CVPR Tutorial

Light Fields

Ramesh Raskar

MIT Media Lab

http:// CameraCulture . info/

Integration of Cameras in Mobile Phones

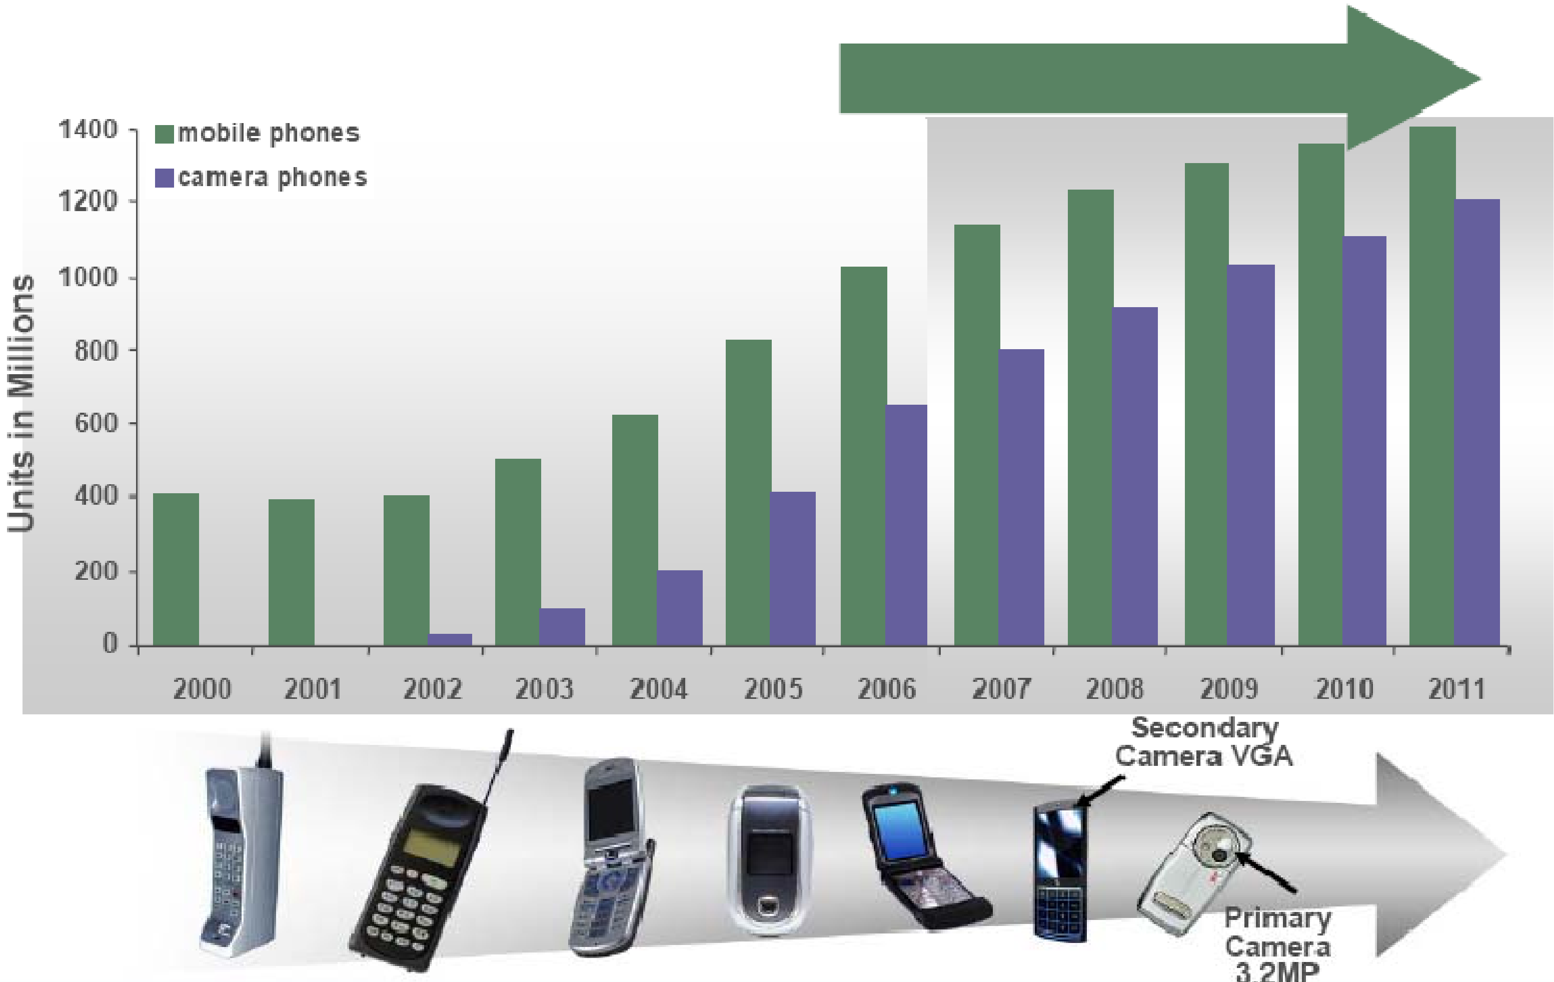

Where are the ‘cameras’?

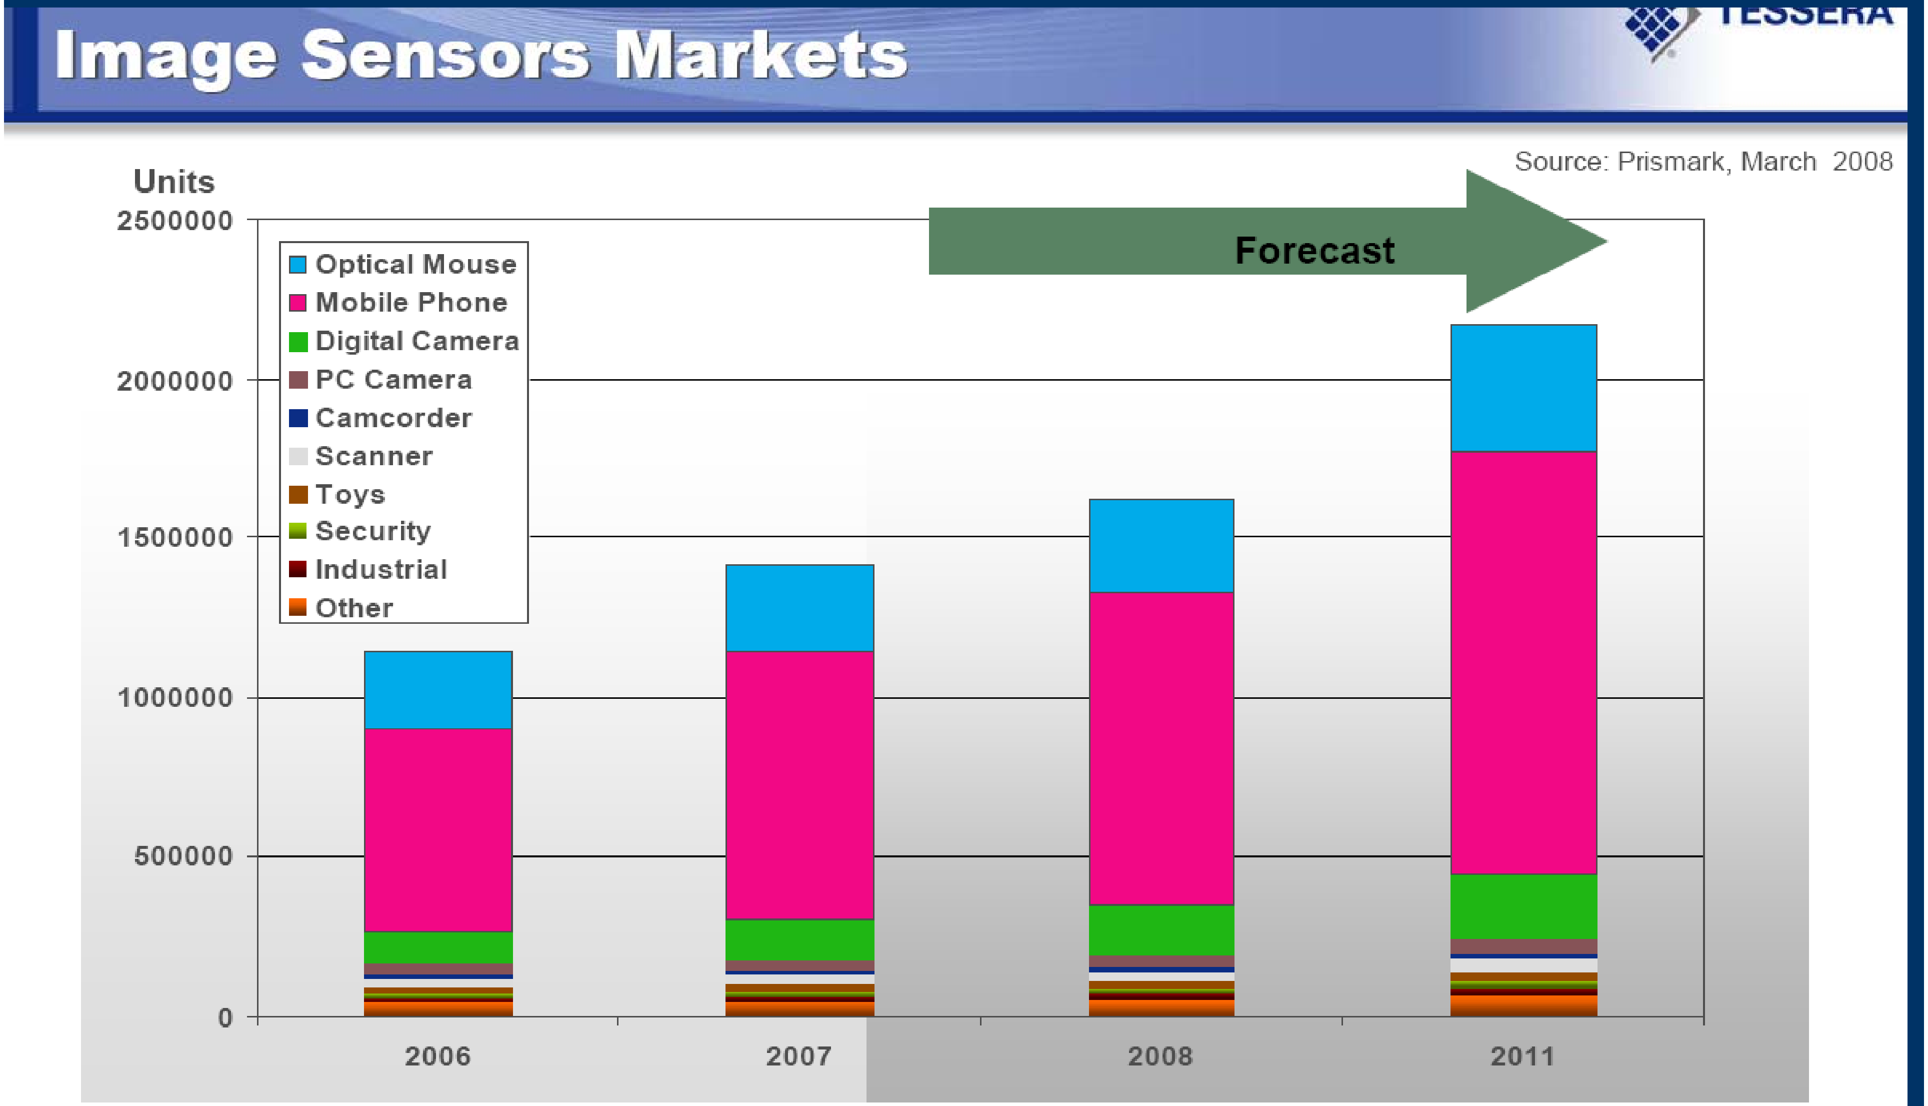

Motivation

- What is the difference between a hologram and a lenticular screen?
- How they capture ‘phase’ of a wavefront for telescope applications?
- What is ‘wavefront coding’ lens for extended depth of field imaging?

Light Fields in Ray and Wave Optics

Introduction to Light Fields:	Ramesh Raskar
Wigner Distribution Function to explain Light Fields:	Zhengyun Zhang
Augmenting LF to explain Wigner Distribution Function:	Se Baek Oh
Q&A	
Break	
Light Fields with Coherent Light:	Anthony Accardi
New Opportunities and Applications:	Raskar and Oh
Q&A:	All

Space of LF representations
Time-frequency representations
Phase space representations
Quasi light field

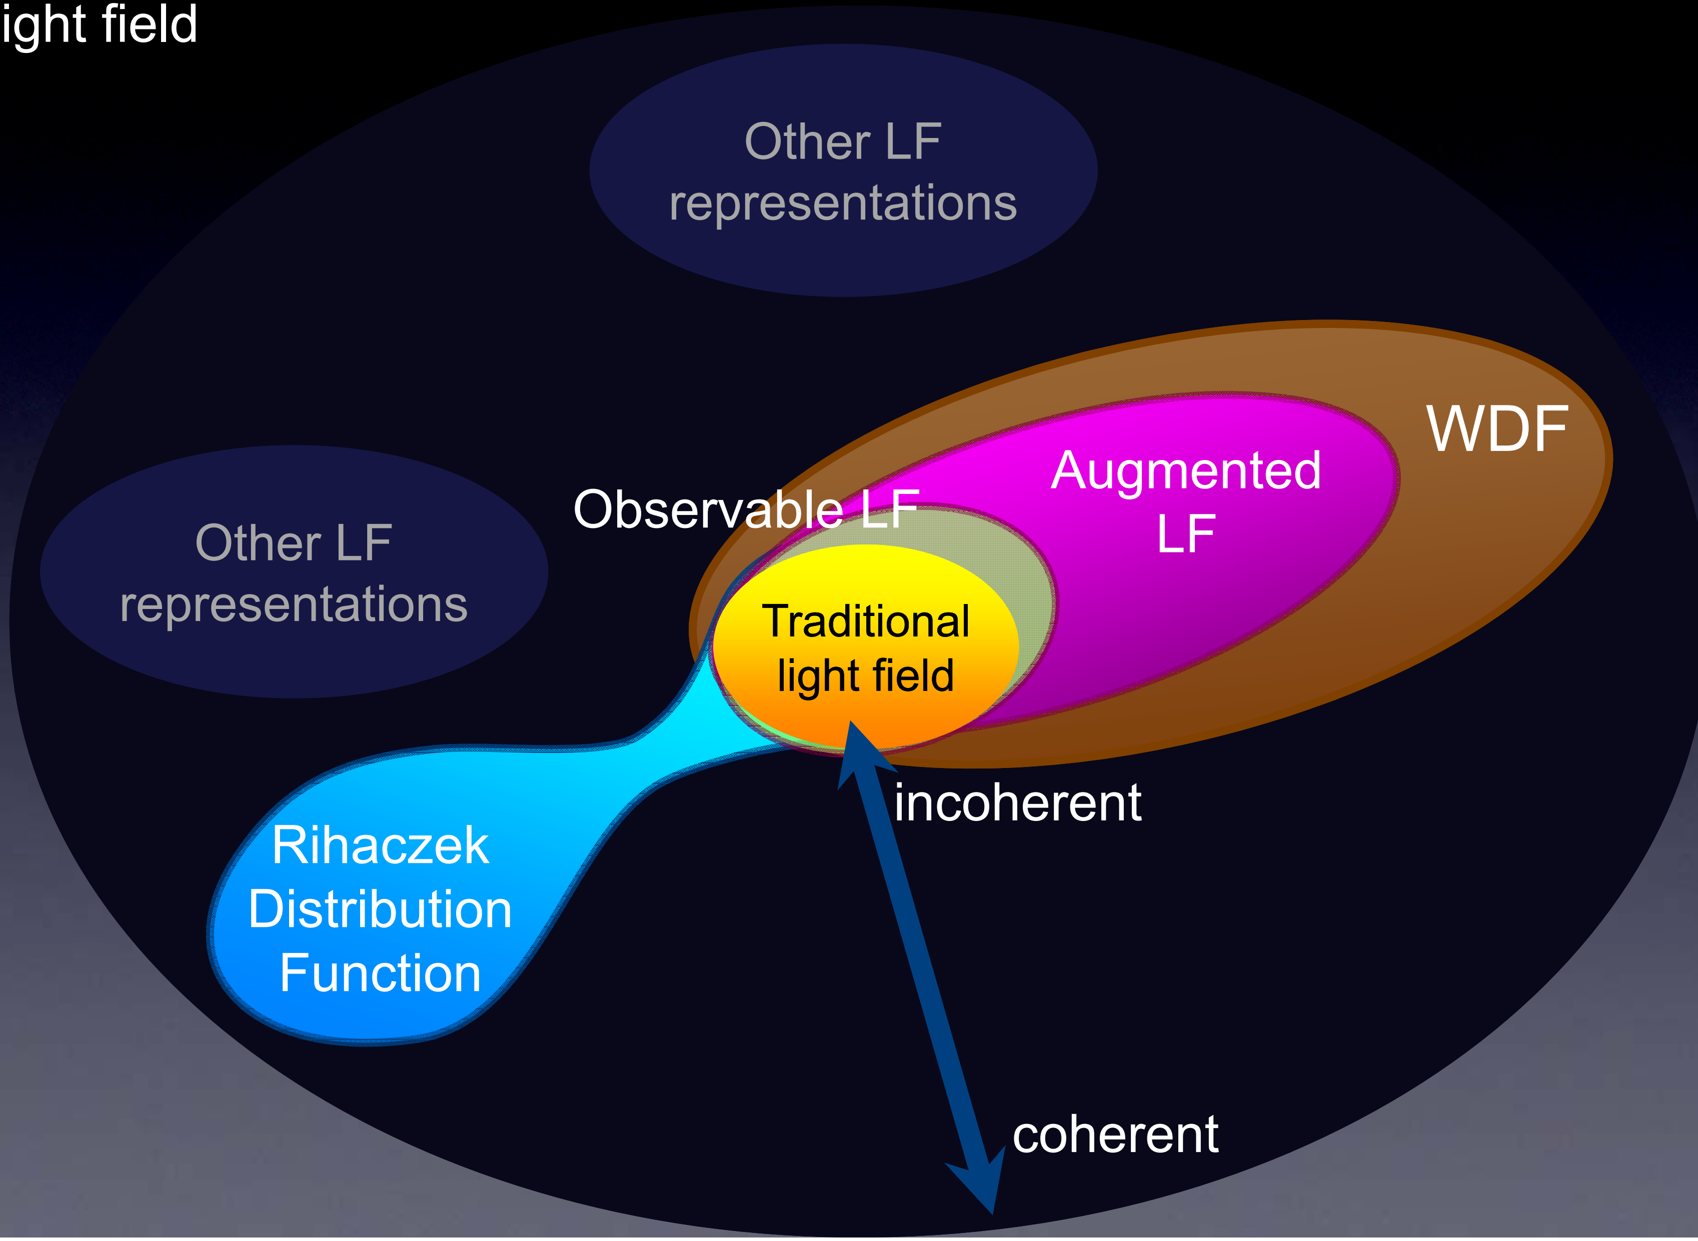

Property of the Representation

	Constant along rays	Non-negativity	Coherence	Wavelength	Interference Cross term
Traditional LF	always constant	always positive	only incoherent	zero	no
Observable LF	nearly constant	always positive	any coherence state	any	yes
Augmented LF	only in the paraxial region	positive and negative	any	any	yes
WDF	only in the paraxial region	positive and negative	any	any	yes
Rihaczek DF	no; linear drift	complex	any	any	reduced

