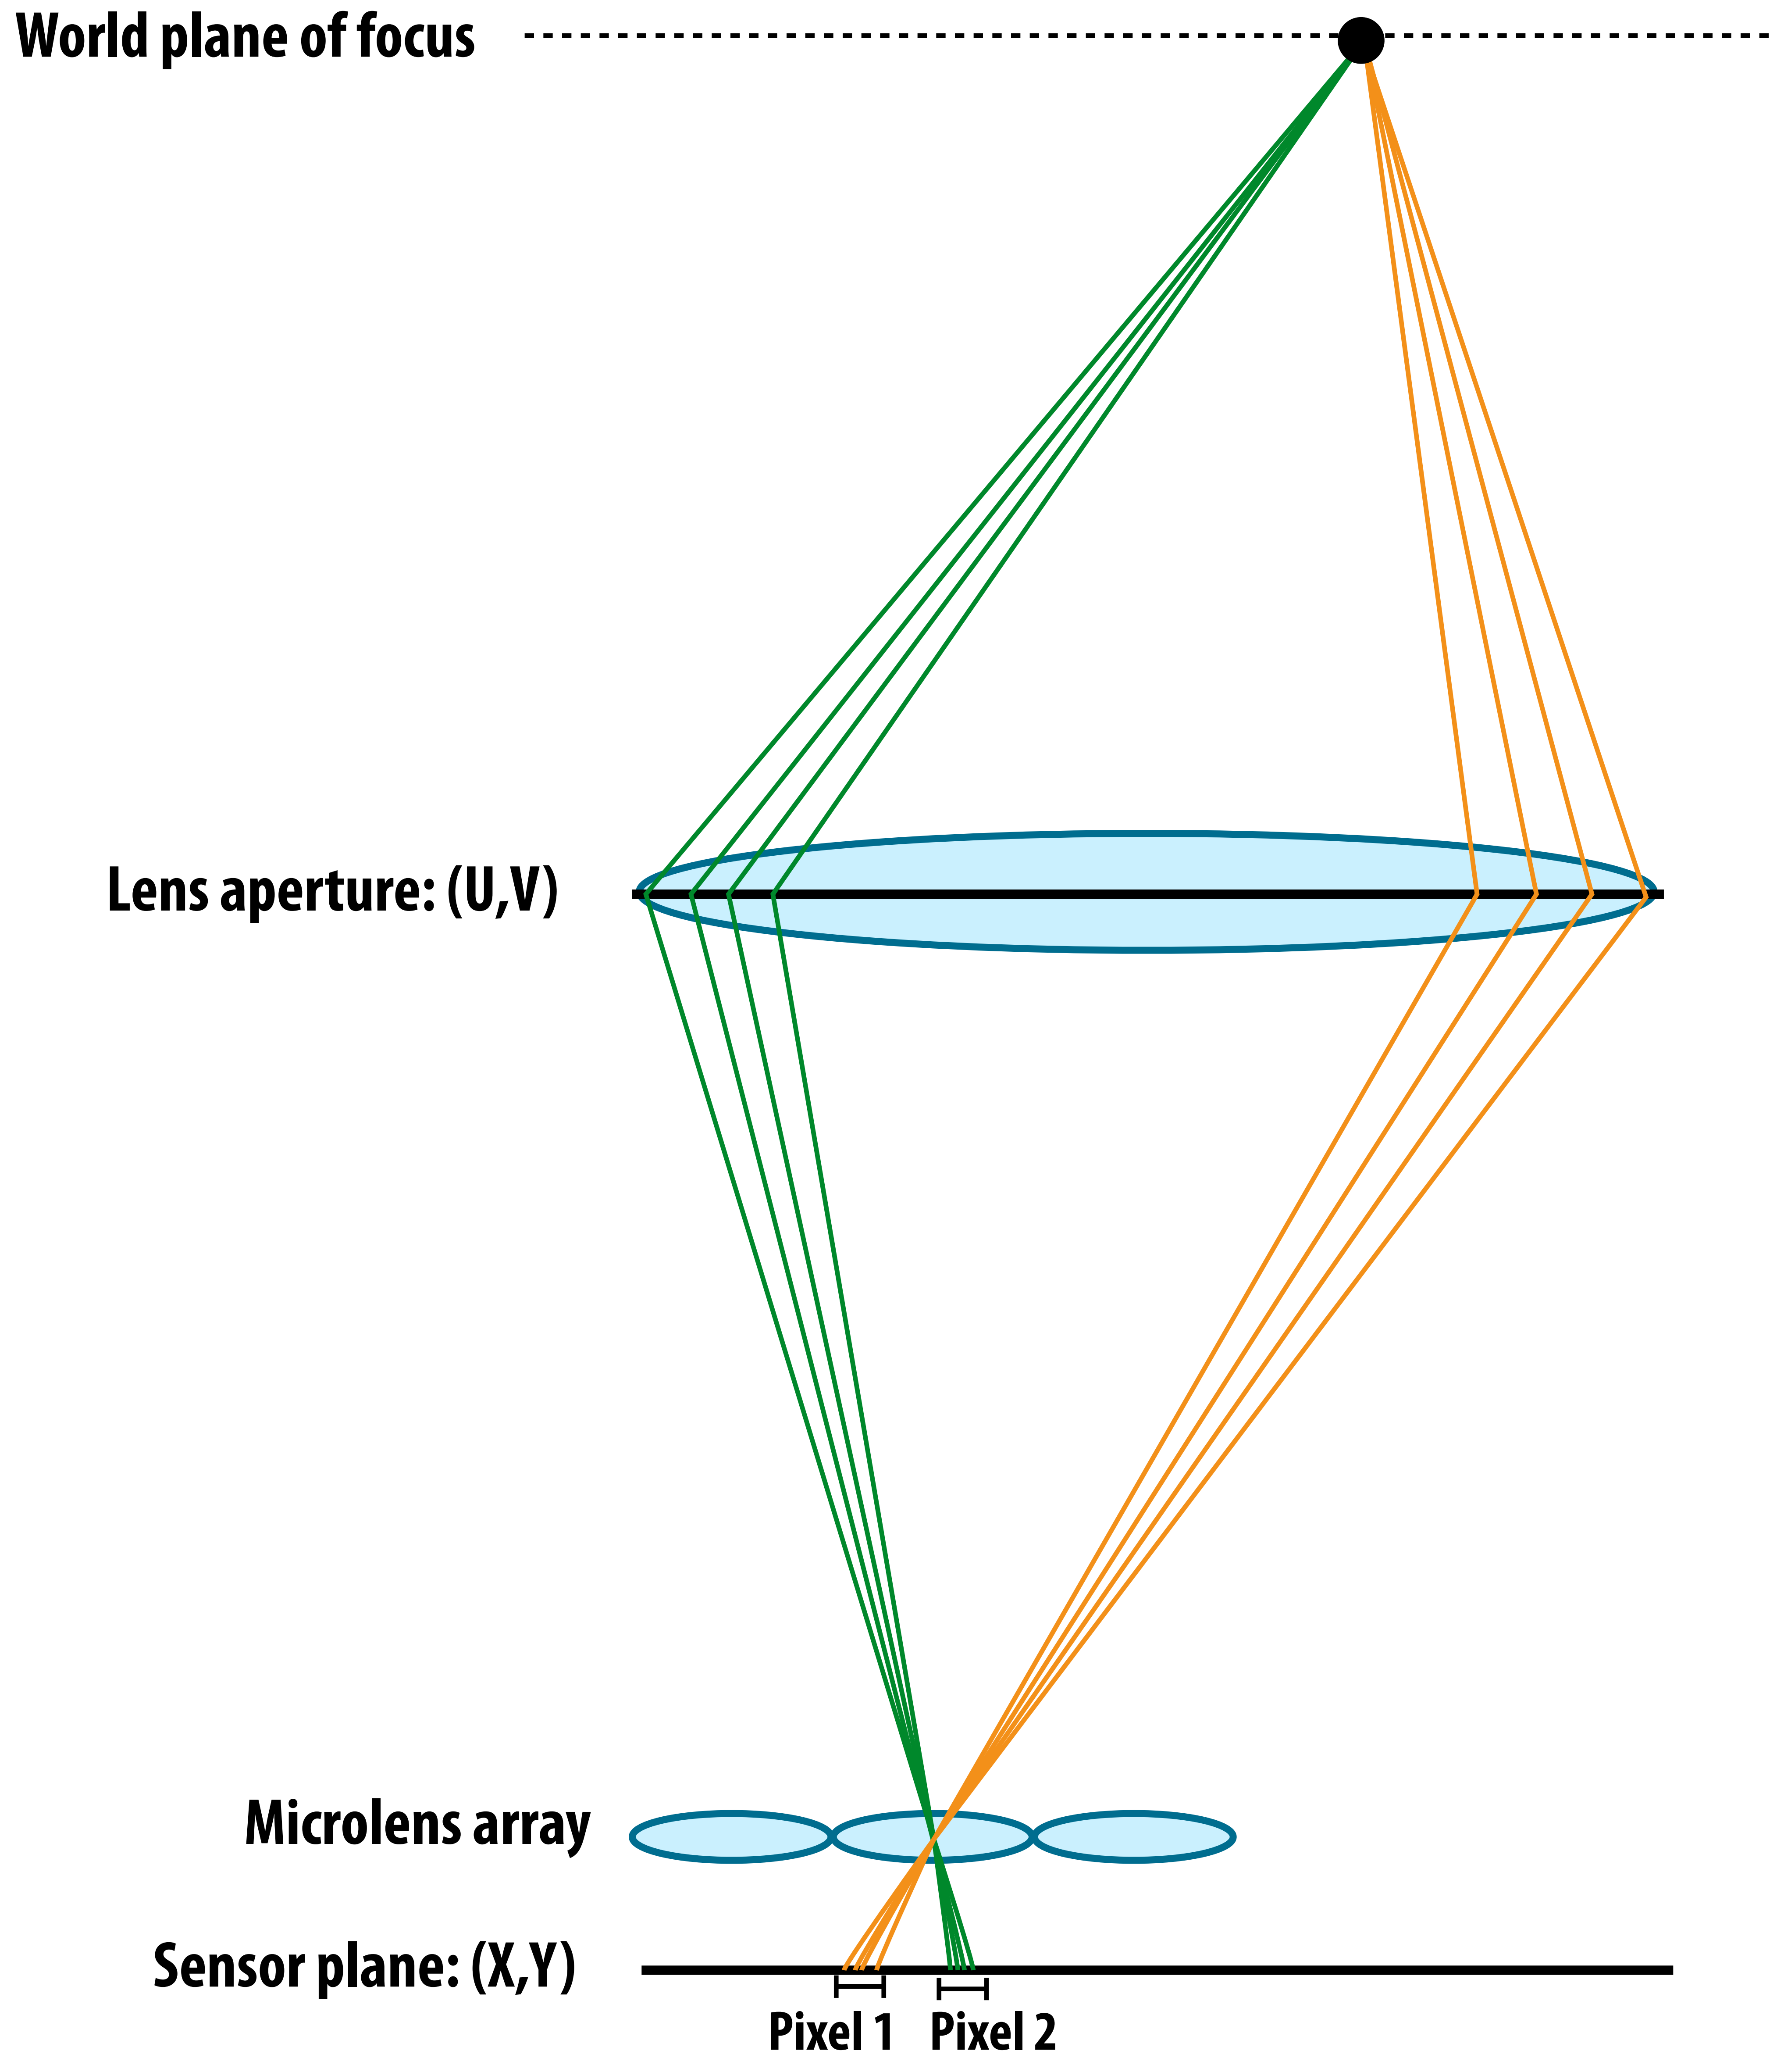

[Ng et al. 2005]

[Adelson and Wang, 1992]

Implementation: microlens array placed just on top of the sensor.