Benefits & Limitations of the Representation

	Ability to propagate	Modeling wave optics	Simplicity of computatio n	Adaptability to current pipe line	Near Field	Far Field
Traditional LF	x-shear	no	very simple	high	no	yes
Observable LF	not x-shear	yes	modest	low	yes	yes
Augmented LF	x-shear	yes	modest	high	no	yes
WDF	x-shear	yes	modest	low	yes	yes
Rihaczek DF	x-shear	yes	better than WDF, not as simple as LF	low	no	yes

Light Fields

Goal: Representing propagation, interaction and image formation of light using purely position and angle parameters

- Radiance per ray
- Ray parameterization:
  - Position : s, x, r
  - Direction : u, θ, s

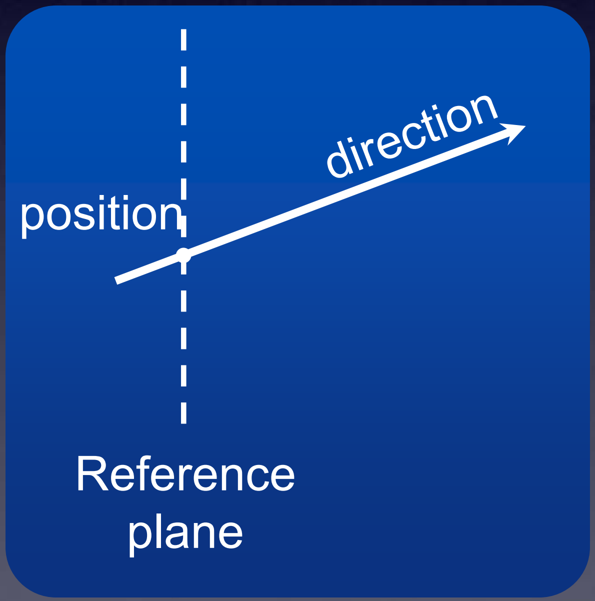

Limitations of Traditional Lightfields

rigorous but cumbersome
wave optics based

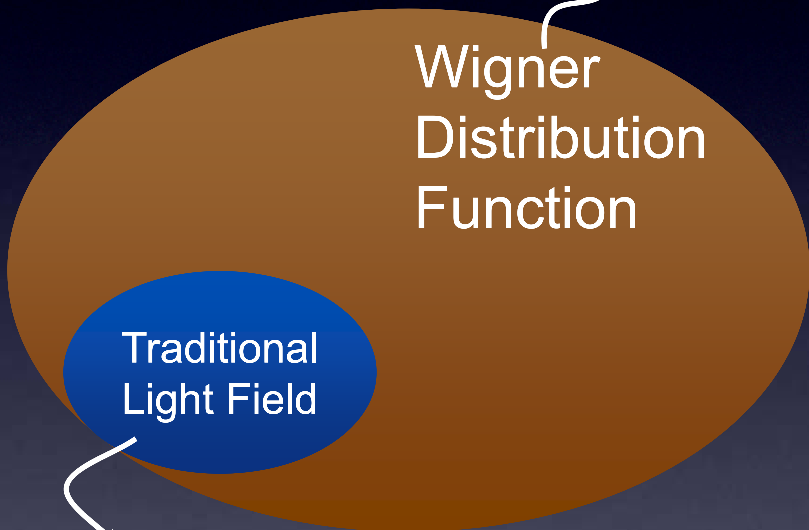

ray optics based
simple and powerful
limited in diffraction & interference

holograms

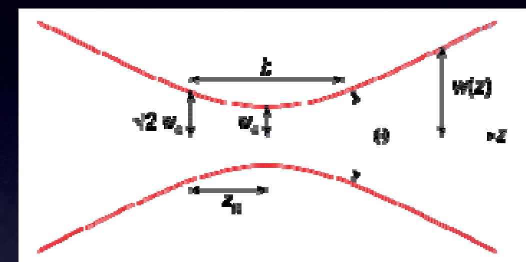

beam shaping

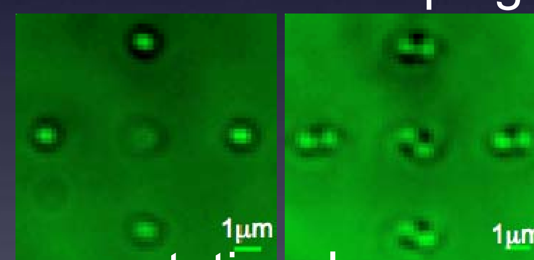

rotational PSF

Example: New Representations Augmented Lightfields

rigorous but cumbersome
wave optics based

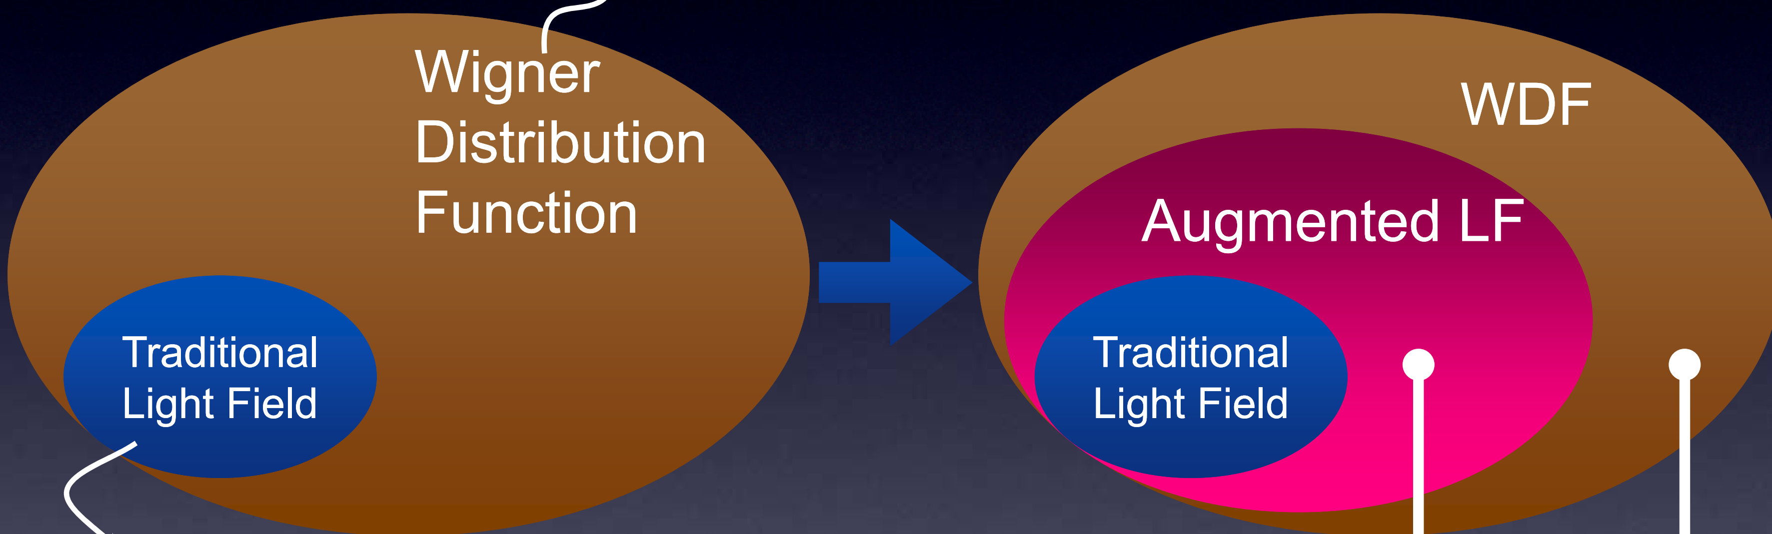

ray optics based
simple and powerful
limited in diffraction & interference

Interference & Diffraction
Interaction w/ optical elements

Non-paraxial
propagation

Introduction to Light Fields

Ramesh Raskar

MIT Media Lab

http:// CameraCulture . info/

Introduction to Light Fields

- Ray Concepts for 4D and 5D Functions
- Propagation of Light Fields
- Interaction with Occluders
- Fourier Domain Analysis and Relationship to Fourier Optics
- Coded Photography: Modern Methods to Capture Light Field
- Wigner and Ambiguity Function for Light Field in Wave Optics
- New Results in Augmenting Light Fields

The Plenoptic Function

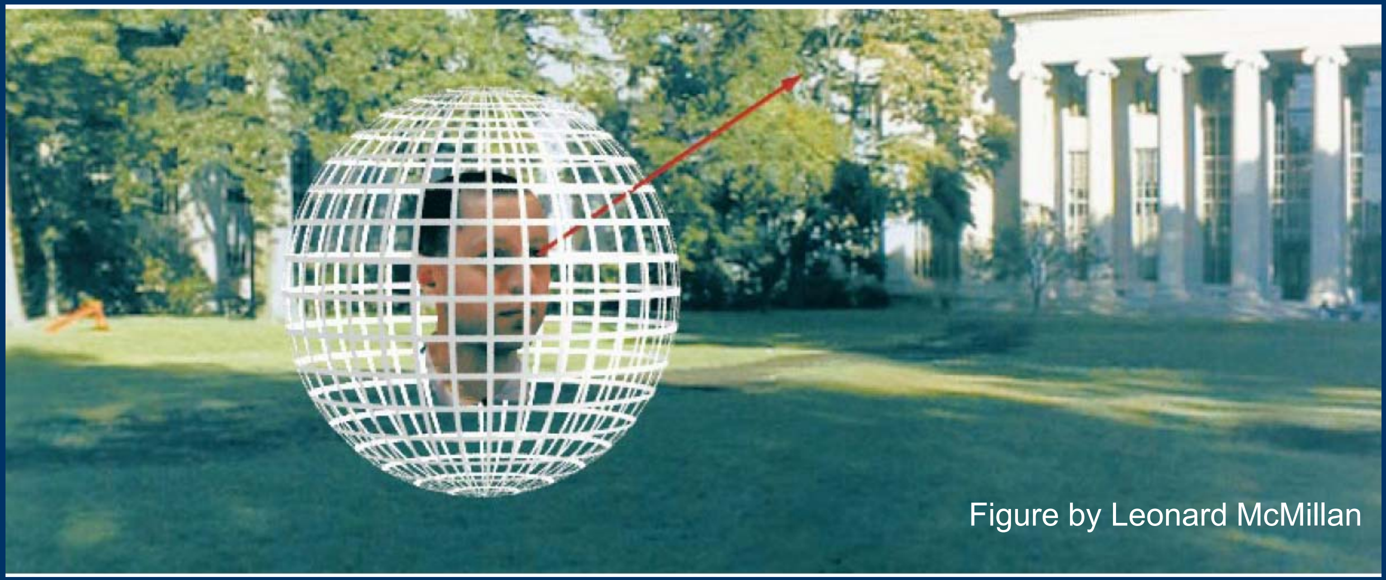

Figure by Leonard McMillan

- Q: What is the set of all things that we can ever see?
- A: The Plenoptic Function (Adelson & Bergen)

- Let's start with a stationary person and try to parameterize everything that he can see…

Grayscale snapshot

$$P(\theta,\phi)$$

- is intensity of light
  - Seen from a single view point
  - At a single time
  - Averaged over the wavelengths of the visible spectrum
- (can also do *P(x,y)*, but spherical coordinate are nicer)

Color snapshot

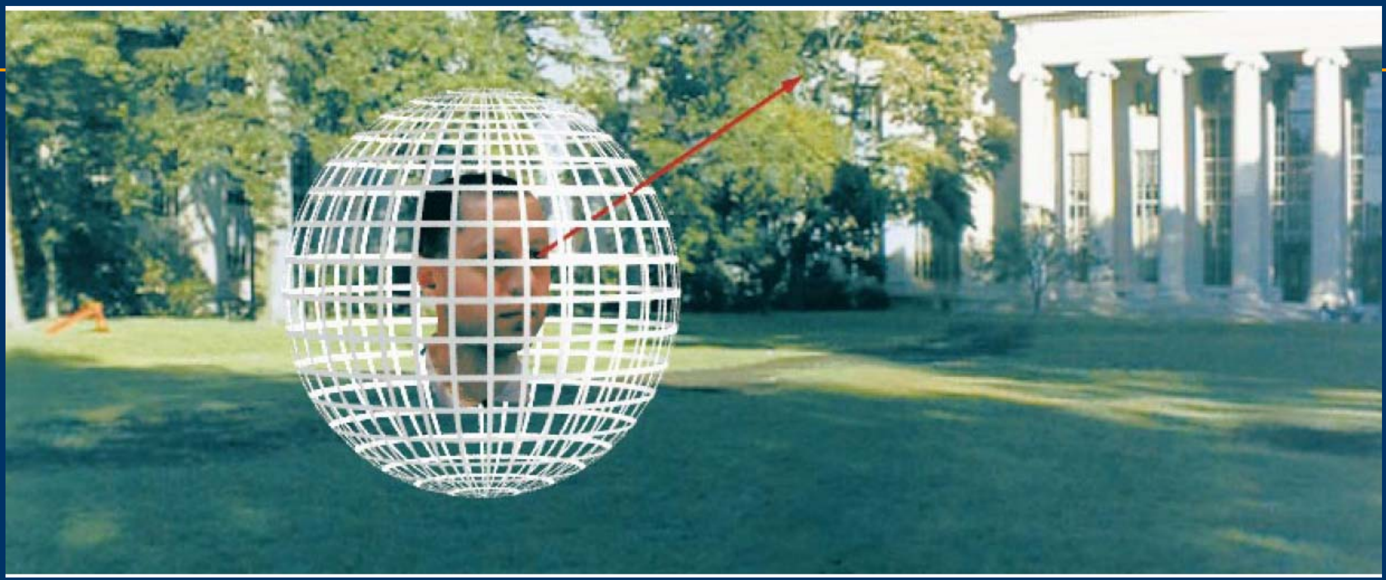

$$P(\theta,\phi,\lambda)$$

- is intensity of light
  - Seen from a single view point
  - At a single time
  - As a function of wavelength

A movie

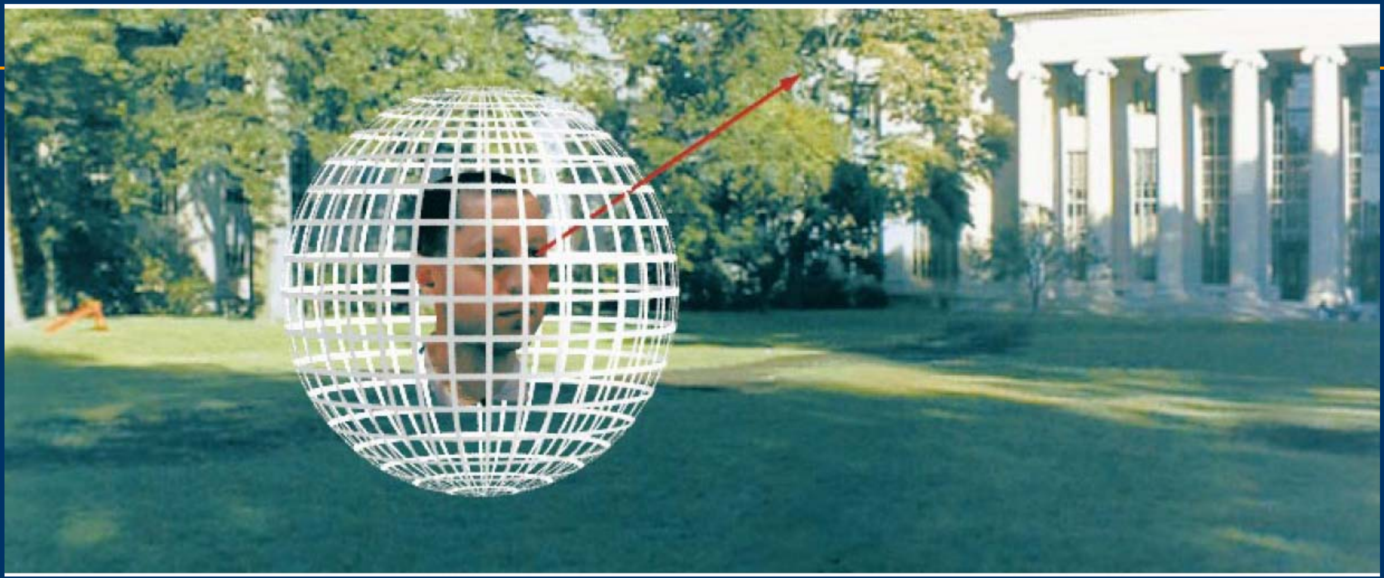

$$P(\theta,\phi,\lambda,t)$$

- is intensity of light
  - Seen from a single view point
  - Over time
  - As a function of wavelength

Holographic movie

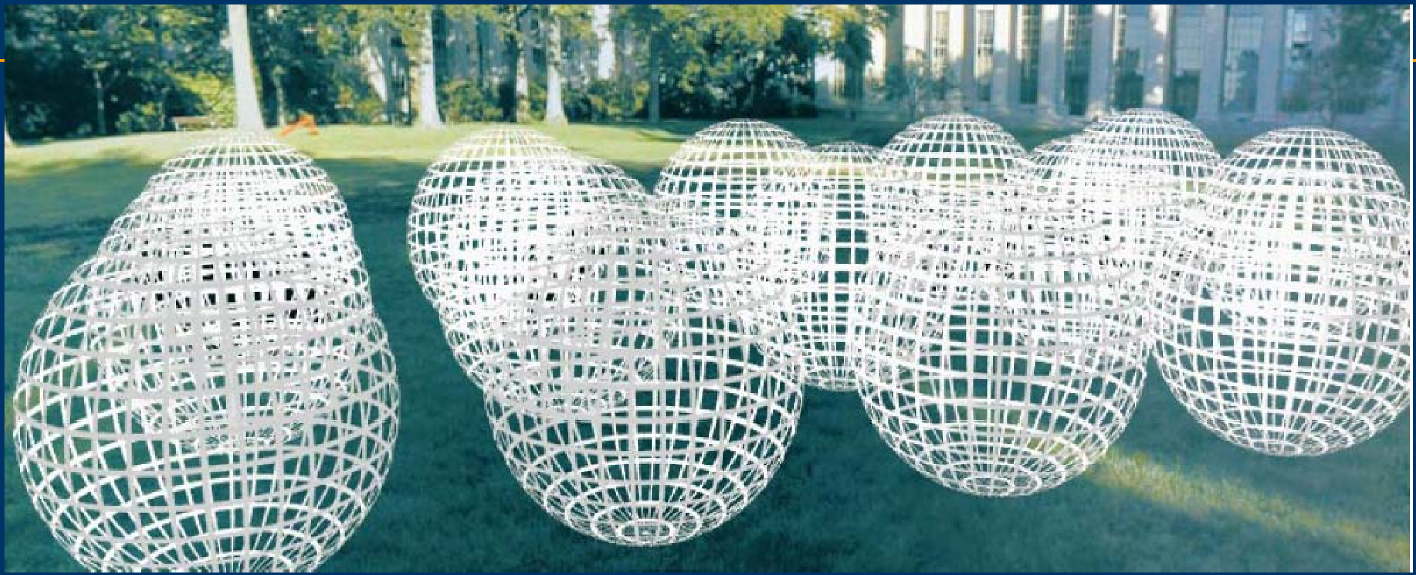

$$P(\theta,\phi,\lambda,t,V_X,V_Y,V_Z)$$

- is intensity of light
  - Seen from ANY viewpoint
  - Over time
  - As a function of wavelength