Each sensor pixel records a small beam of light inside the camera

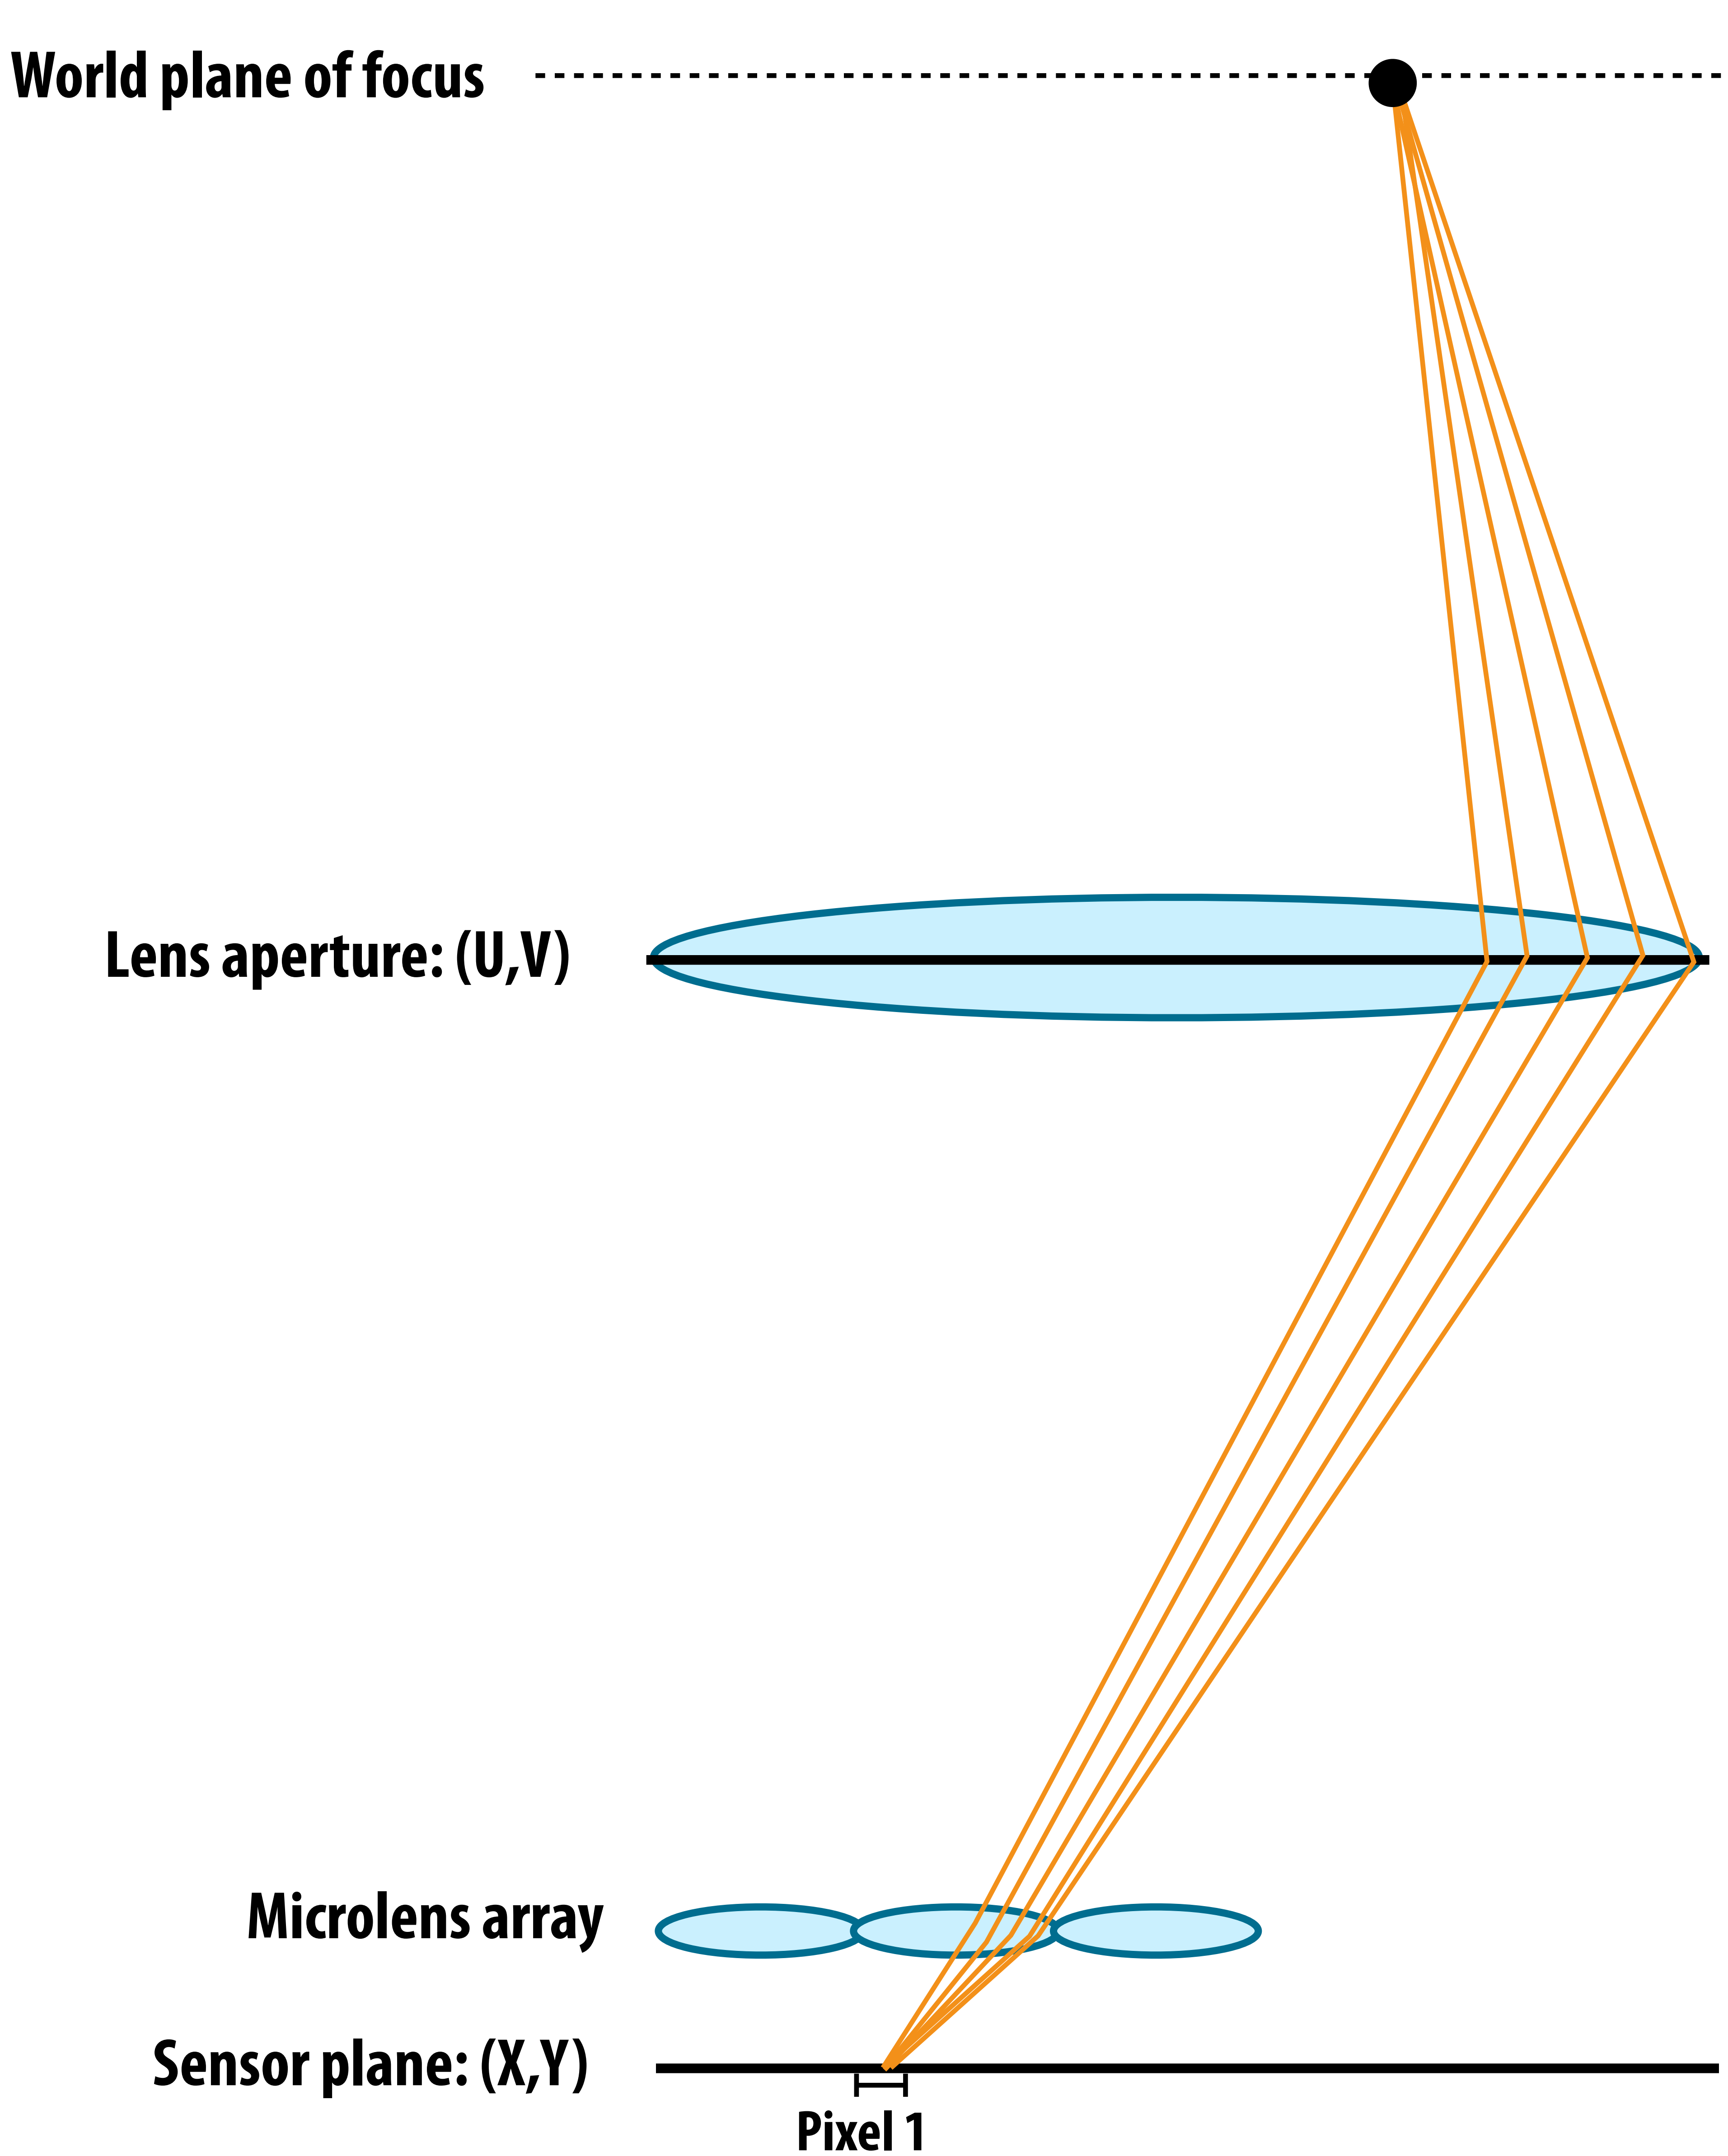

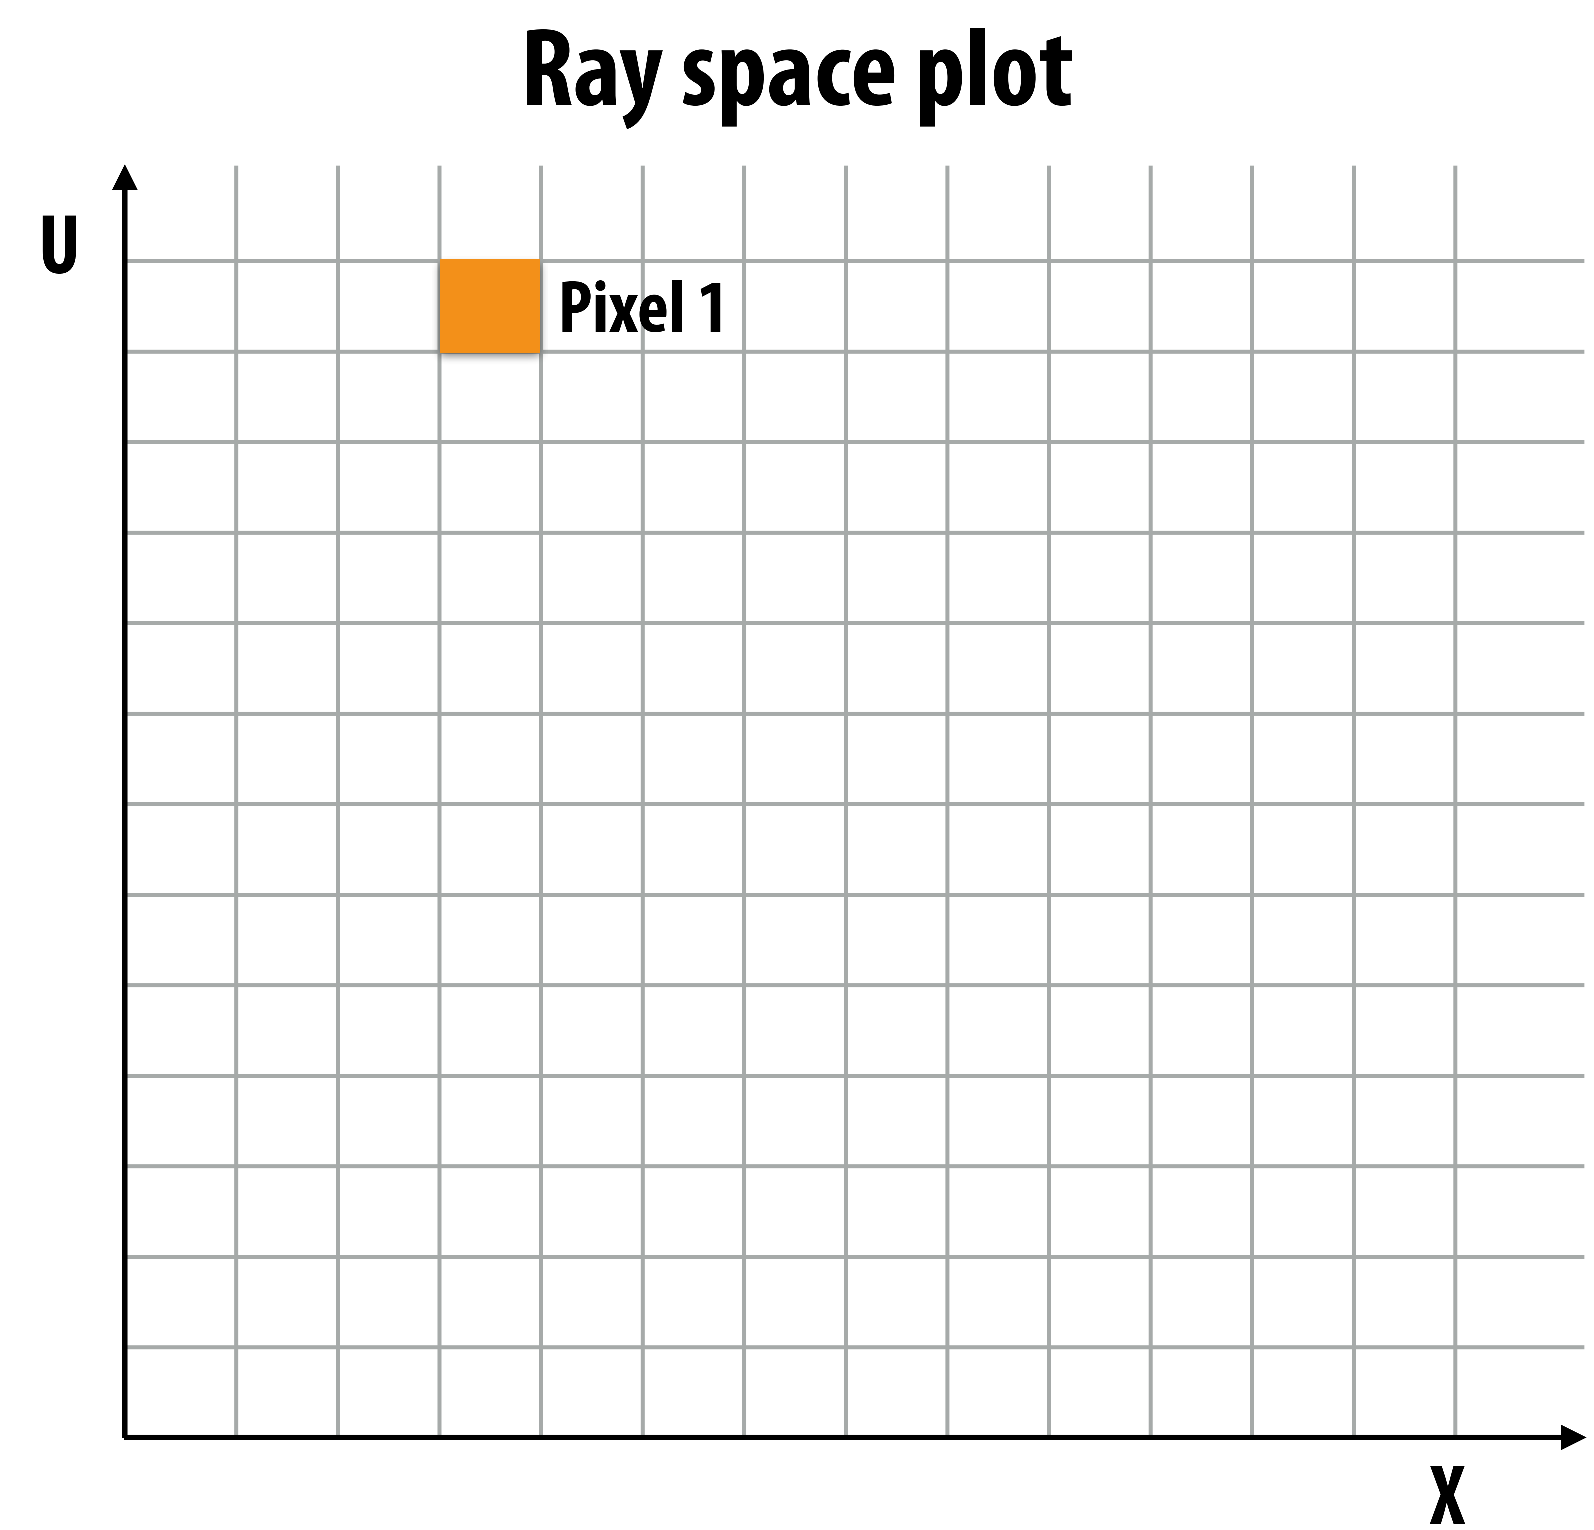

Each sensor pixel records a small beam of light inside the camera

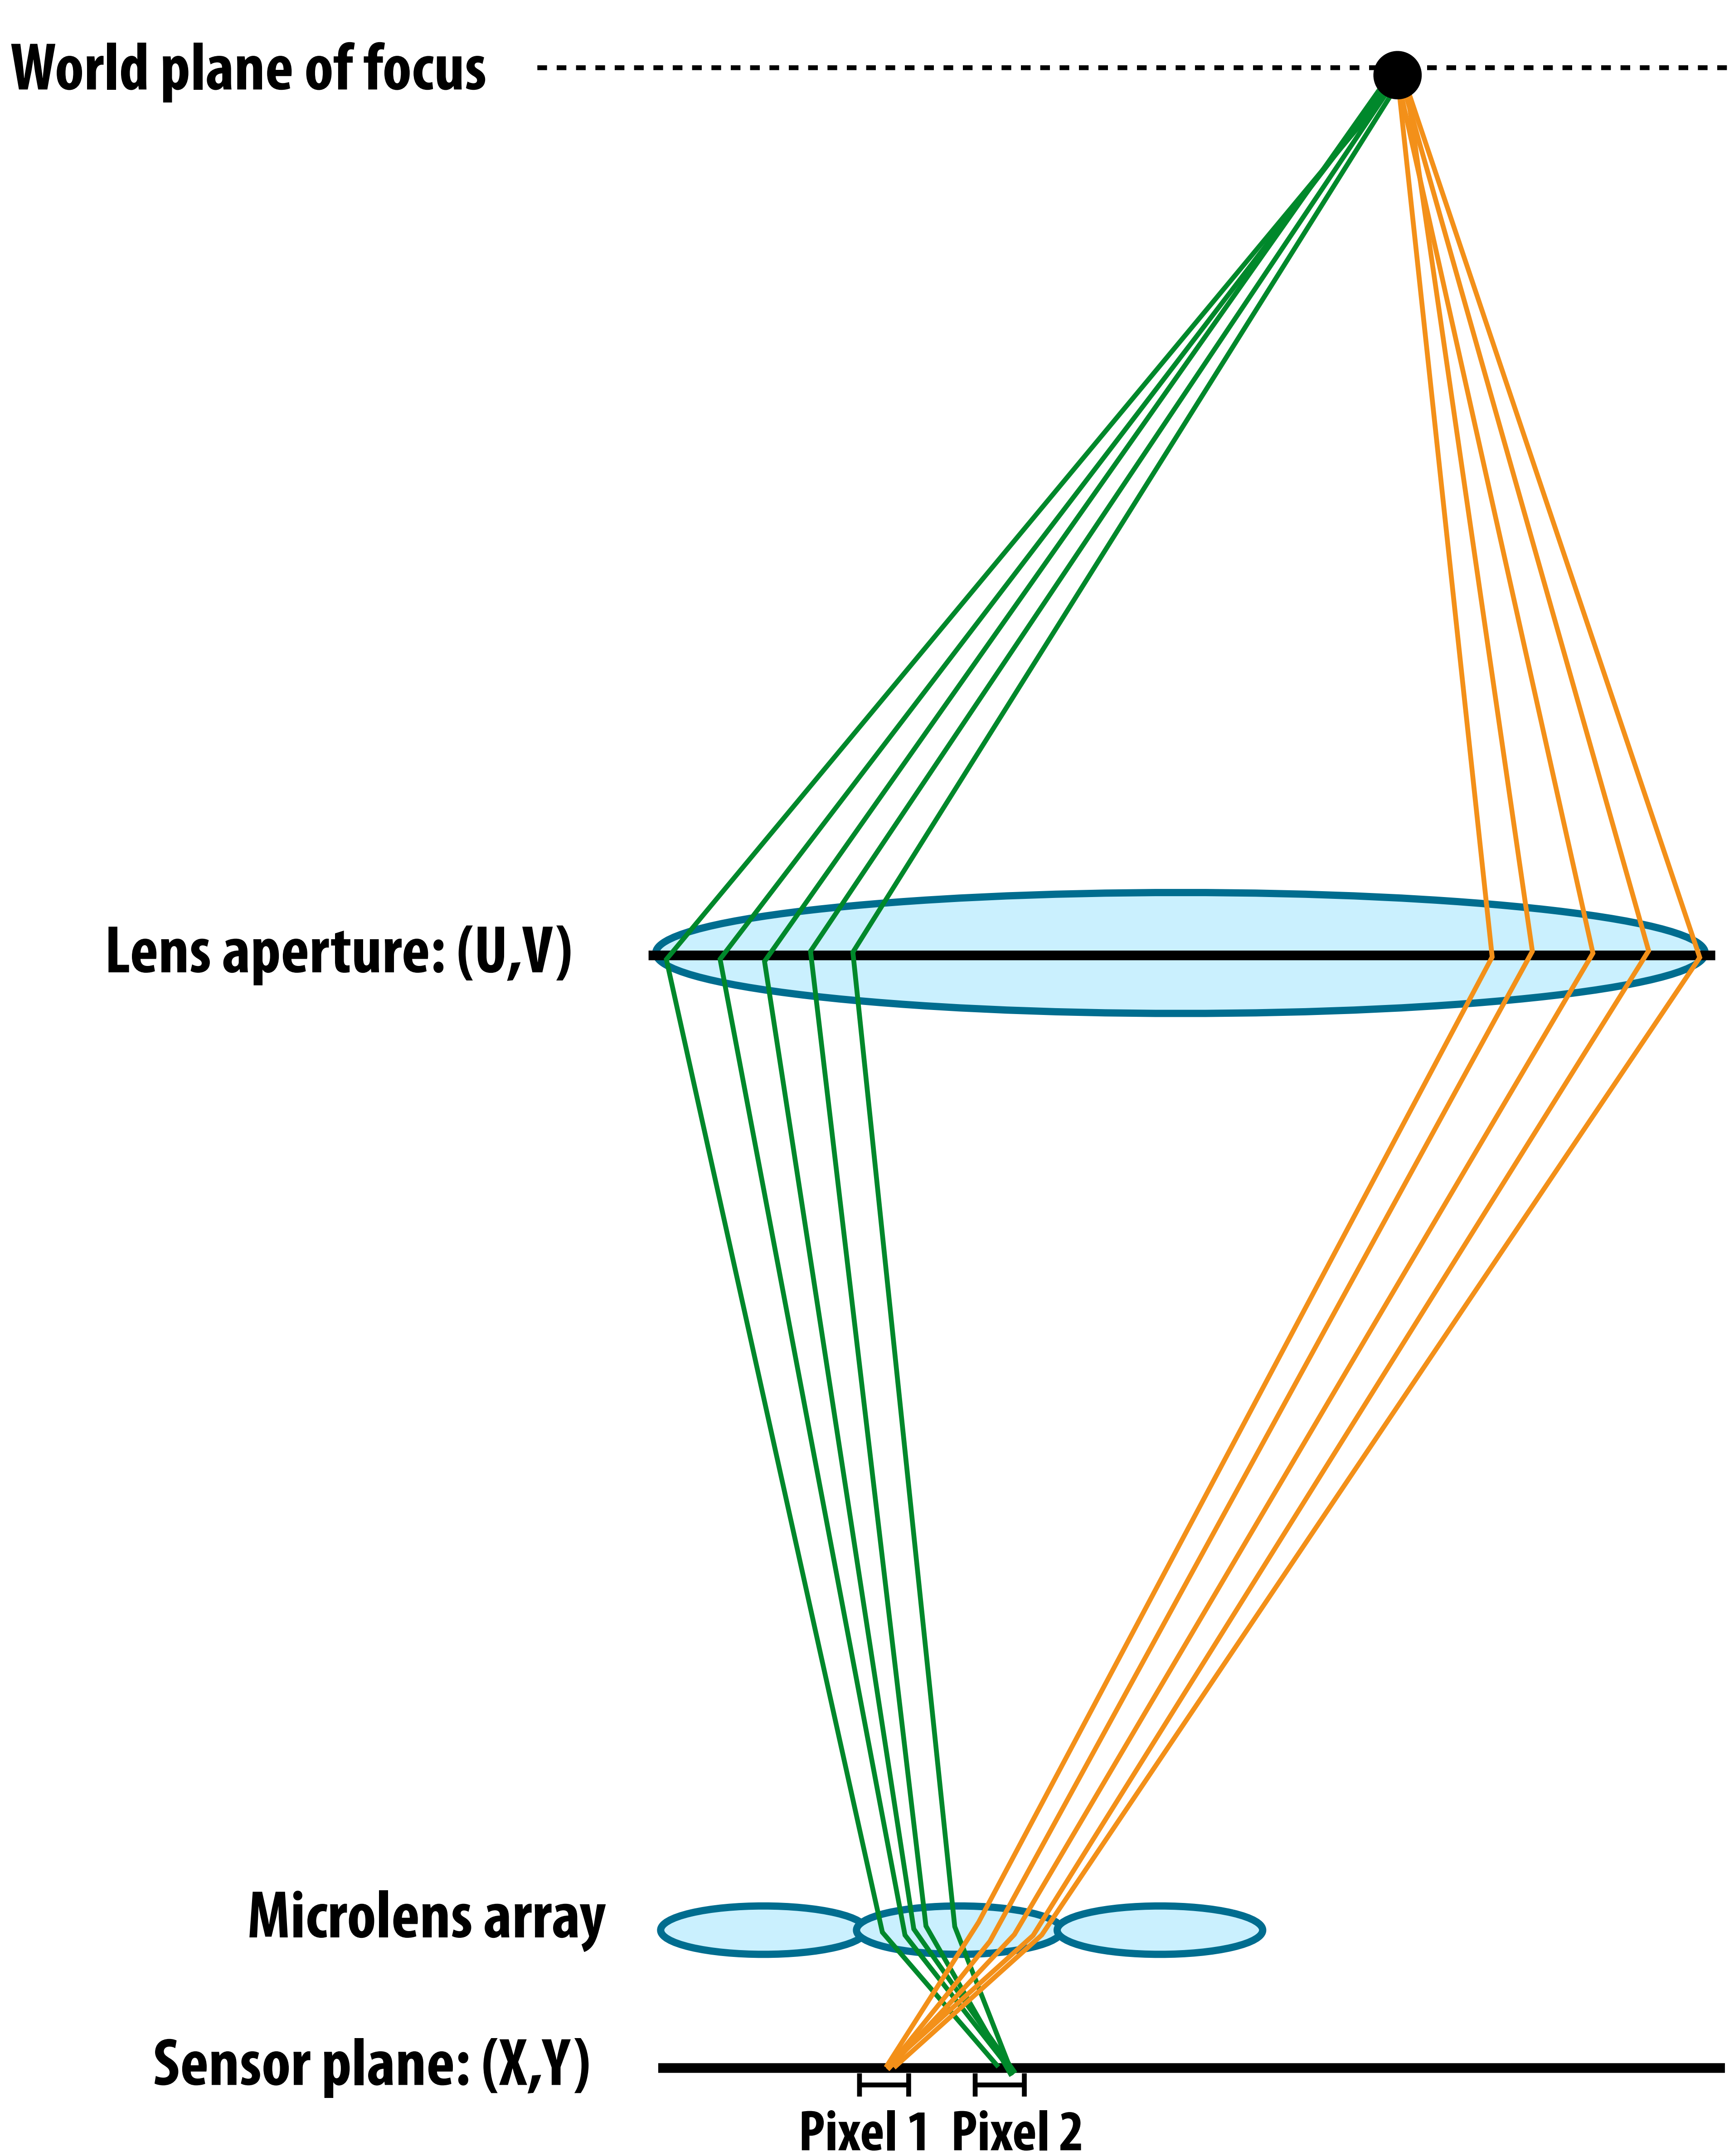

Ray space plot

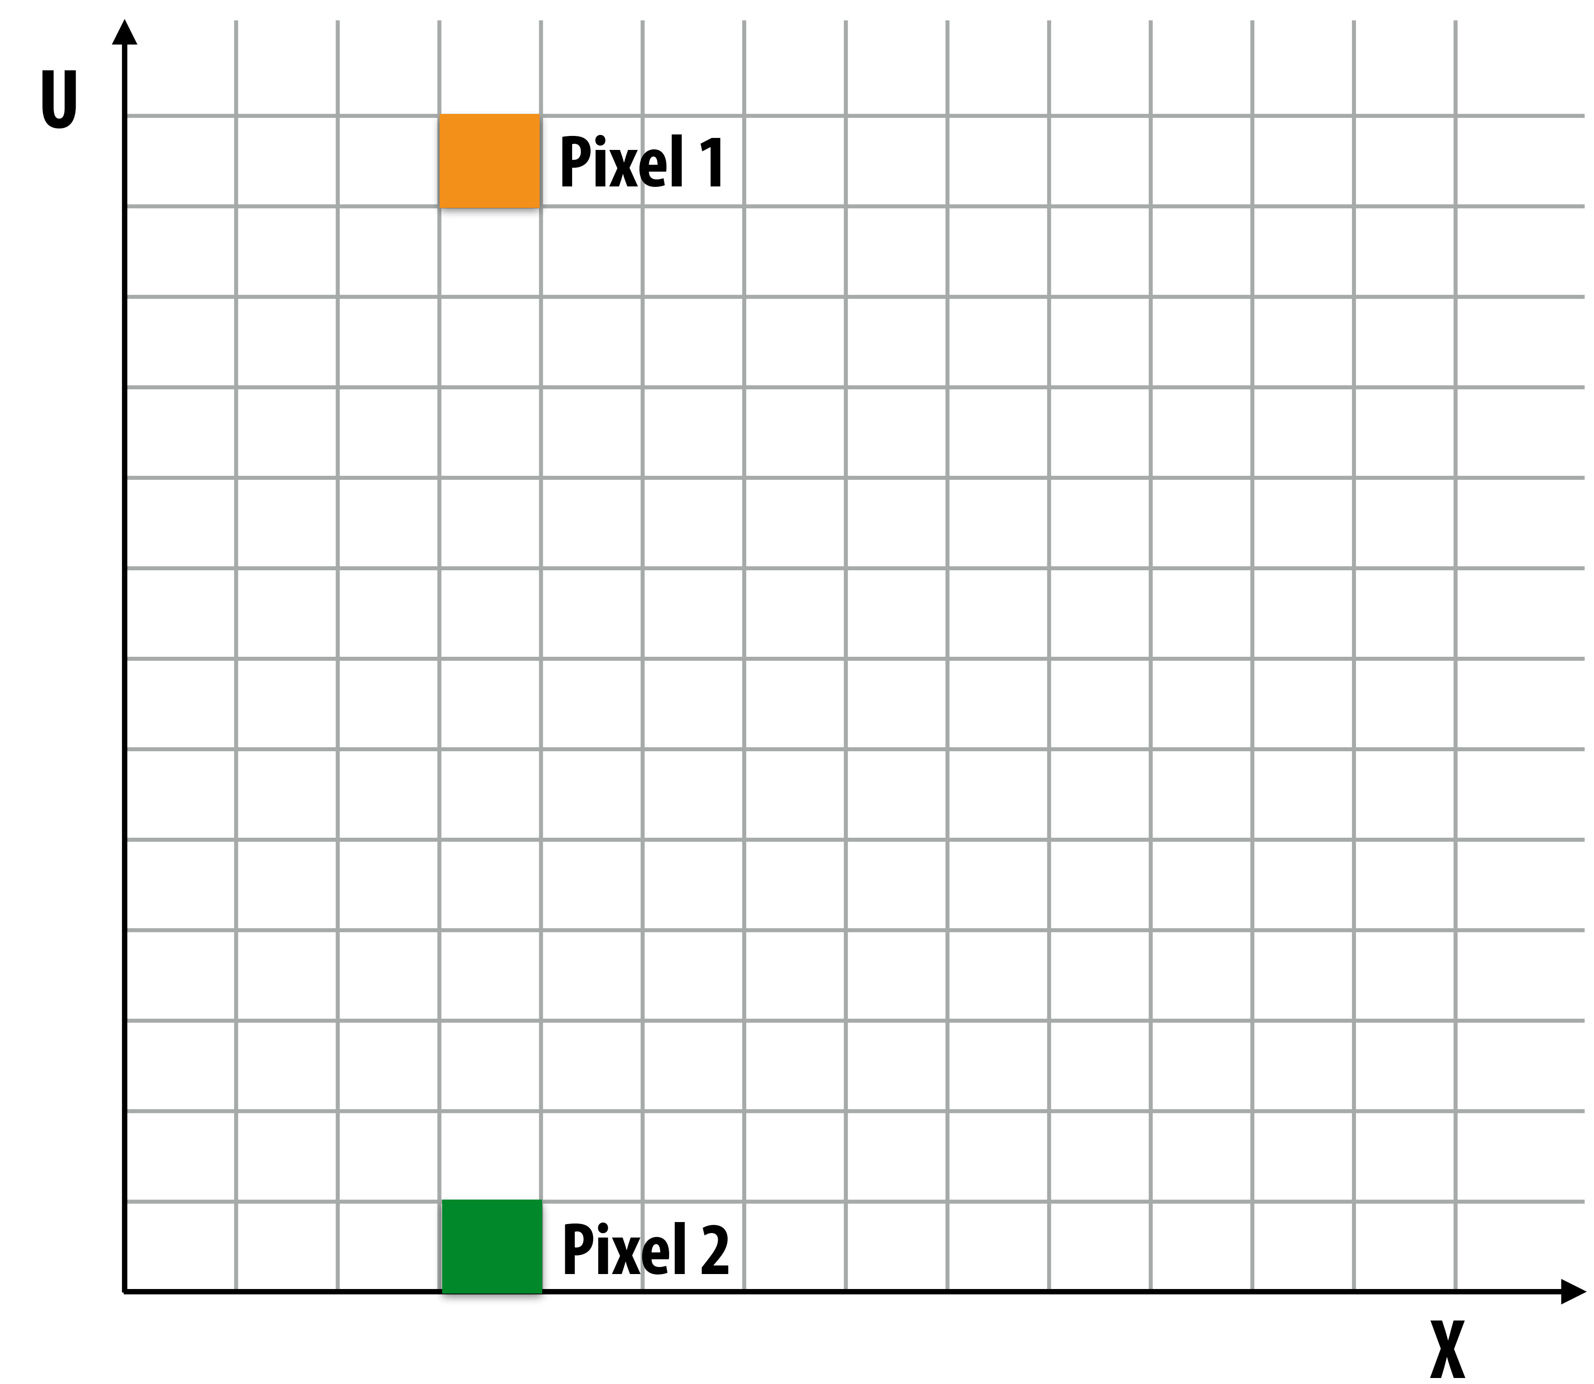

Microlens Array

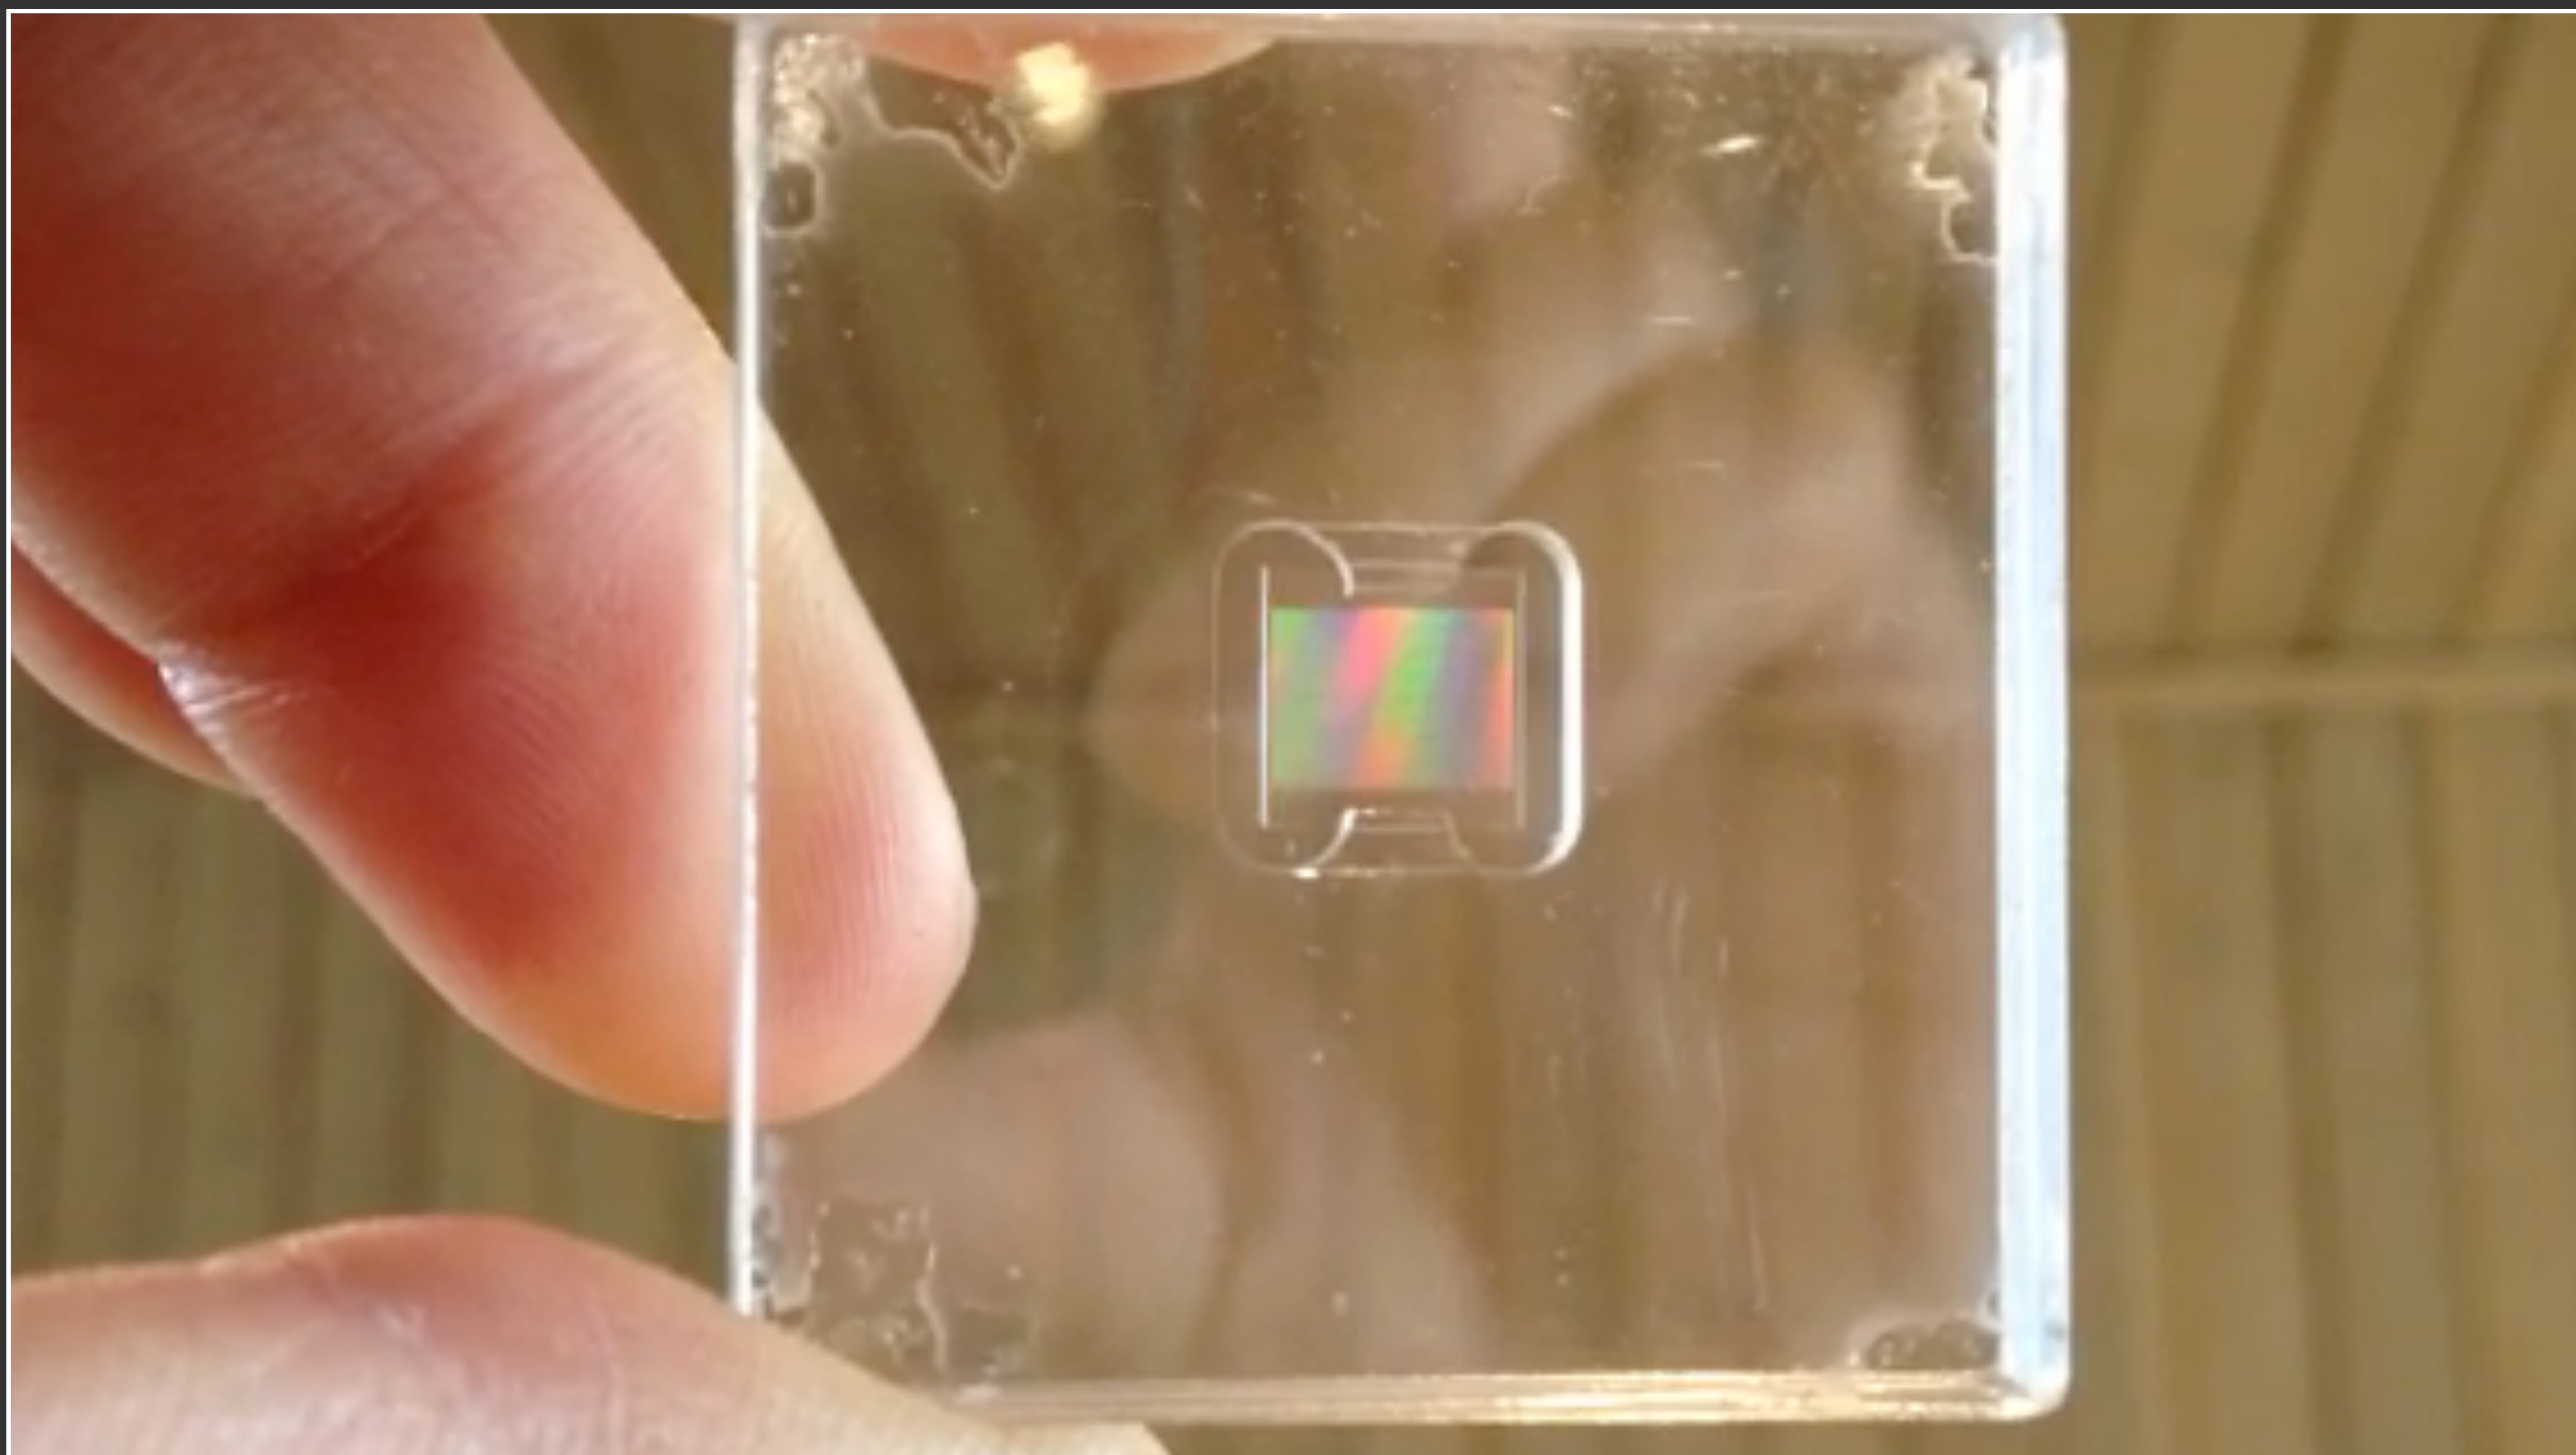

Where Microlenses Go Inside Camera

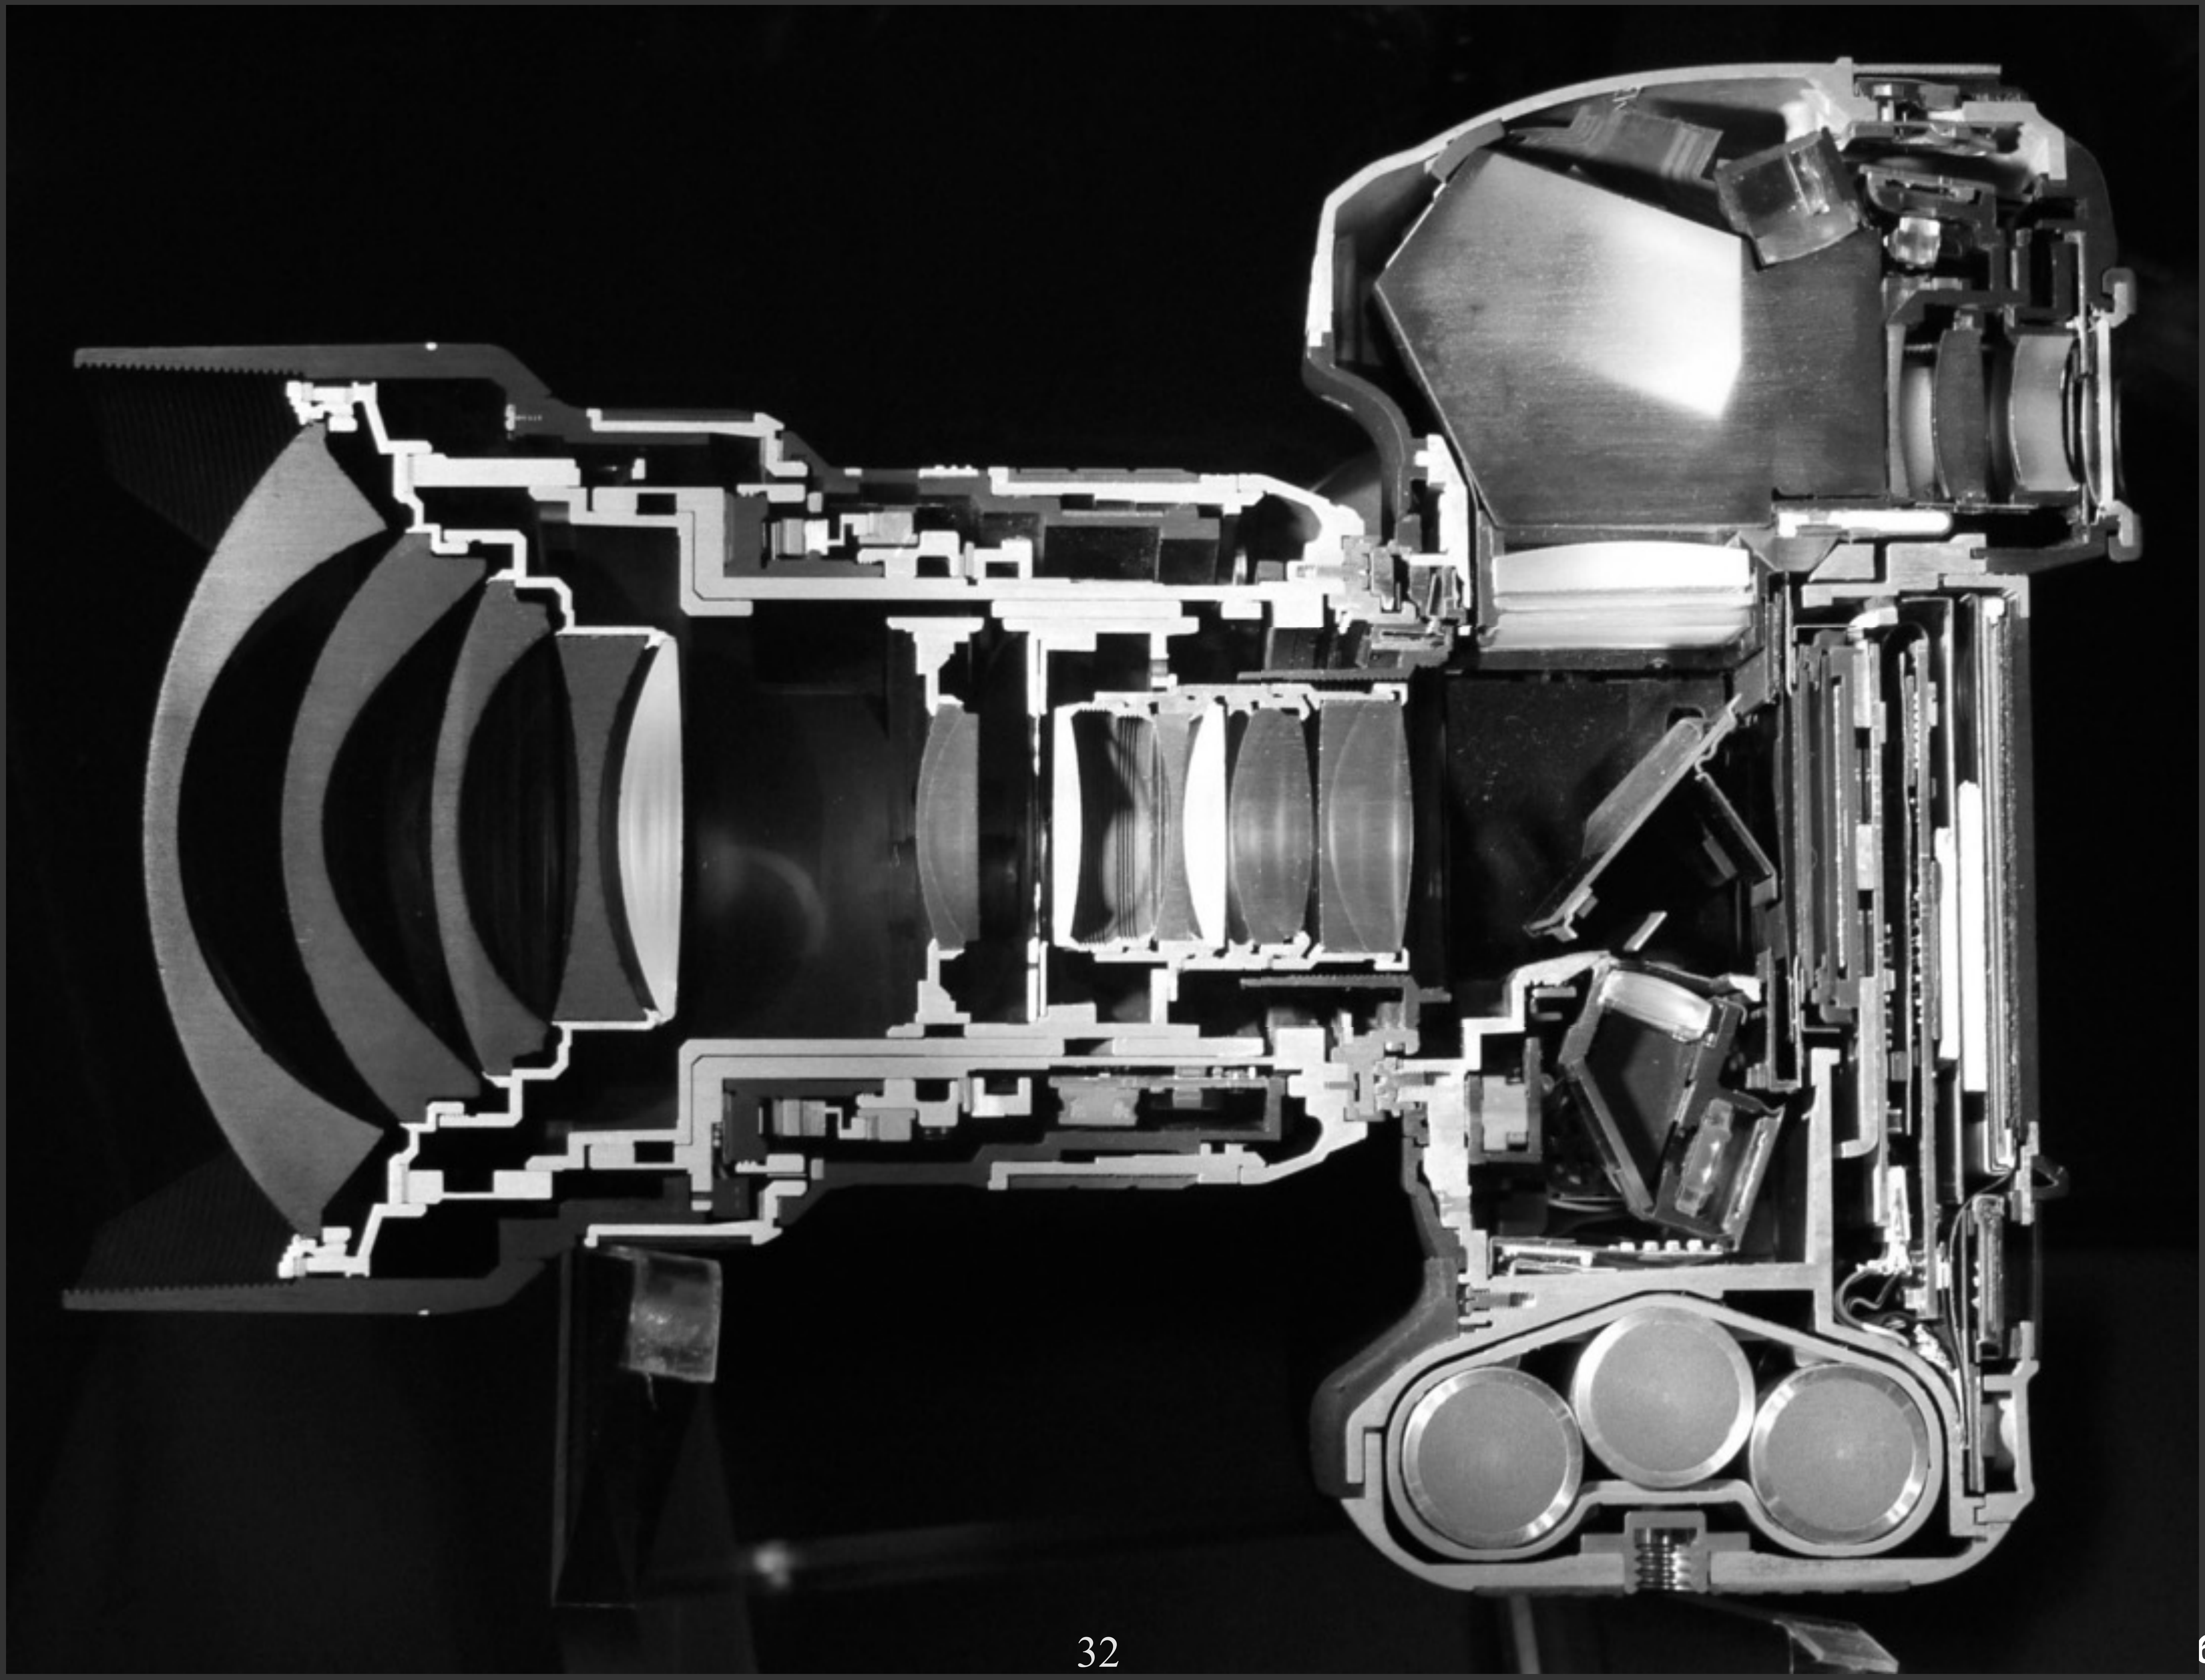

Cross-section of Nikon D3, 14-24mm F28 lens, Credit: ??