The Plenoptic Function

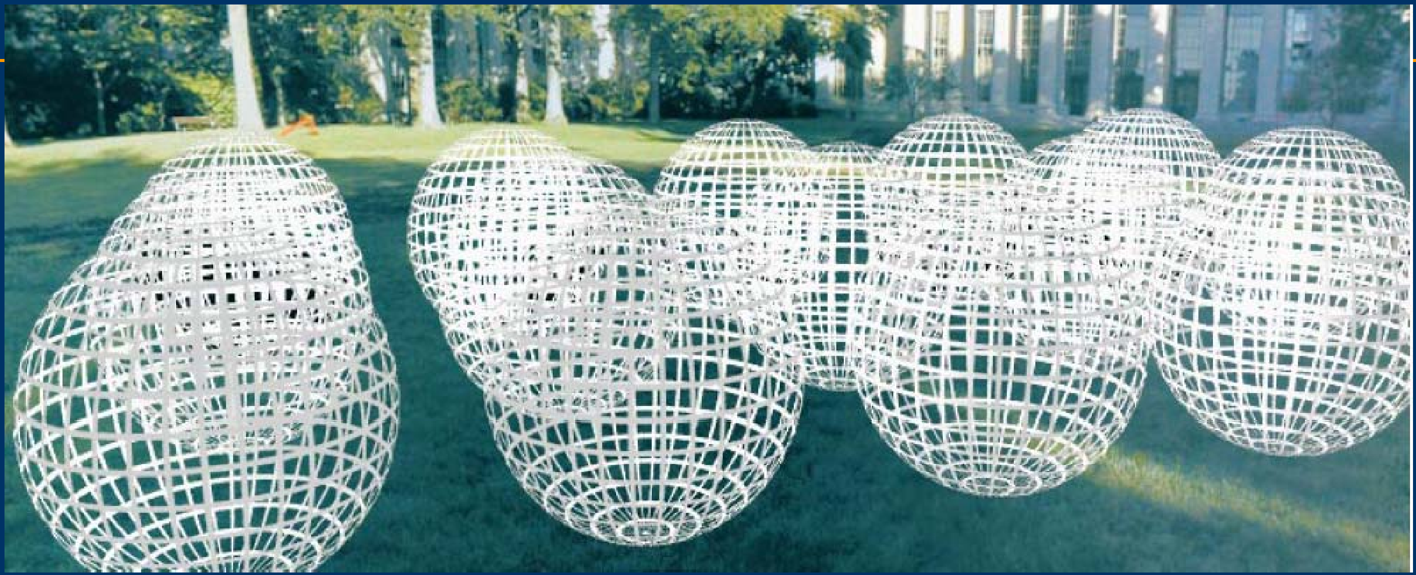

$$P(\theta,\phi,\lambda,t,V_X,V_Y,V_Z)$$

– Can reconstruct every possible view, at every moment, from every position, at every wavelength

– Contains every photograph, every movie, everything that anyone has ever seen.

Sampling Plenoptic Function (top view)

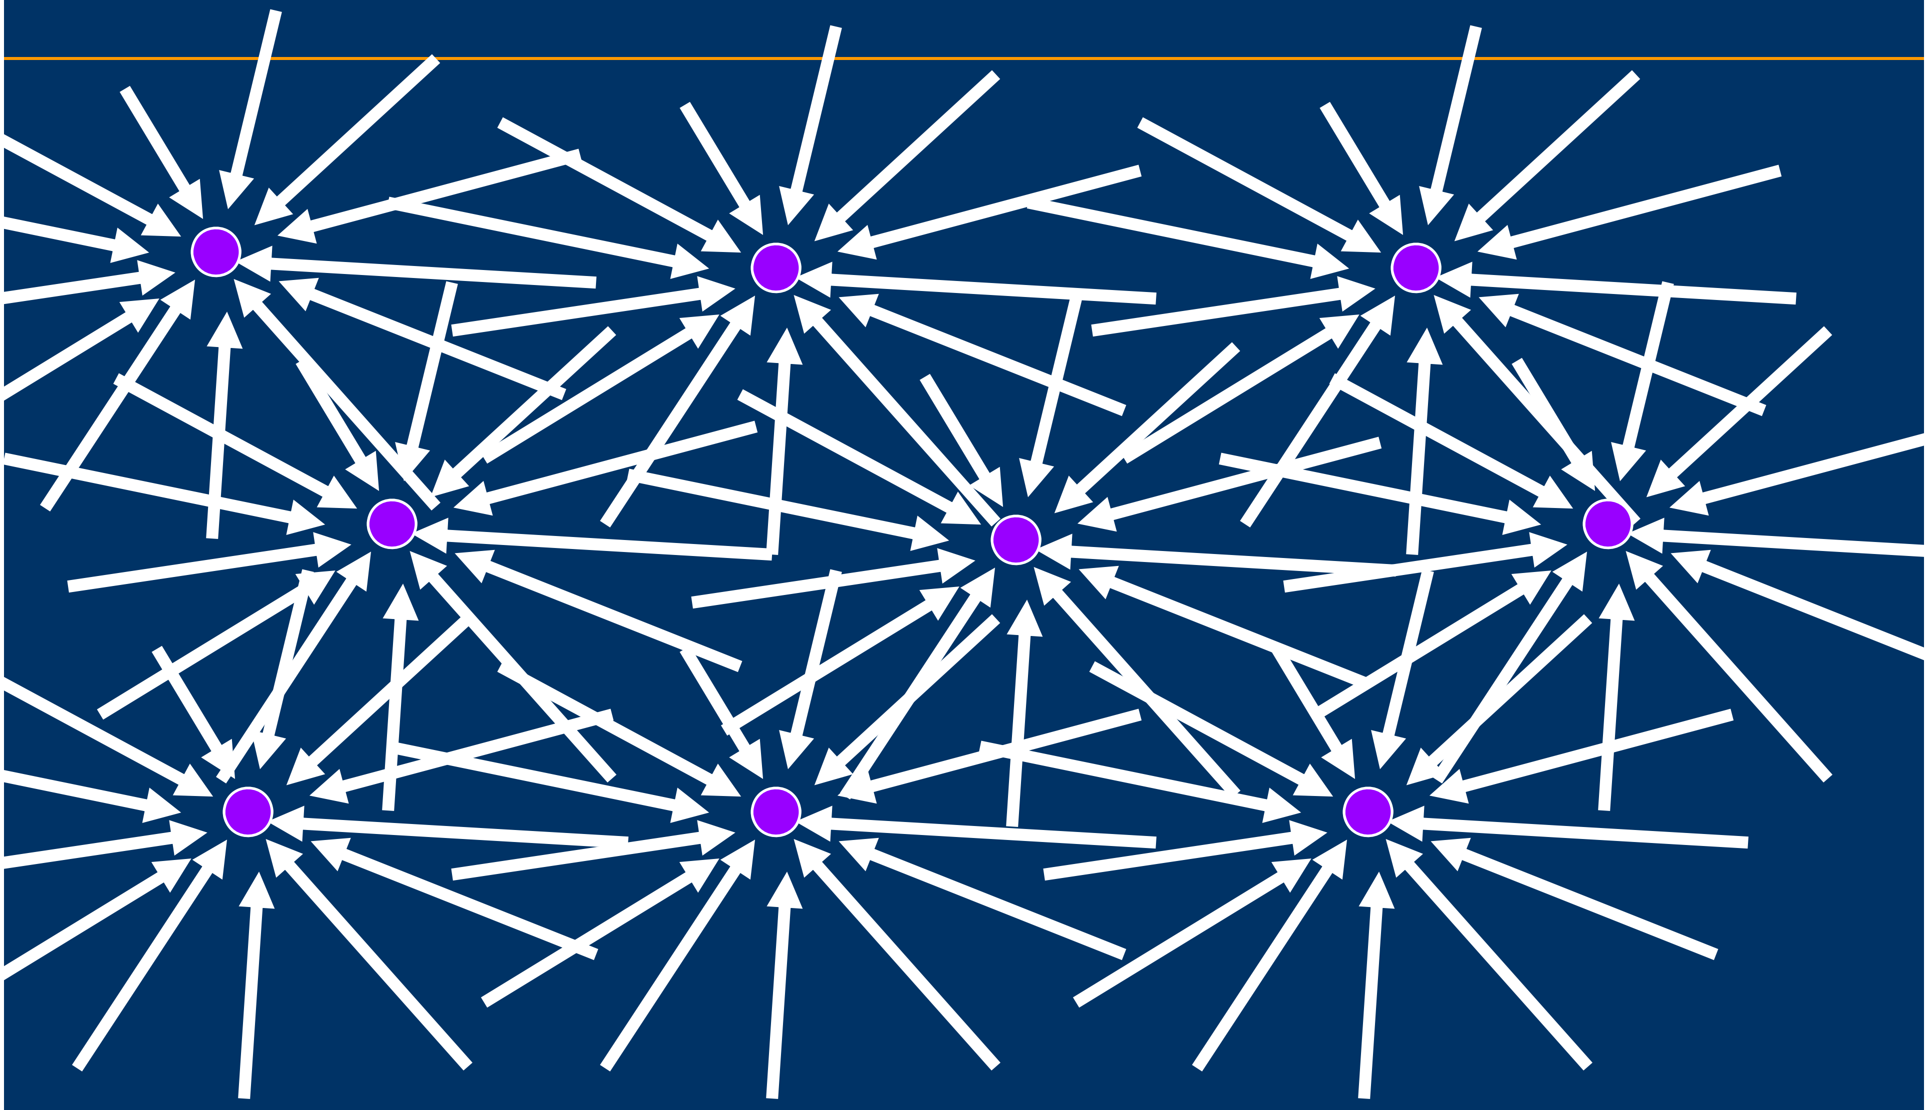

Ray

- Let's not worry about time and color:

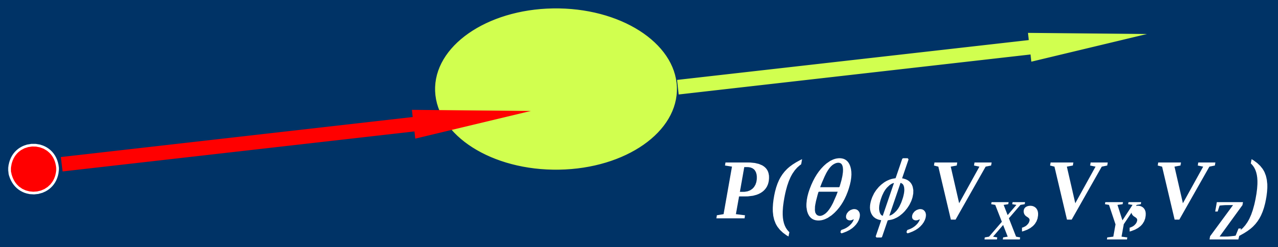

- 5D
  - 3D position
  - 2D direction

Ray

- No Occluding Objects

$P(\theta,\phi,V_X,V_Y,V_Z)$

- 4D
  - 2D position
  - 2D direction

- The space of all lines in 3-D space is 4D.