Where Microlenses Go Inside Camera

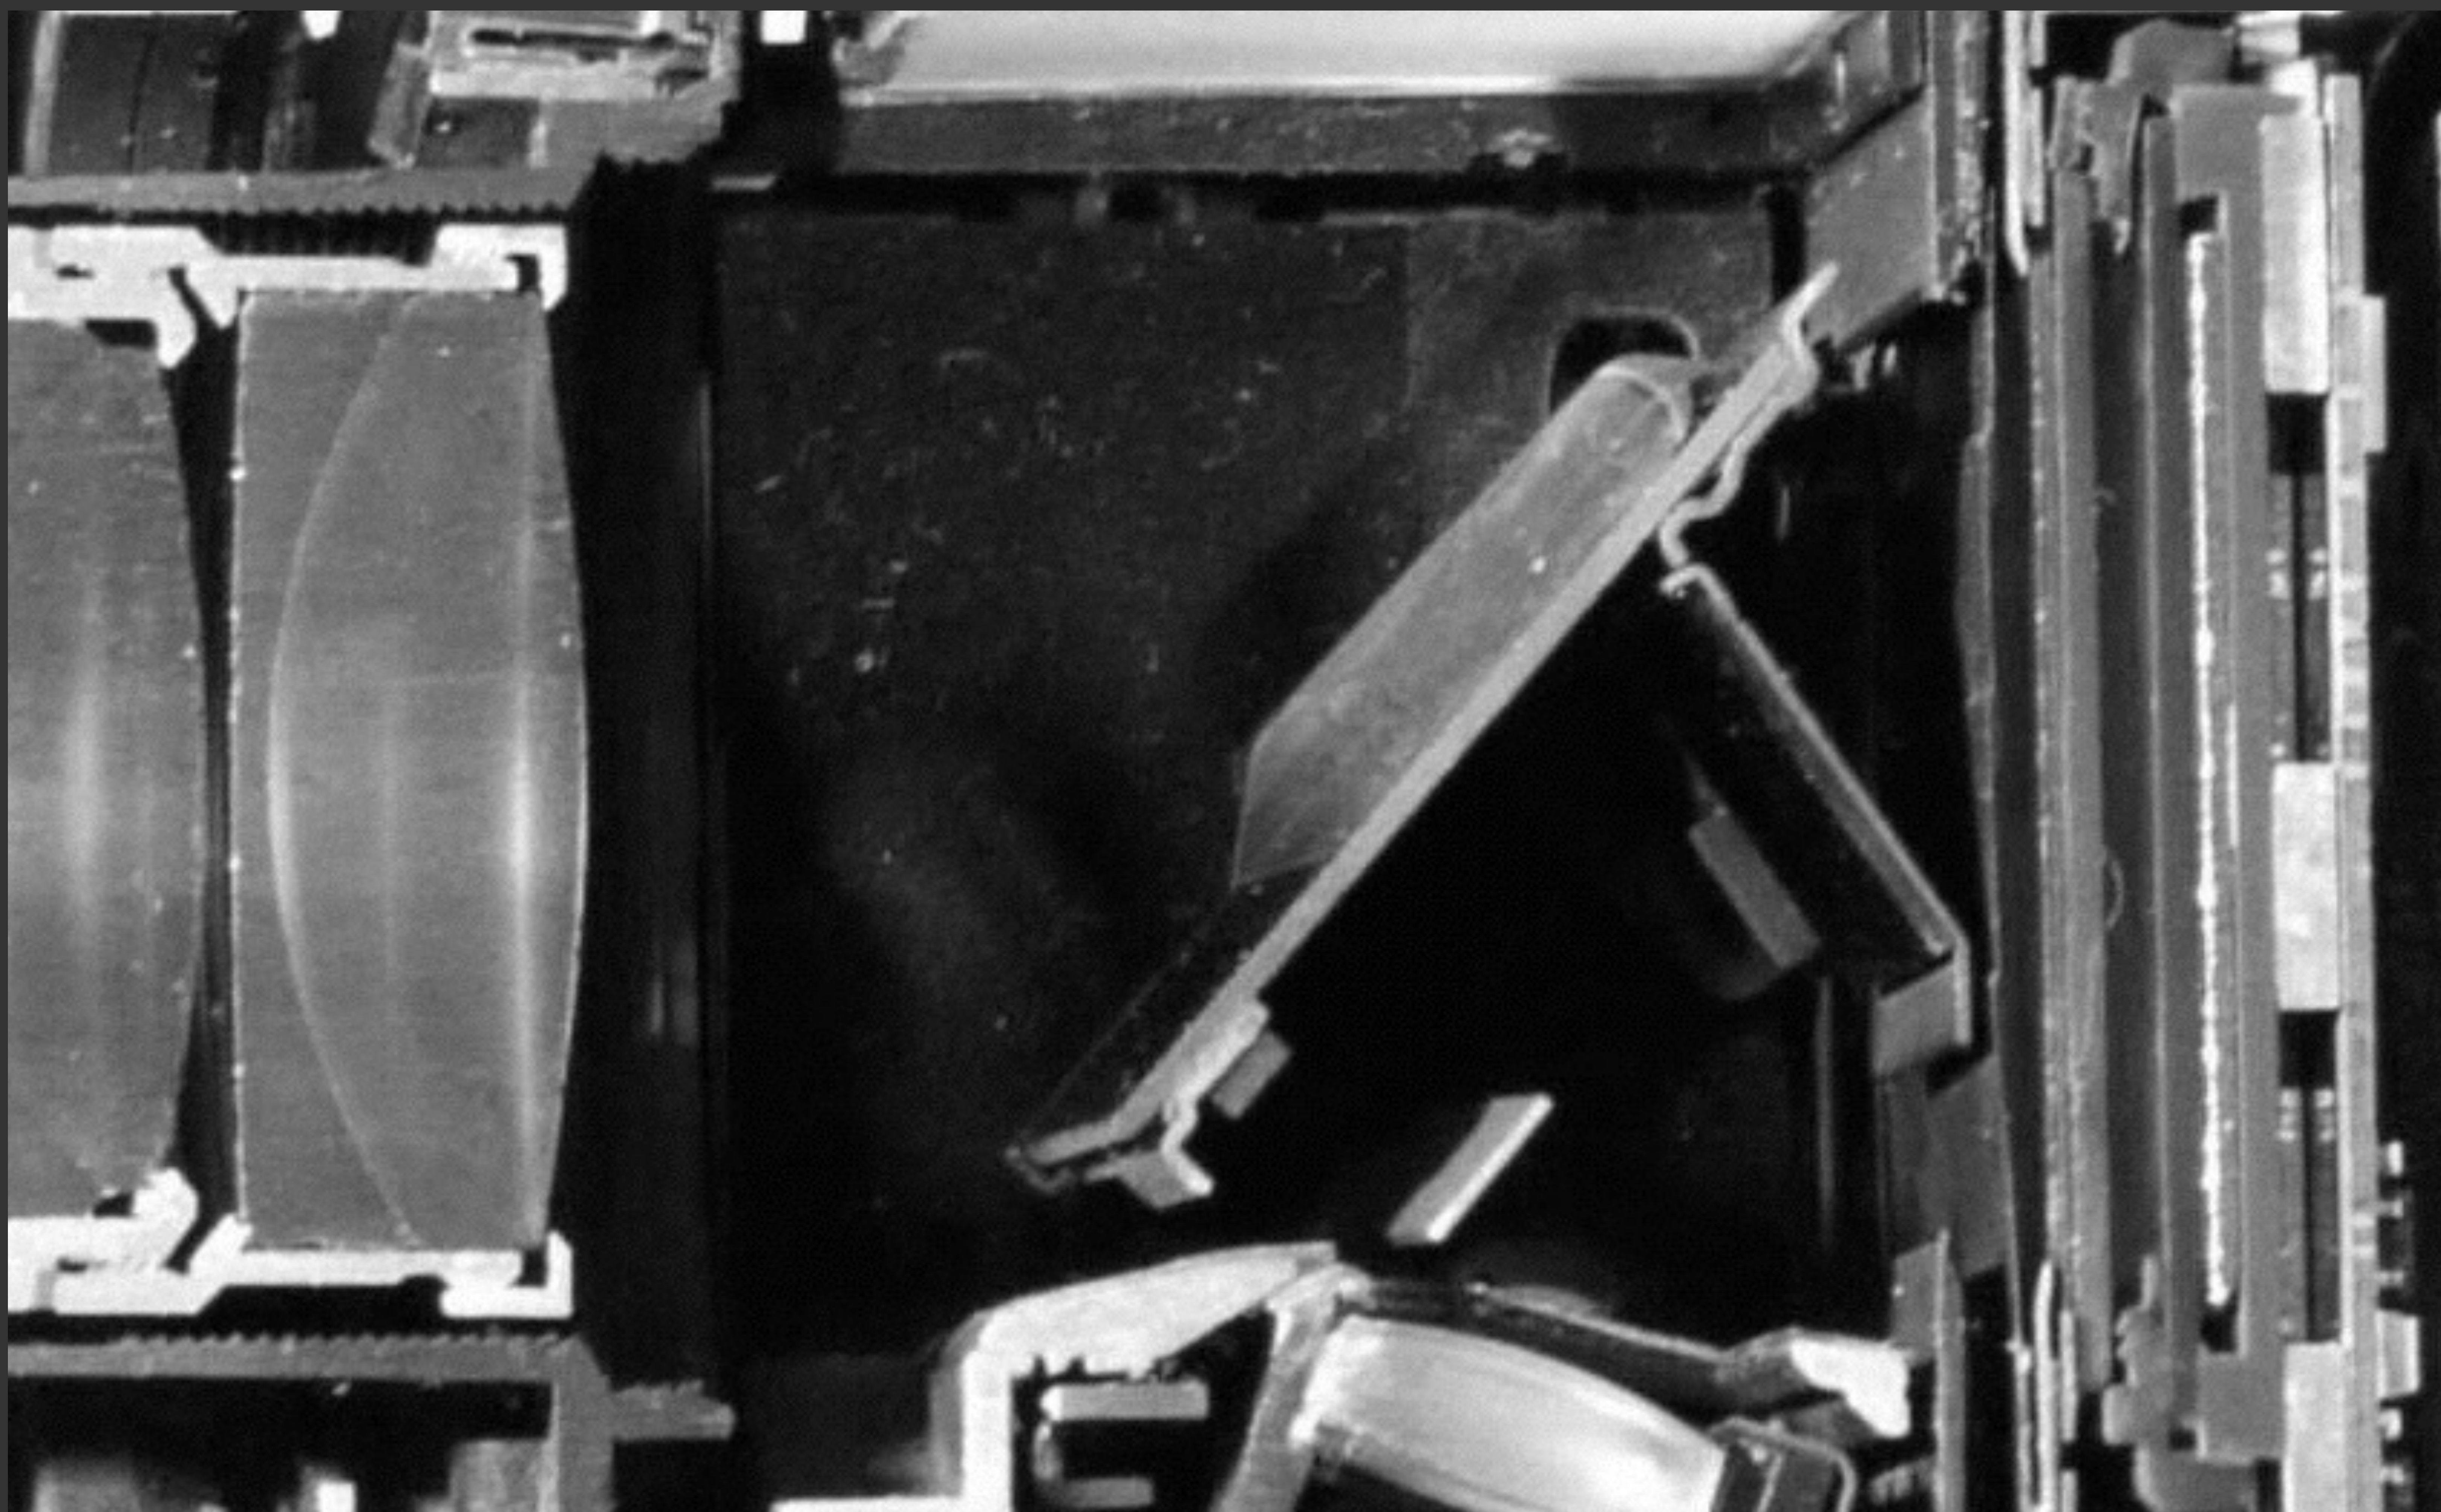

Where Microlenses Go Inside Camera

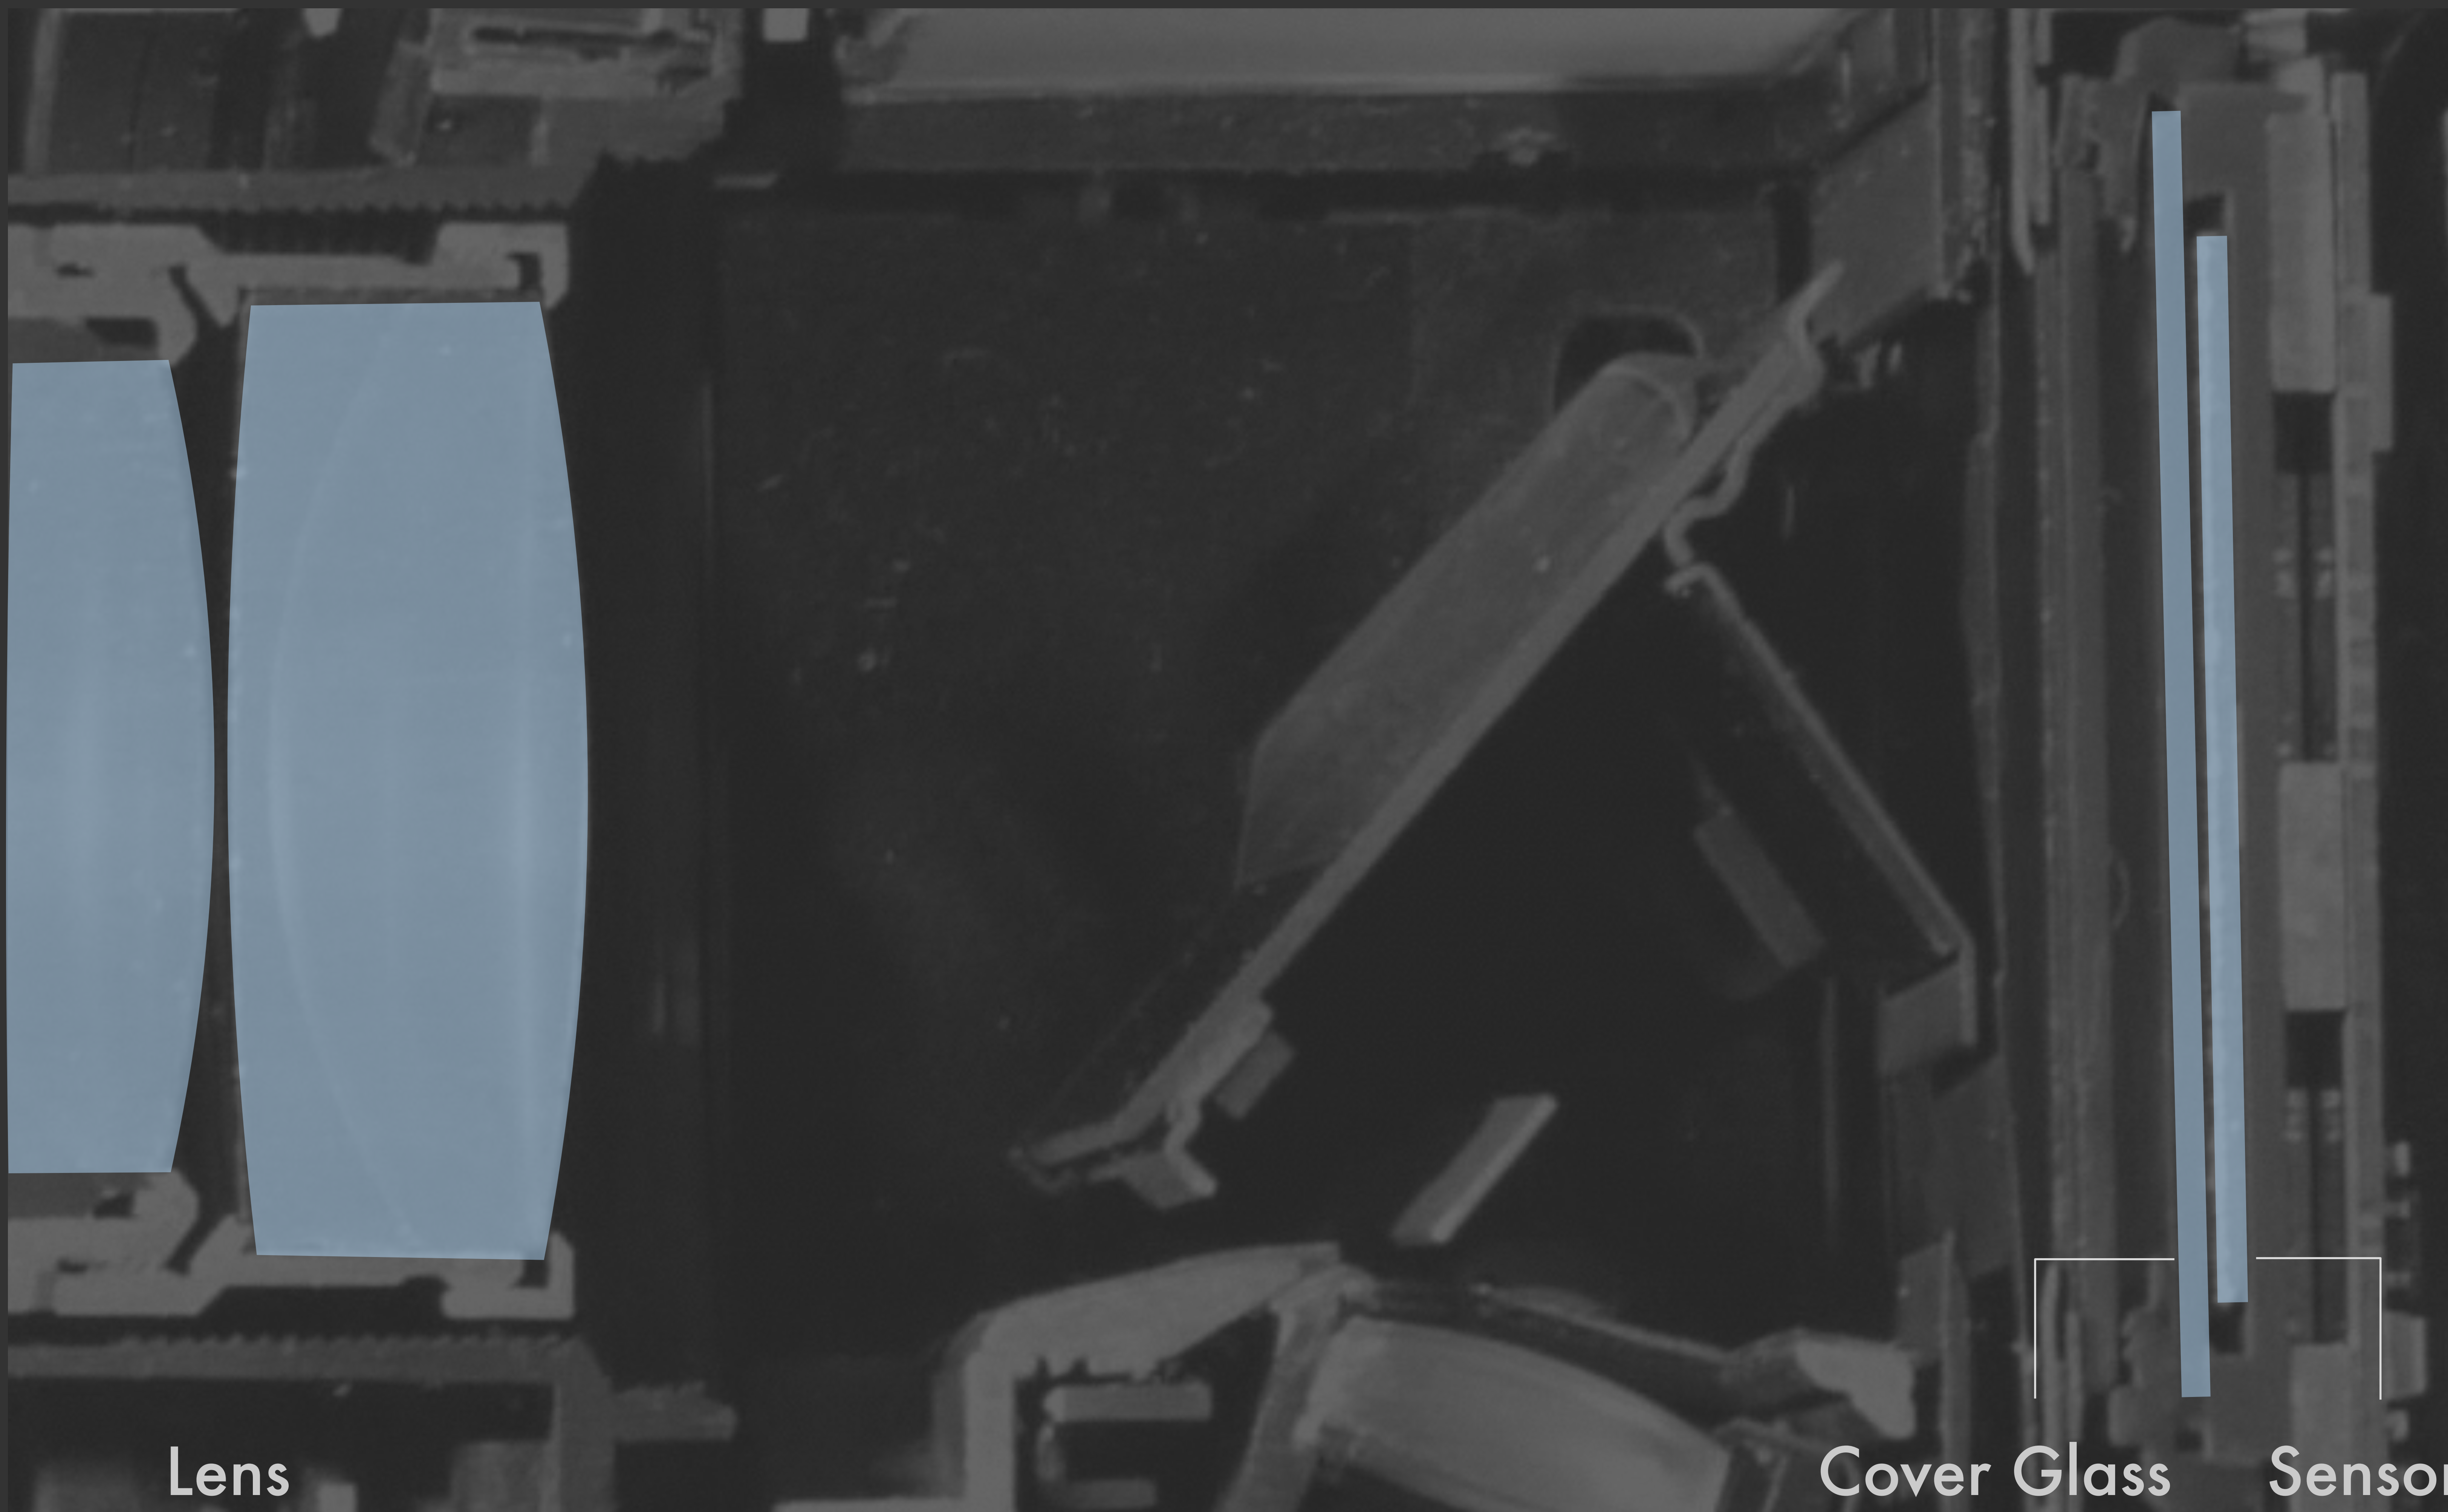

Where Microlenses Go Inside Camera

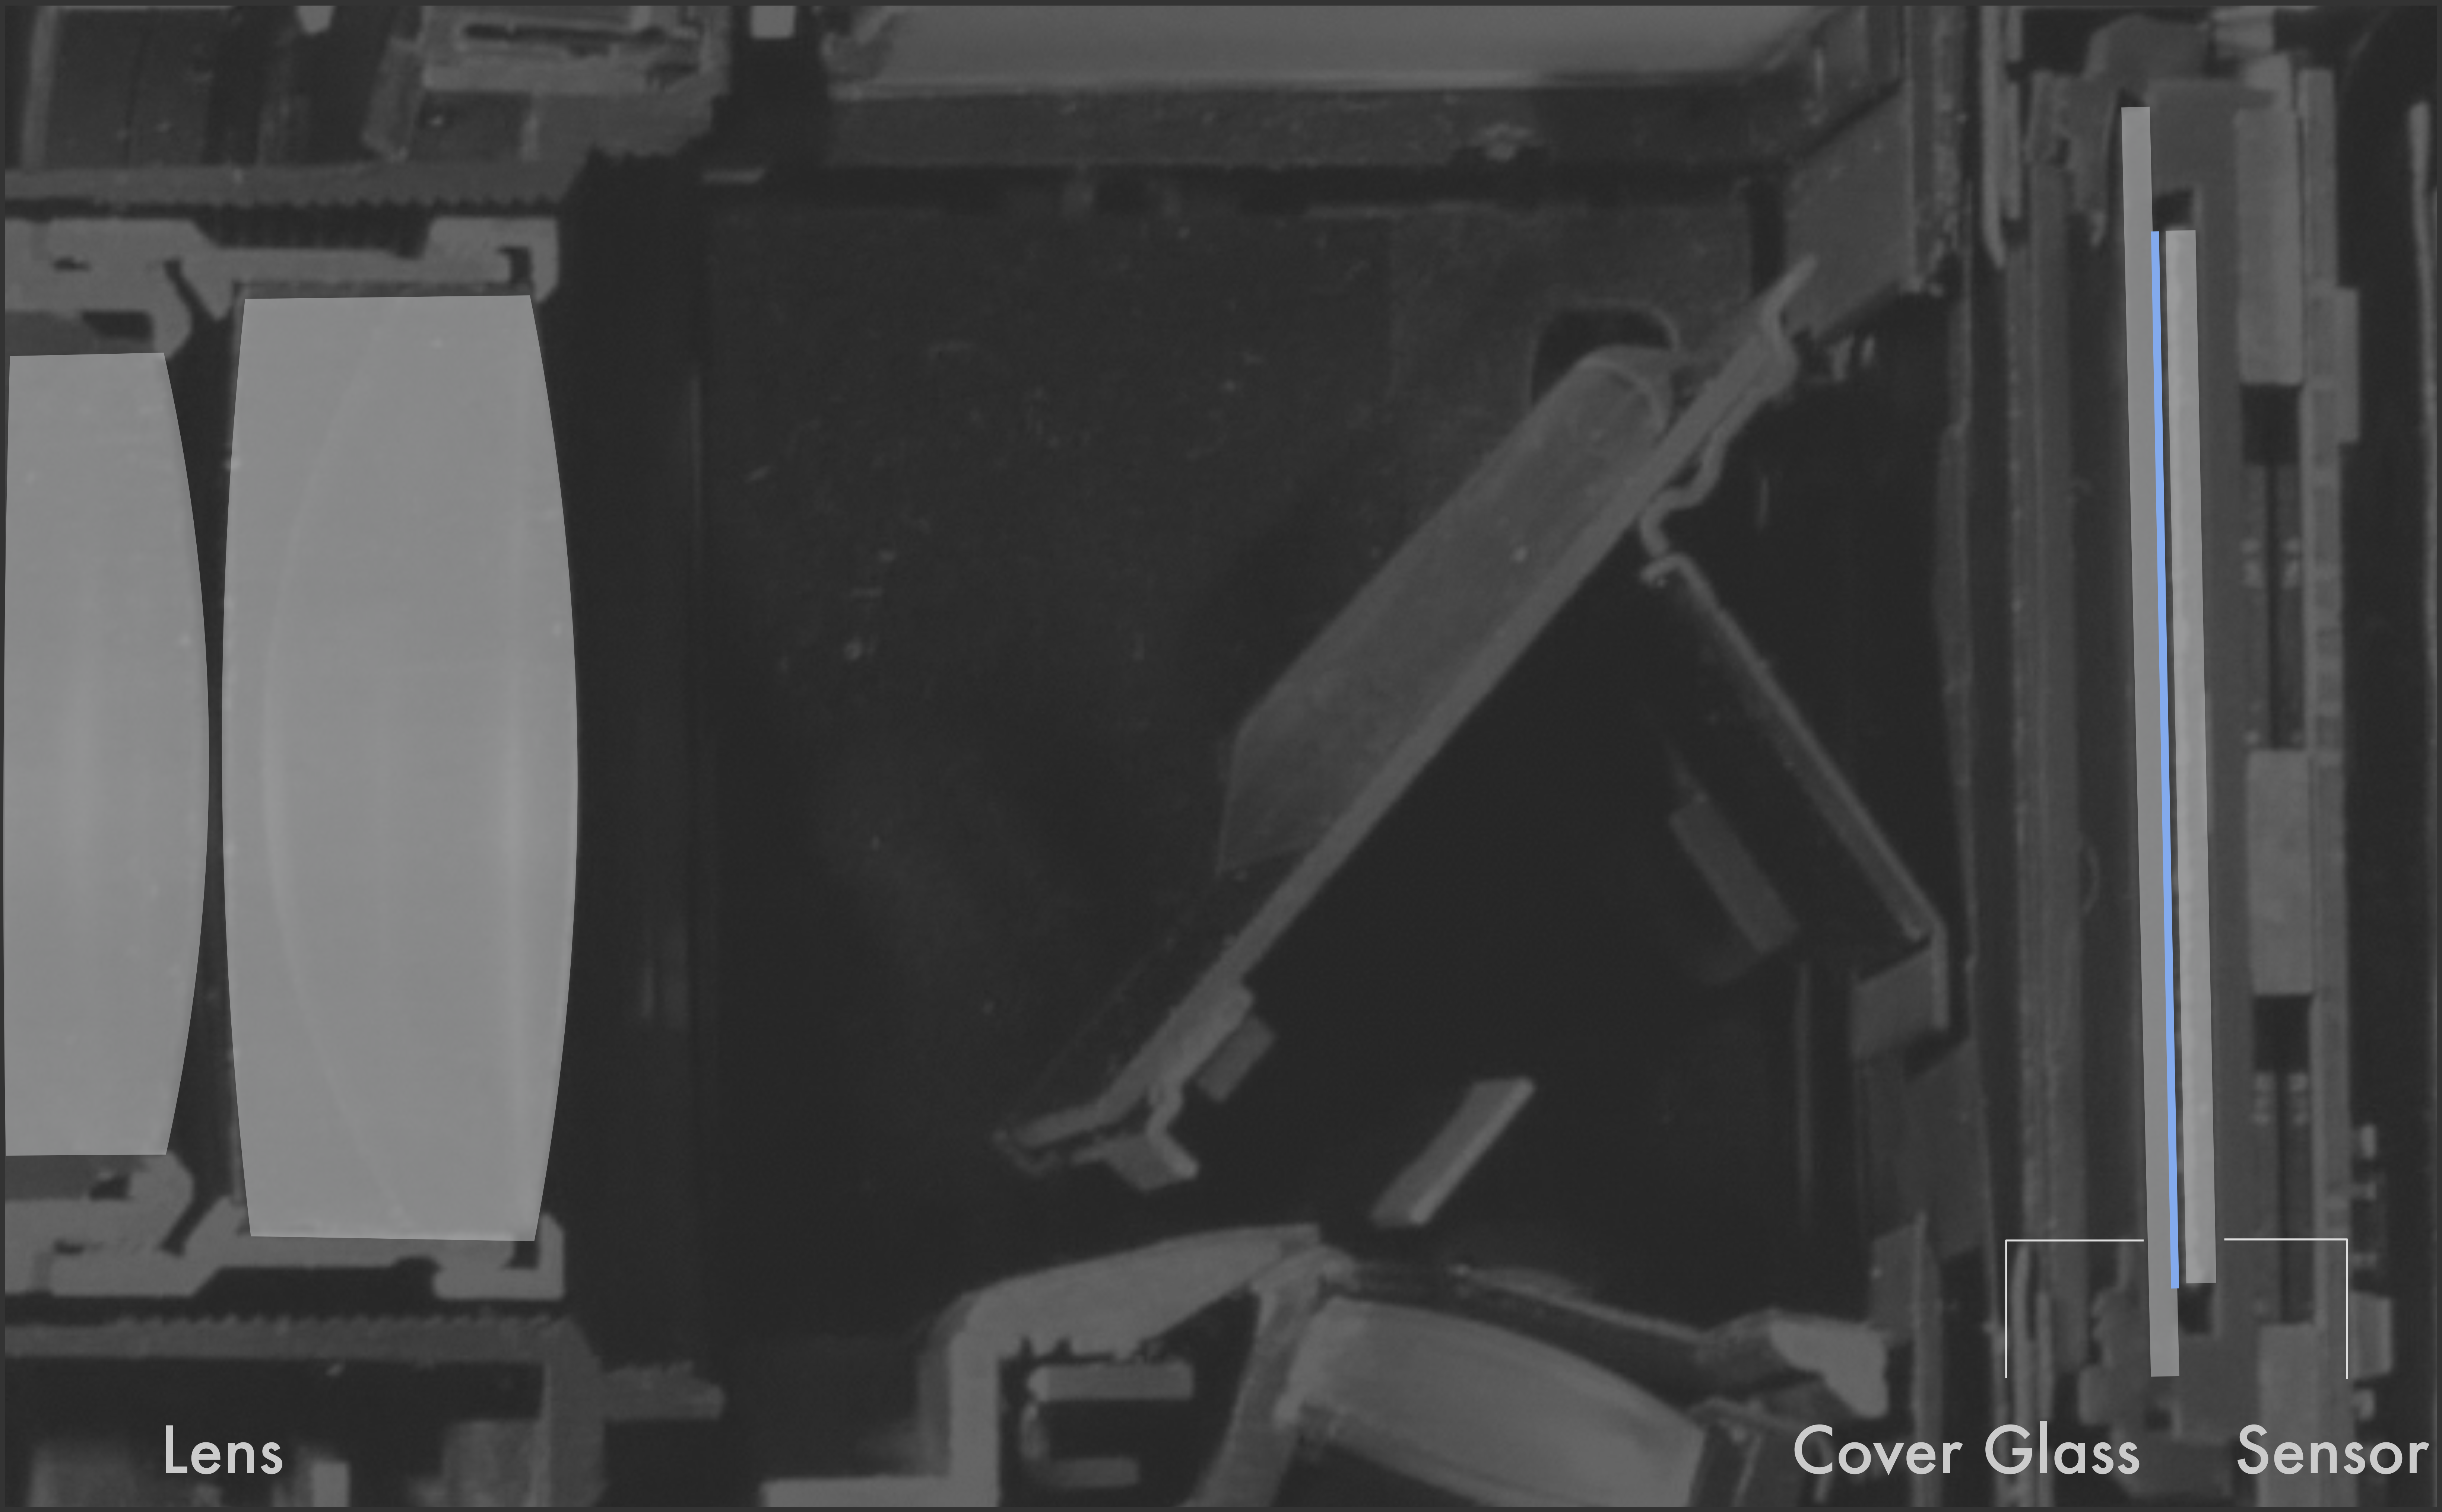

Raw Data From Light Field Sensor

Raw Data From Light Field Sensor

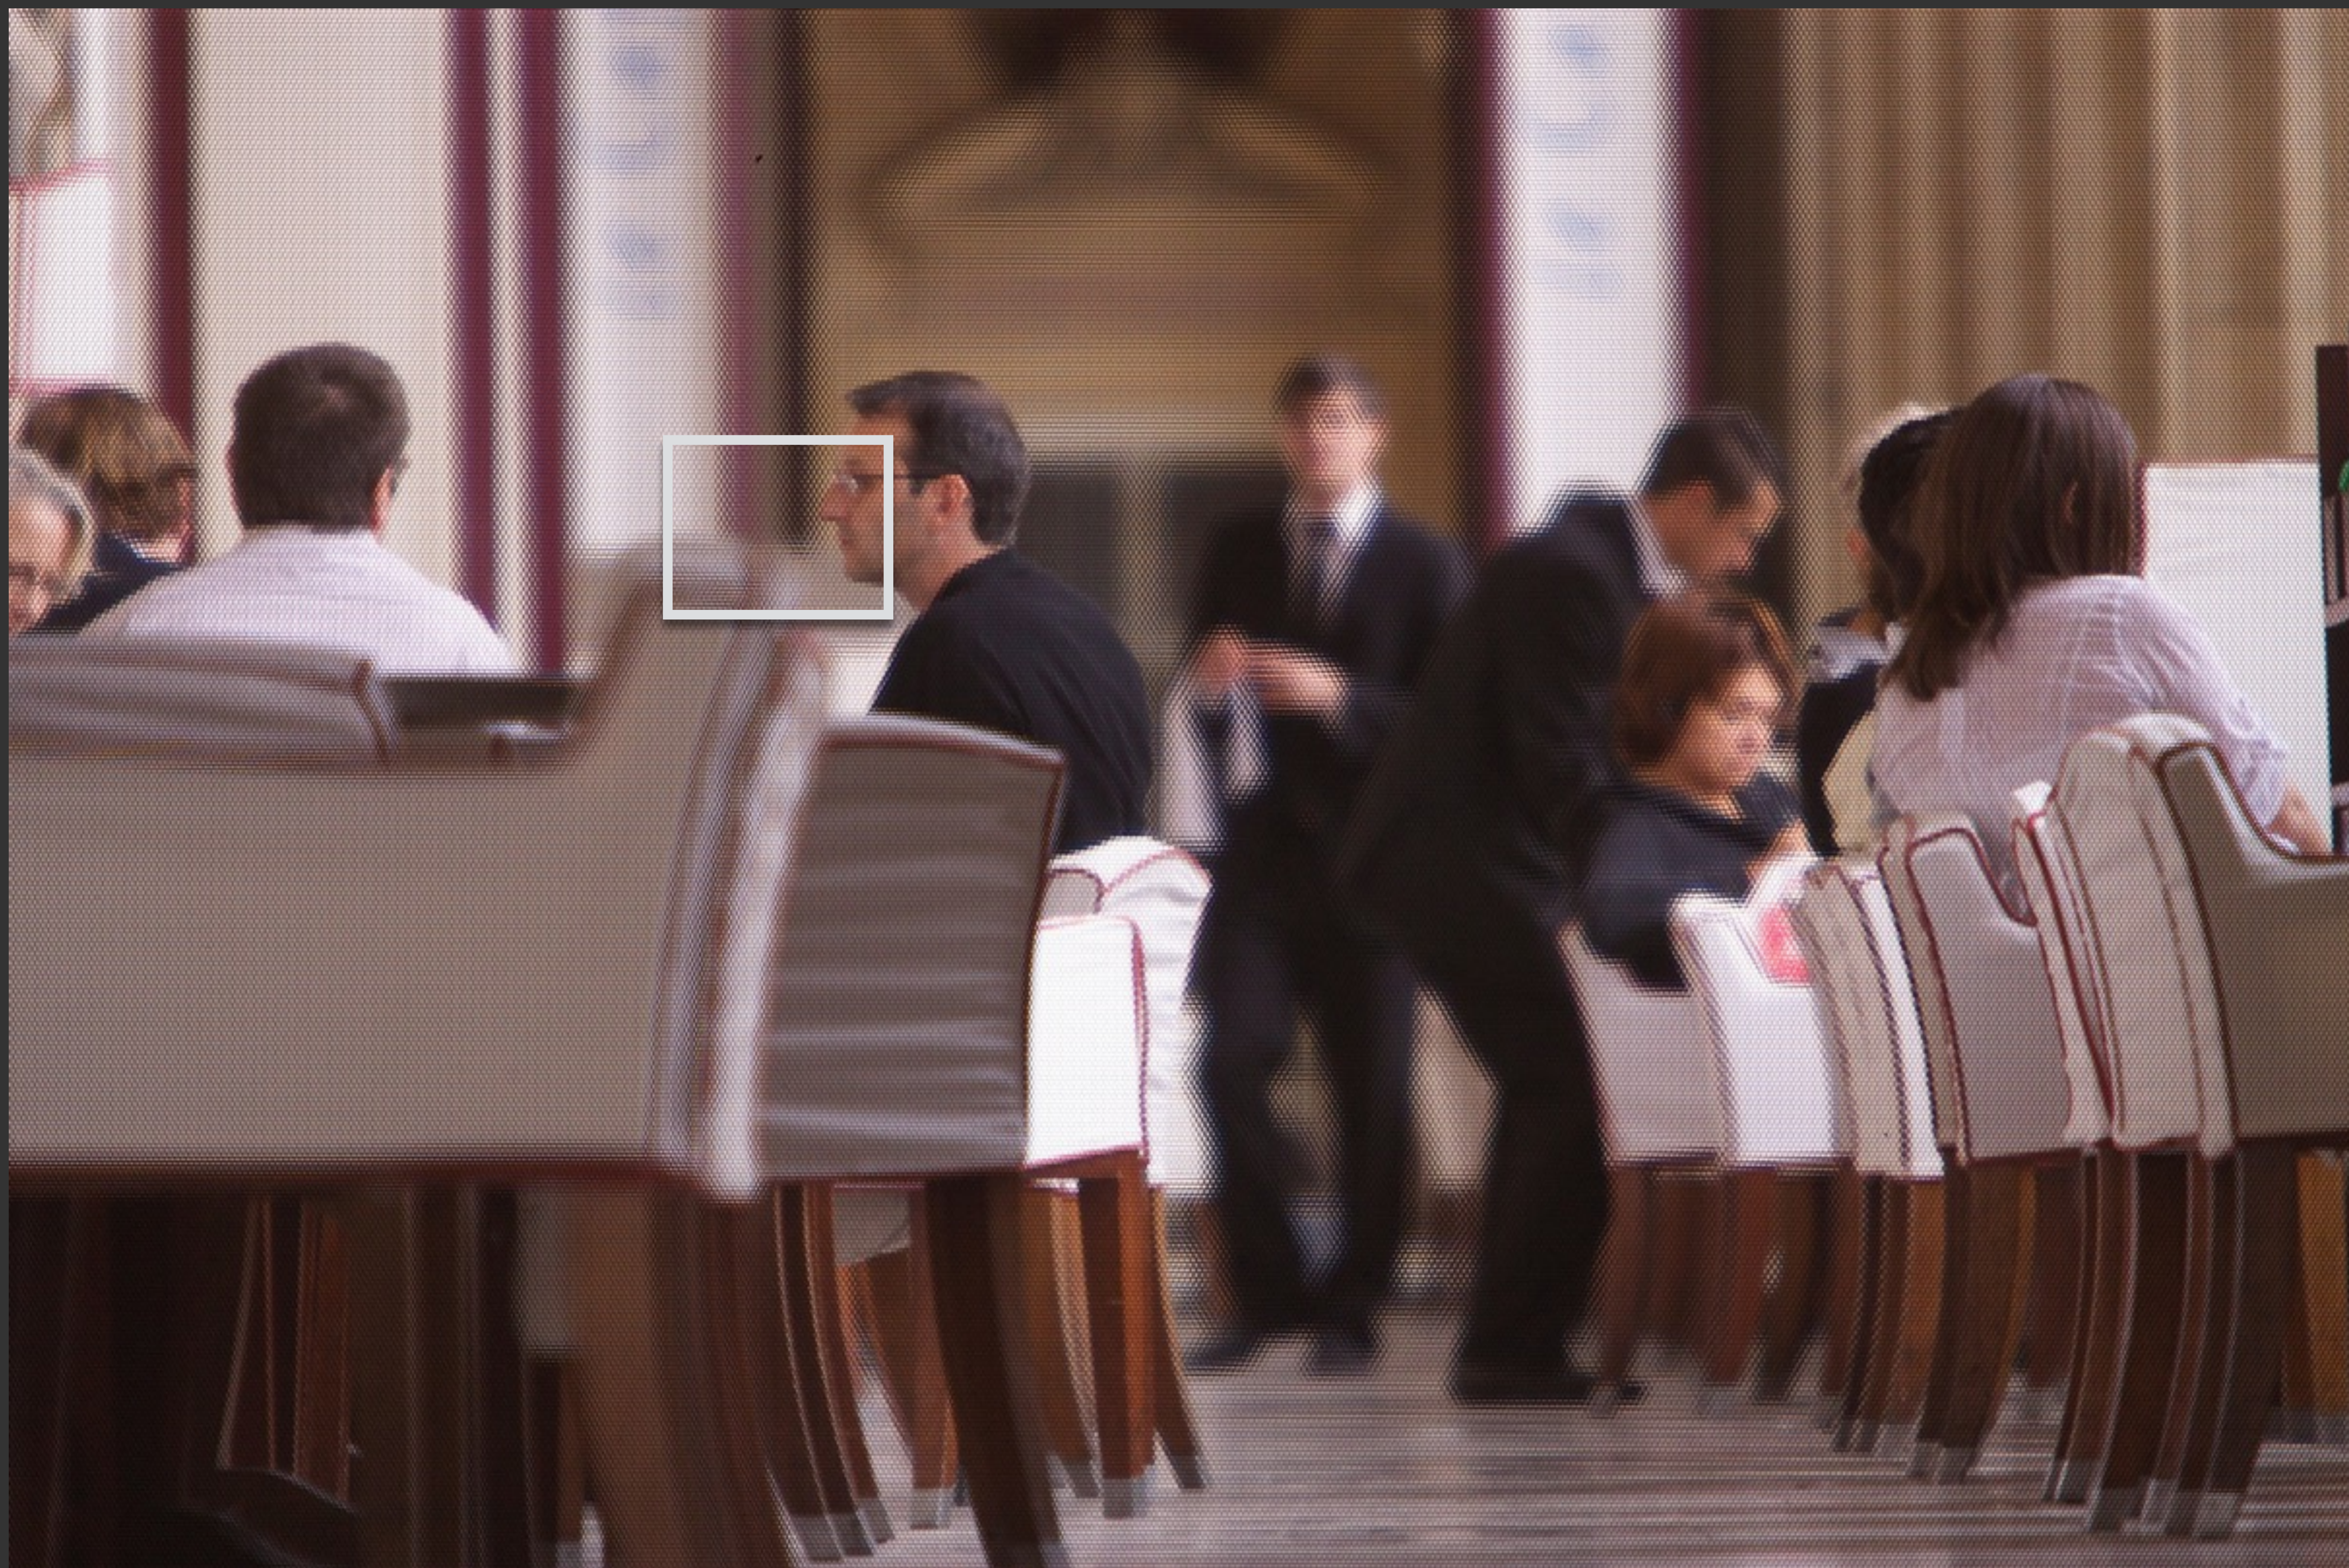

Raw Data From Light Field Sensor

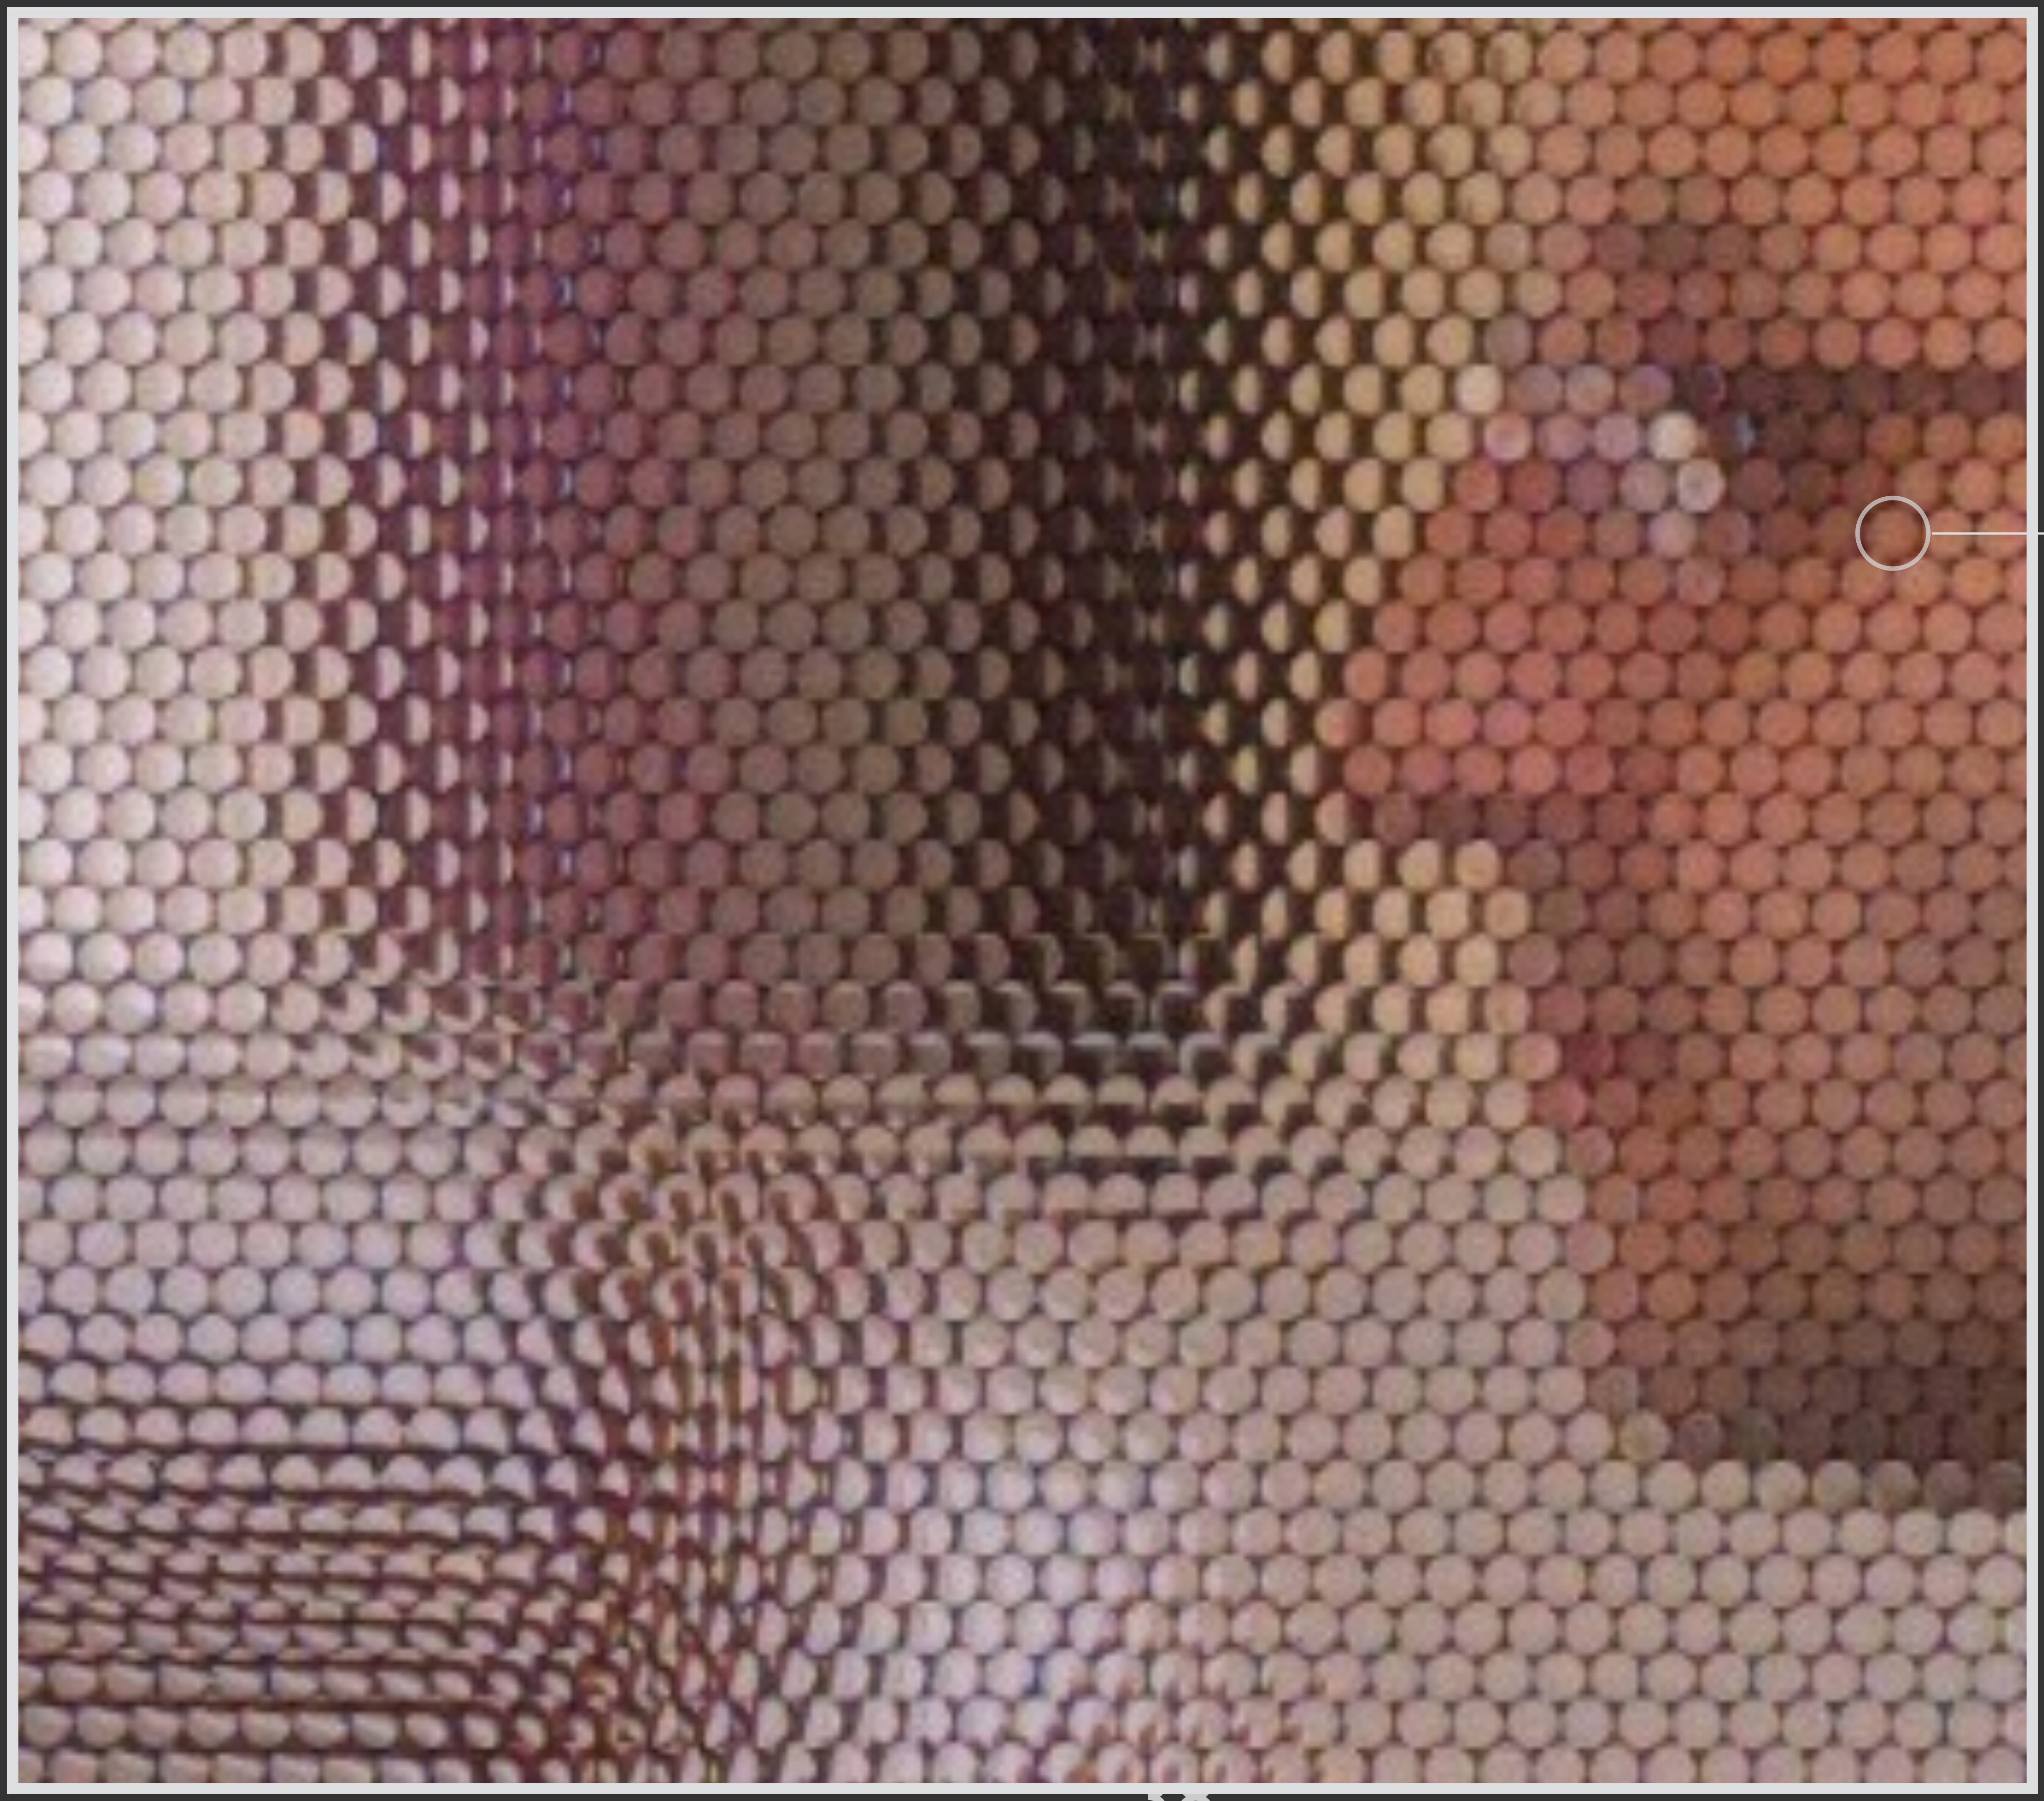

One disk image

Raw Data From Light Field Sensor

Mapping Sensor Pixels to (x,y,u,v) Rays

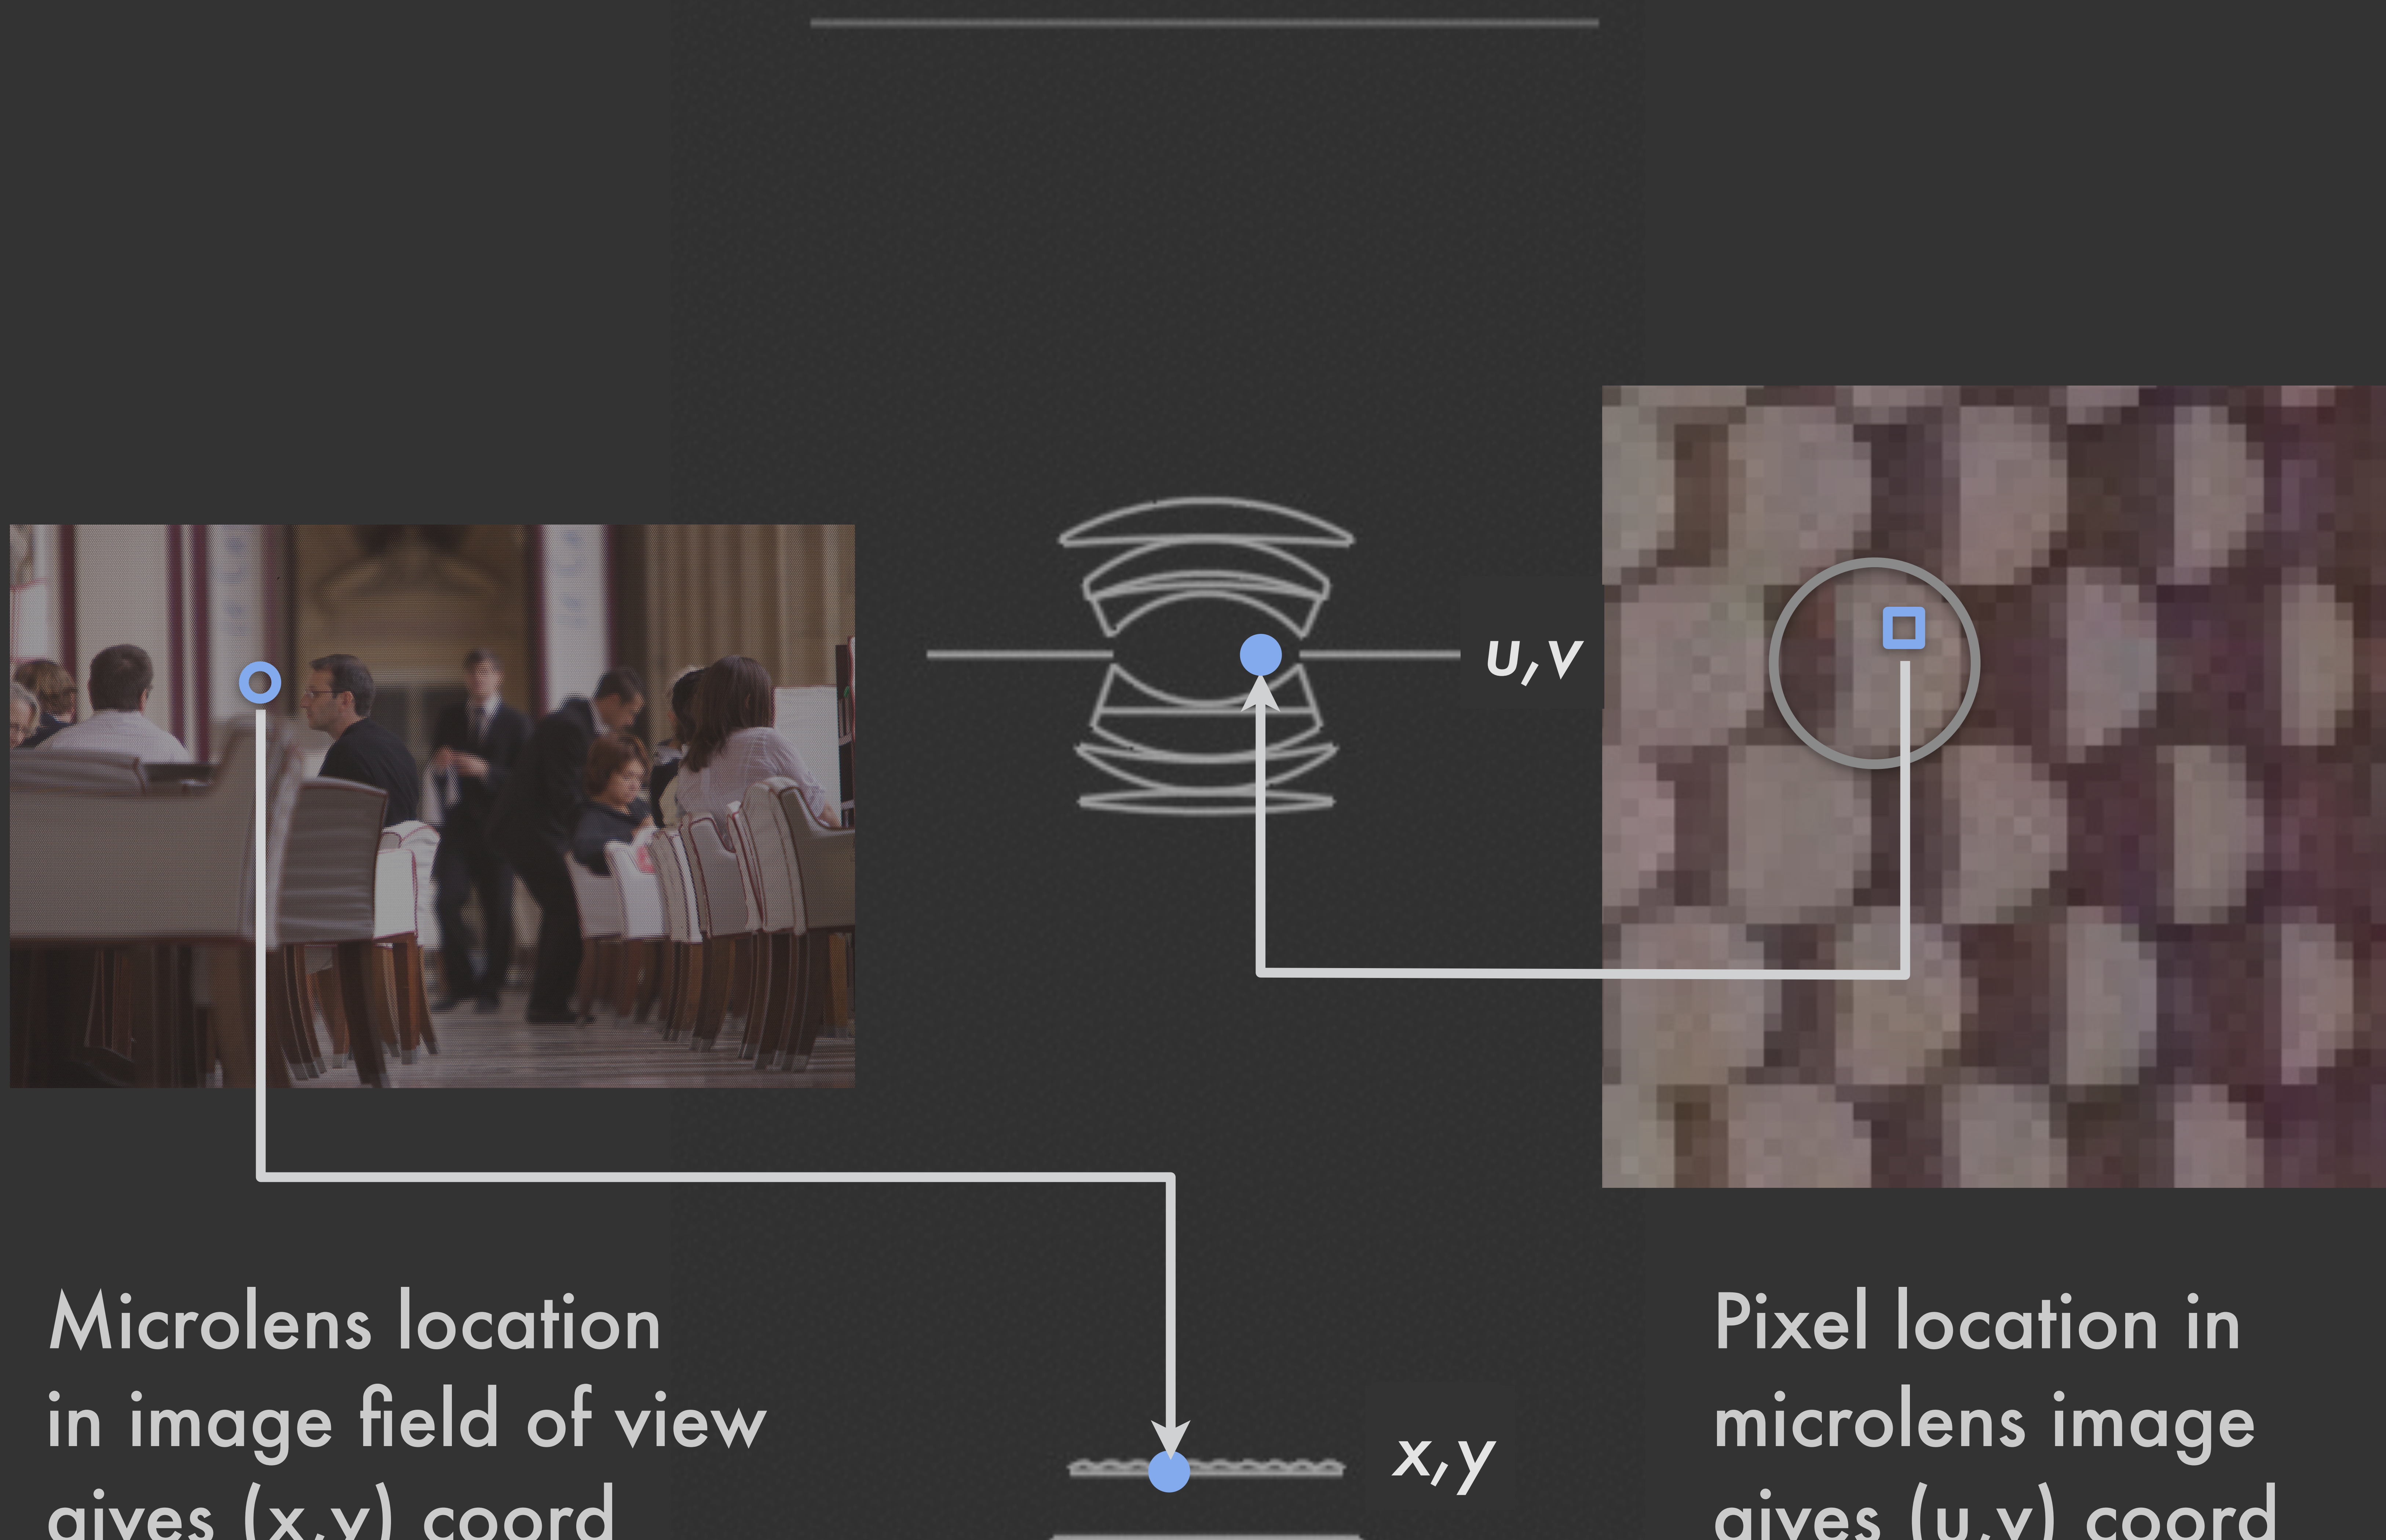

Microlens location in image field of view gives (x,y) coord

Pixel location in microlens image gives (u,v) coord

Mapping Sensor Pixels to (x,y,u,v) Rays

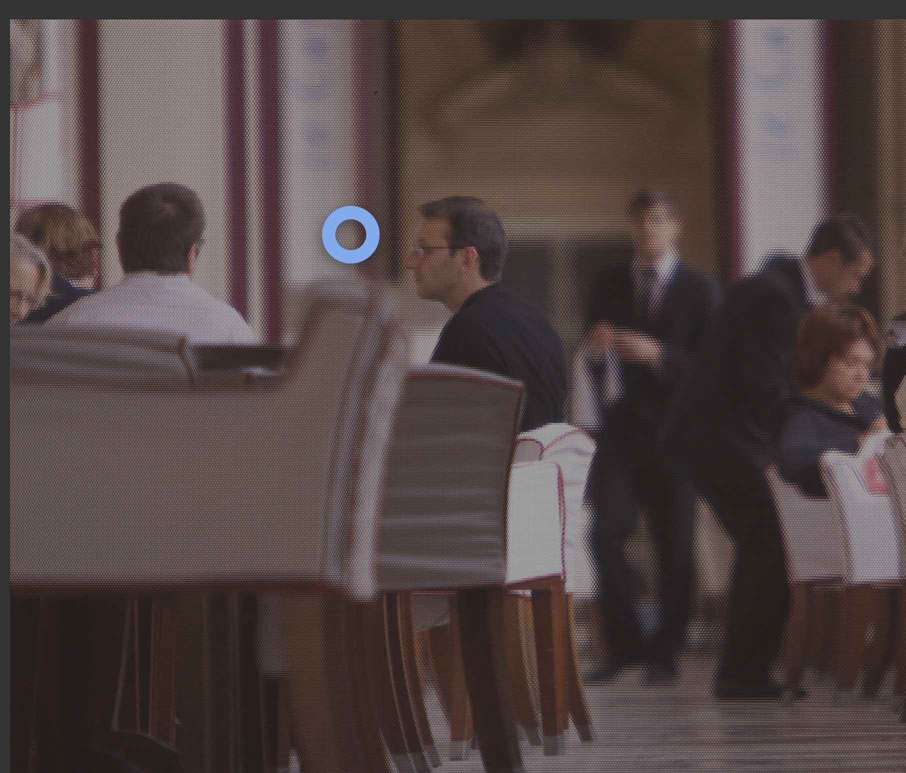

Microlens location in image field of view gives (x,y) coord

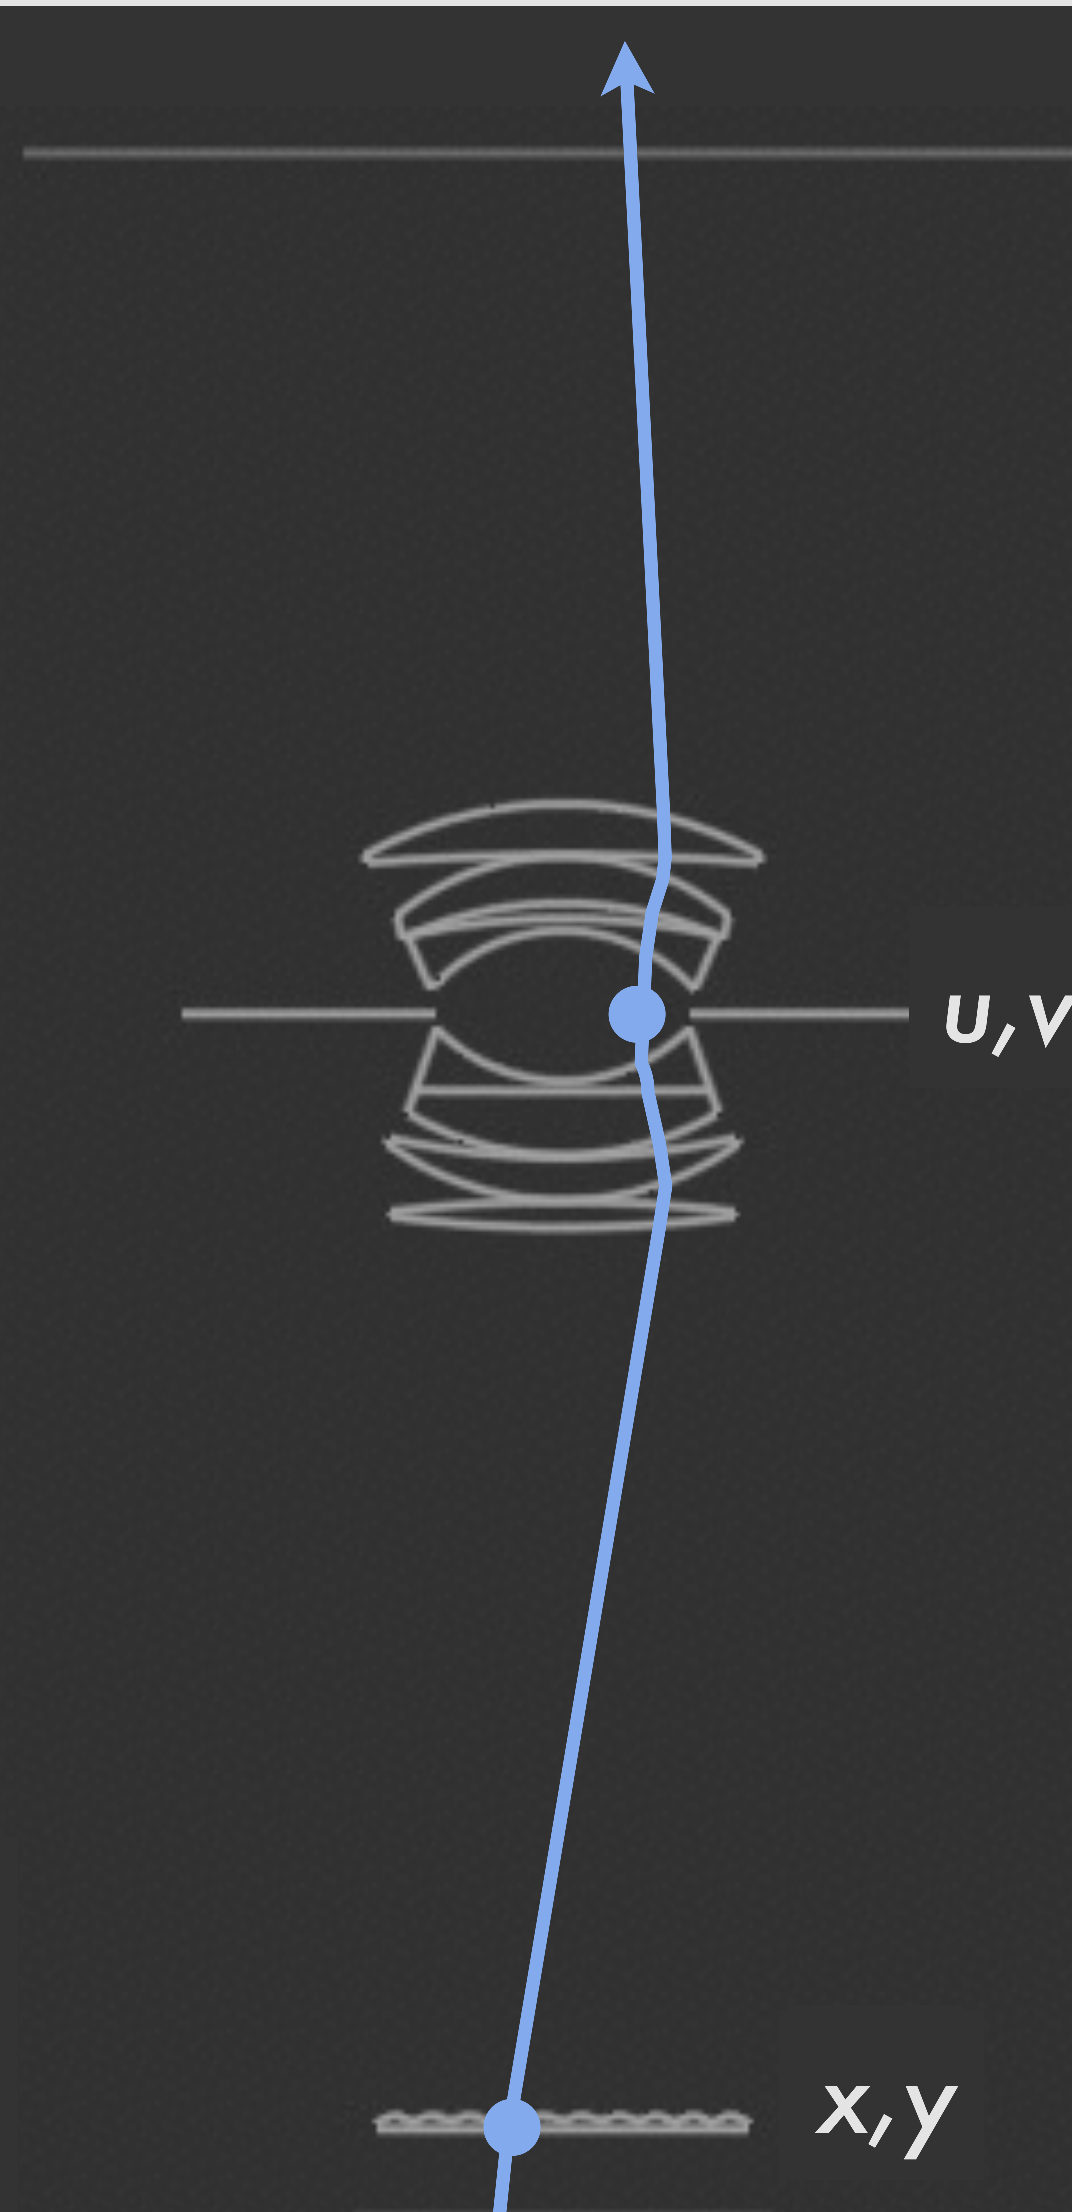

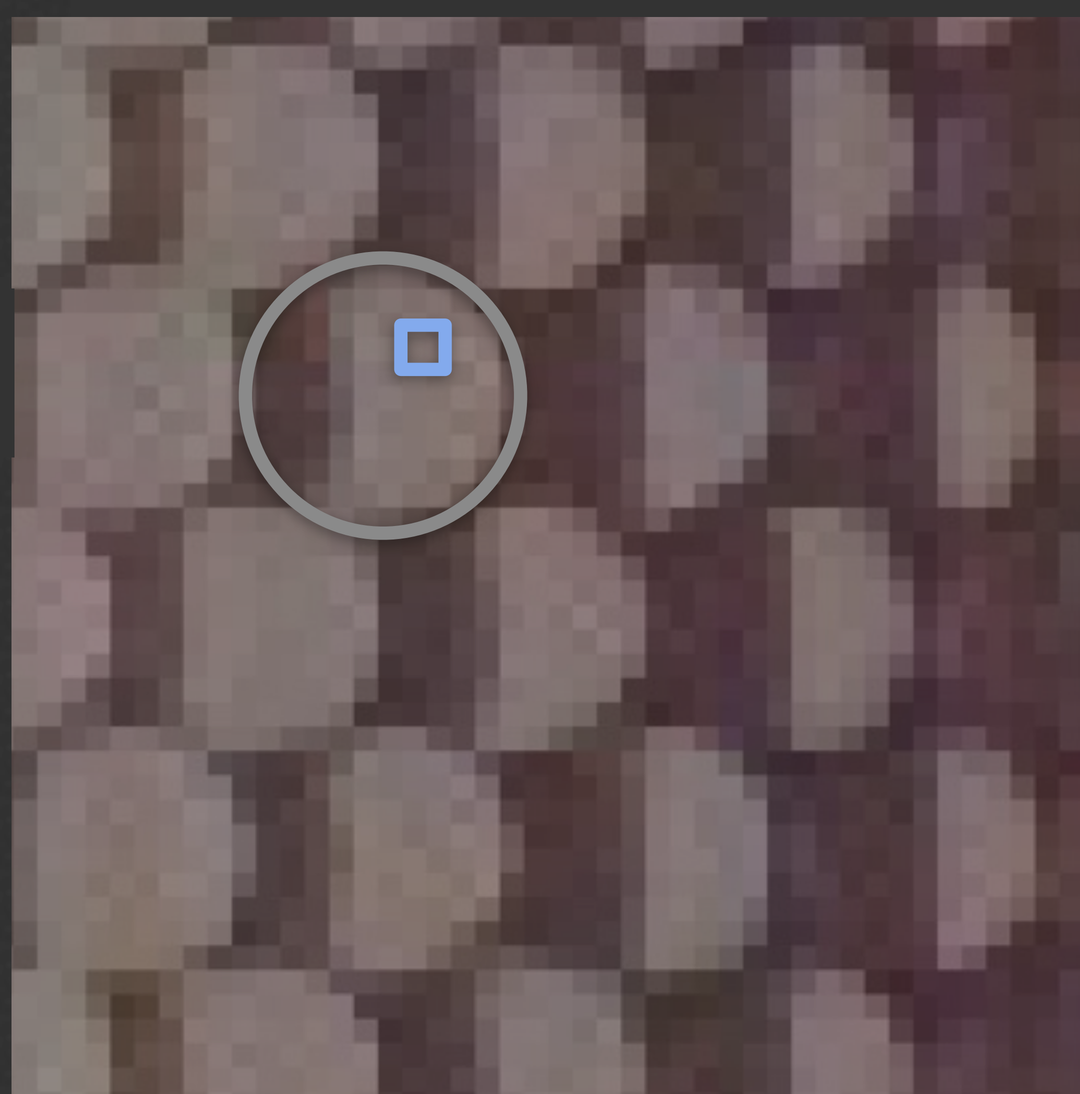

Pixel location in microlens image gives (u,v) coord