Lumigraph/Lightfield - Organization

- 2D position
- 2D direction

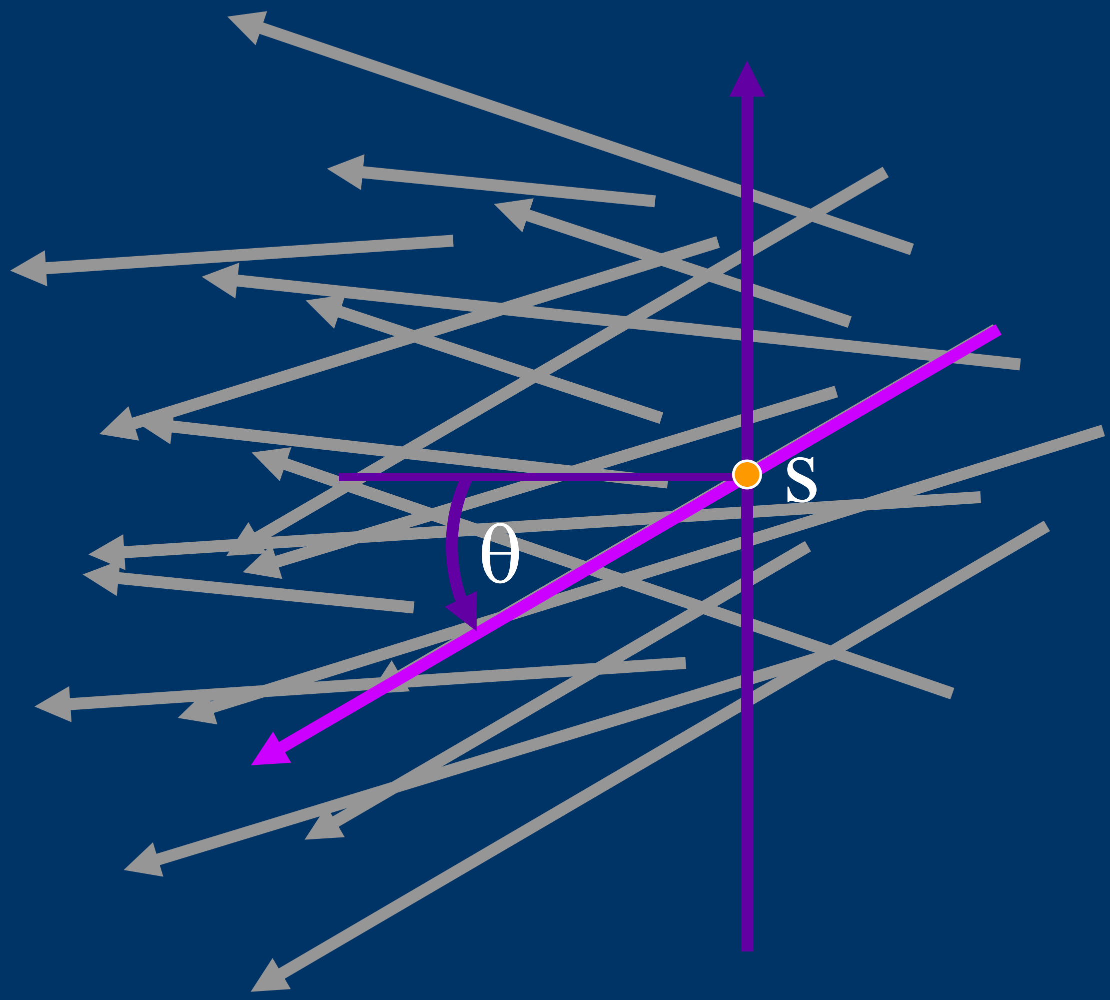

Slide by Rick Szeliski and Michael Cohen

- 2D position
- 2D position

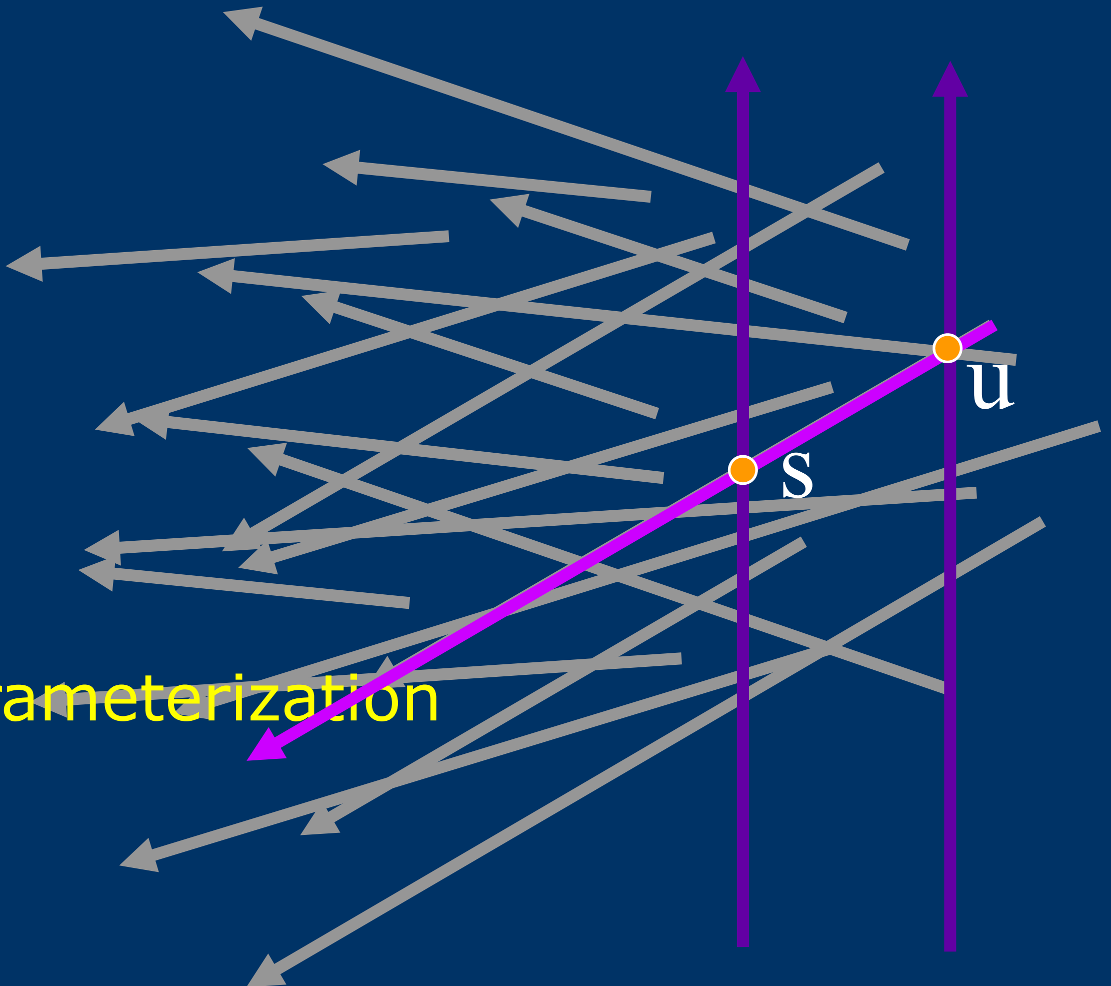

- 2 plane parameterization

Slide by Rick Szeliski and Michael Cohen

- 2D position
- 2D position

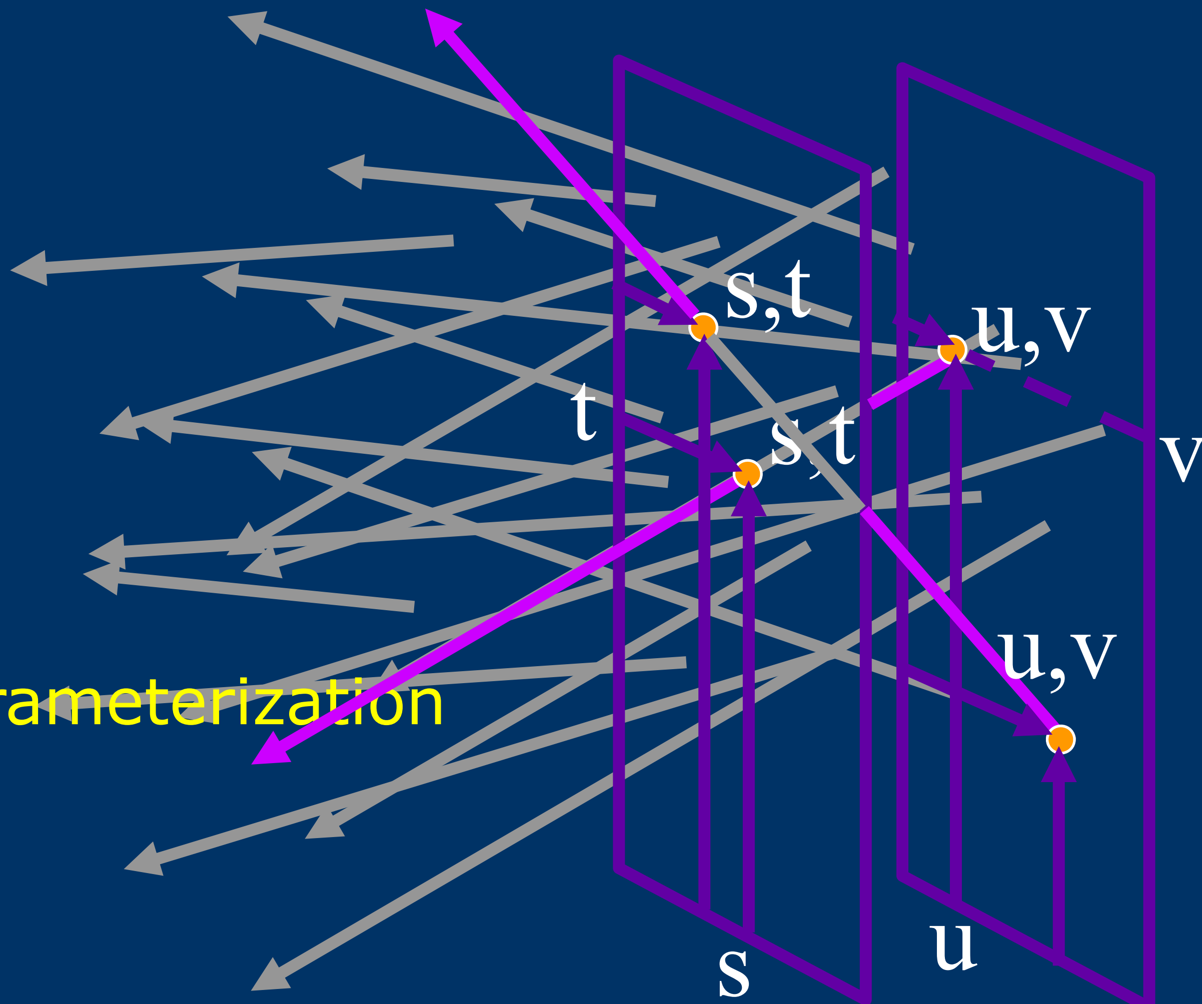

- 2 plane parameterization

Slide by Rick Szeliski and Michael Cohen

Light Field = Array of (virtual) Cameras

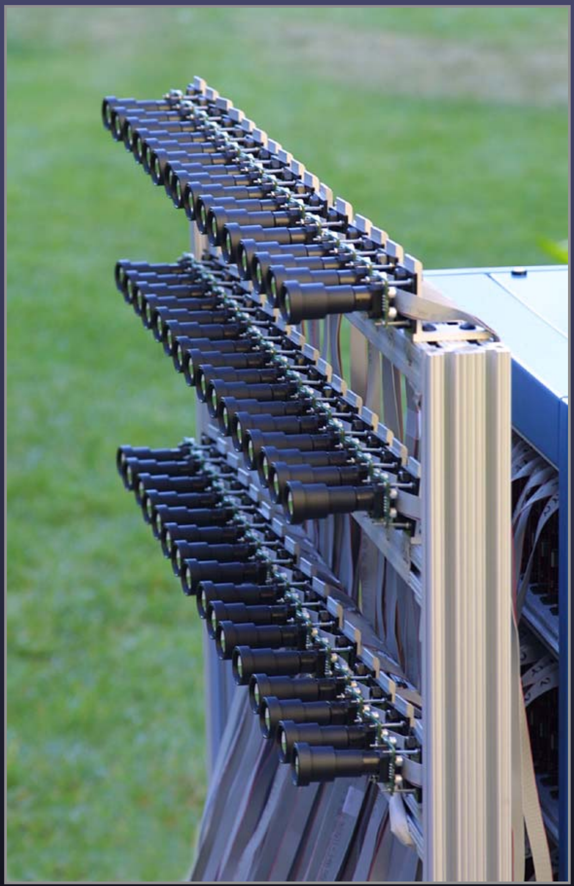

Pixel= (s,t)   Camera = (u,v)

Light Field = Array of (virtual) Cameras

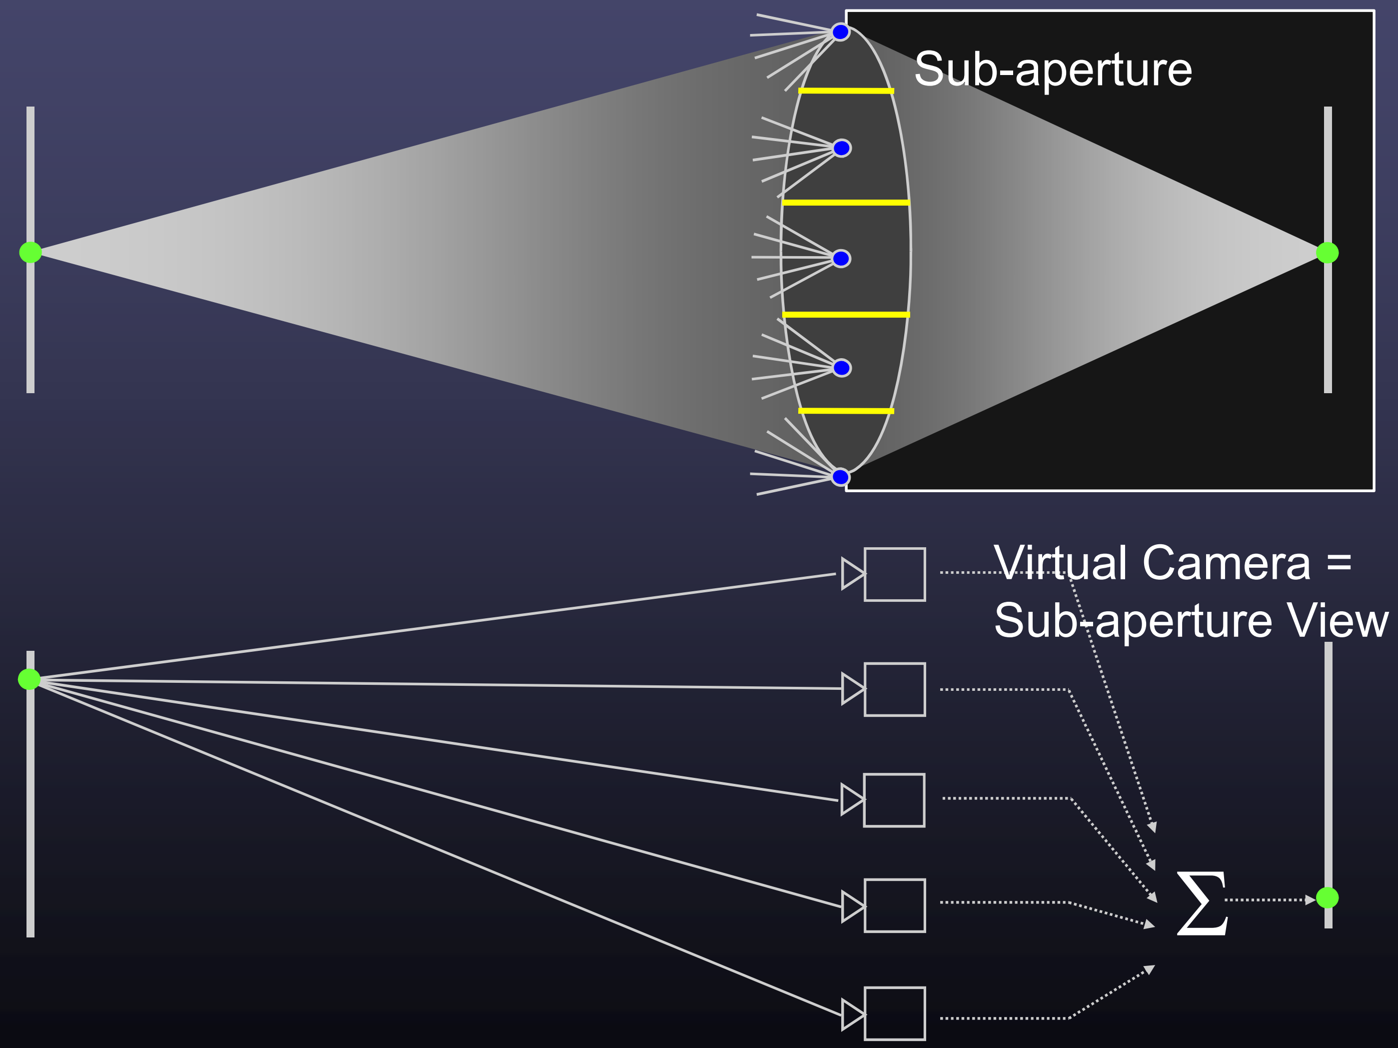

Conventional versus plenoptic camera

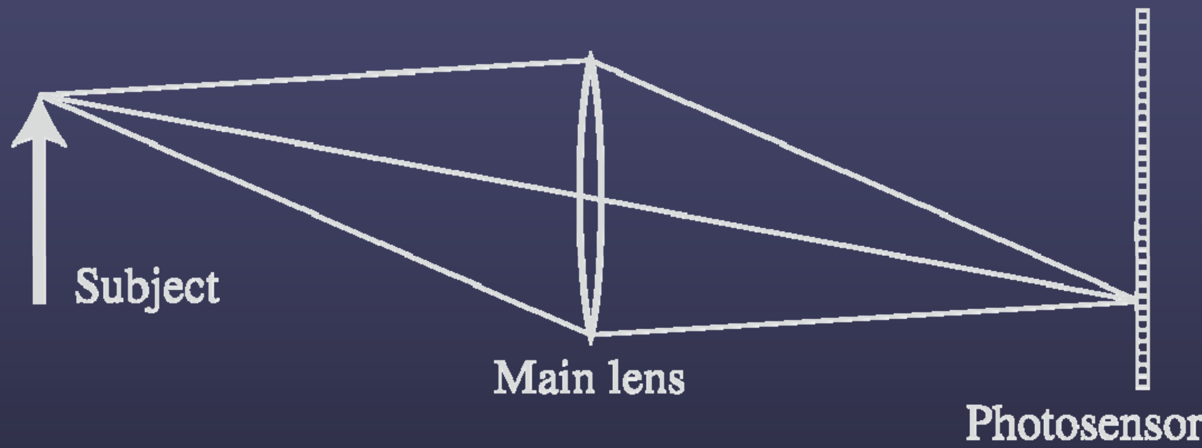

Scene Pixel = (s,t)

Virtual Camera = (u,v)

Pixel = (s,t)

uv-plane

st-plane

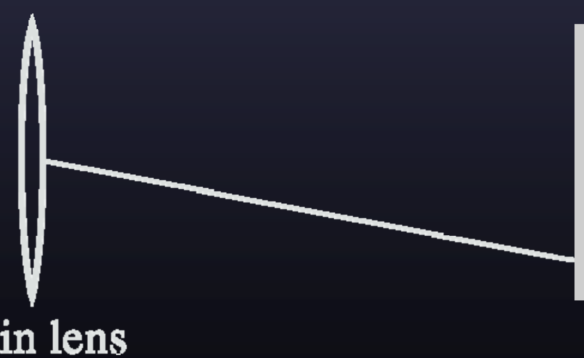

Light Field = Array of (virtual) Cameras

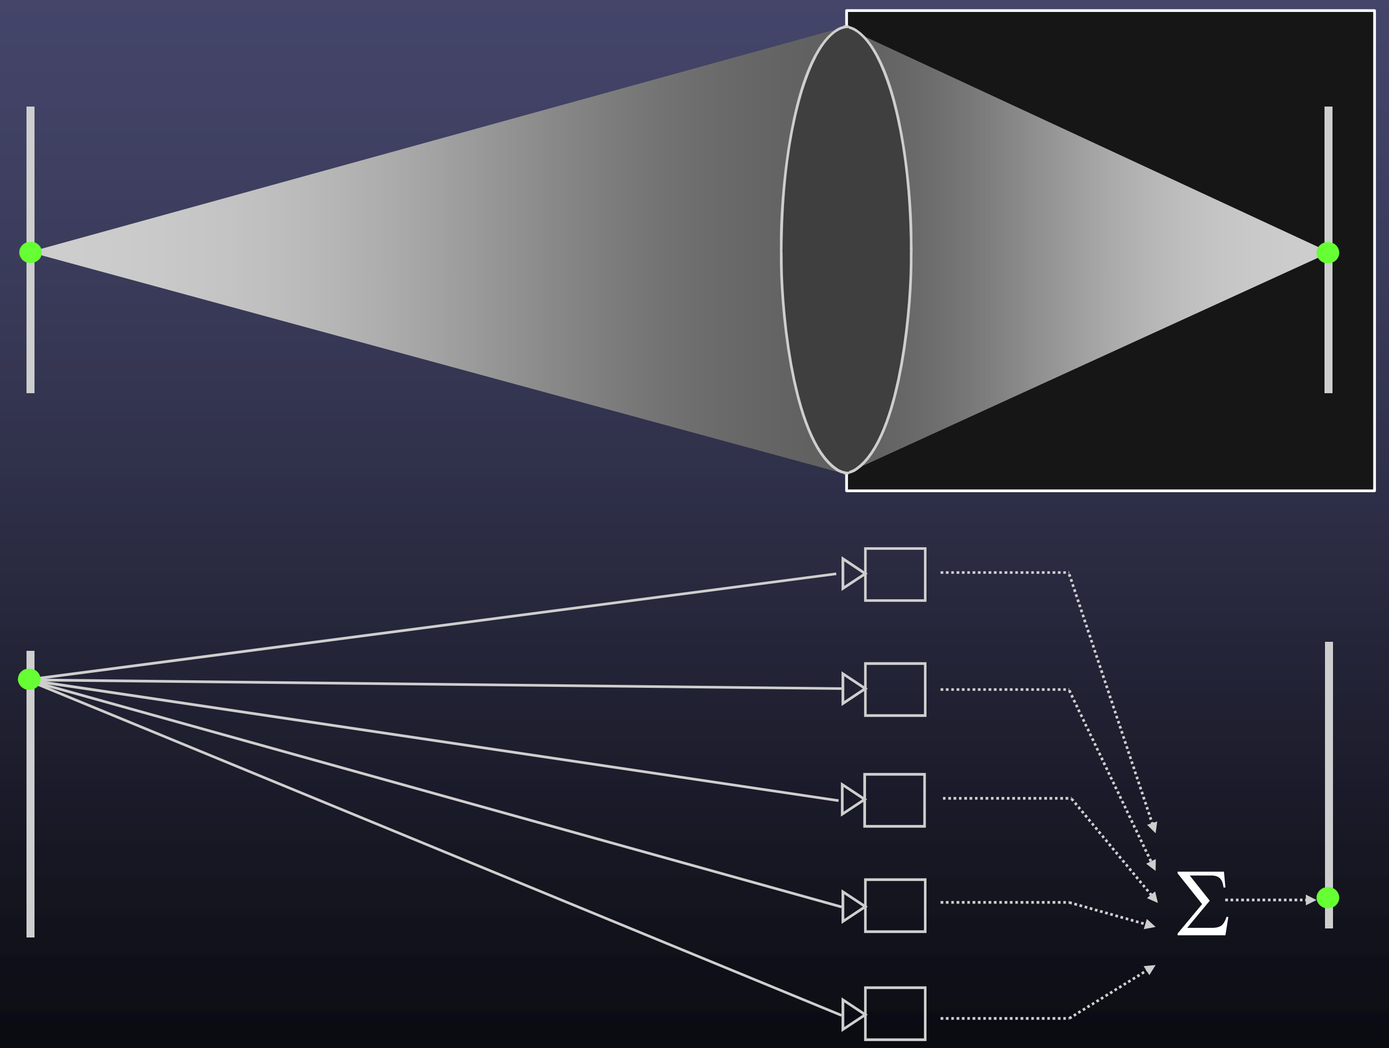

Light Field = Array of (virtual) Cameras

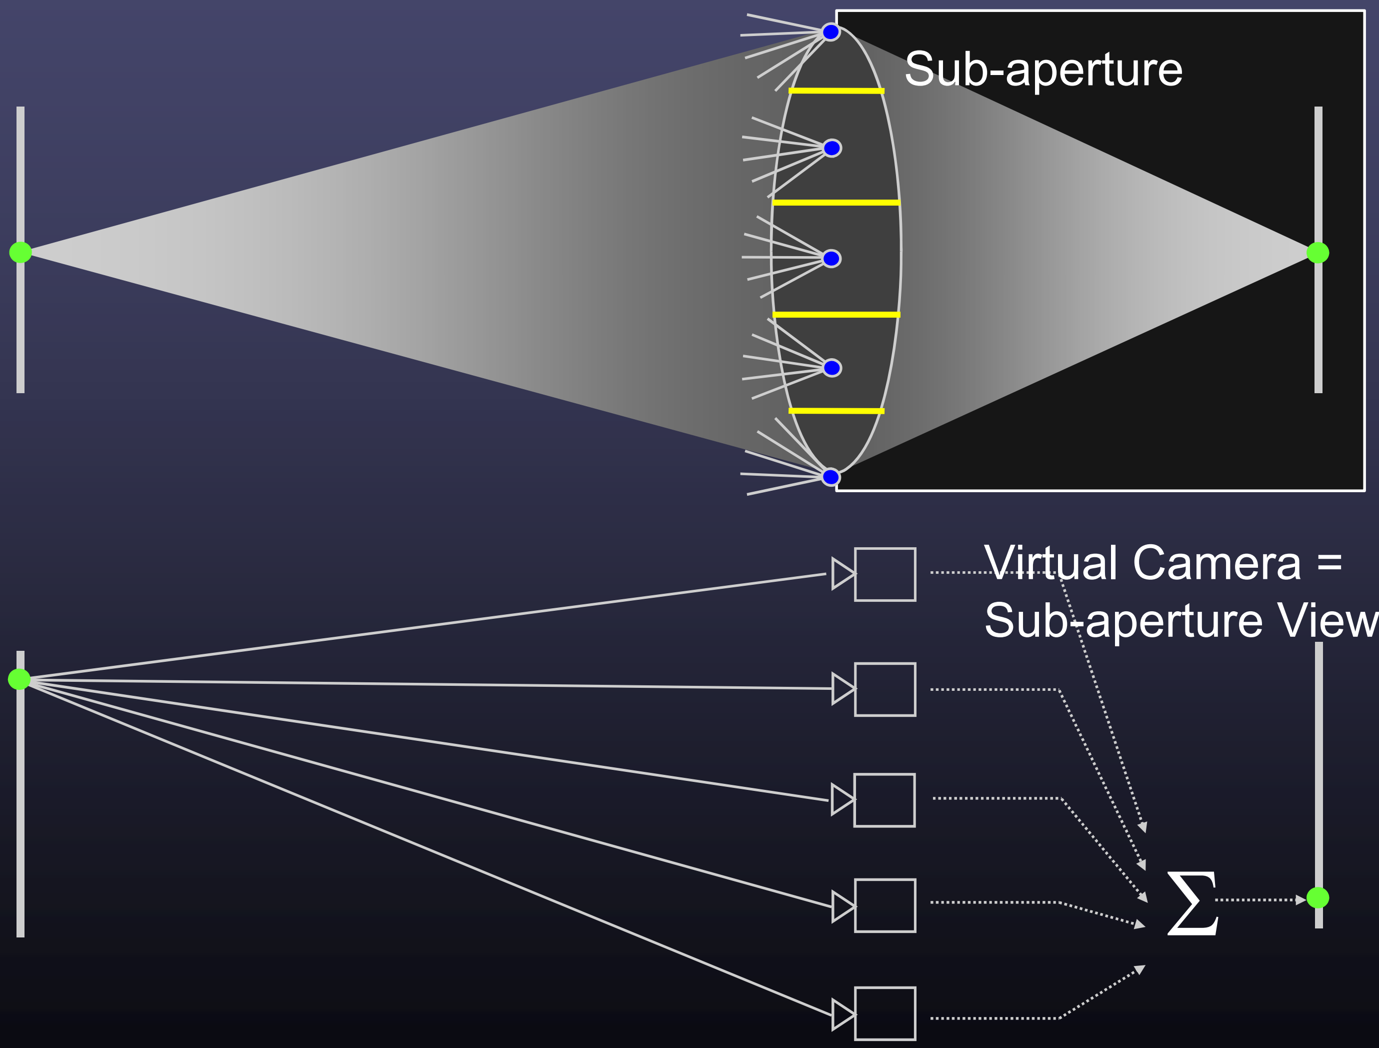

Light Field = Array of (virtual) Cameras

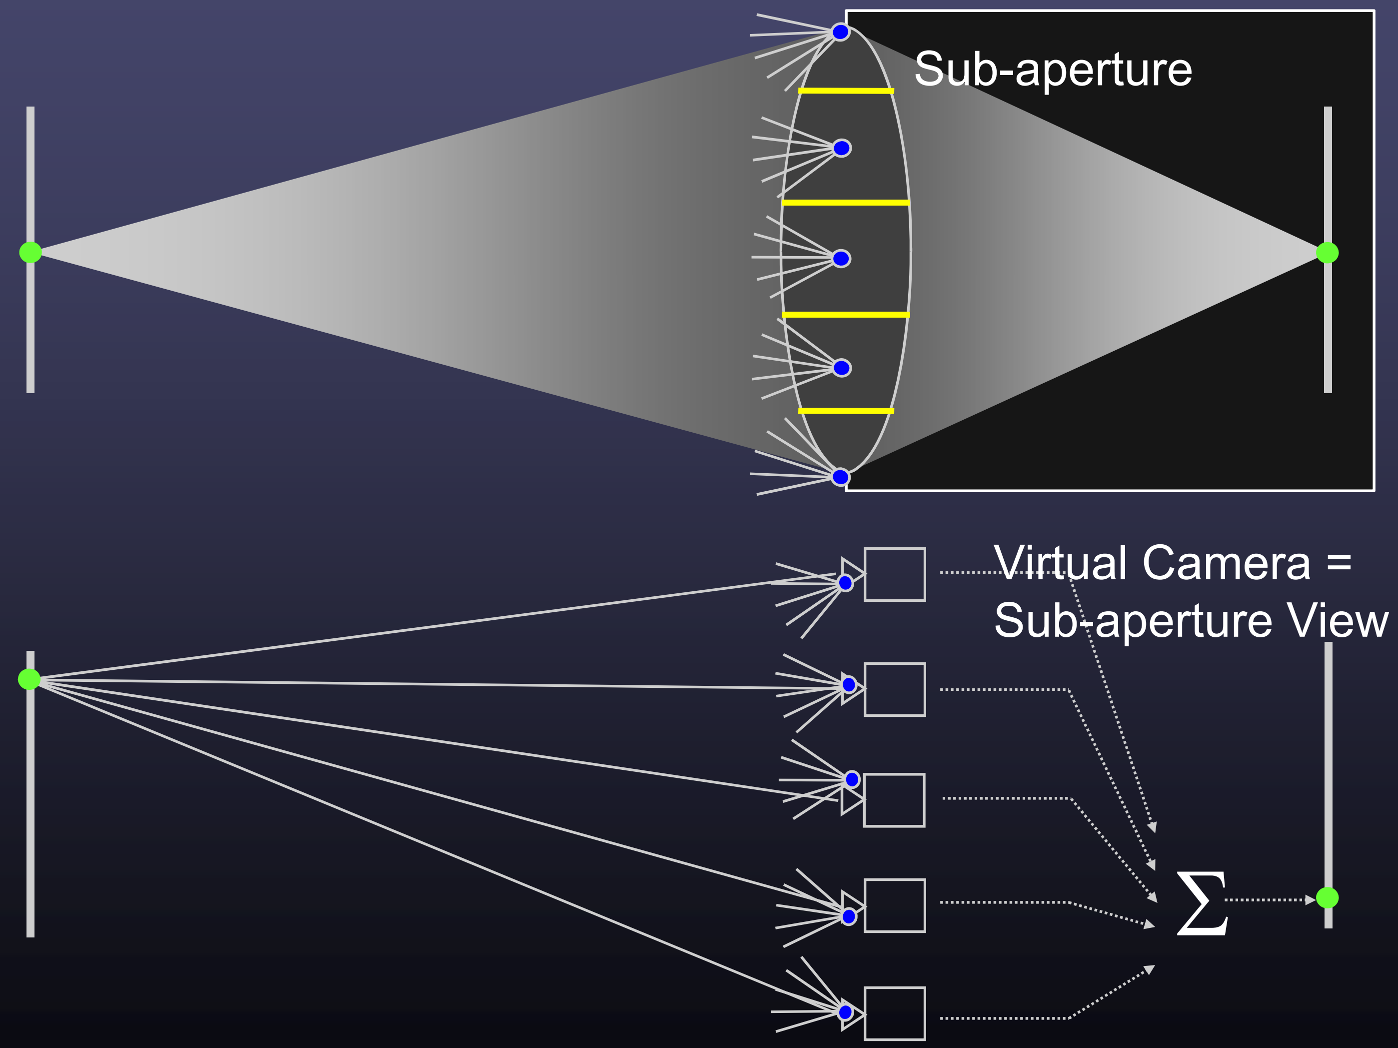

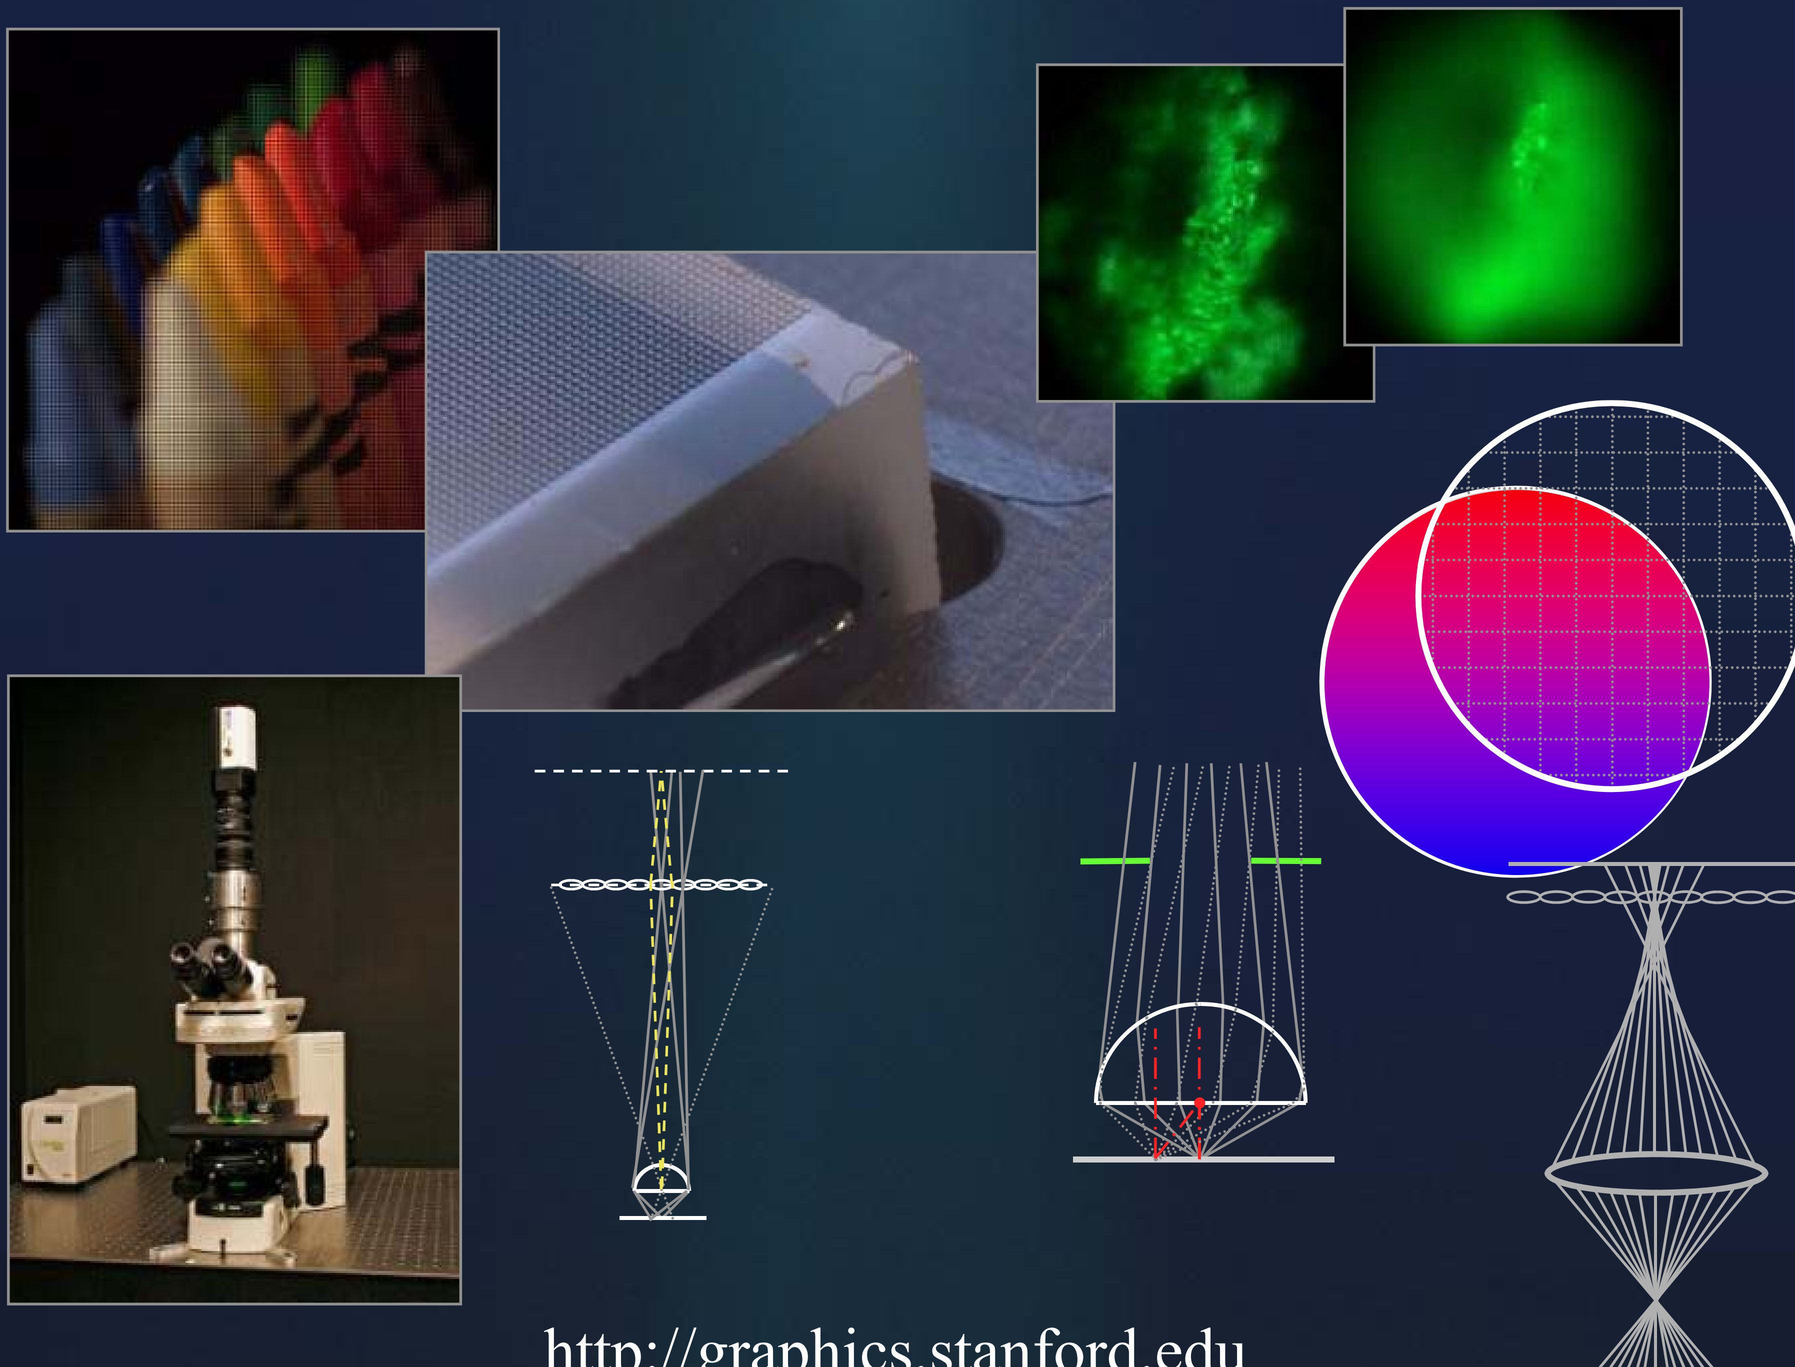

http://graphics.stanford.edu

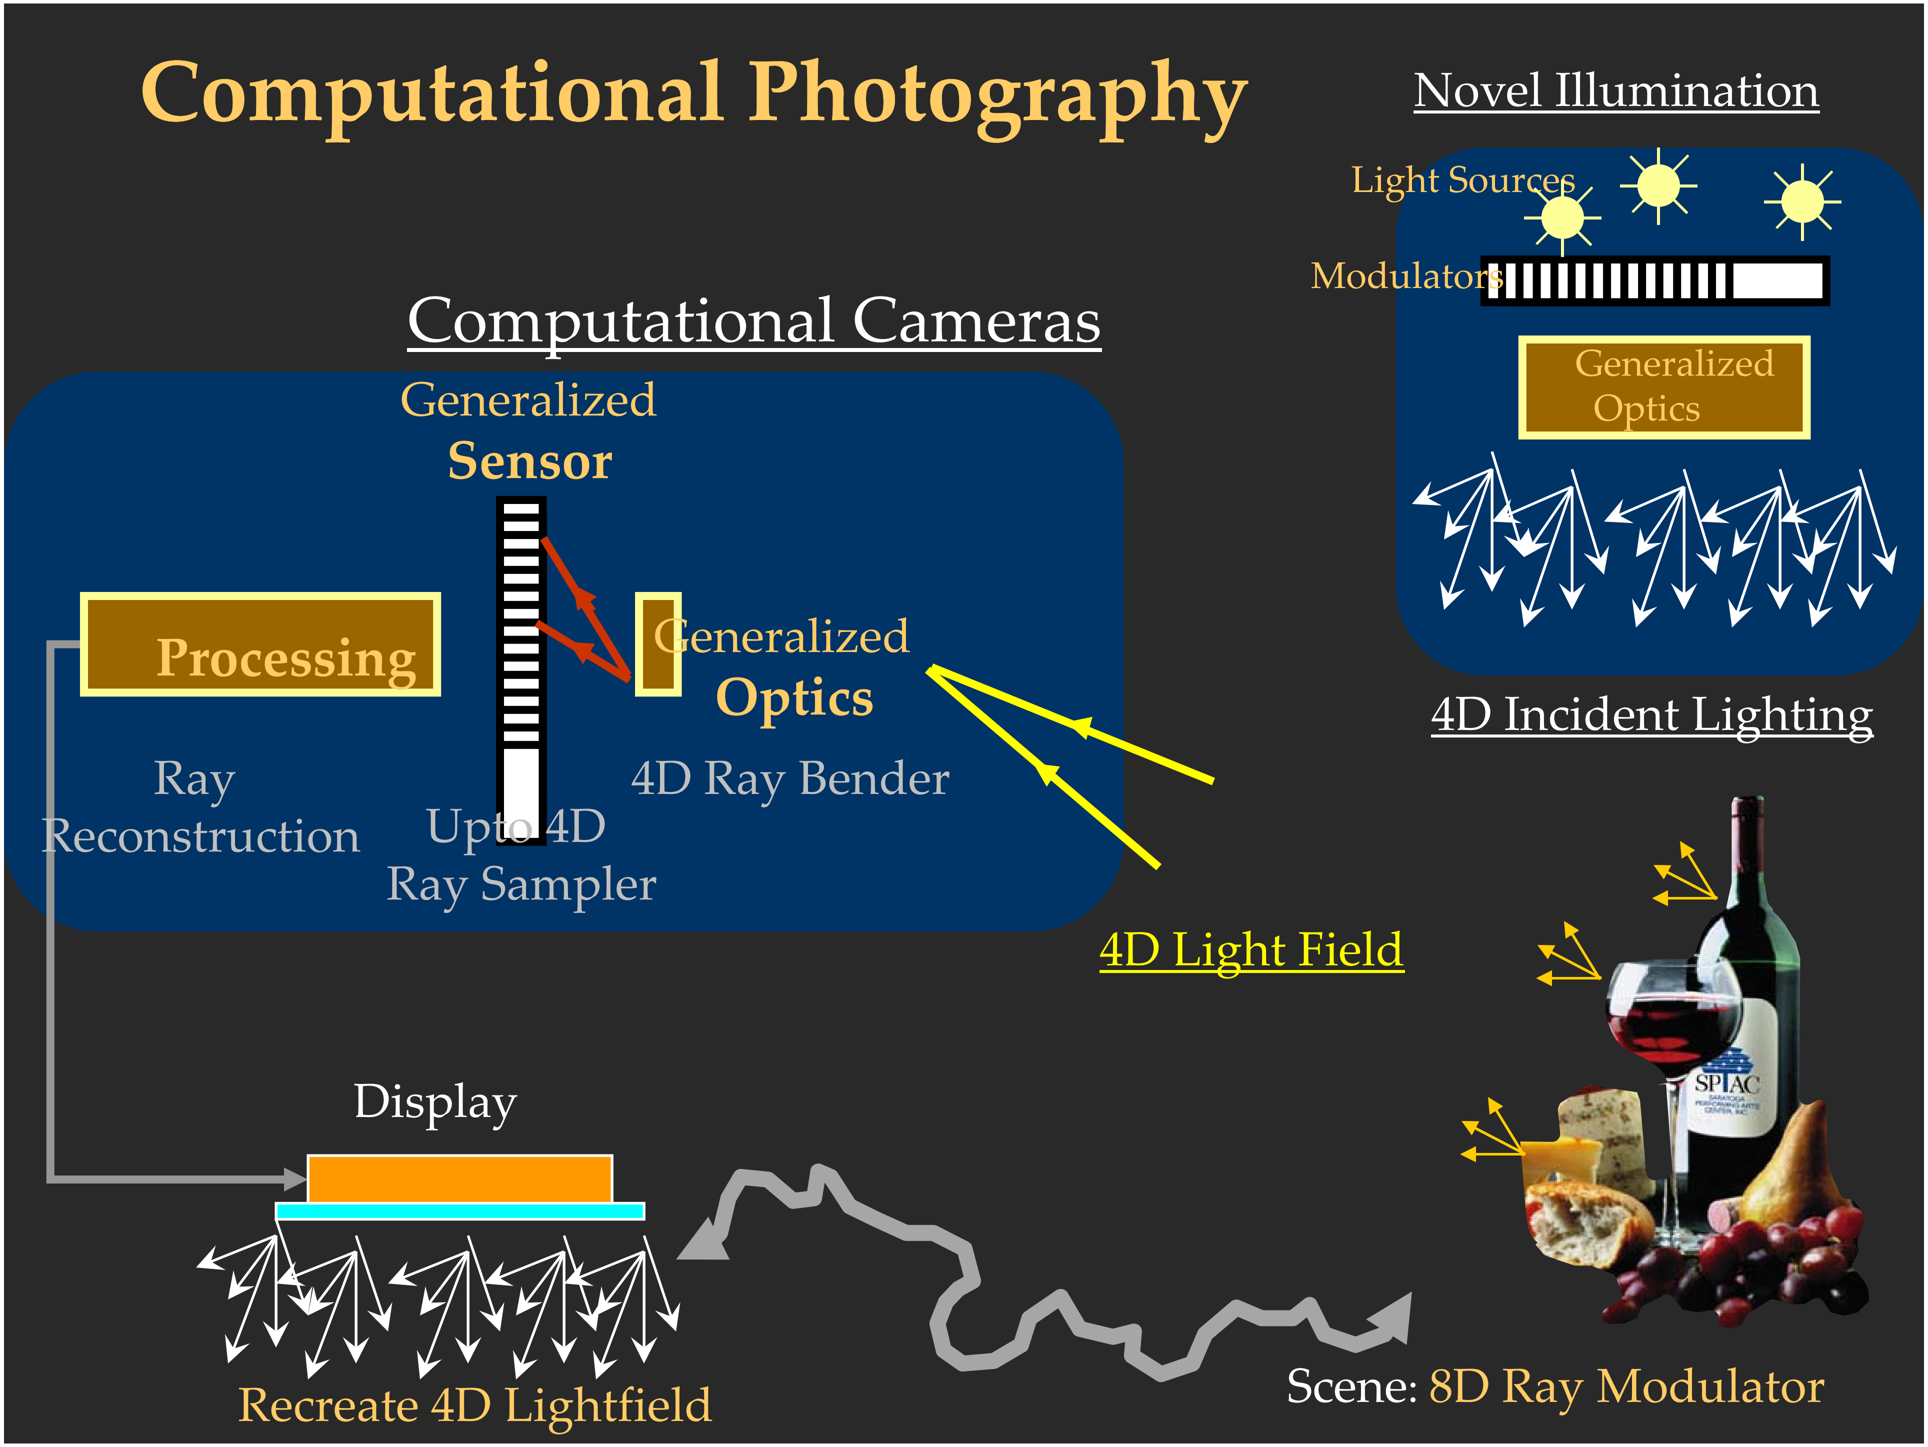

Computational Photography
Novel Illumination
Light Sources
Modulators
Generalized Optics
4D Incident Lighting
Computational Cameras
Generalized
Sensor
Processing
Generalized
Optics
Ray
Reconstruction
Upto 4D
Ray Sampler
4D Ray Bender
4D Light Field
Display
Recreate 4D Lightfield
Scene: 8D Ray Modulator

The light field
[Gershun 1936]

Radiance as a function of position and direction

- for general scenes
  - the "plenoptic function"
  - five-dimensional
  - $L(x, y, z, \theta, \phi)$ (w/m^{2}sr)

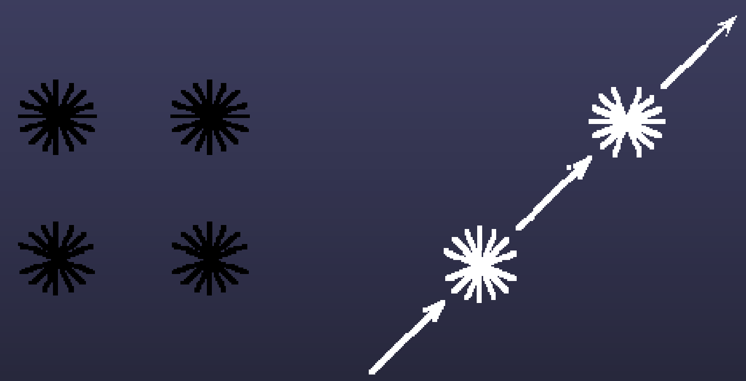

- in free space
  - four-dimensional
  - $L(u, v, s, t)$

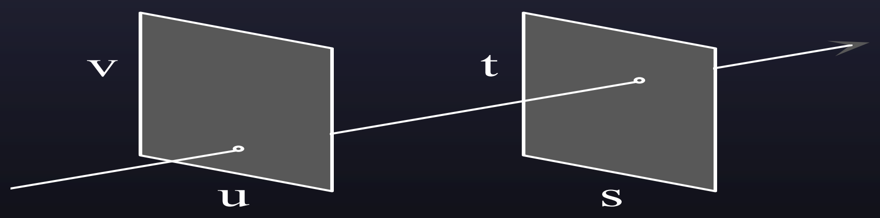