Analogy: Shooting RAW vs JPG (2D Photos)

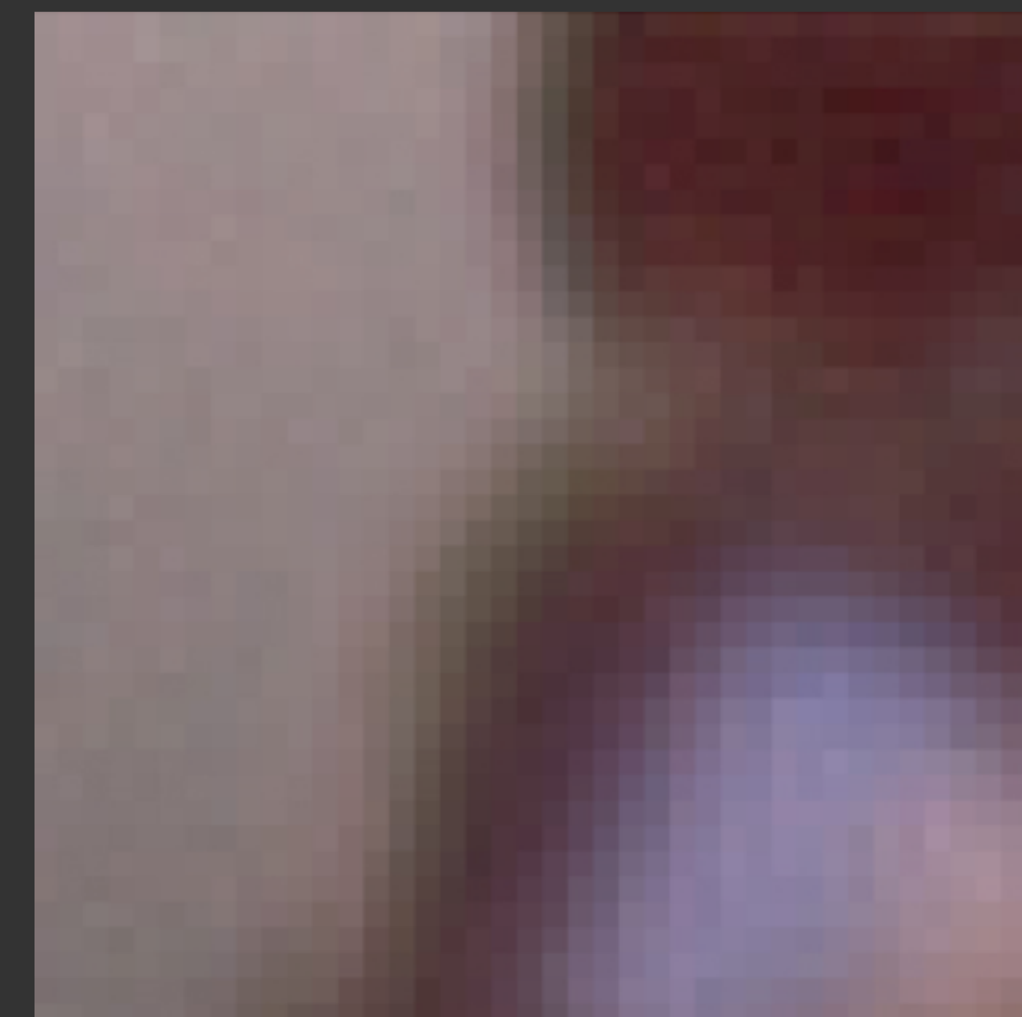

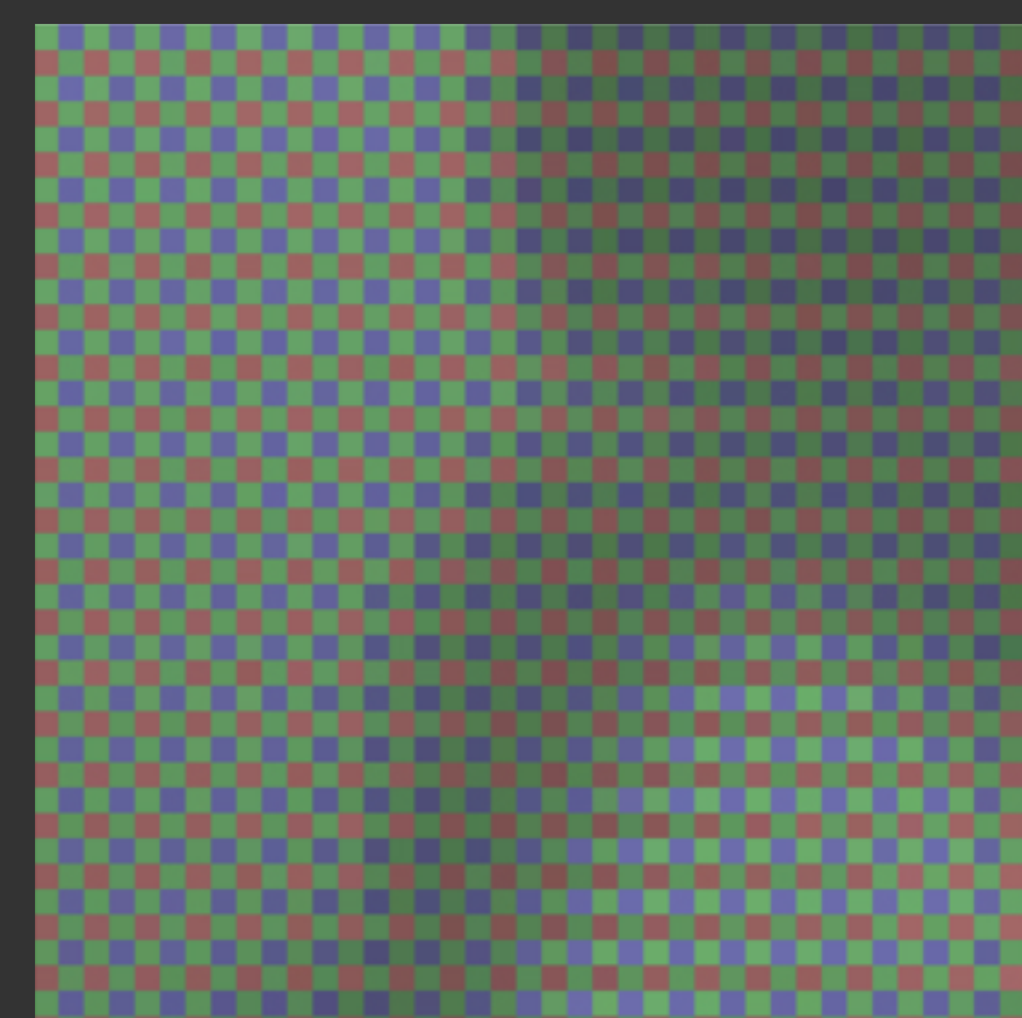

RAW Bayer 14-bit data preserves intensity information, and provides flexibility over exposure and color.

Shooting 4D Light Field vs 2D Photographs

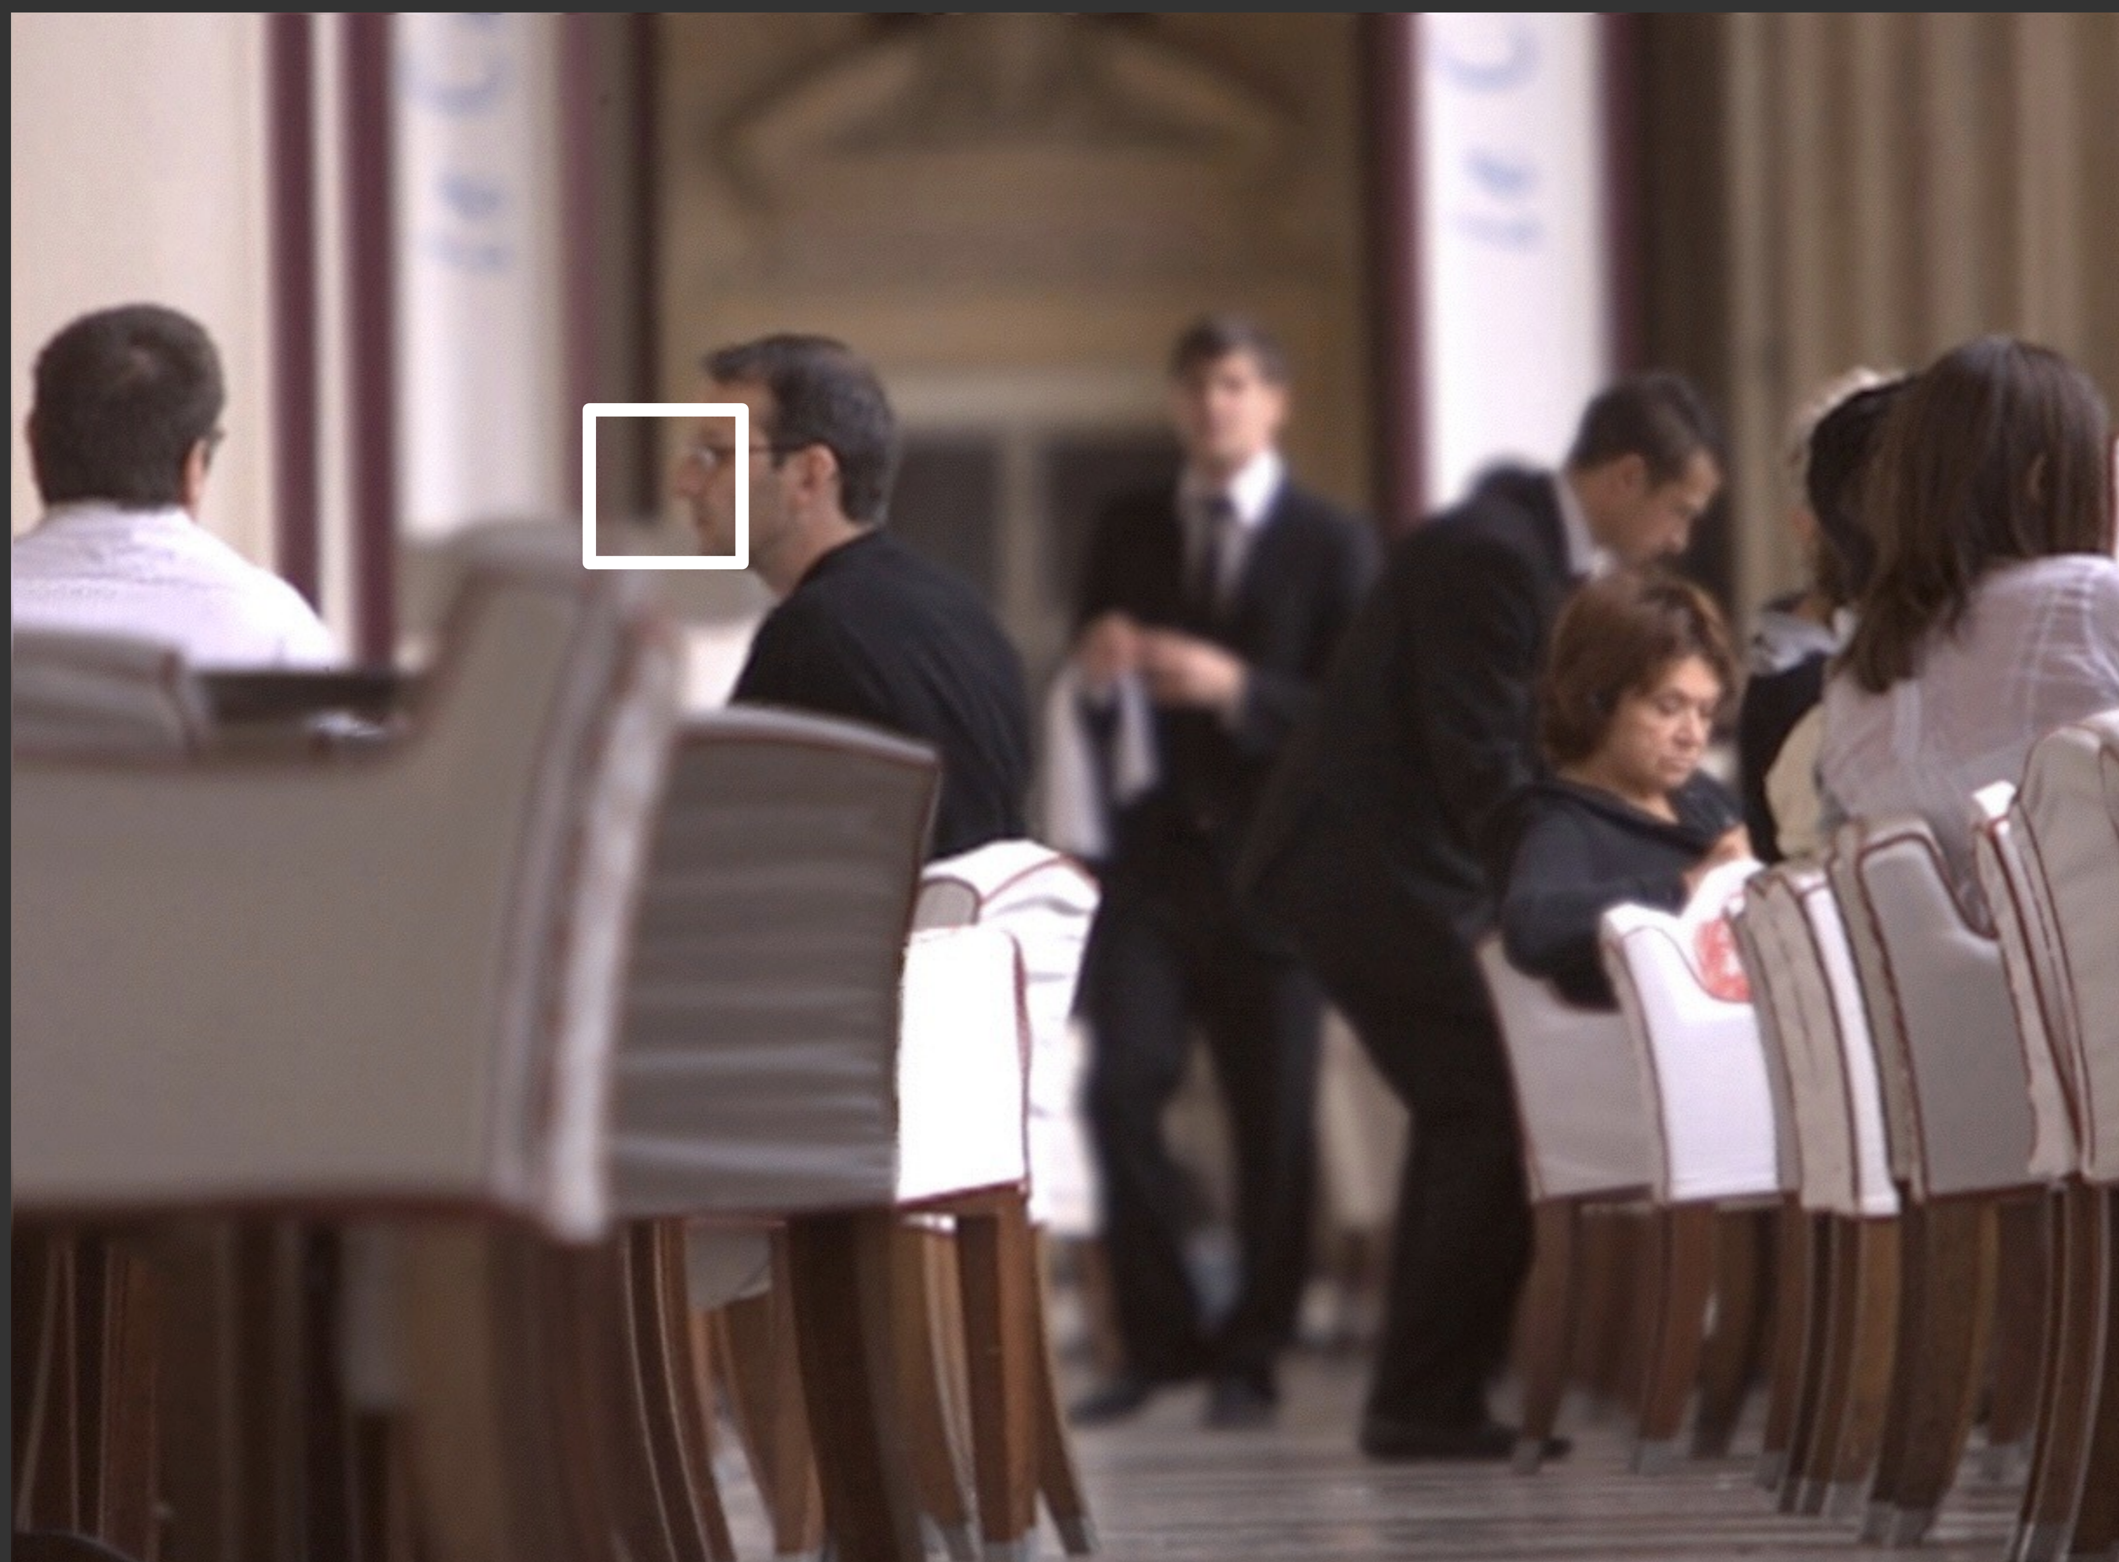

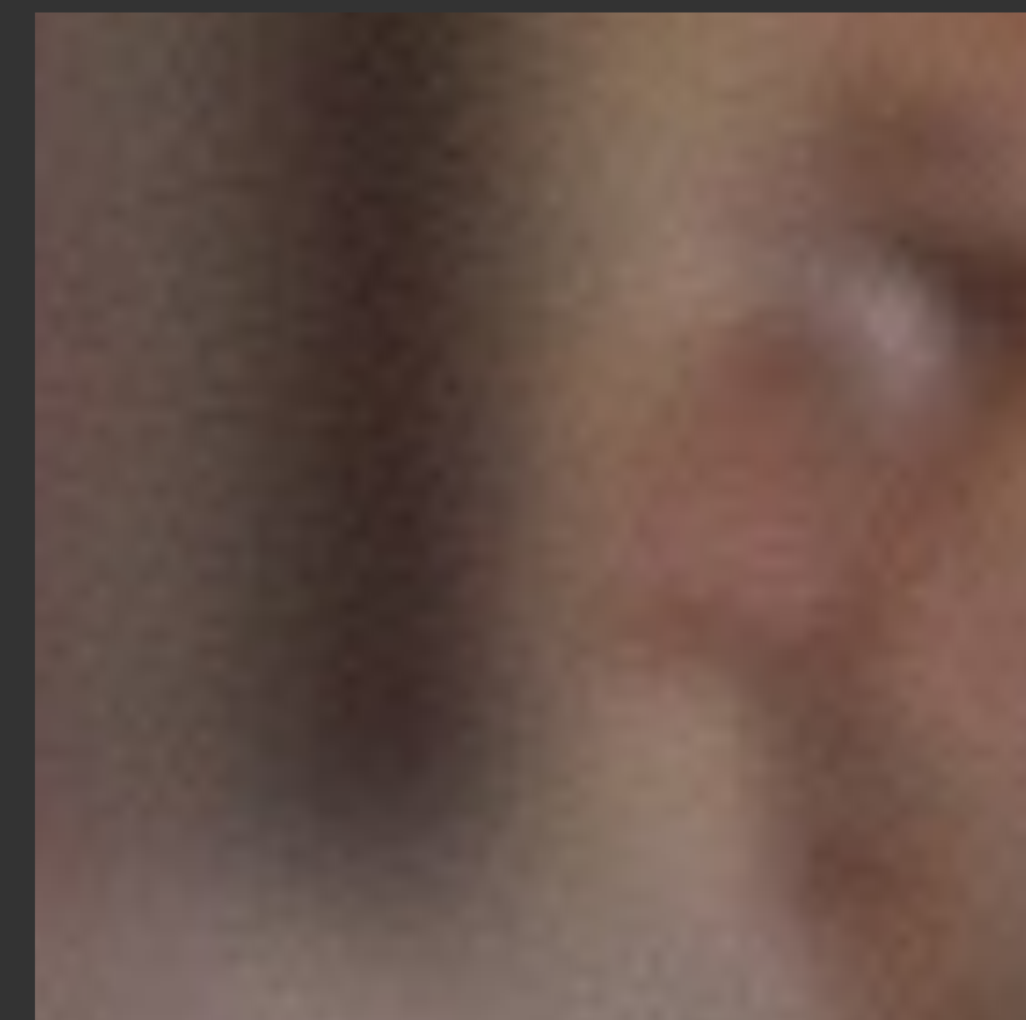

Light field data preserves directional intensity information, and provides flexibility over focus and aberrations.