two-plane parameterization
[Levoy and Hanrahan 1996]

Radiance 'along a ray'

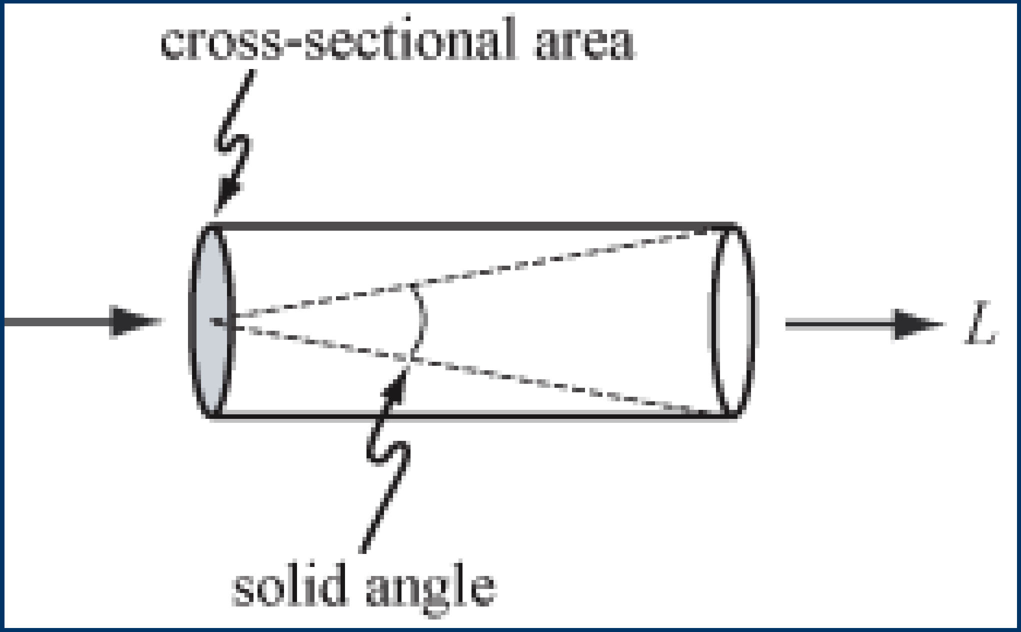

Radiance *L* along a ray can be thought of as the amount of light traveling along all possible straight lines through a tube whose size is determined by its solid angle and cross-sectional area.

measured in watts (W) per steradian (sr) per meter squared (m^2).

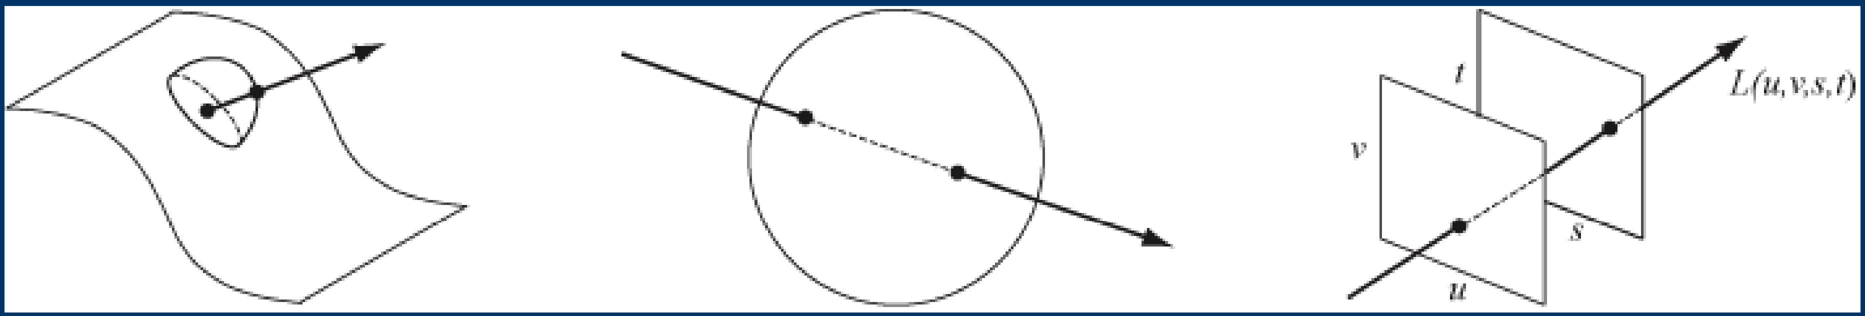

Some alternative parameterizations of the 4D light field, which represents the flow of light through an empty region of three-dimensional space. *Left:* points on a plane or curved surface and directions leaving each point. *Center:* pairs of points on the surface of a sphere. *Right:* pairs of points on two planes in general (meaning any) position.

Light Field Inside a Camera

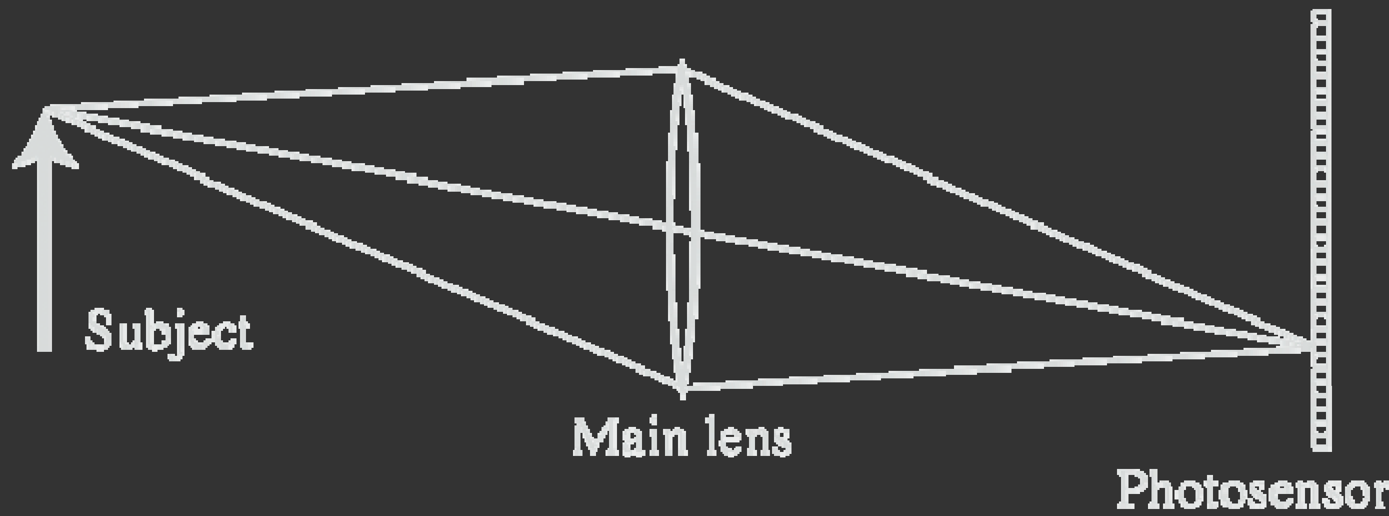

Light Field Inside a Camera

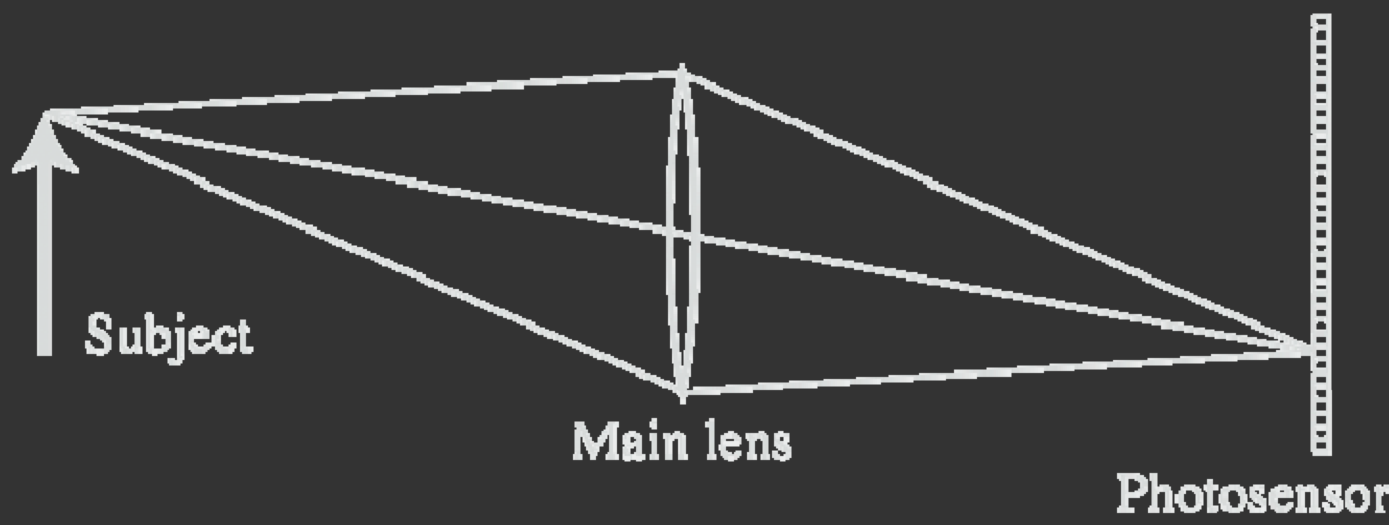

Lenslet-based Light Field camera

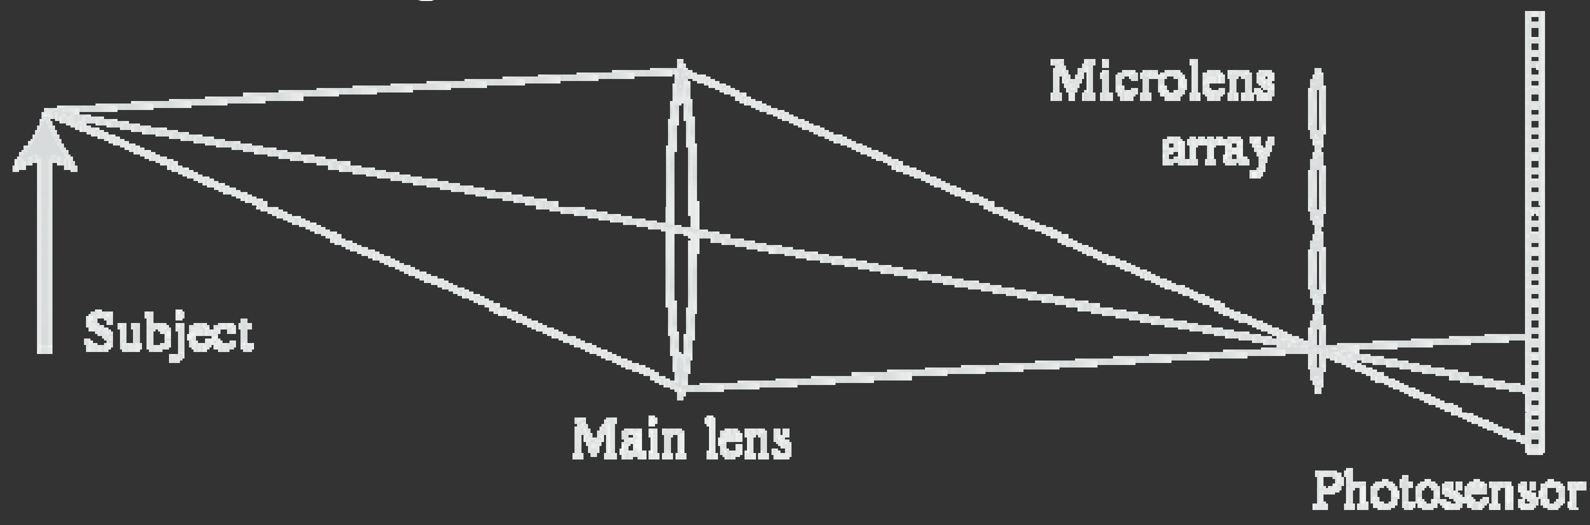

[Adelson and Wang, 1992, Ng et al. 2005]