Check your understanding

Sub-Aperture Images

Image from selecting the same pixel under every microlens

Sub-Aperture Images

Image from selecting same pixel under every microlens

Sub-Aperture Images

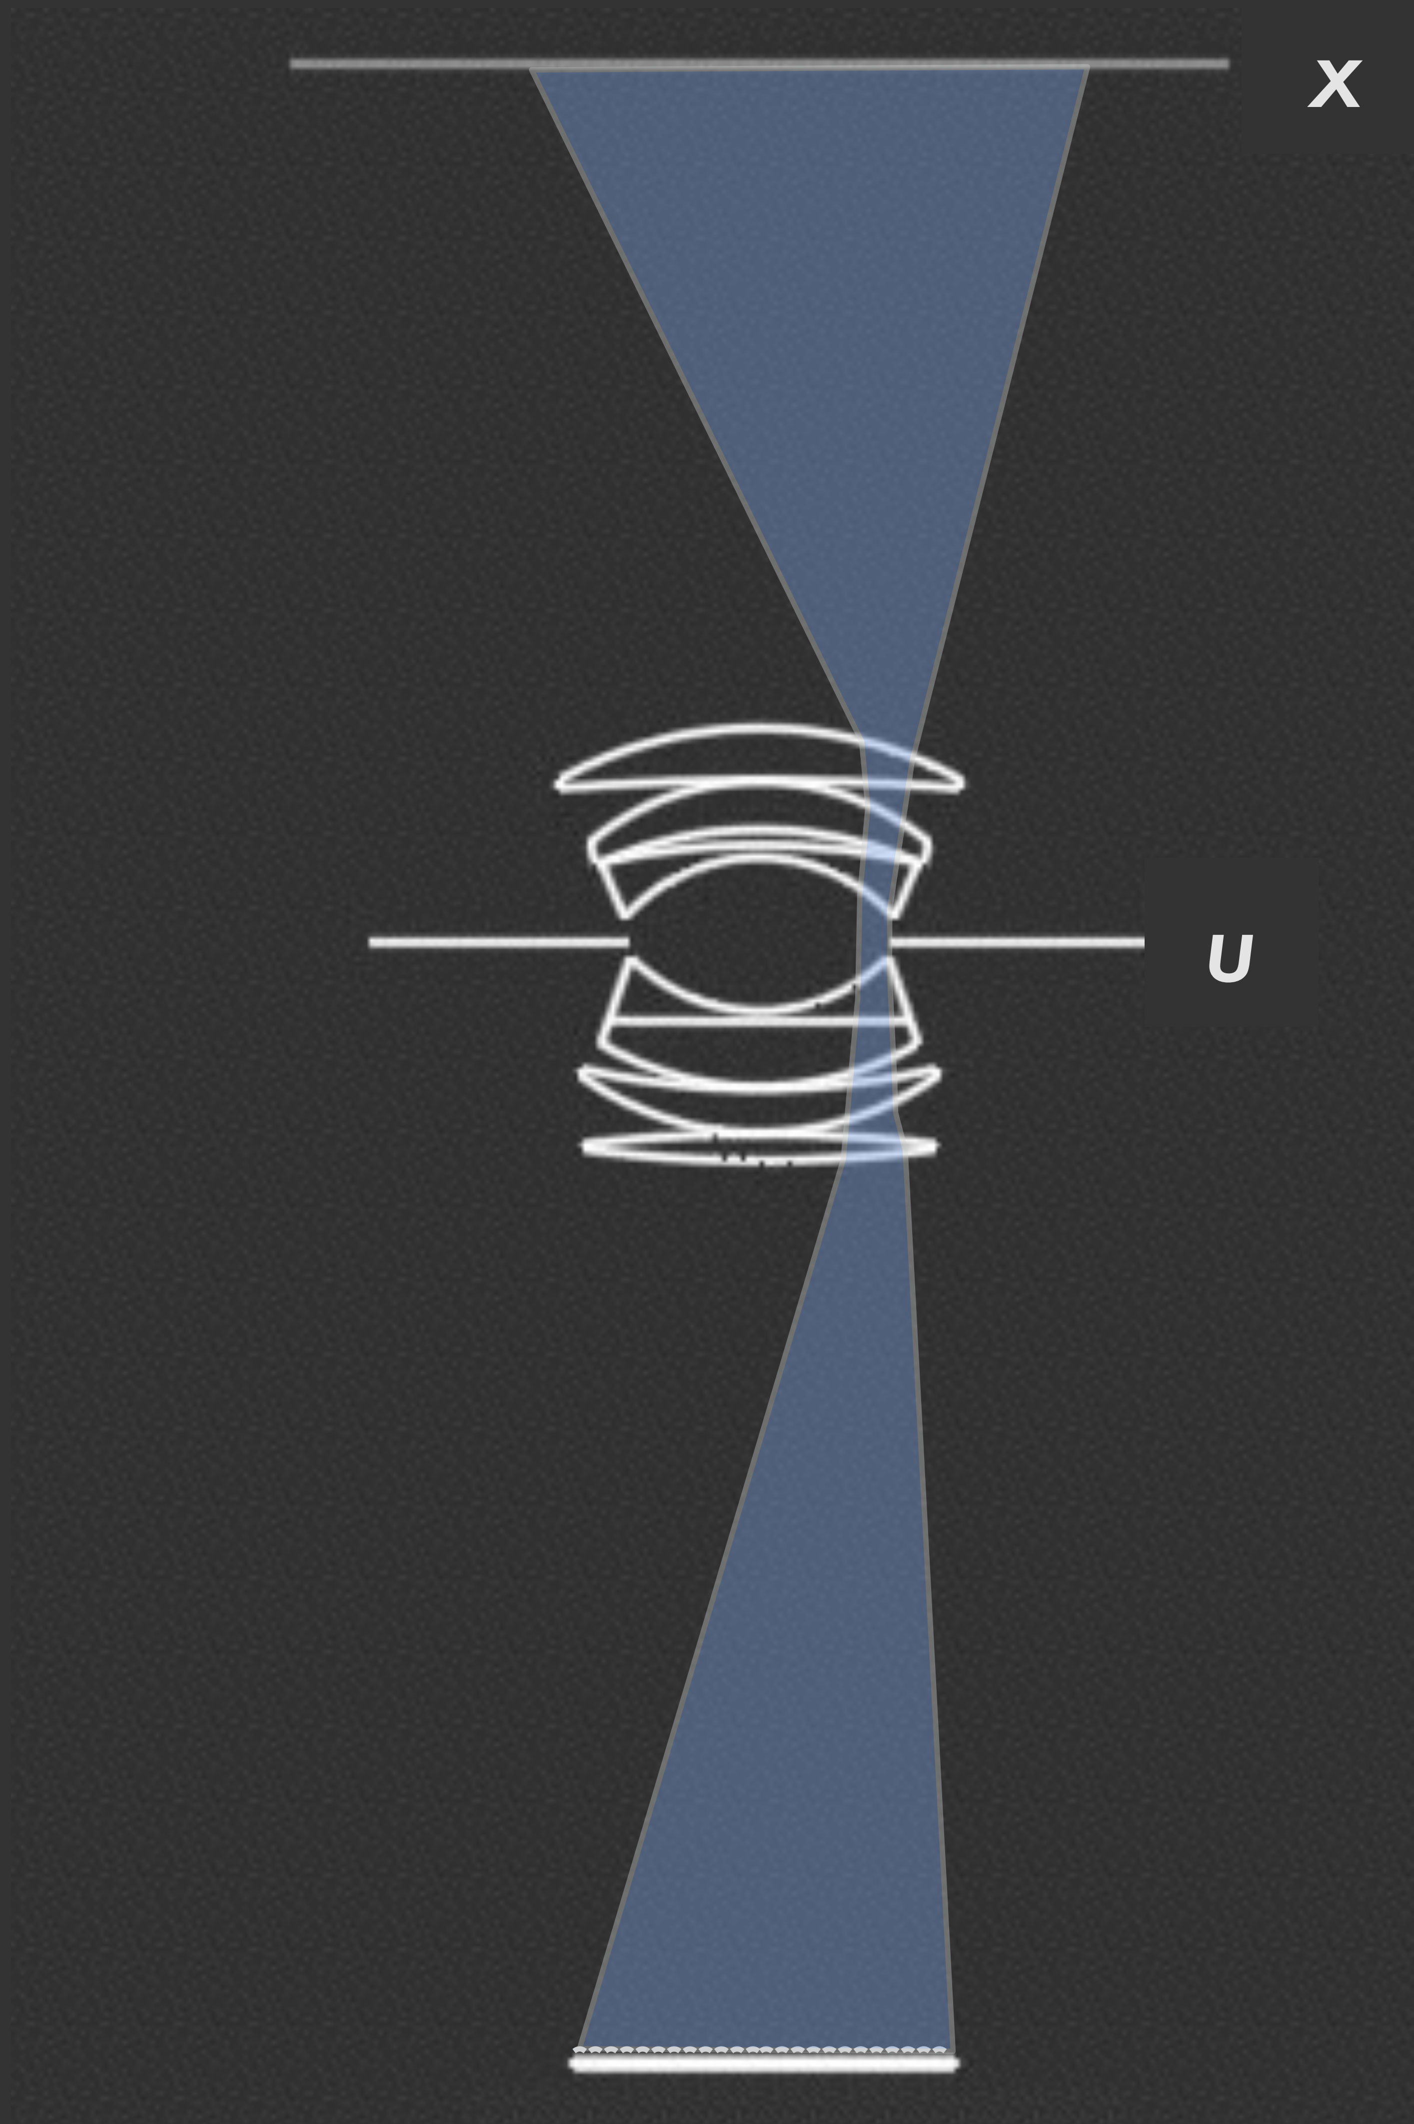

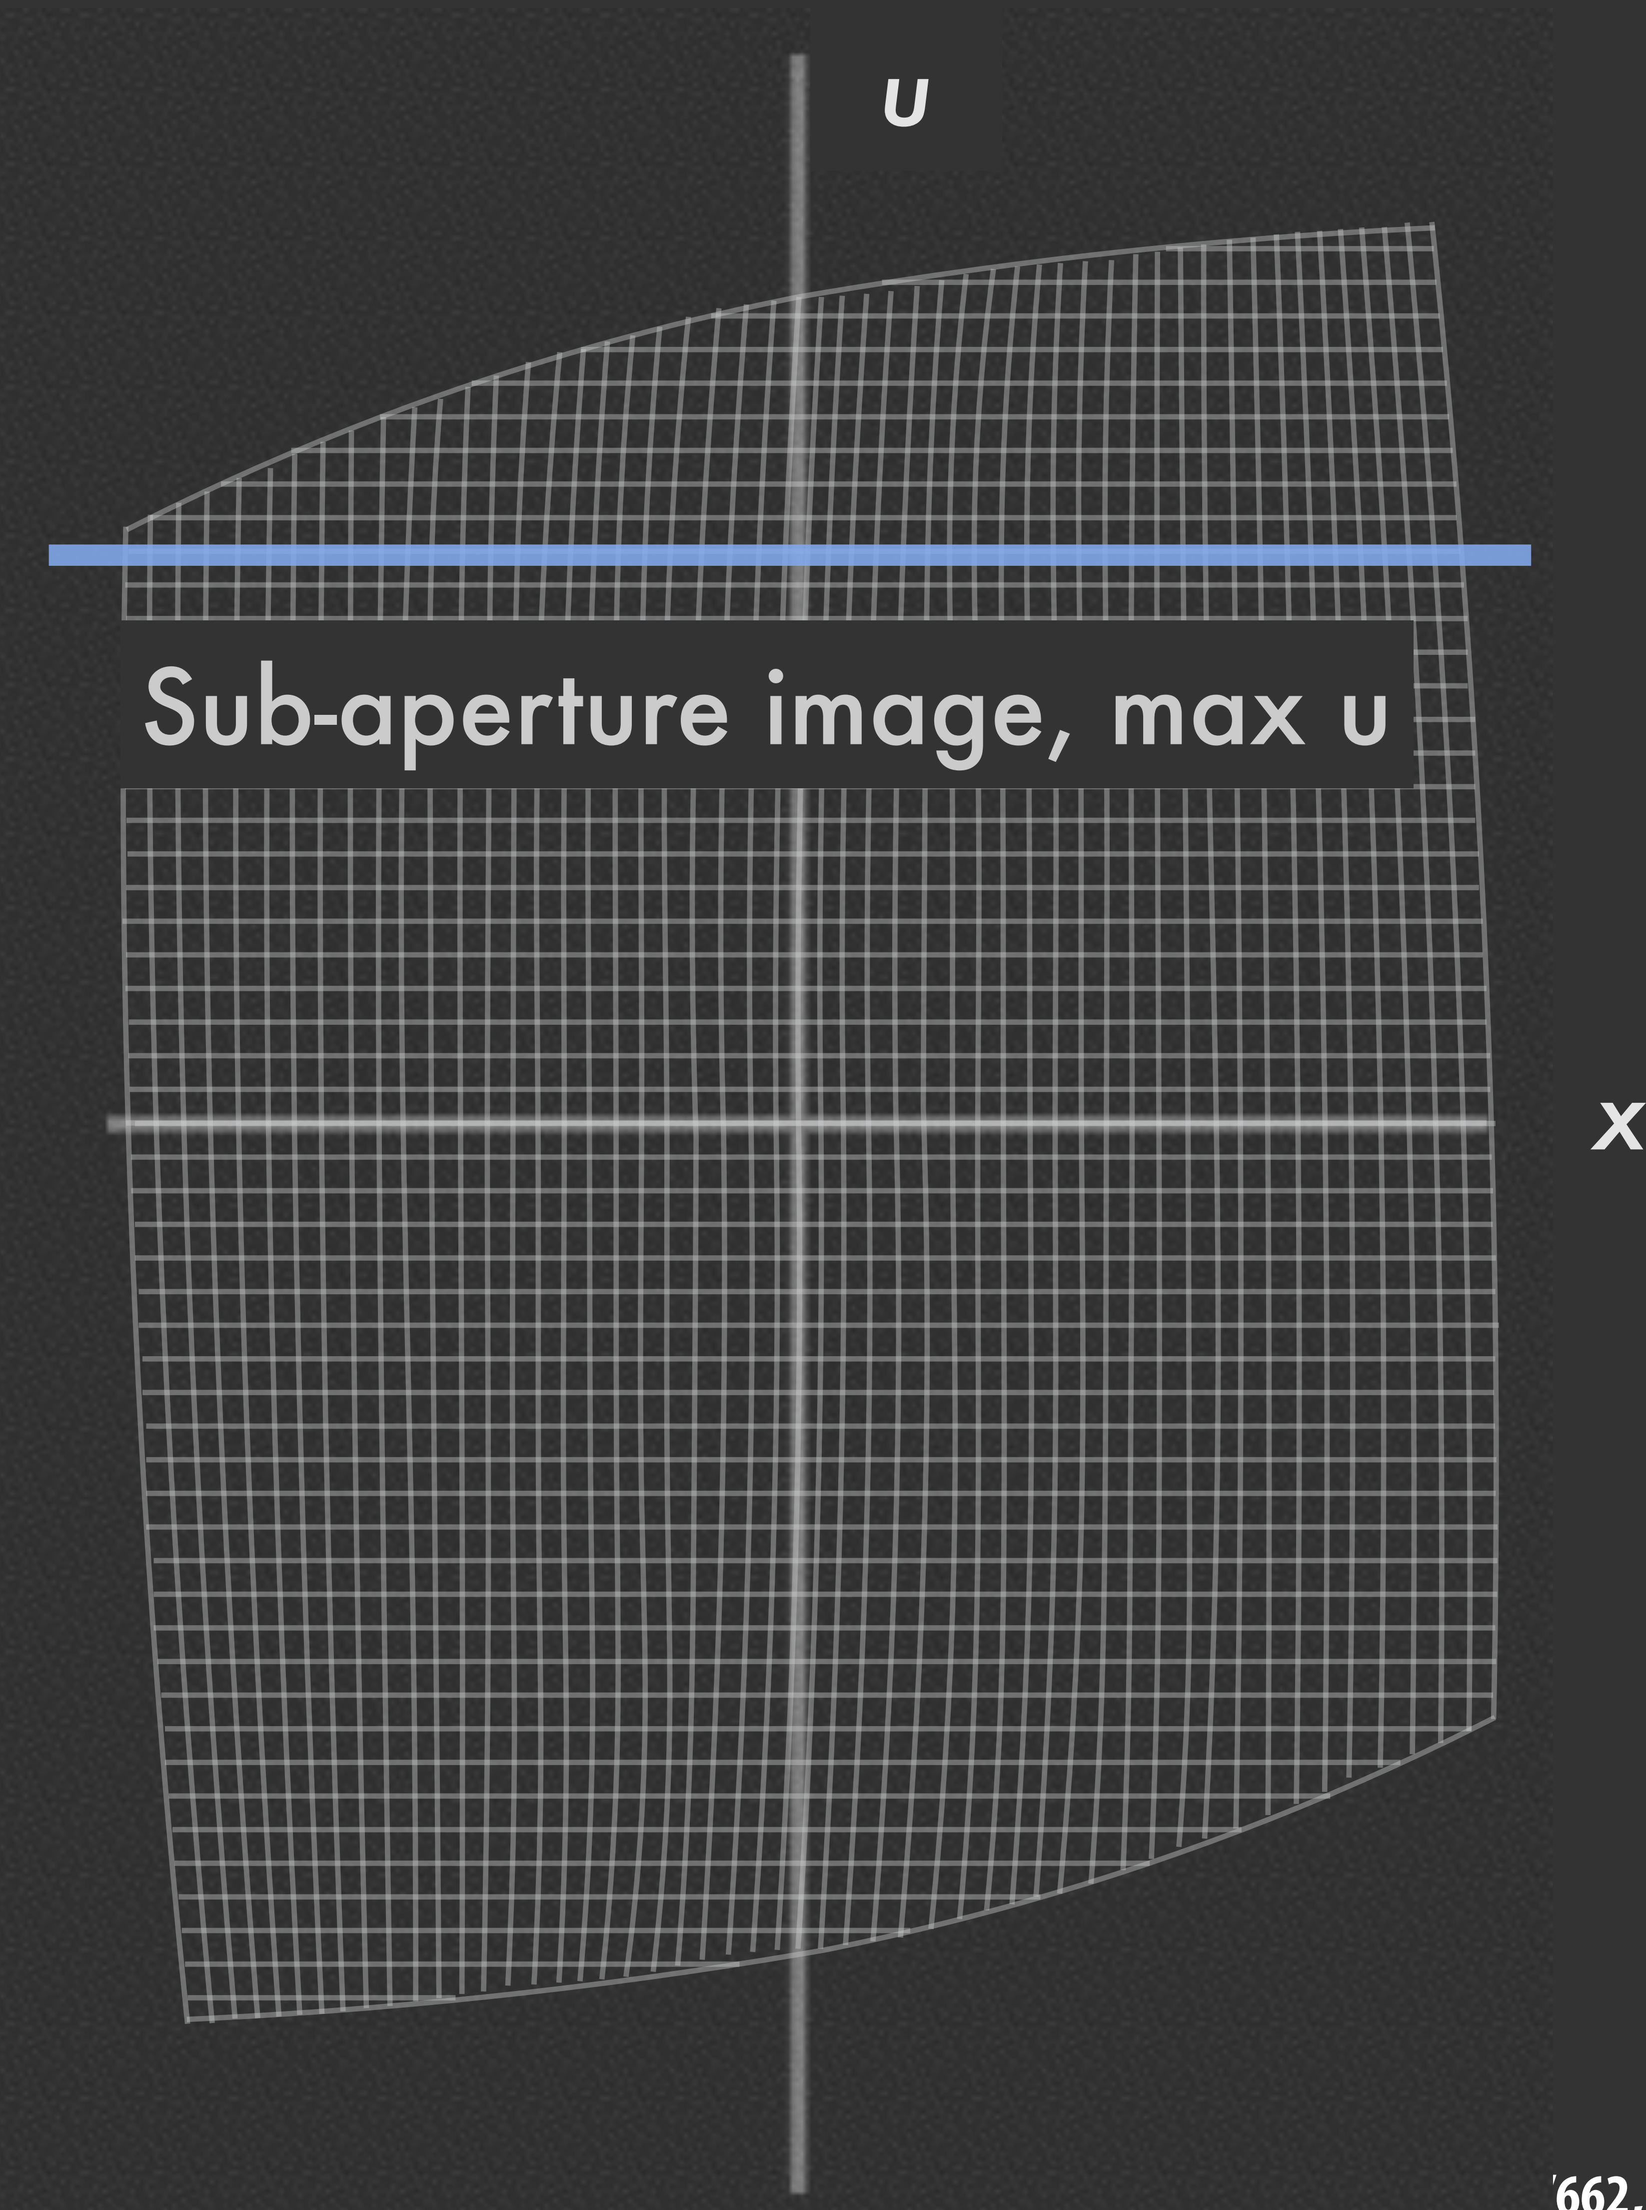

Sub-Aperture Images

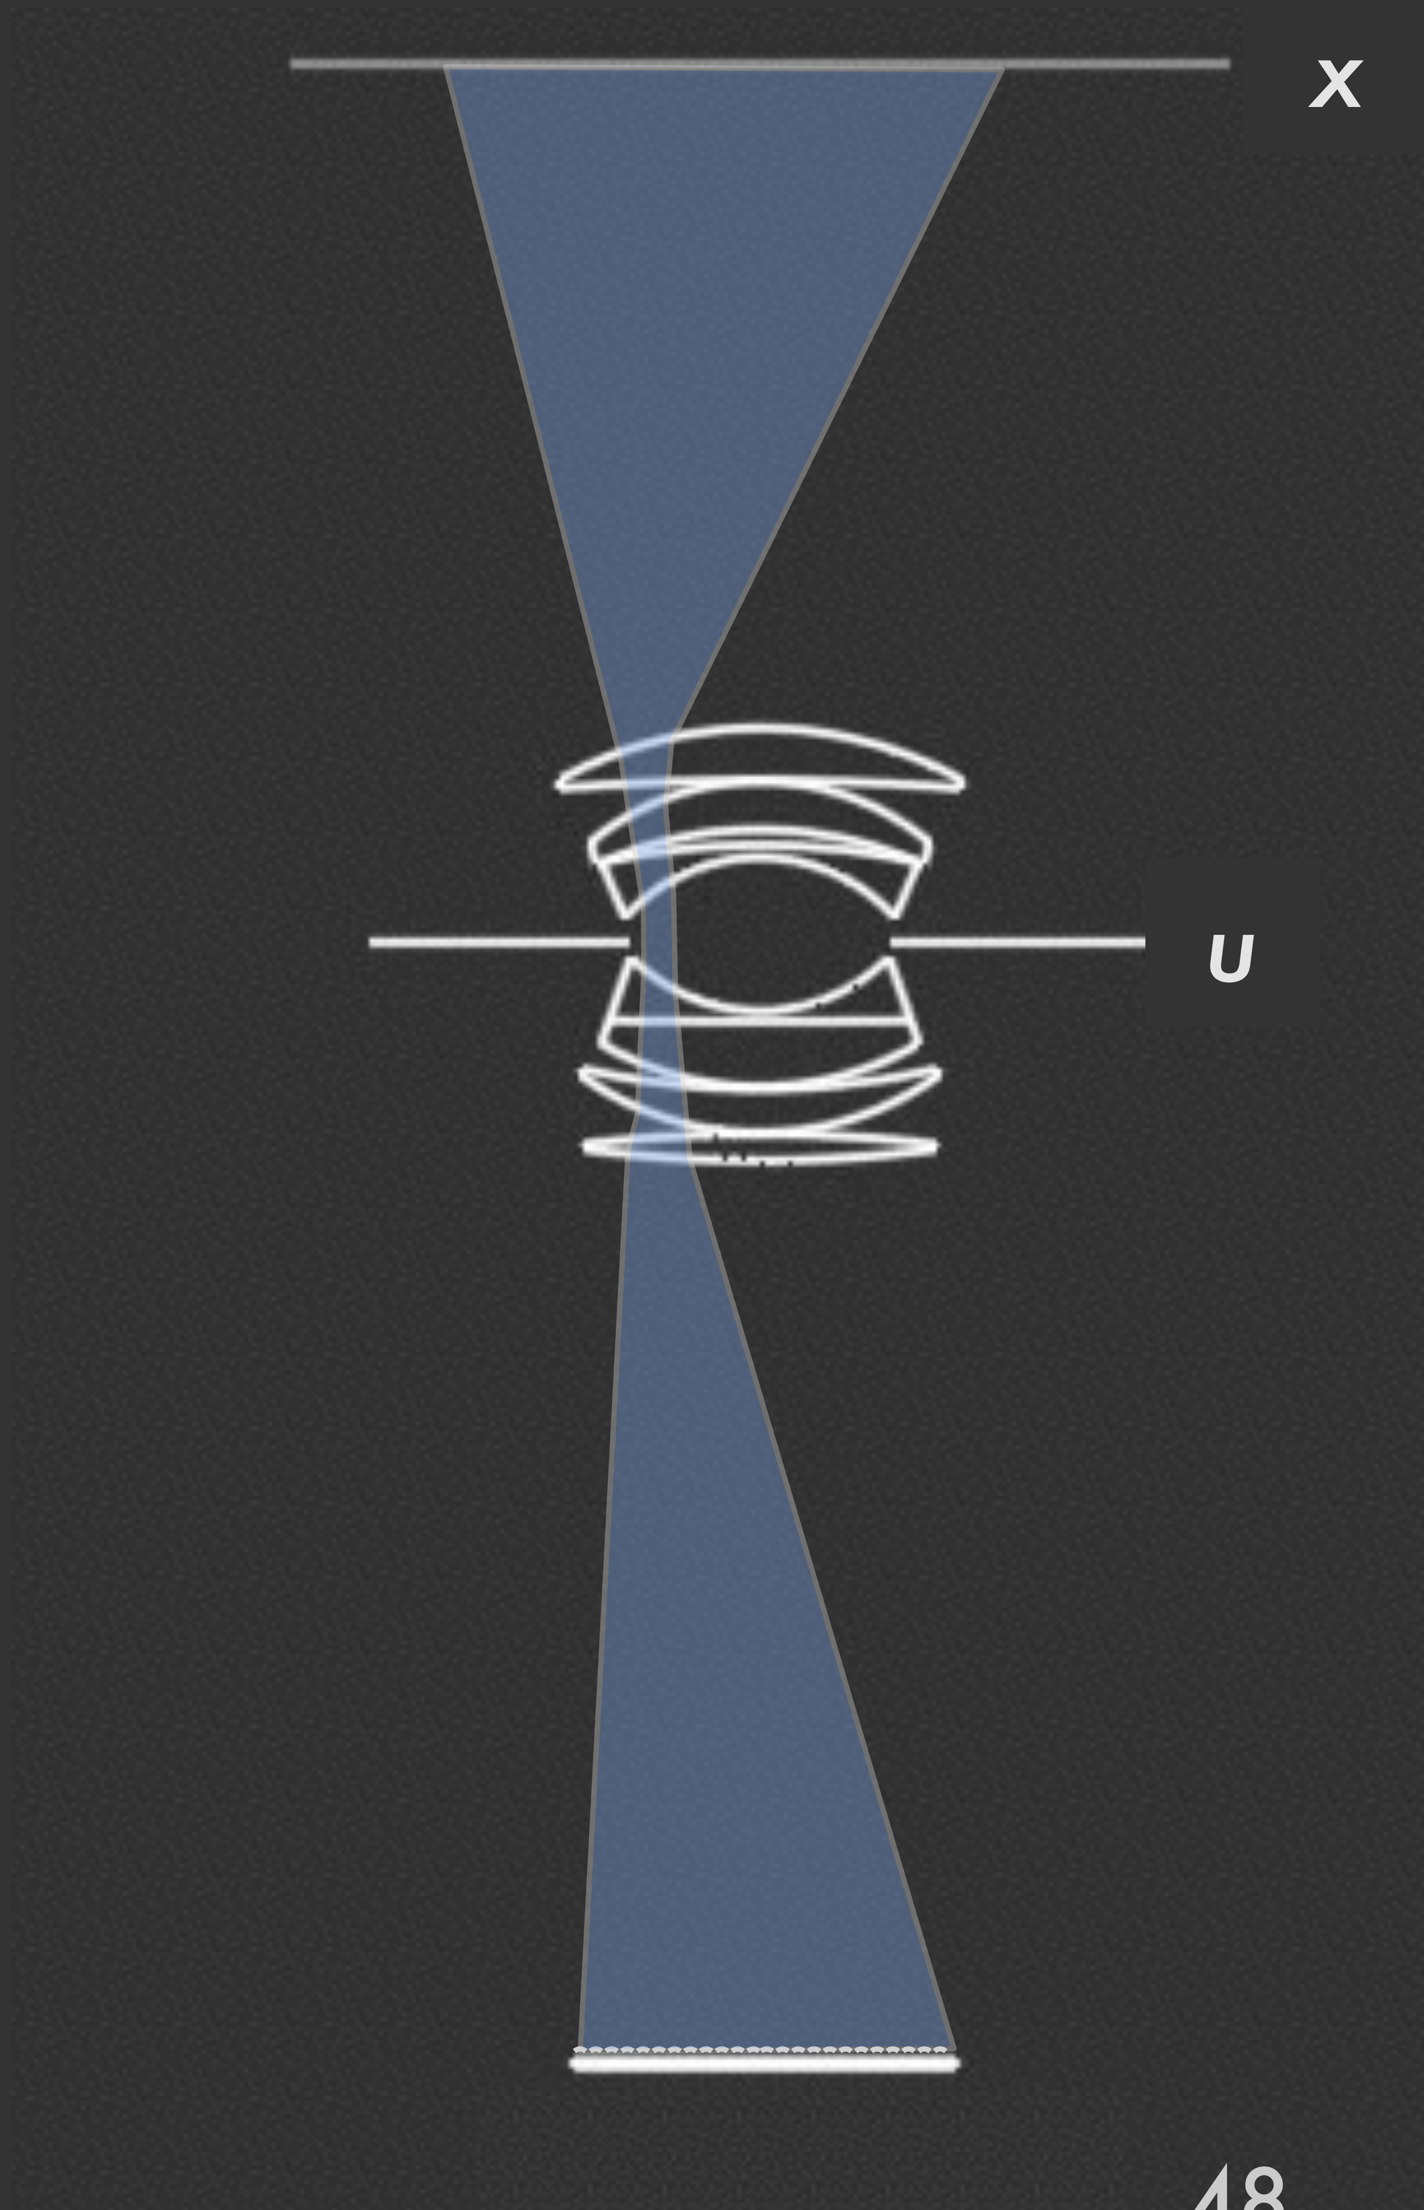

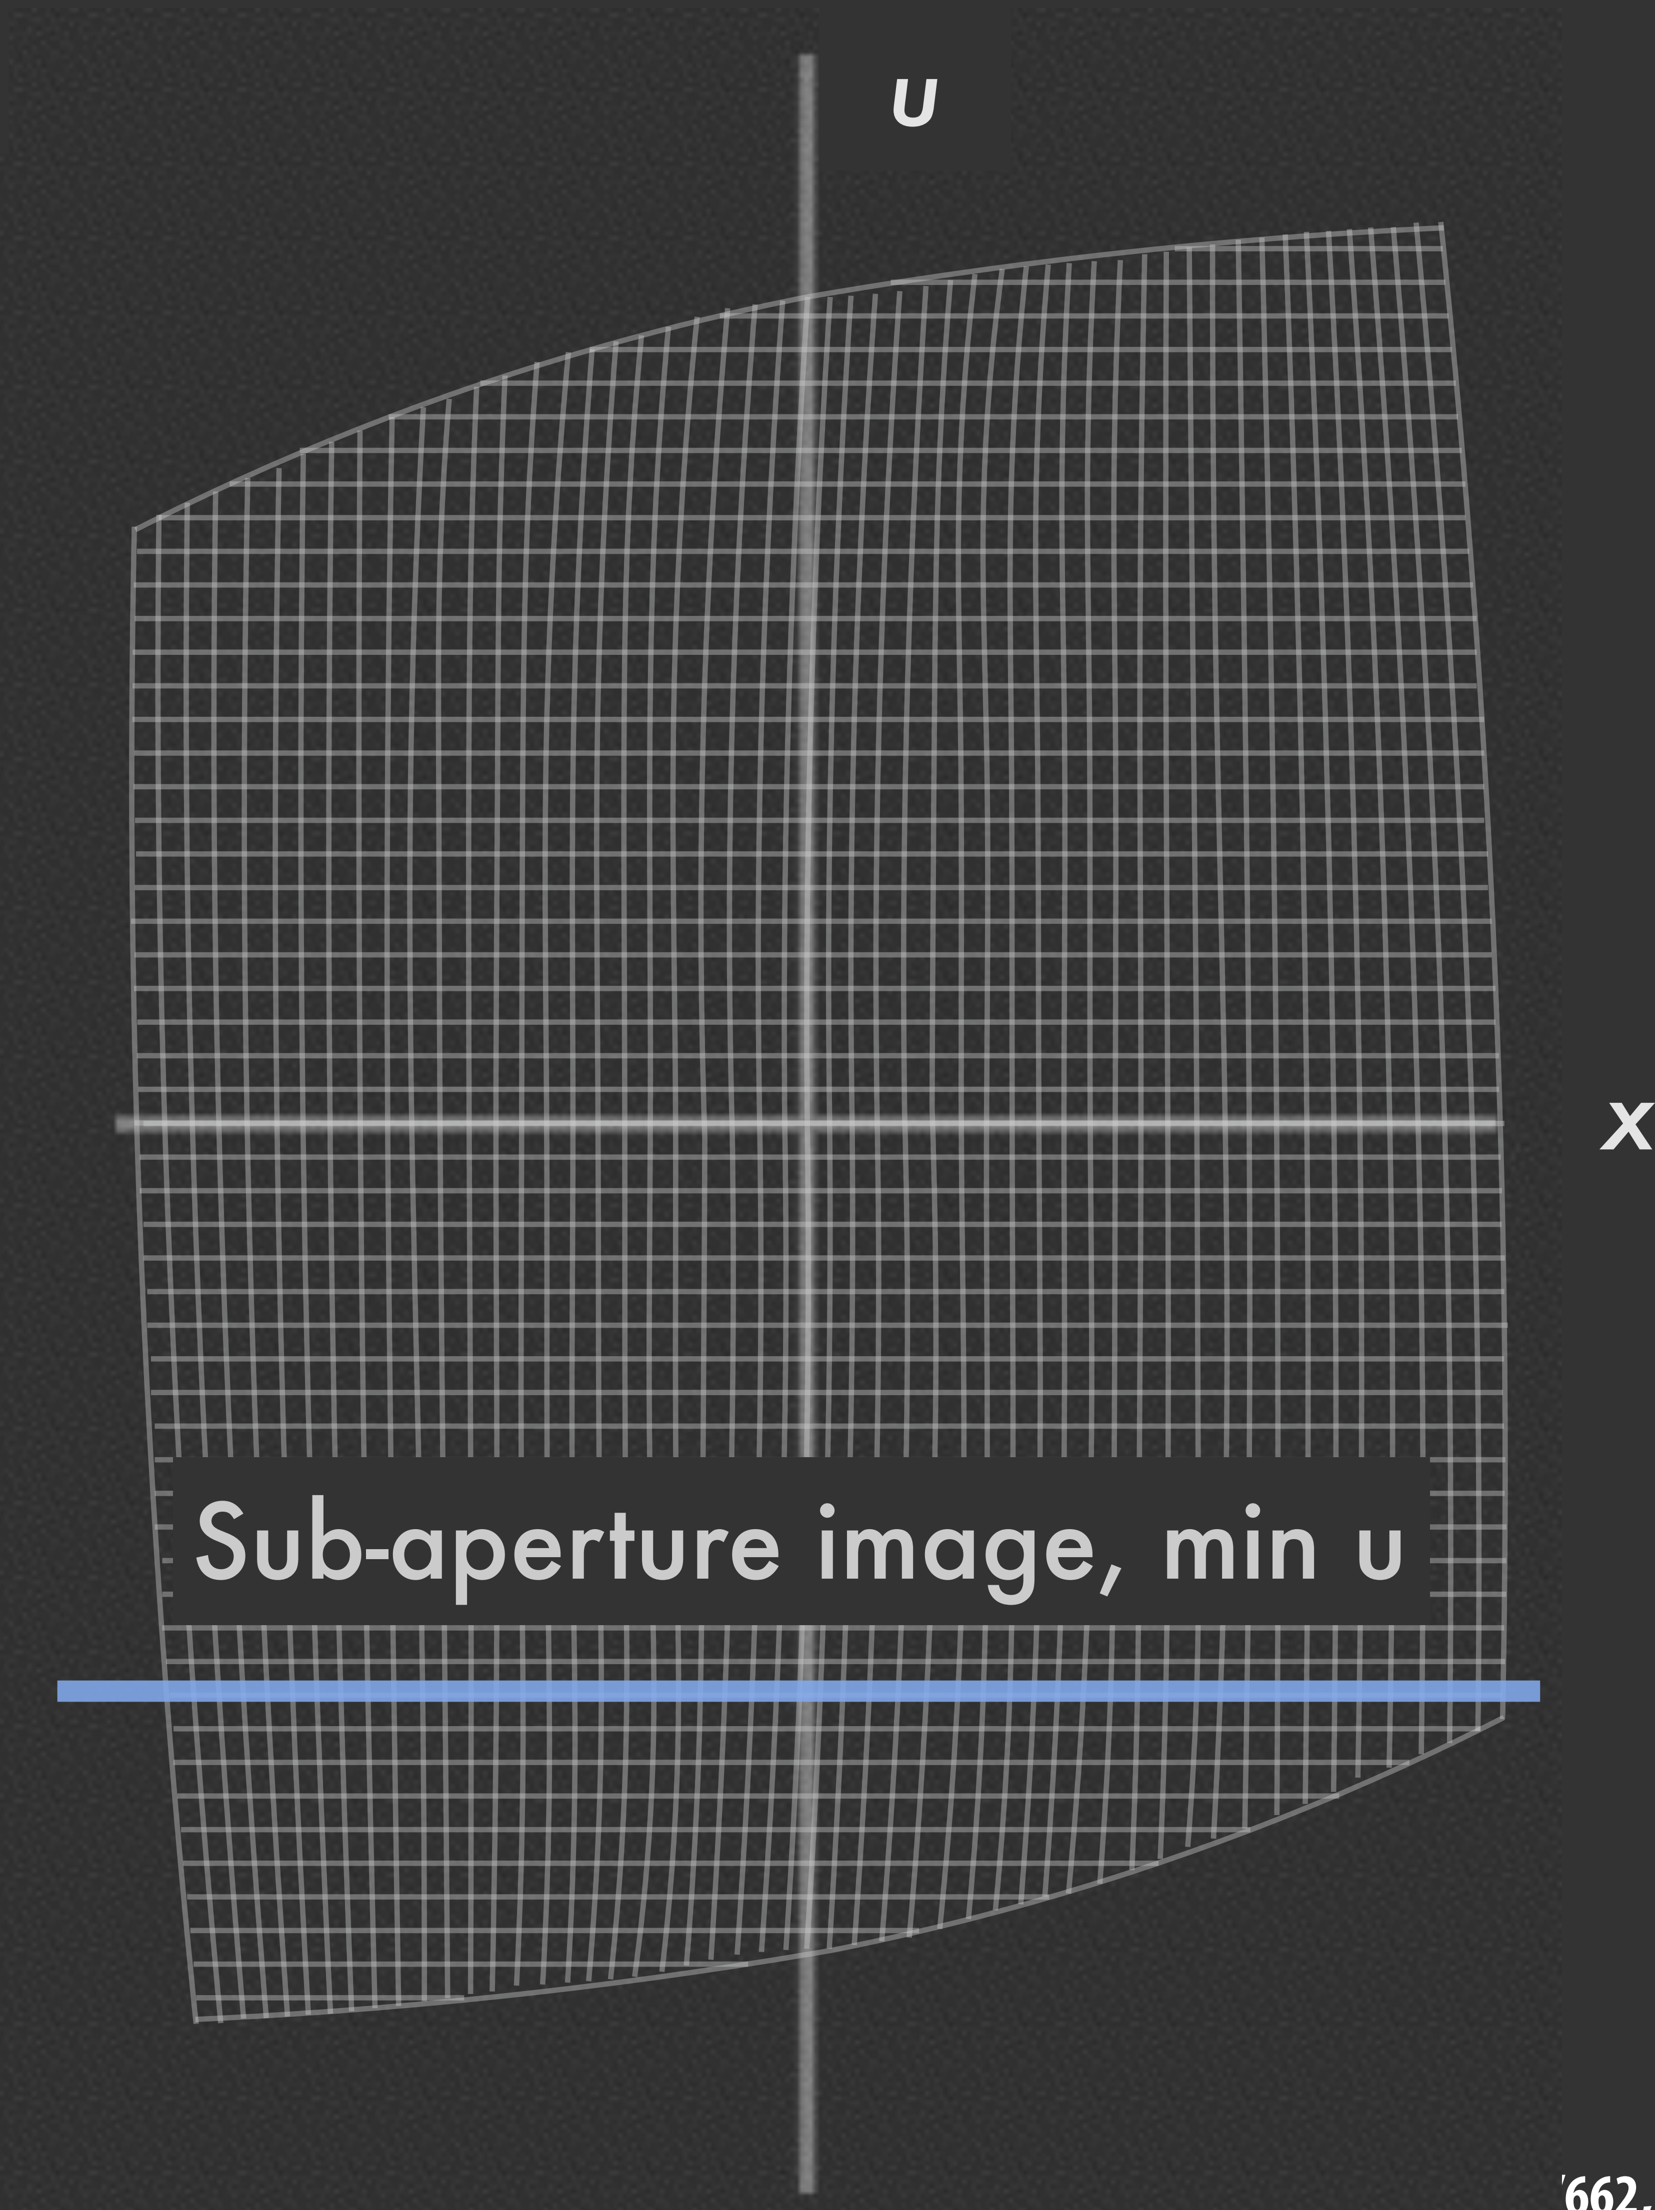

Application: computational Change of Viewpoint

Lateral movement (left)

Application: computational Change of Viewpoint

Lateral movement (right)

How do you compute a photograph from a light field?