Stanford Plenoptic Camera [Ng et al 2005]

Contax medium format camera

Kodak 16-megapixel sensor

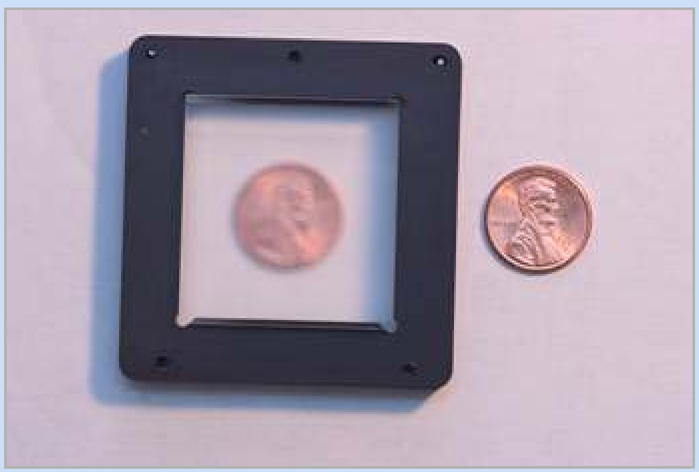

Adaptive Optics microlens array

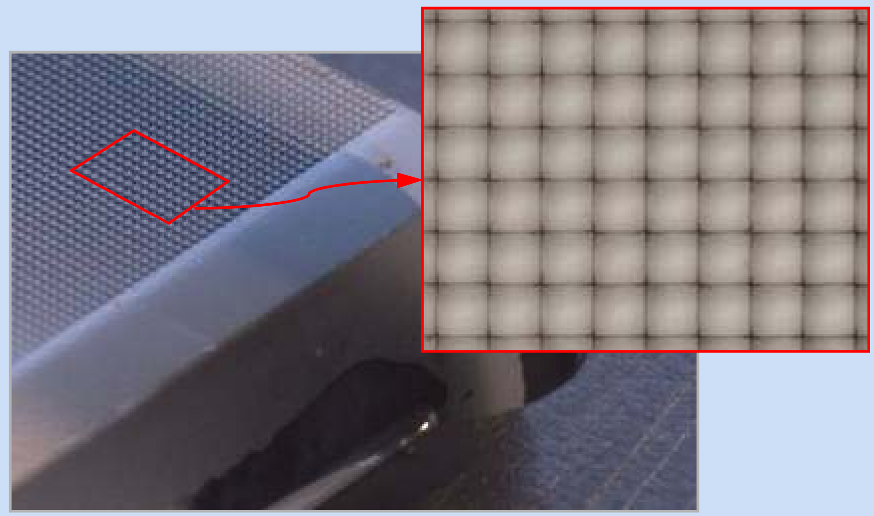

125μ square-sided microlenses

4000 × 4000 pixels ÷ 292 × 292 lenses = 14 × 14 pixels per lens

Digital Refocusing

[Ng et al 2005]

Adaptive Optics

- A deformable mirror can be used to correct wavefront errors in an astronomical telescope

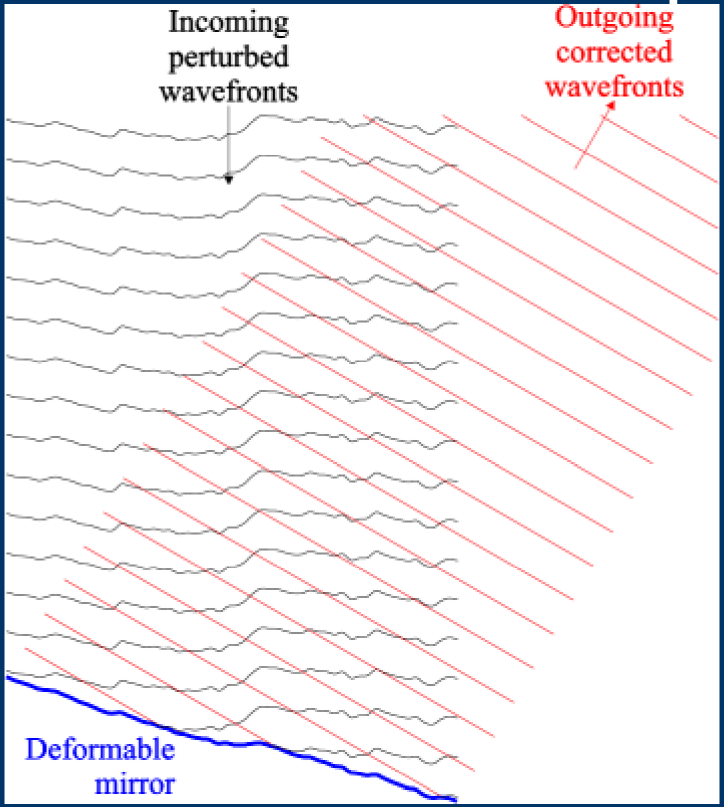

http://en.wikipedia.org/wiki/Image:Adaptive_optics_correct.png

Shack Hartmann wavefront sensor (commonly used in Adaptive optics).

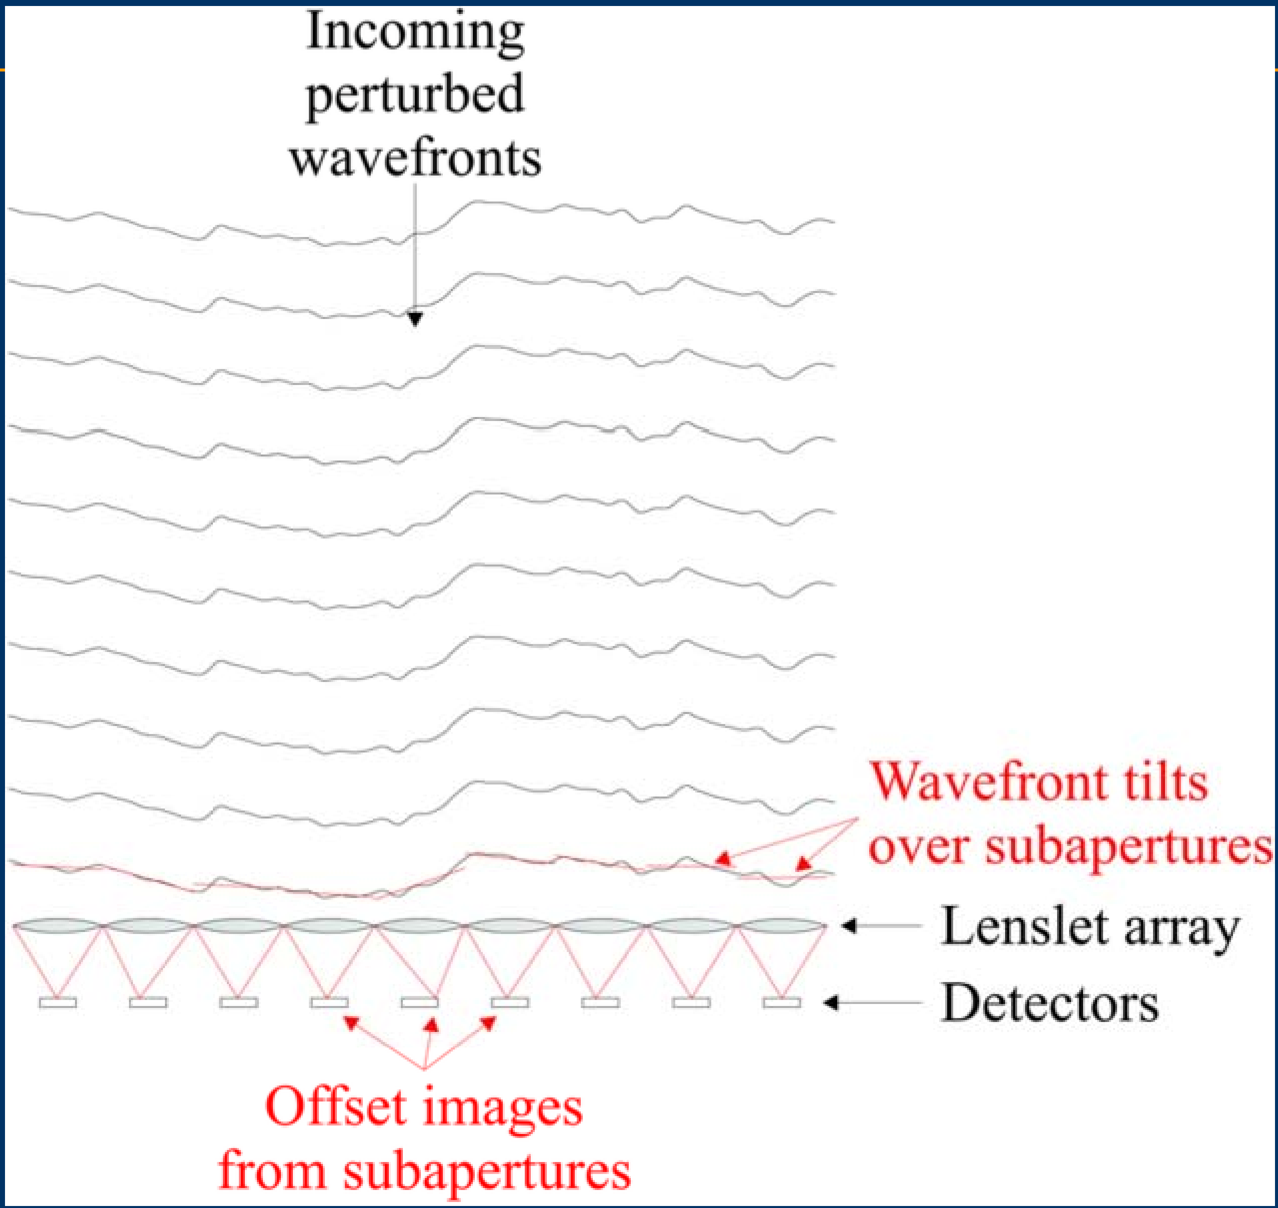

http://en.wikipedia.org/wiki/Image:Shack_hartmann.png

Measuring shape (phase) of wavefront ~ Lightfield Capture

- http://www.cvs.rochester.edu/williamslab/r_shackhartmann.html

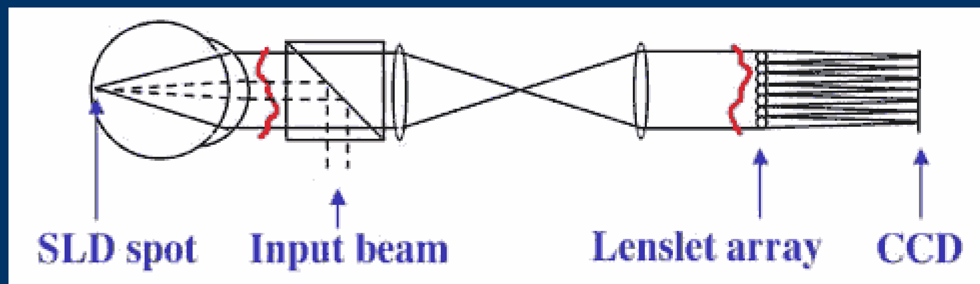

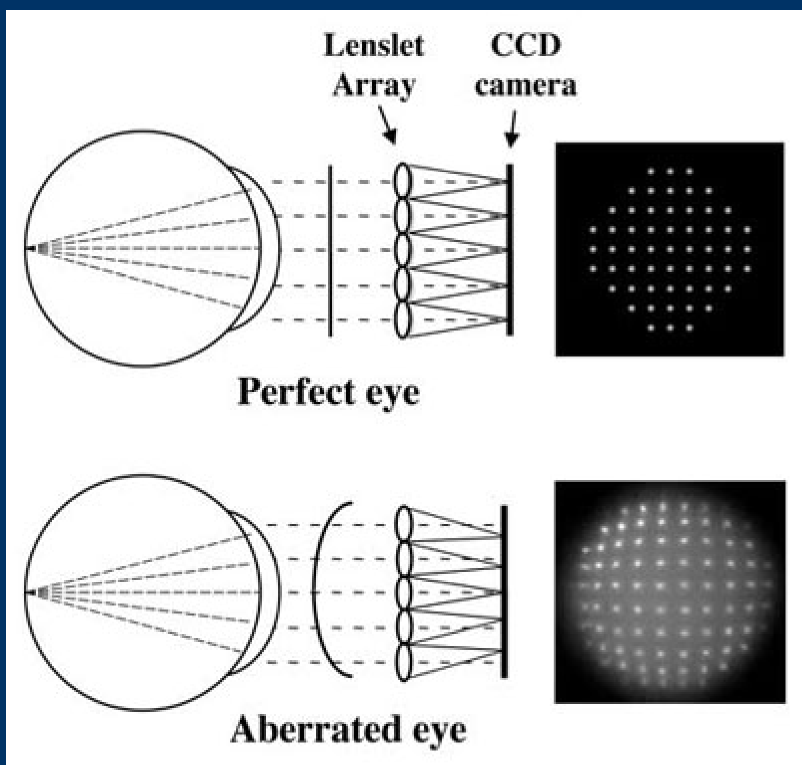

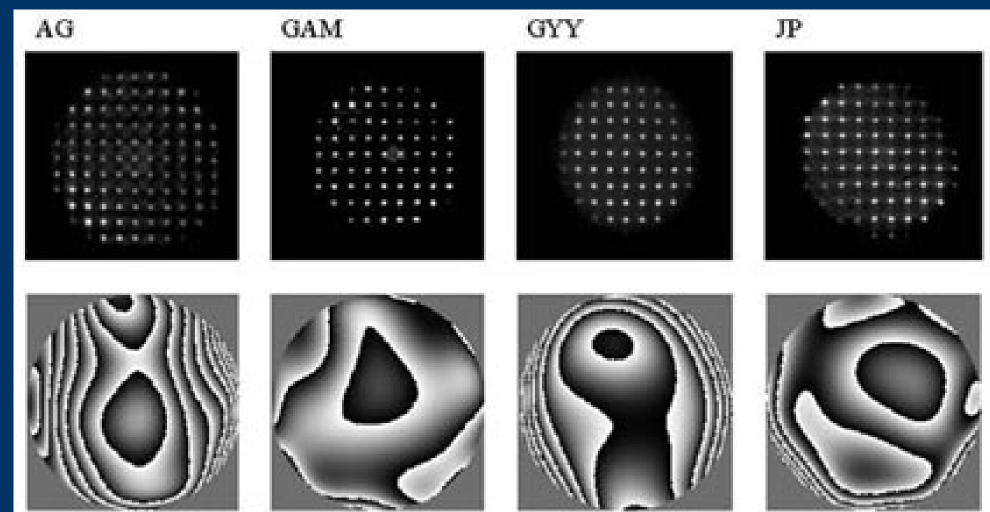

The spots formed on the CCD chip for the eye will be displaced because the wavefront will hit each lenslet at an angle rather than straight on.

Example using 45 cameras
[Vaish CVPR 2004]