Computing a photograph from a light field

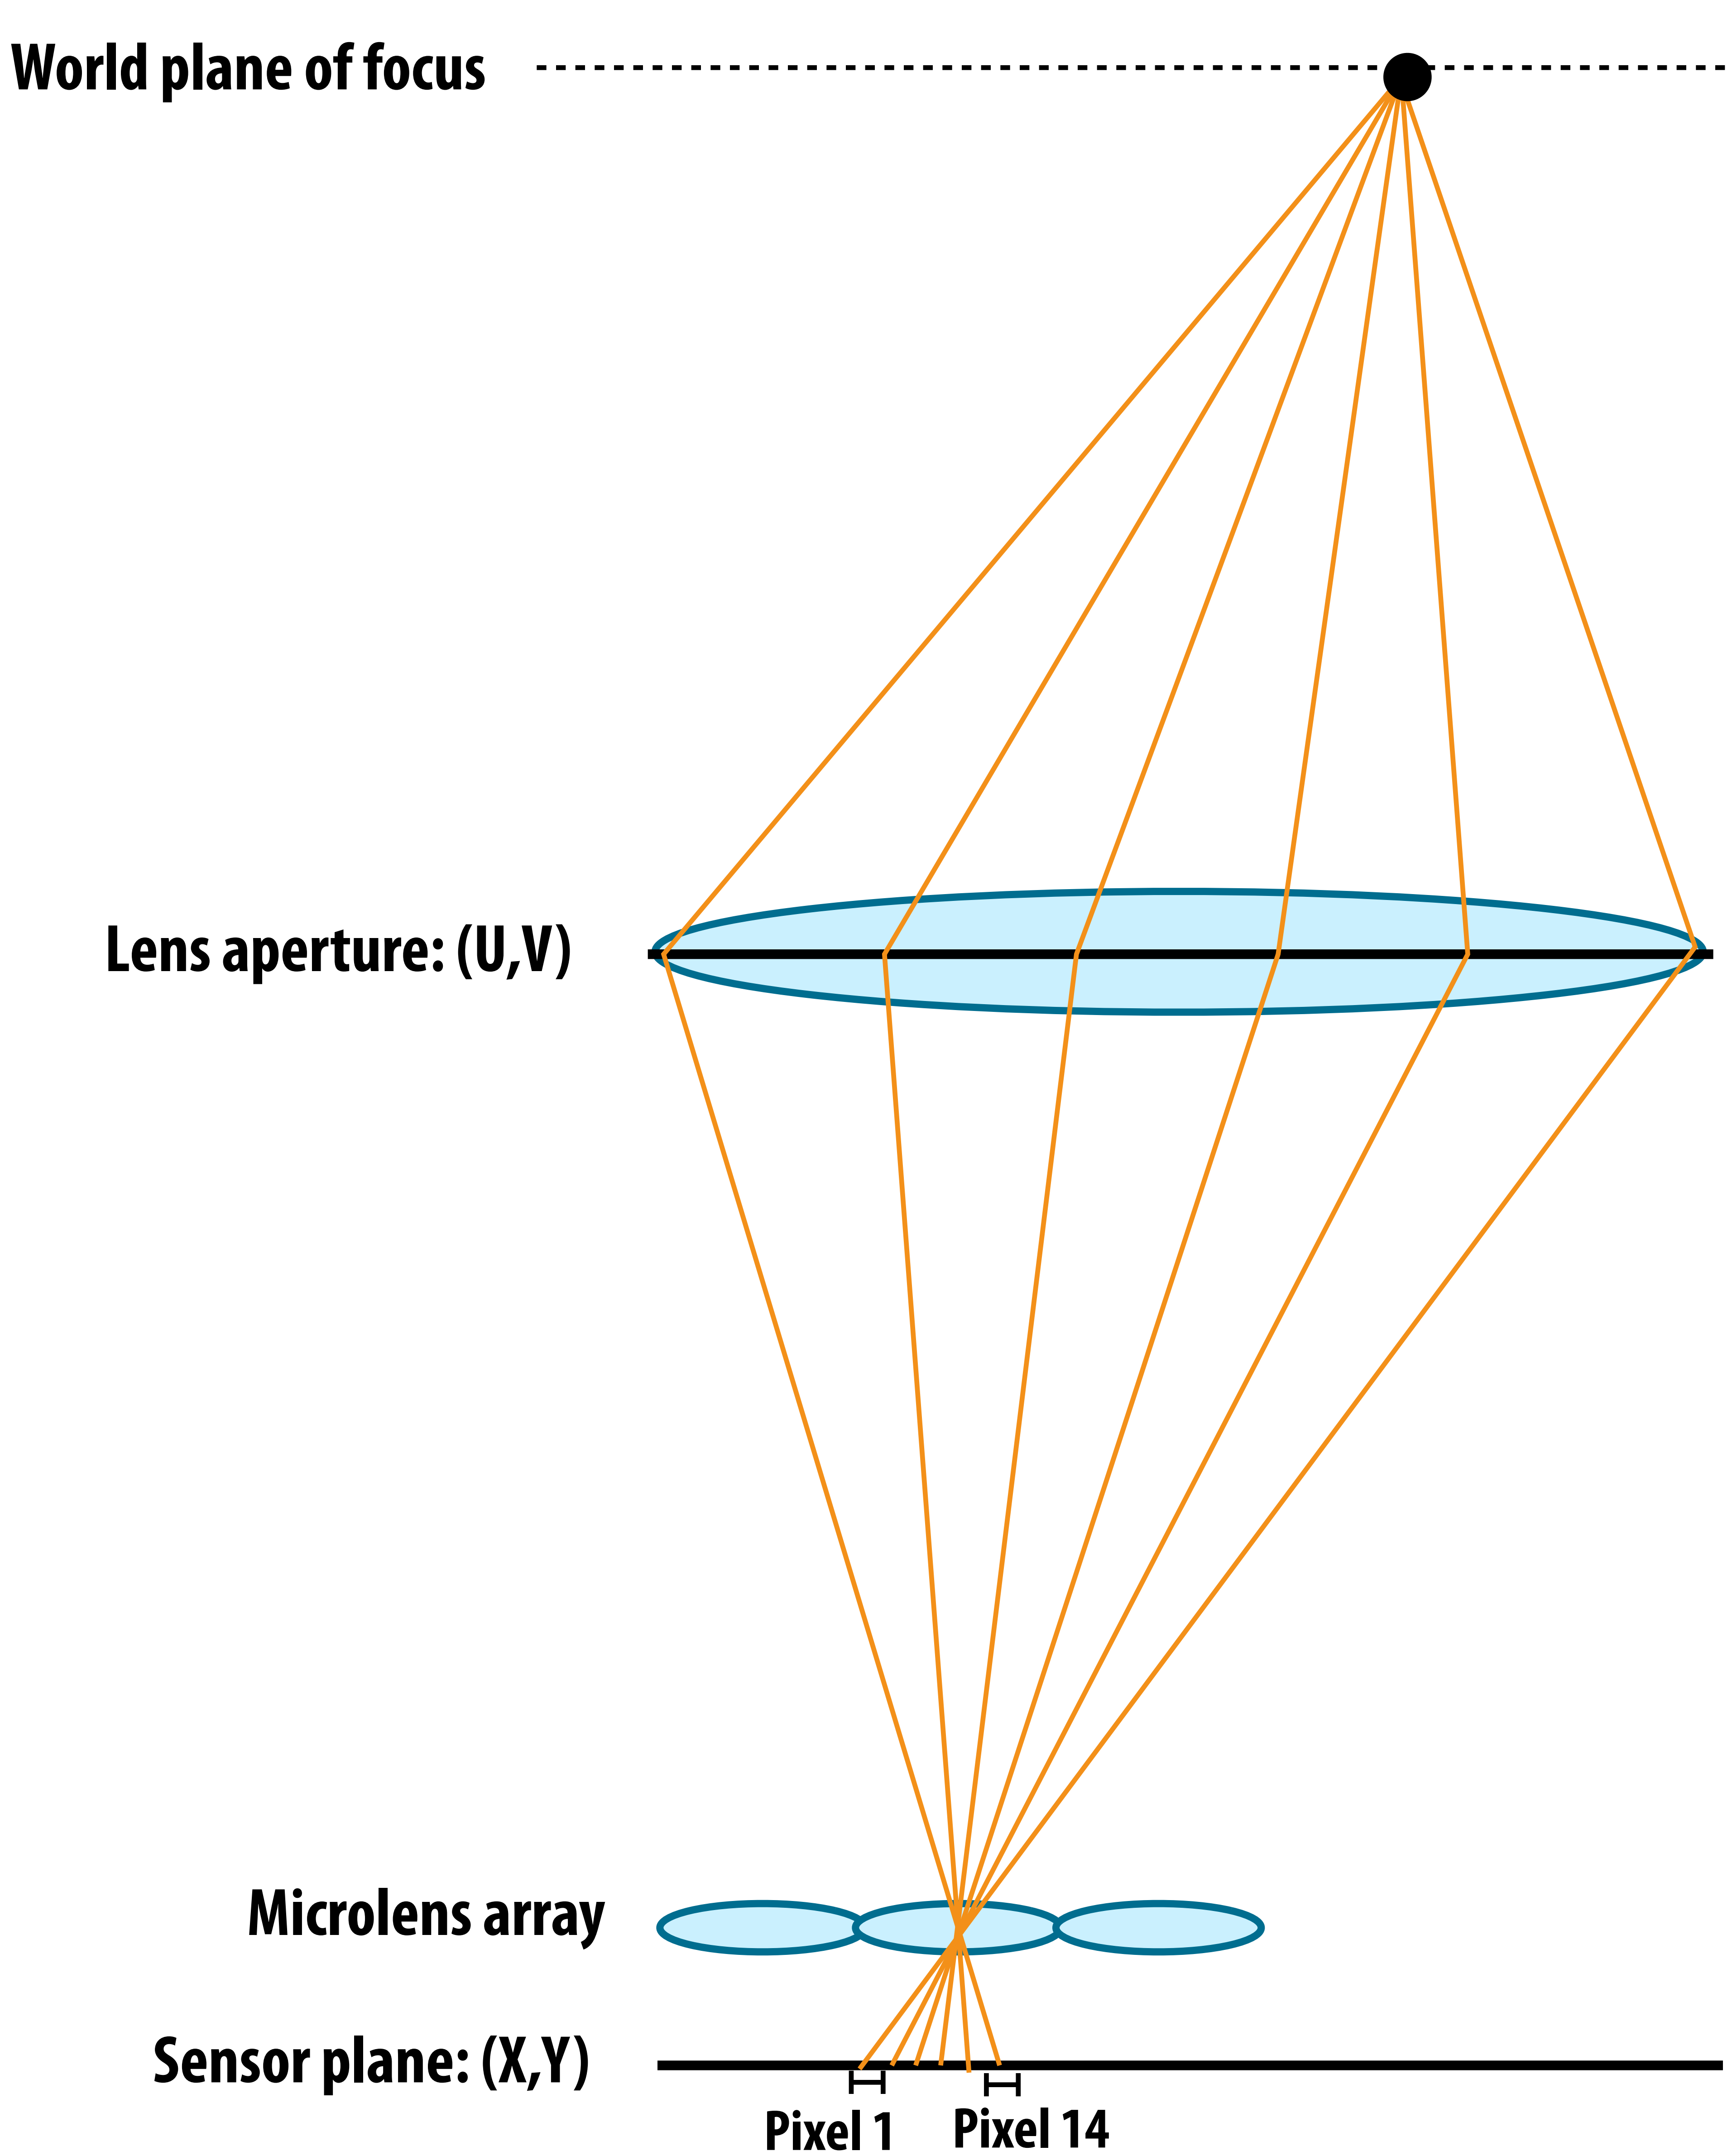

Ray space plot

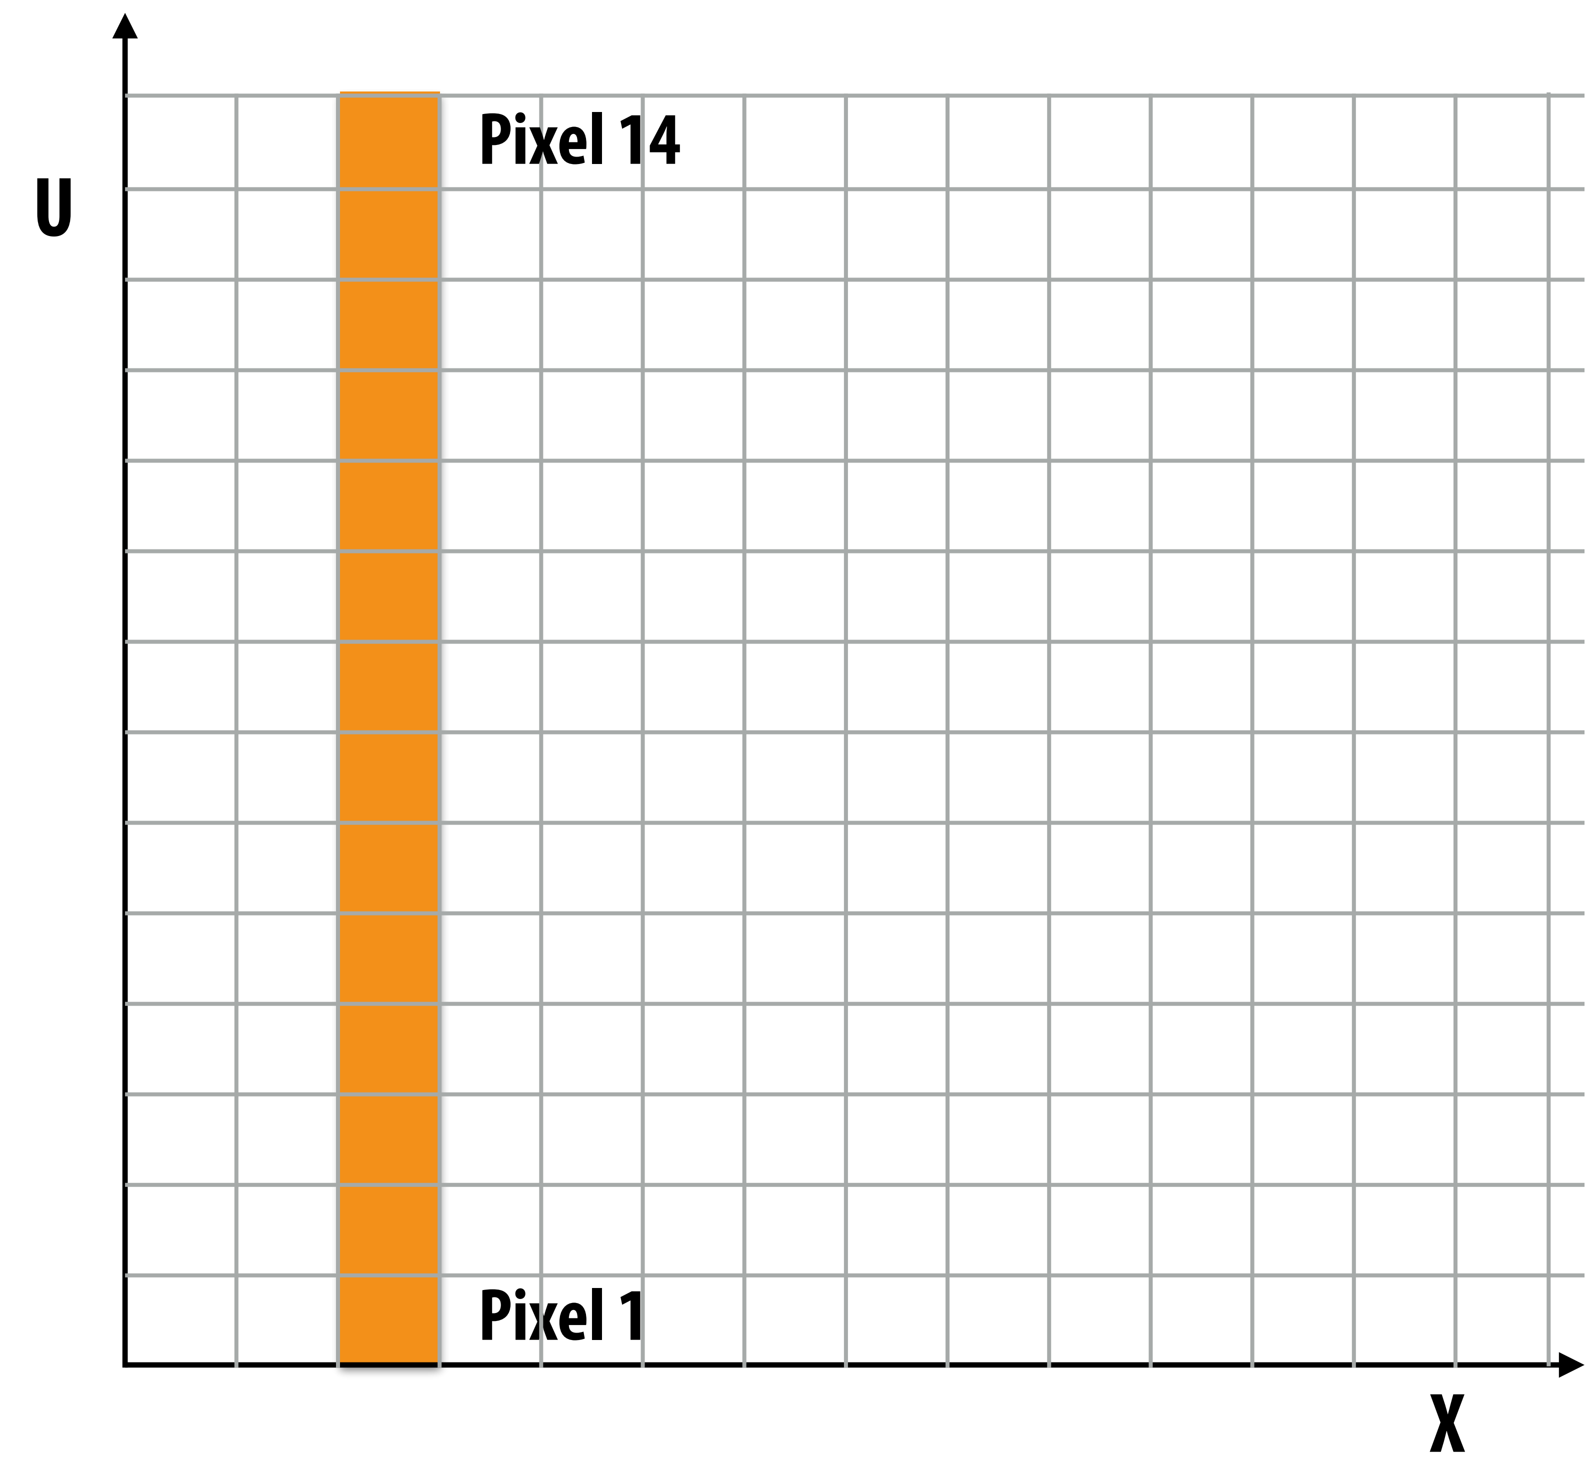

Computing photograph is integral projection (Output image pixel is sum of highlighted light-field sensor pixels)

Computing a photograph at a new focal plane

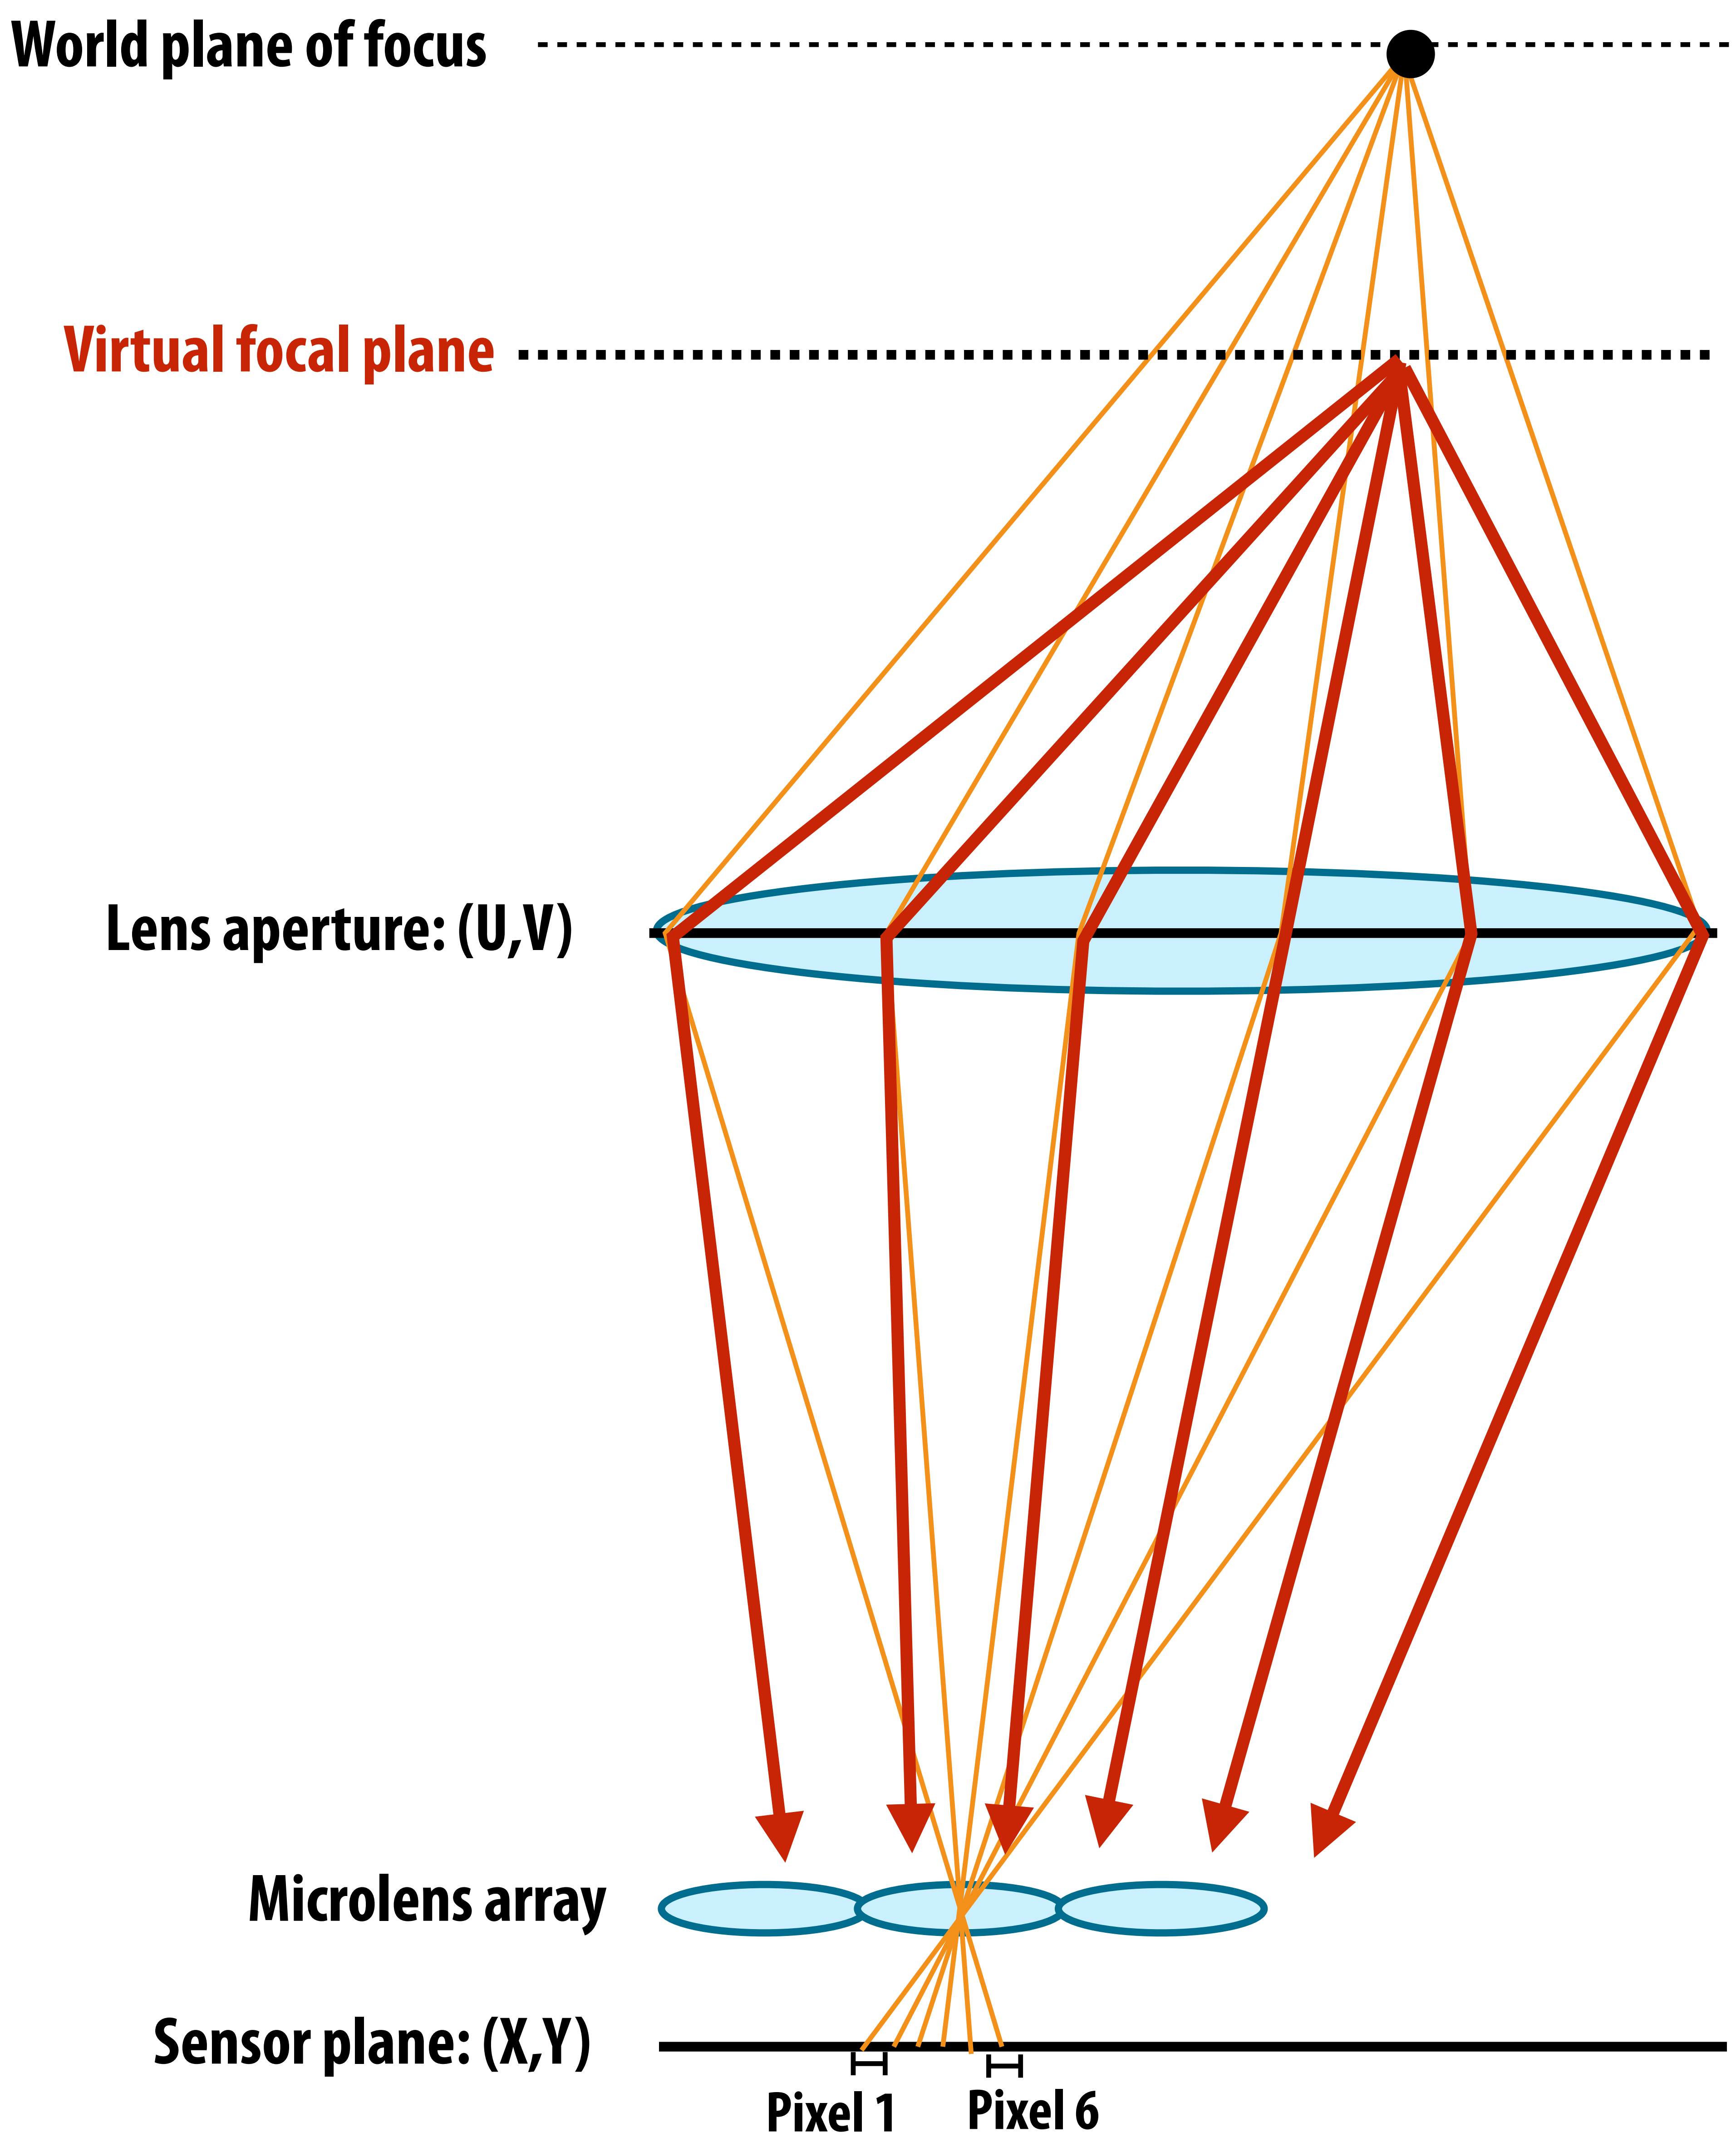

Ray space plot

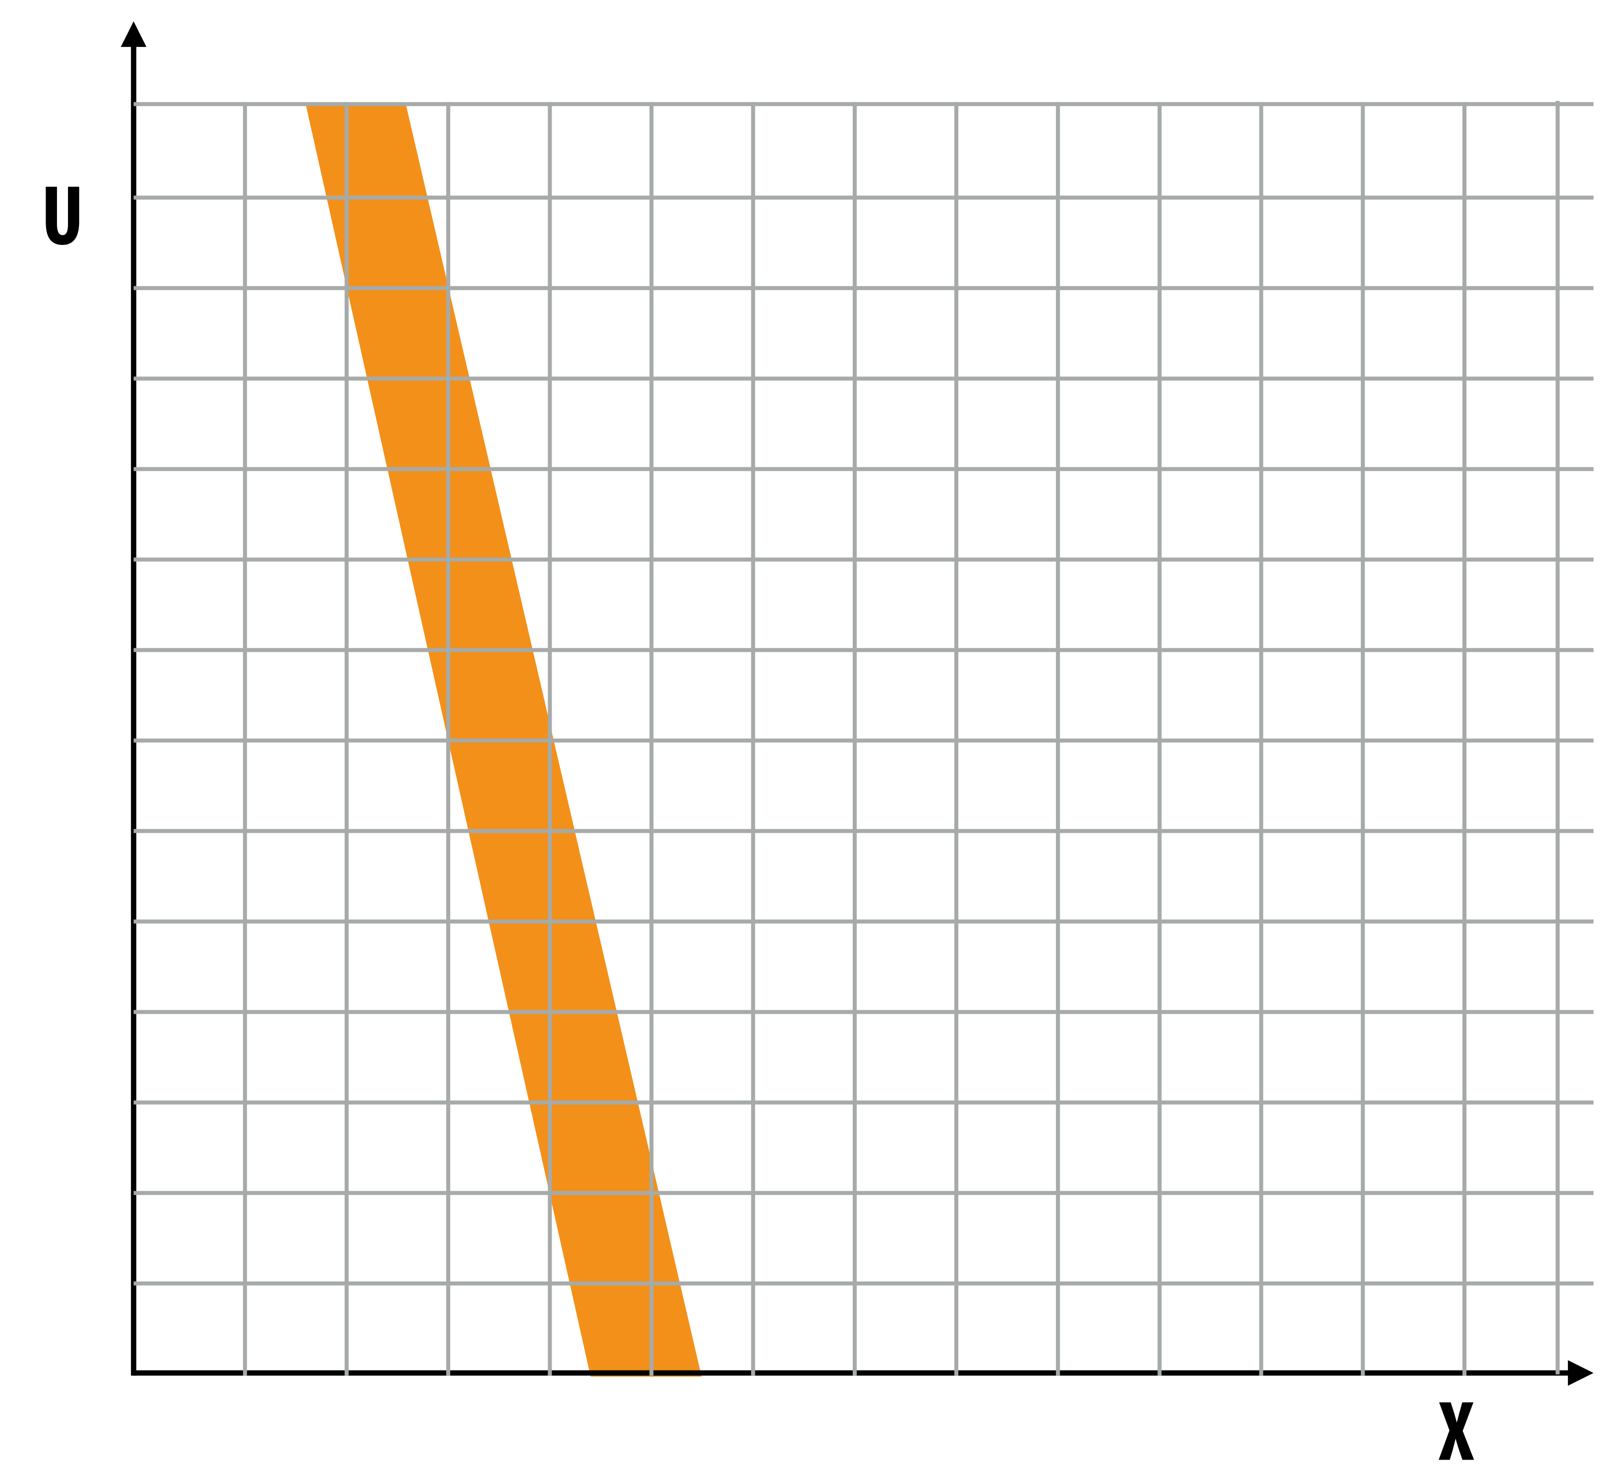

Computing photograph is integral projection (Output image pixel is integral over highlighted region: resample)

Output Image Pixel is Sum of Many Sensor Pixels

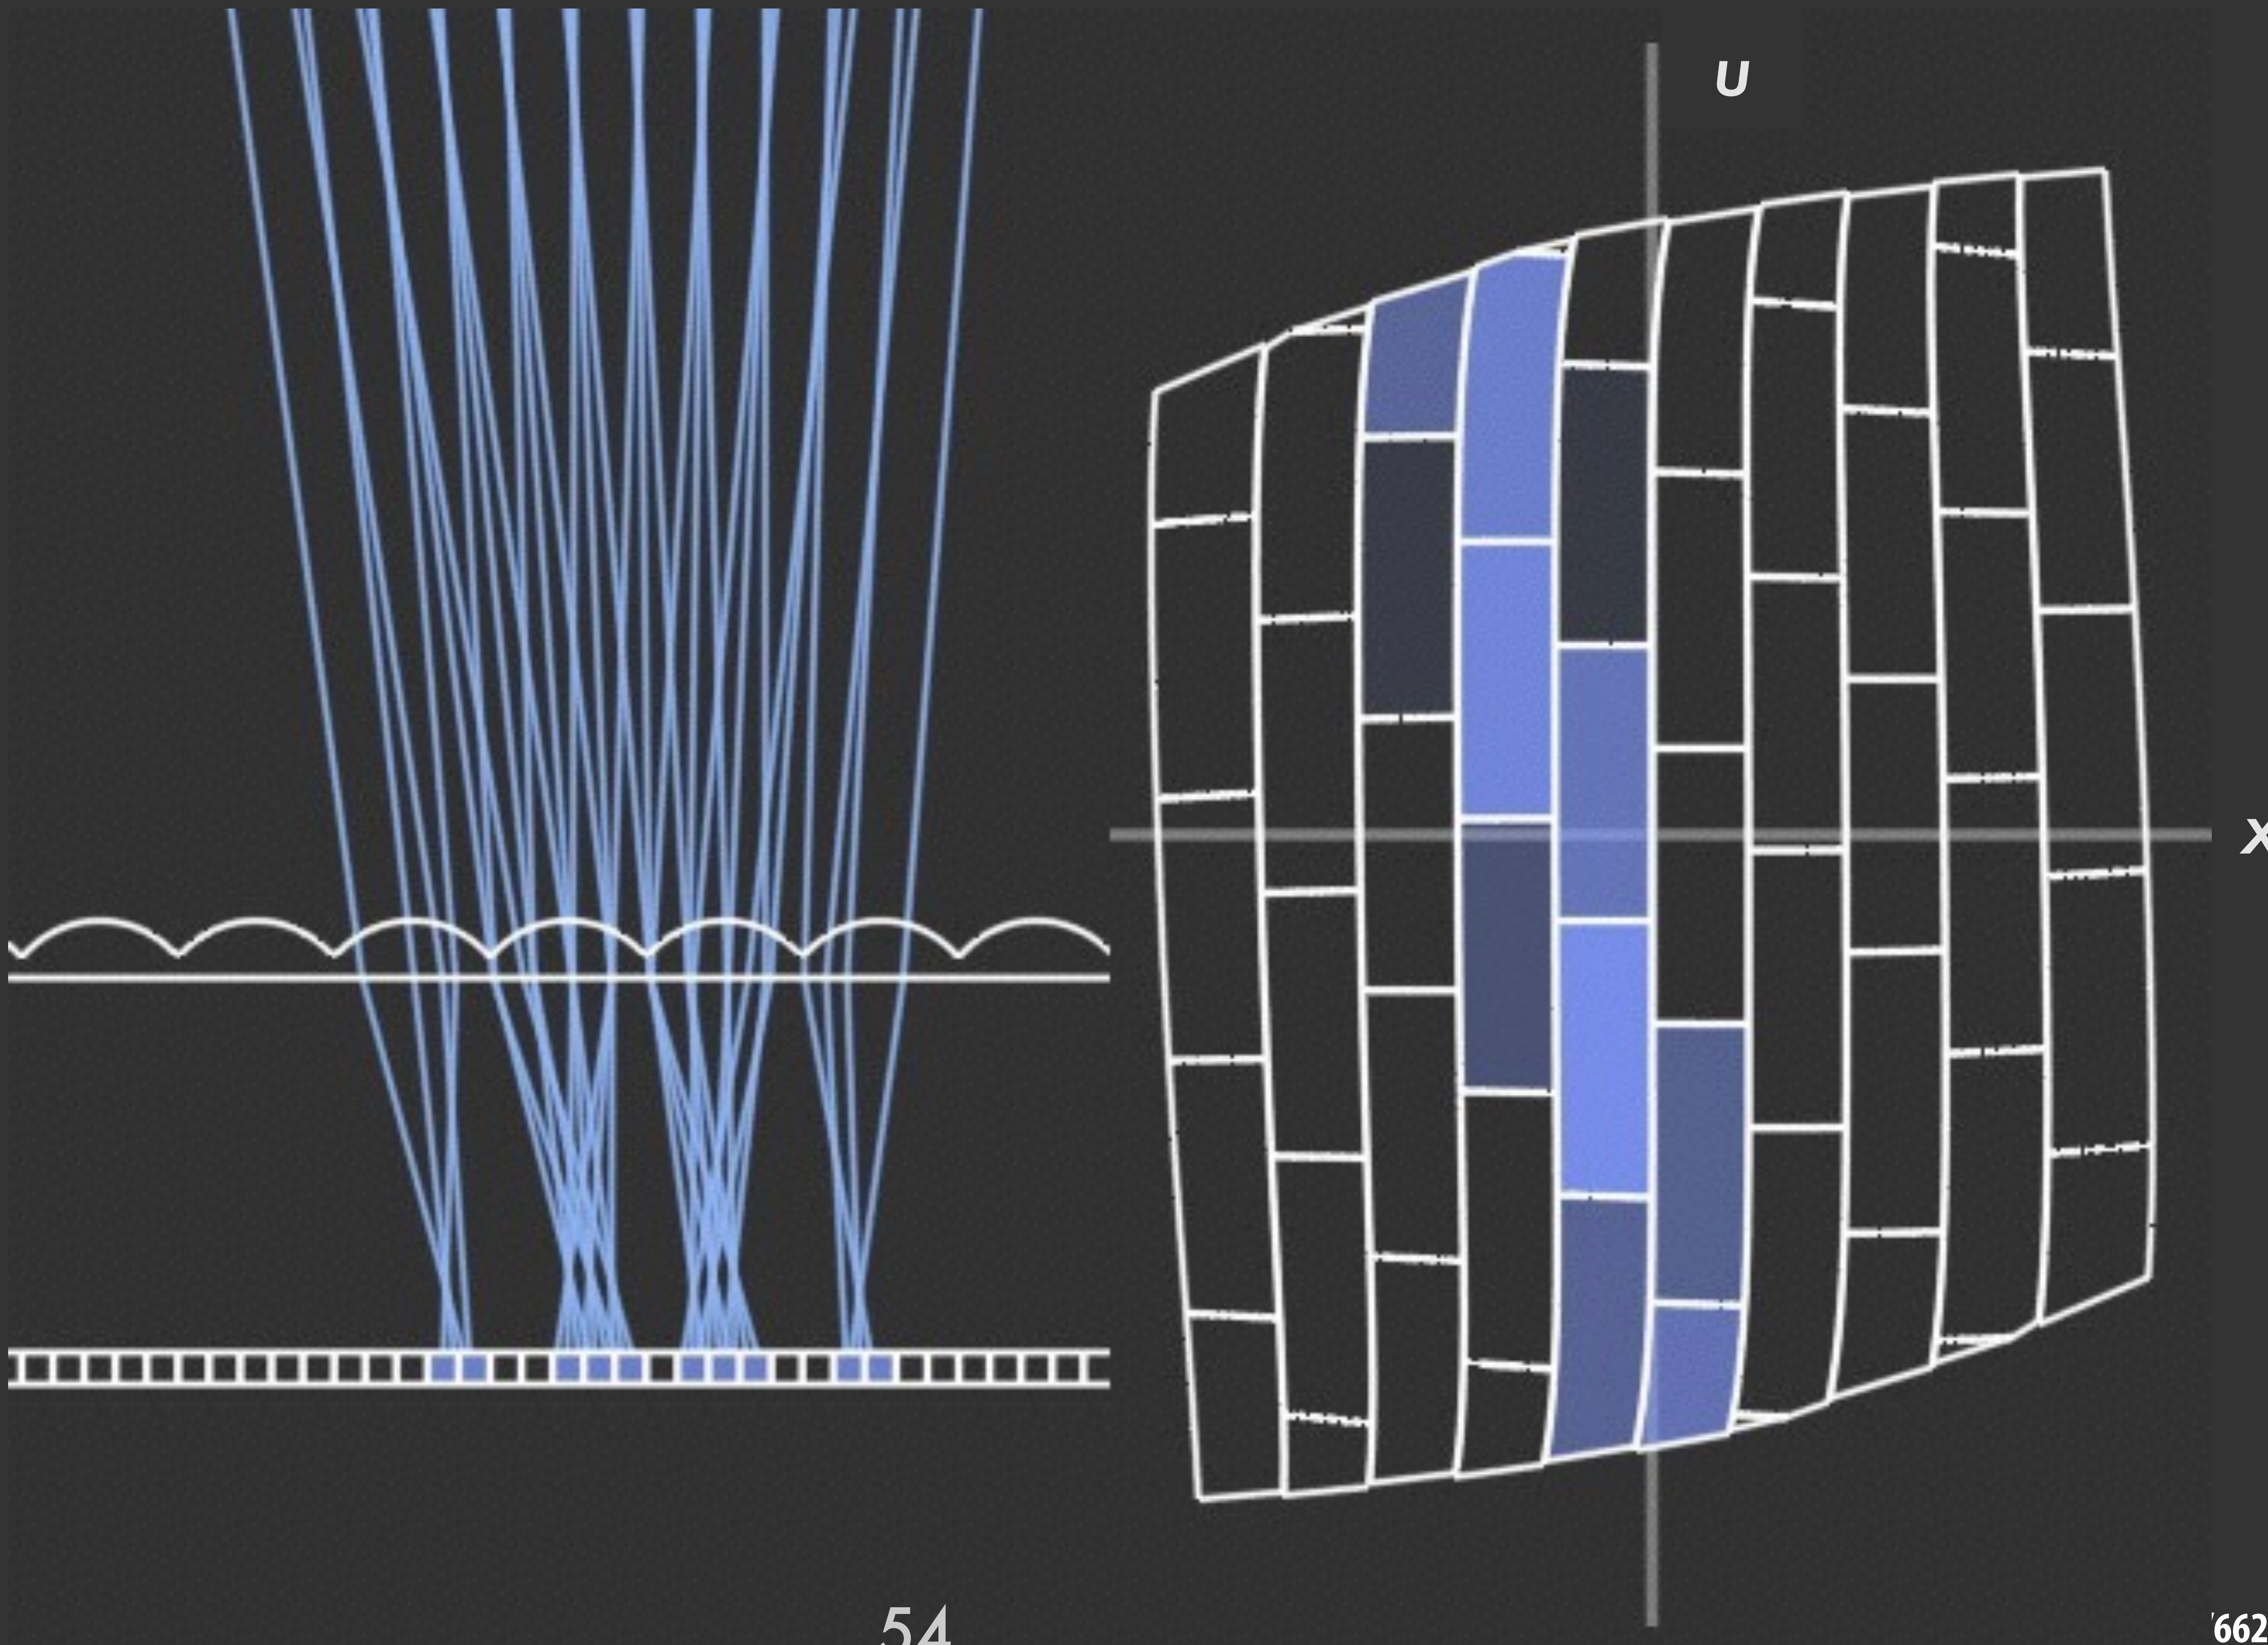

Consumer Light Field Resolutions Today

Lytro (2012)
10 MegaRay
~10 pixels / microlens

Lytro ILLUM (2014)
40 MegaRay
~14 pixels / microlens

Sensor Industry Has Large Untapped Resolution

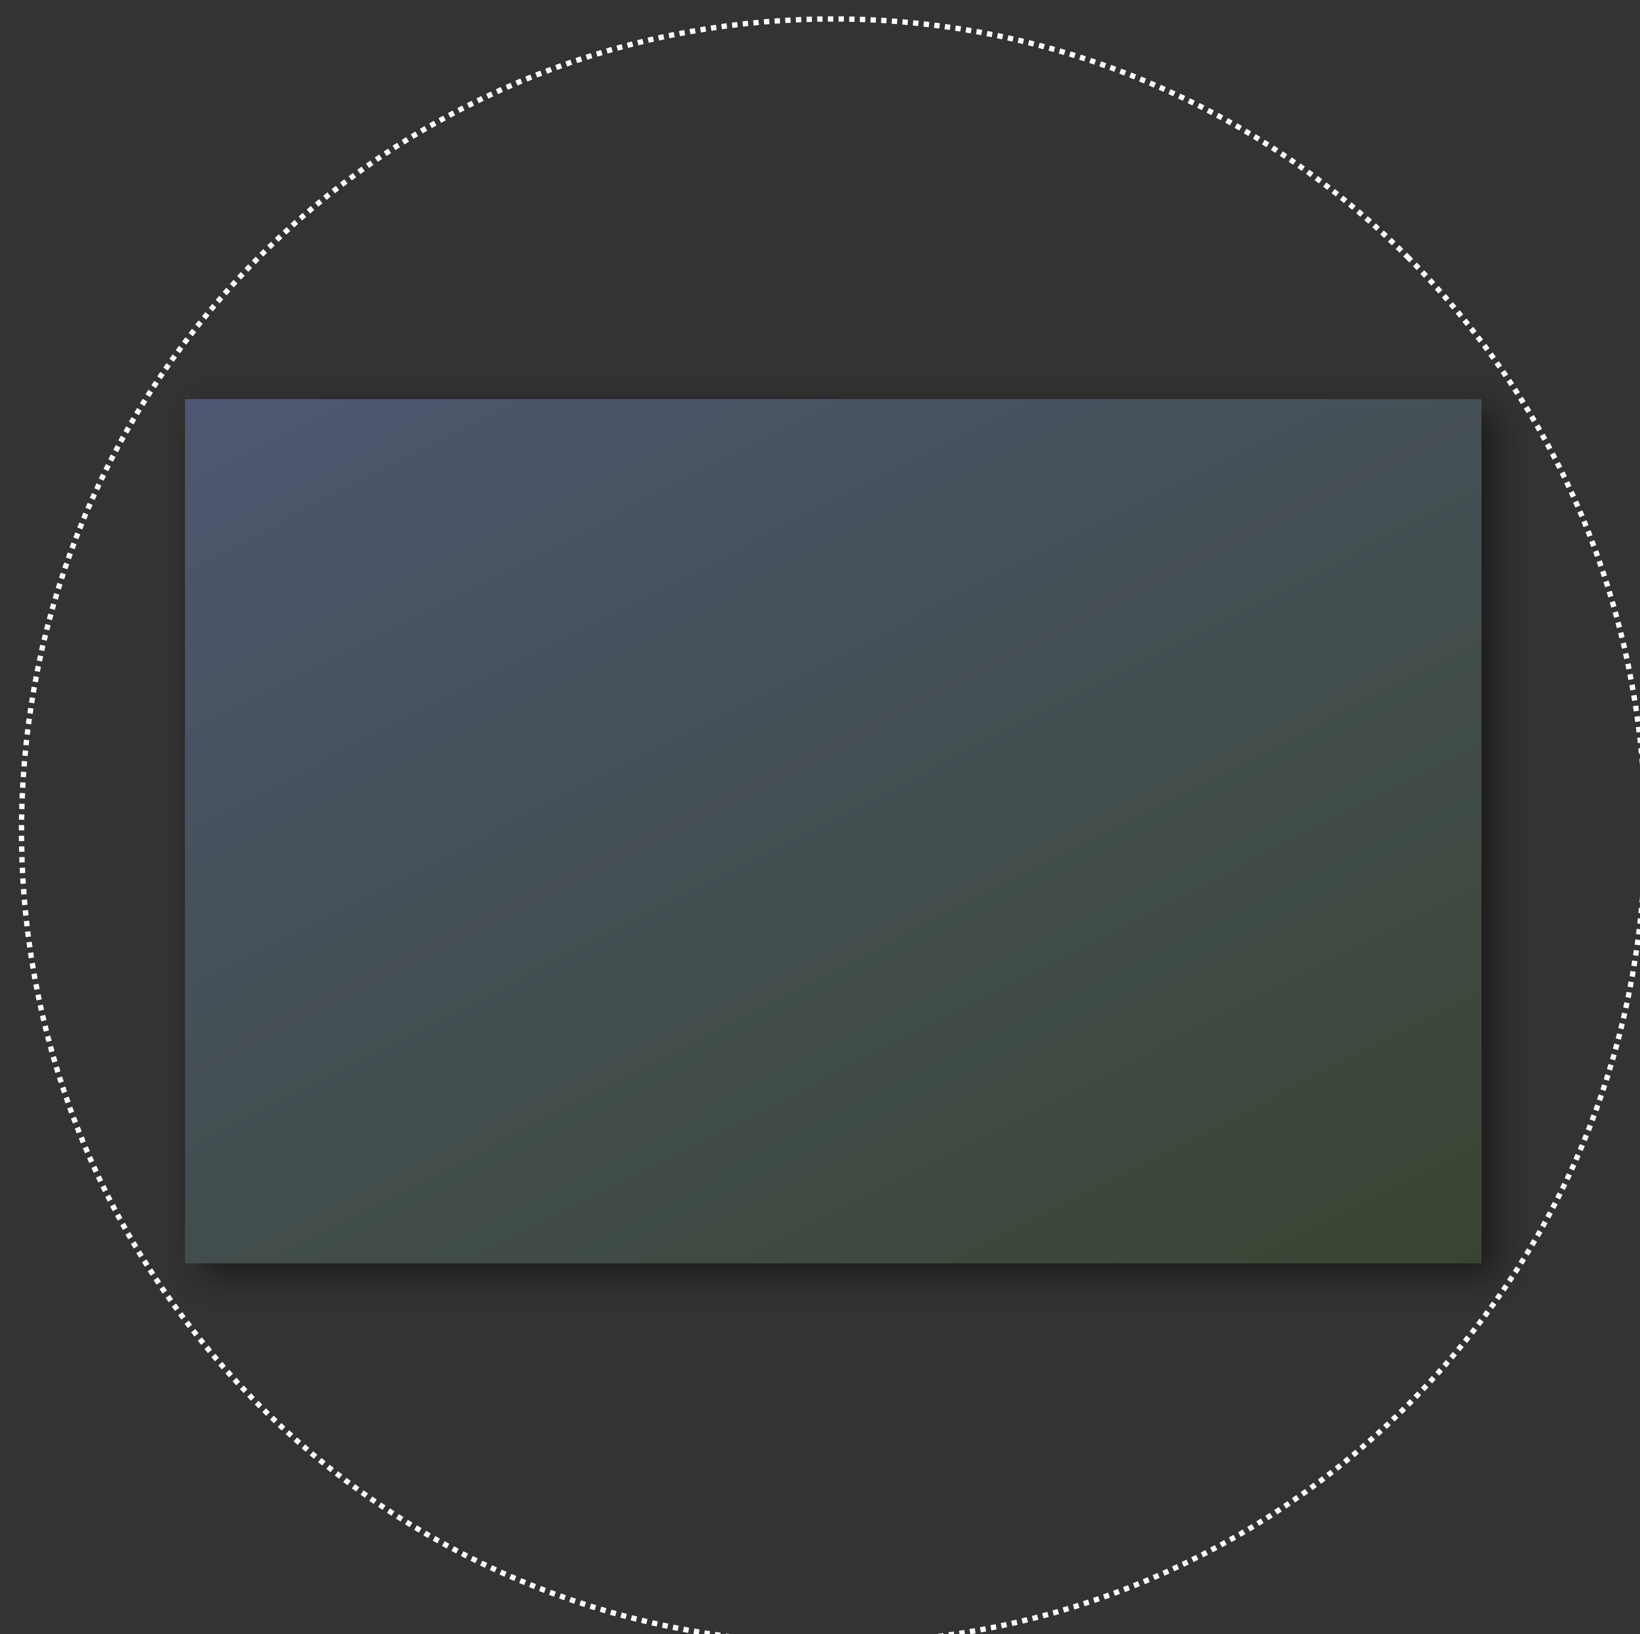

Full-Frame Sensor
36 x 24 mm
Up to 36 MP
4.9 micron pixel

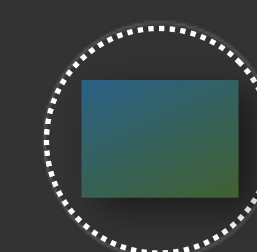

1/3″ Sensor
4.8 x 3.6 mm
Up to 13 MP
1.12 micron pixel

Sensor Industry Has Large Untapped Capability

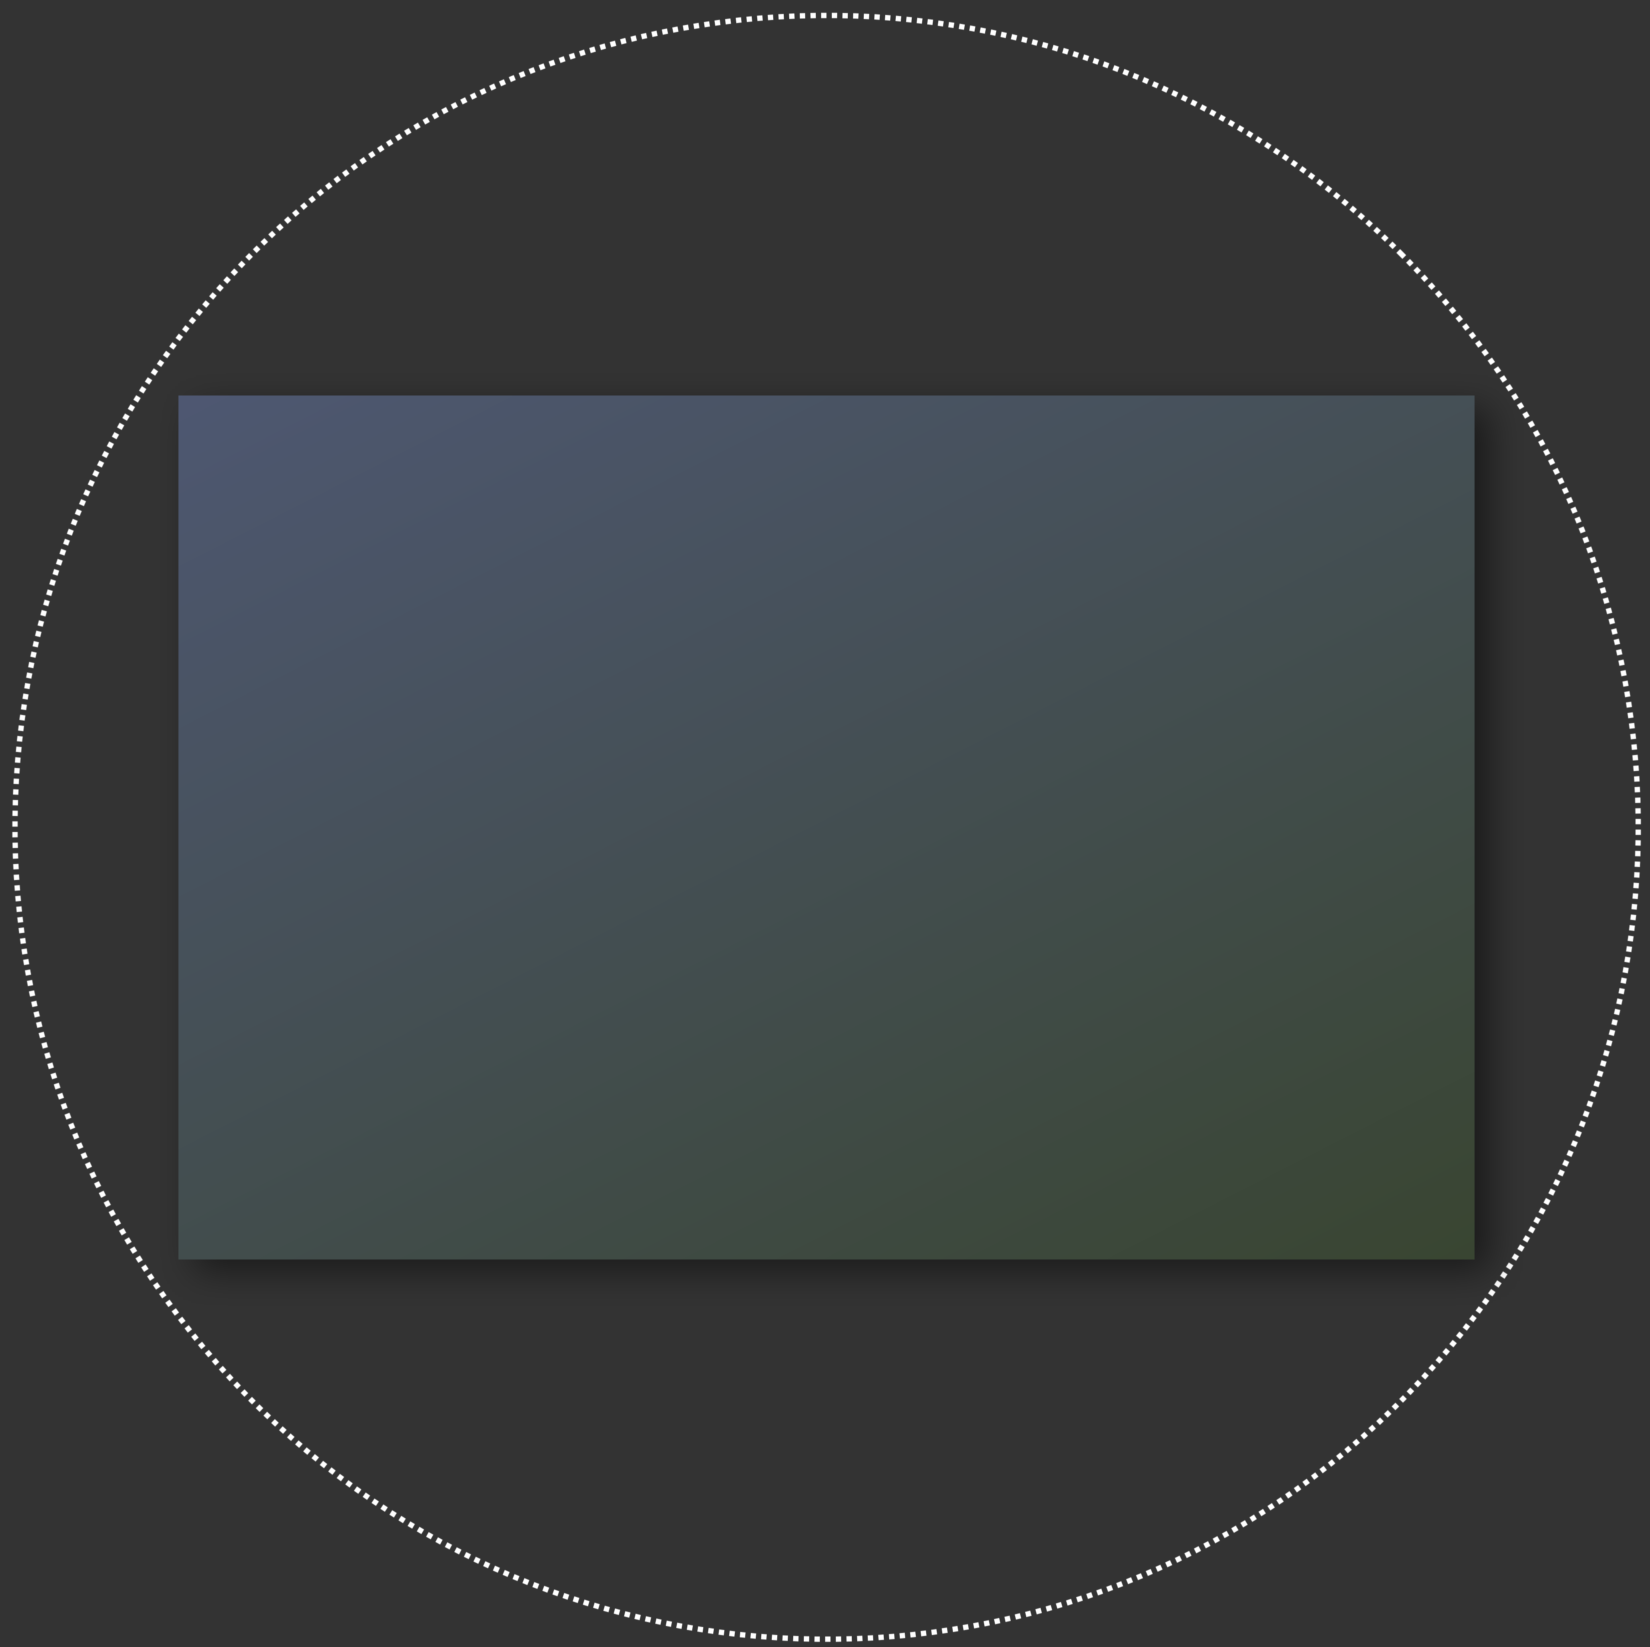

Full-Frame Sensor
36 x 24 mm
Up to 36 MP
4.9 micron pixel

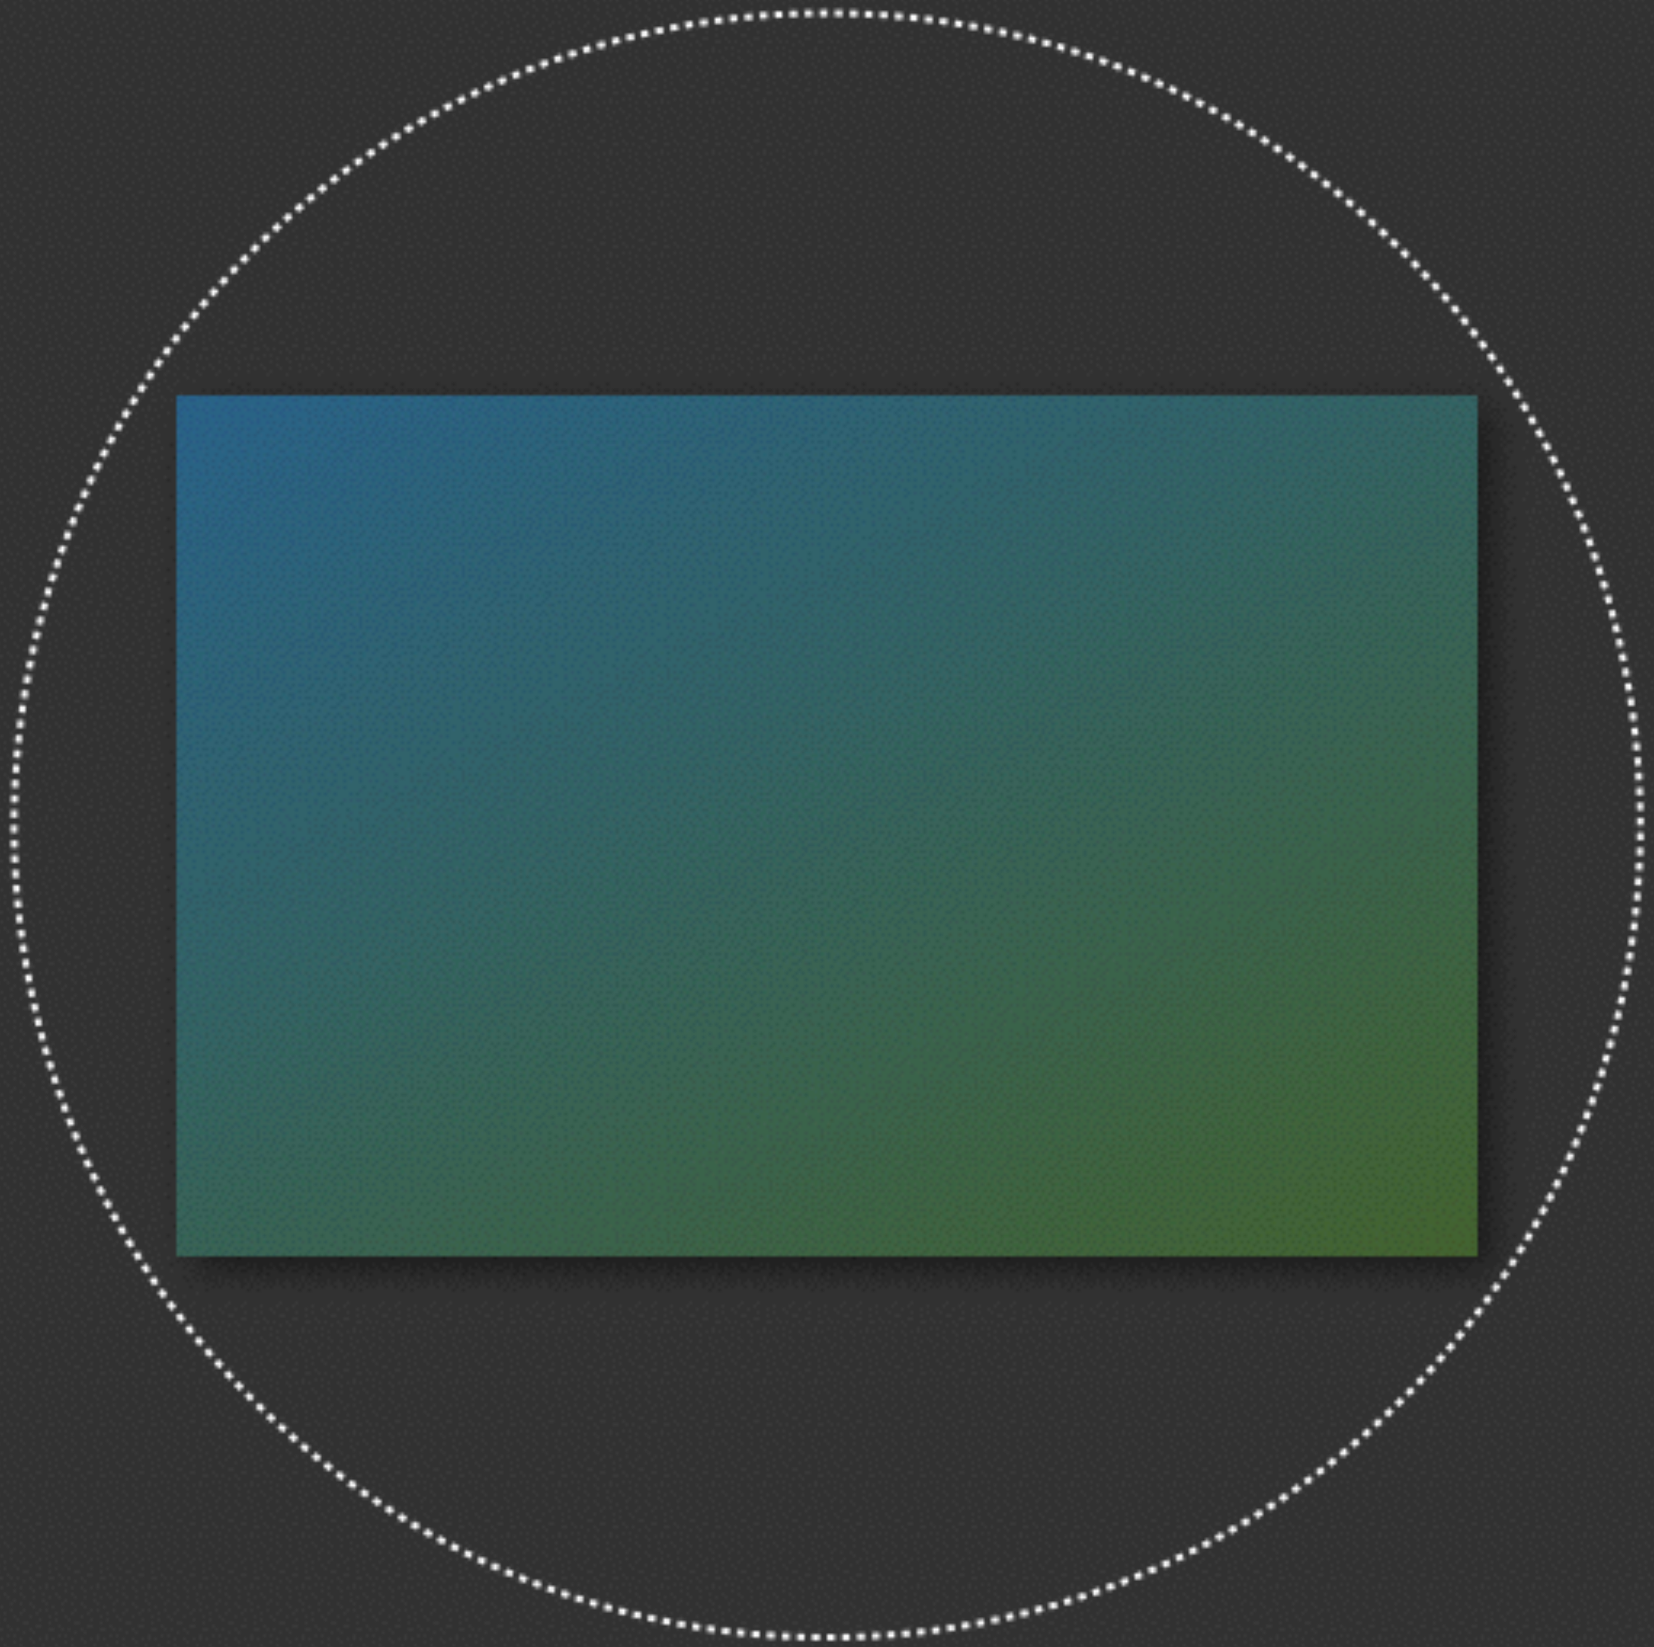

Full-Frame Sensor
36 x 24 mm
688 MP
1.12 micron pixel

New photography applications

- <u>Single shot</u> capture of information that can be used to generate many different photographs
  - Digital (post shot) refocusing
  - Parallax (computational change of viewpoint)
  - Extended depth of field (put entire image in focus)
  - Stereo images

- **Better camera performance:**
  - Reduced shutter lag: in the limit, no need for autofocus
  - Potential for better low-light performance (always shoot with wide aperture, since misfocus due to shallow depth-of-field can be corrected after the shot)
  - Correction of lens aberrations

Trends

- No free lunch: measure directional information at cost of spatial resolution
  - Ng's original prototype: 16 MP sensor, but output was 300x300 images
  - Original Lytro camera: 11MP sensor, ~1MP output images

- Light field cameras can make use of increasing sensor pixel densities
  - More directional resolution = increased refocusing capability
  - More spatial resolution at fixed directional resolution
  - Few motivations high-pixel-count sensors for traditional cameras today

- High-resolution cameras introduce computational challenges
  - Processing challenges
  - Storage challenges
  - Data-transfer challenges

What you should know:

- Give examples of the types of computation that may be done within the camera itself to assemble and improve the final generated image.
- What are sources of noise in capturing / measuring images in a typical digital camera?
- How is color handled in a typical digital camera (e.g., explain the Bayer mosaic)?
- What does it mean to "demosaic" an image?
- Give some examples of how an image may be improved using outputs from multiple cameras (e.g., as in the Light L16 camera).
- Give a 4D lightfield parameterization of rays entering a camera using coordinates on the image plane and on the camera lens.
- How does the microlens array of a lightfield camera allow this 4D space to be captured within a single 2D array? Draw sketches to illustrate your answer.
- What are the advantages to capturing a 4D representation of the light energy in a scene vs. a standard 2D image representation?
- What are the disadvantages?