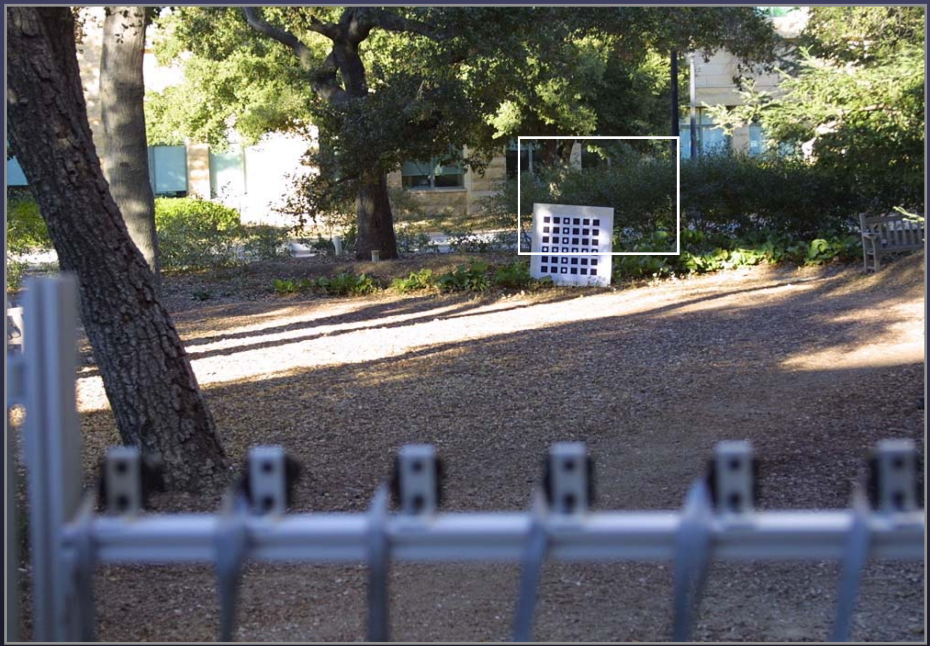

Synthetic aperture videography


A Virtual Optical Bench

- Understanding rays during defocus

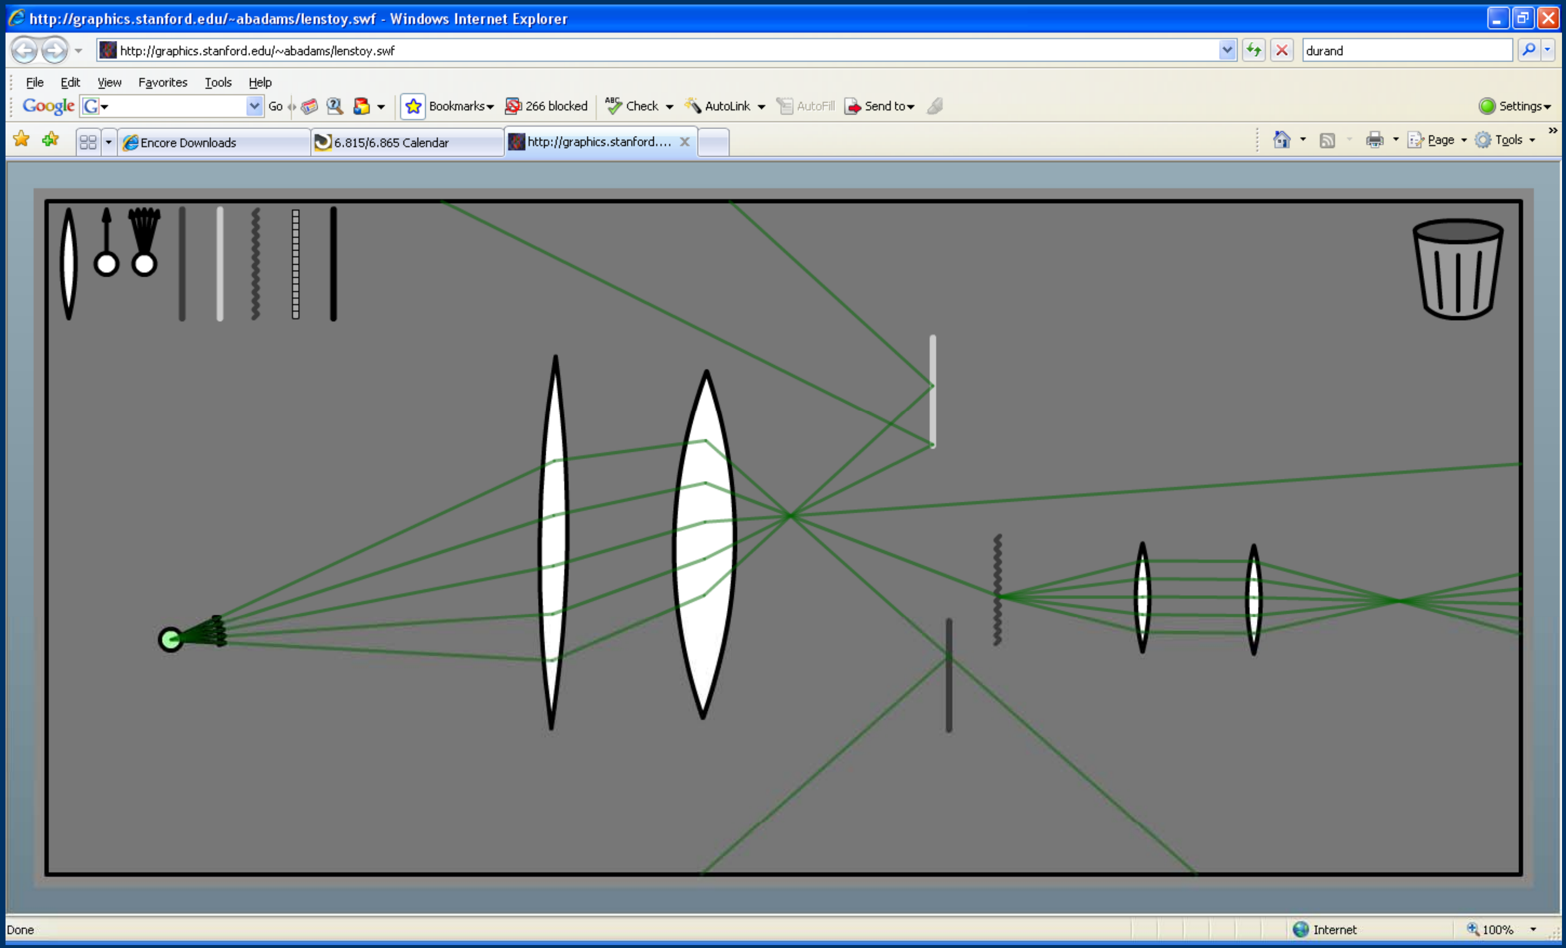

Visualizing Lightfield

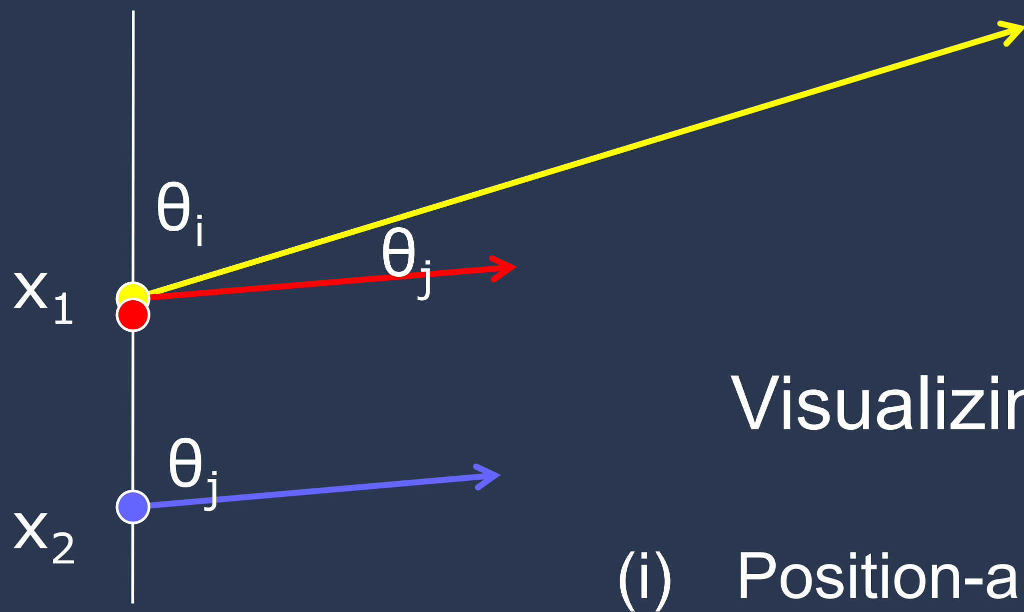

(i) Position-angle space
(ii) Phase-space
(iii) Space- Spatial Frequency
(iv) Spectrogram

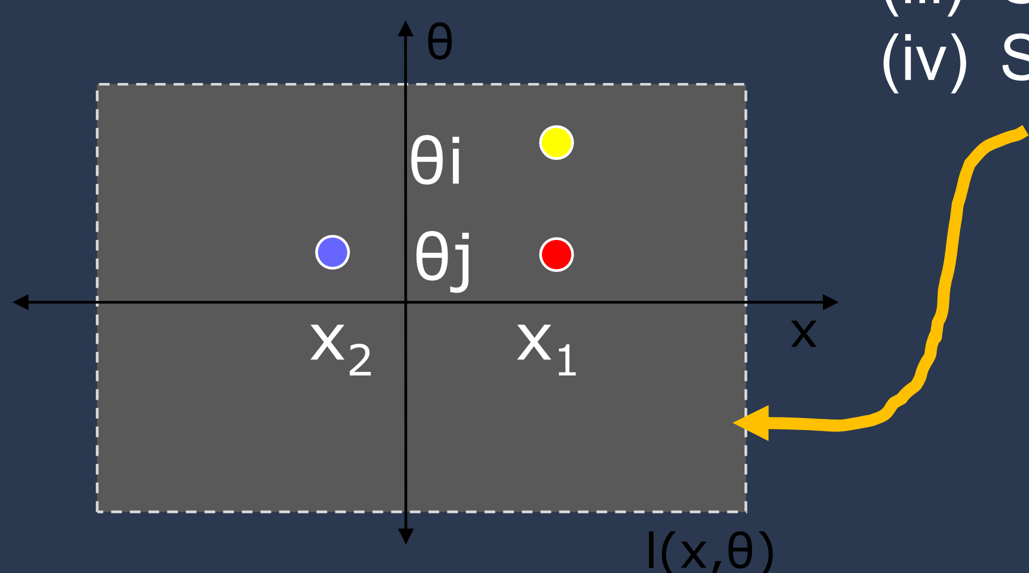

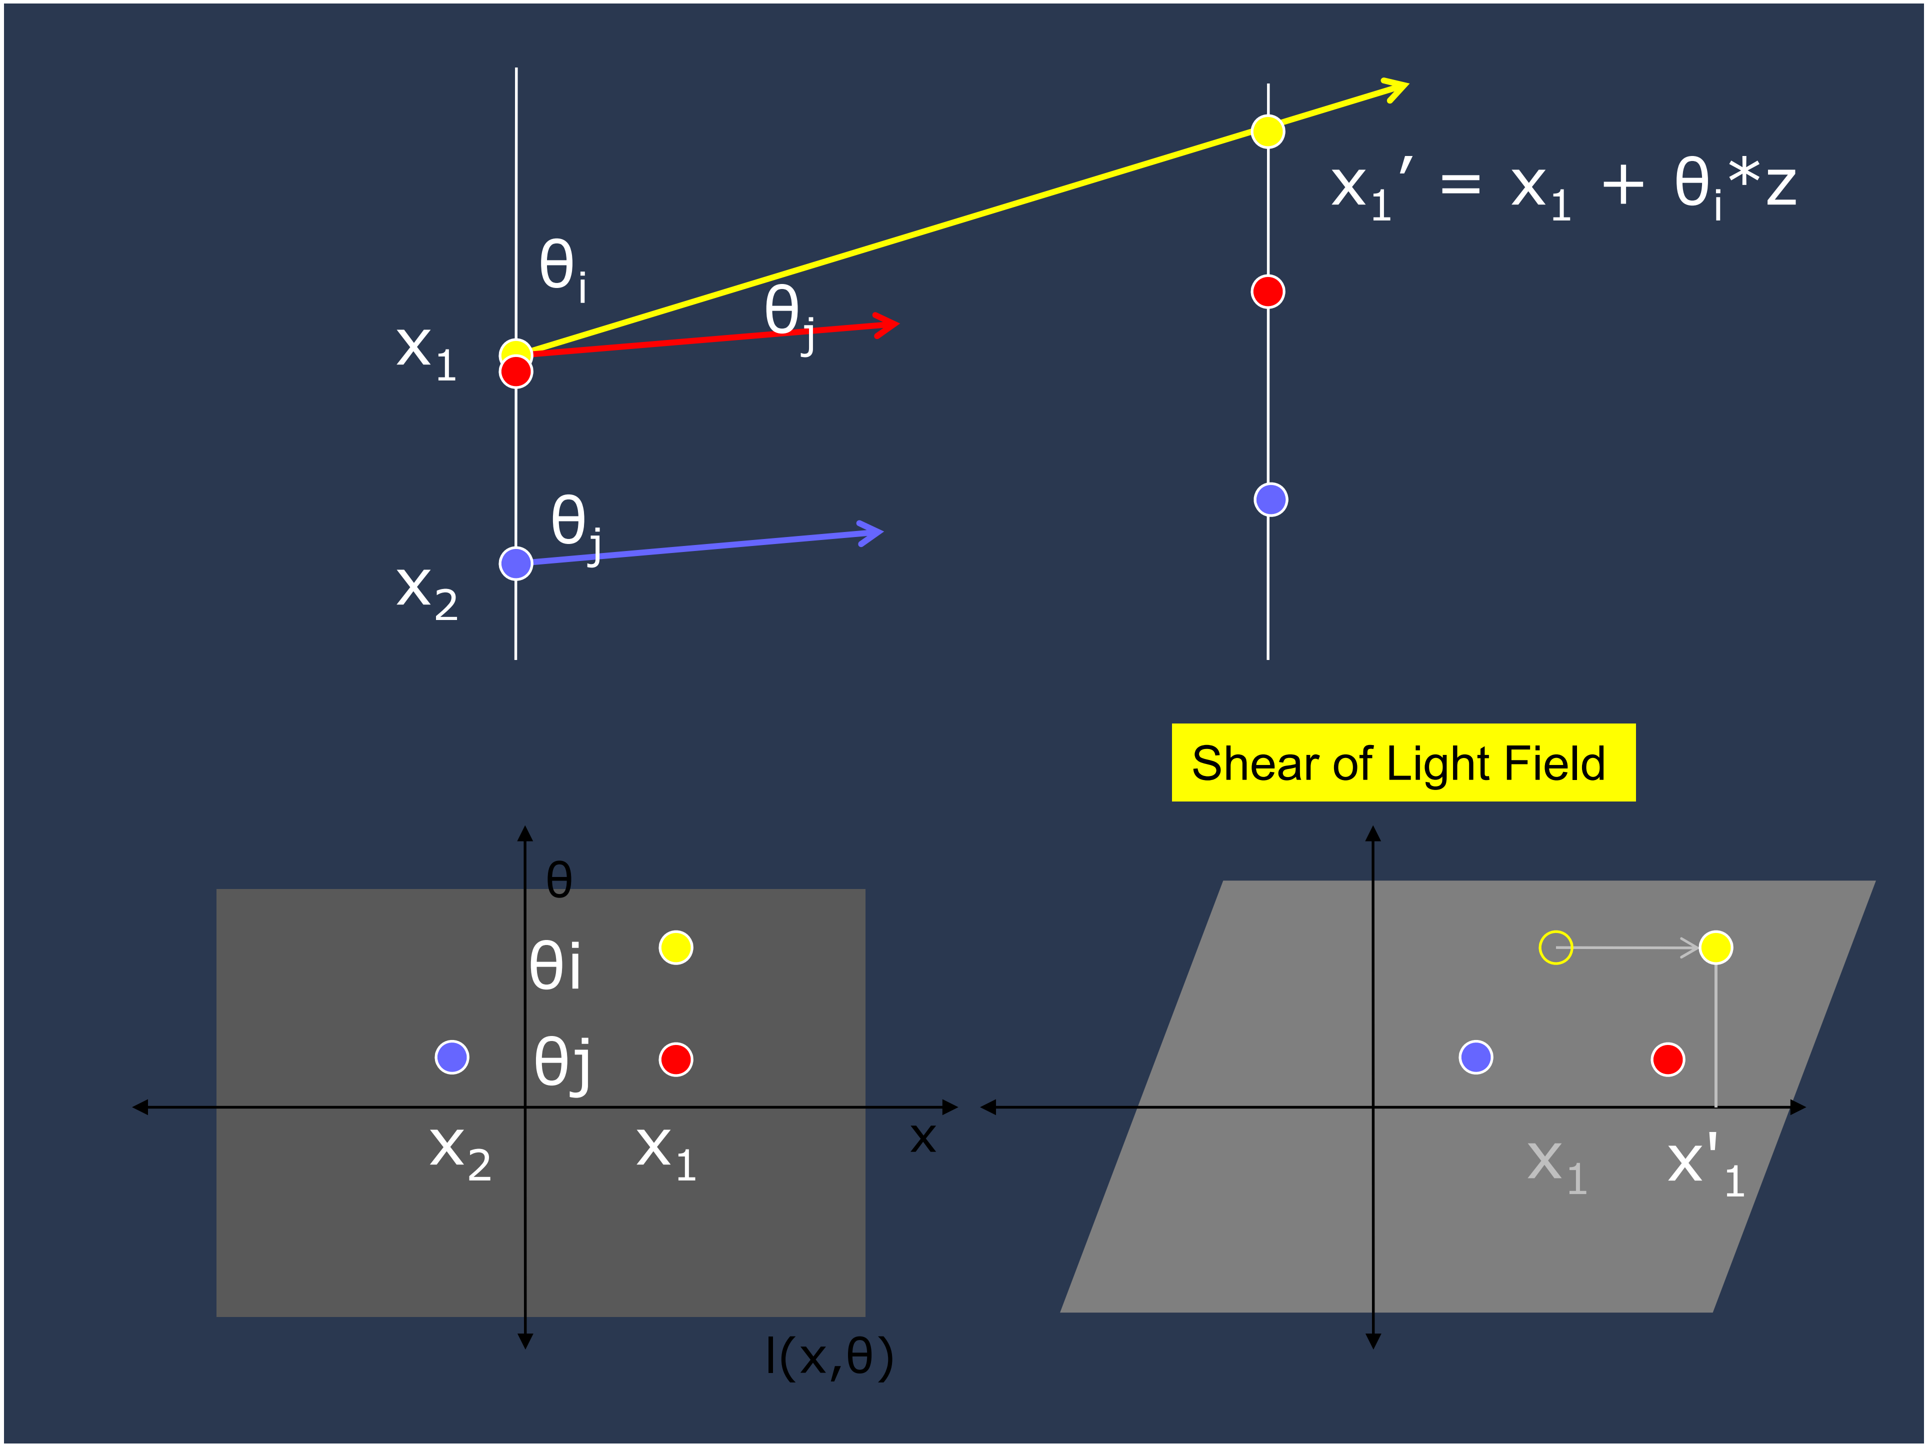

$x_1' = x_1 + \theta_i * z$
θ_i
θ_j
x_1
θ_j
x_2
Shear of Light Field
θ
θi
θj
x_2
x_1
x
$I(x,\theta)$
x_1
x'_1

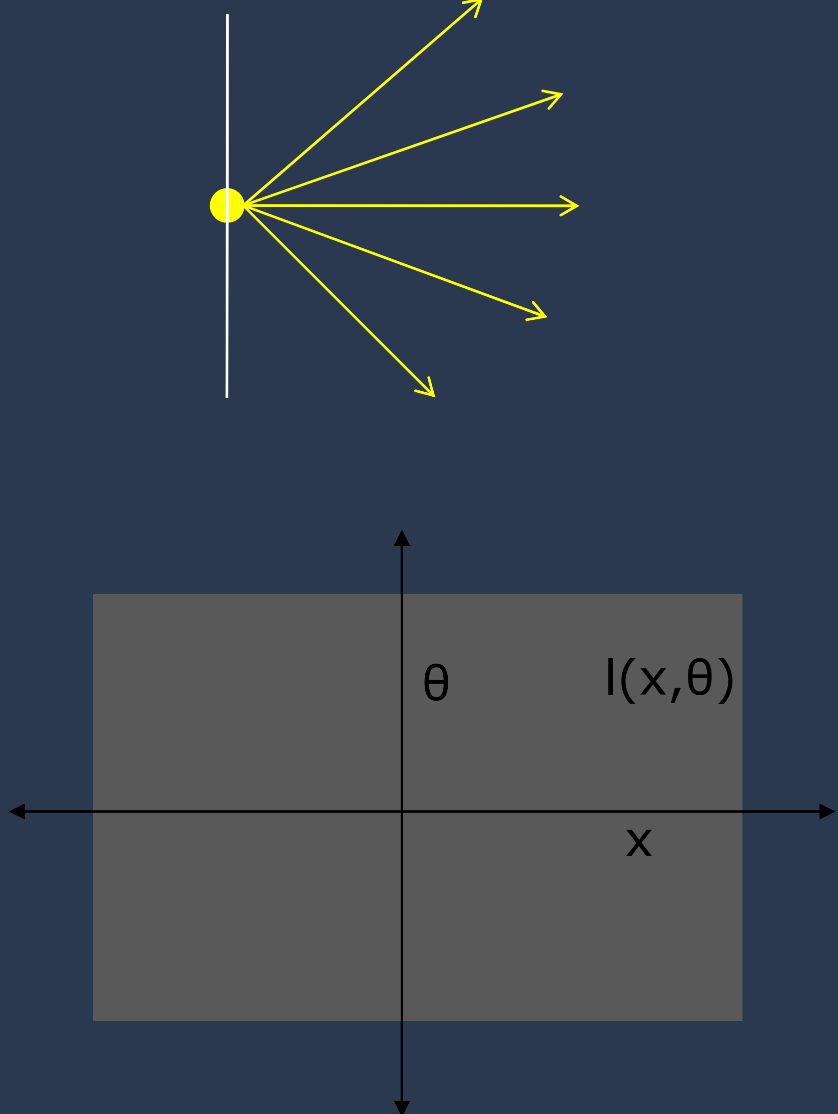
θ
I(x,θ)
x

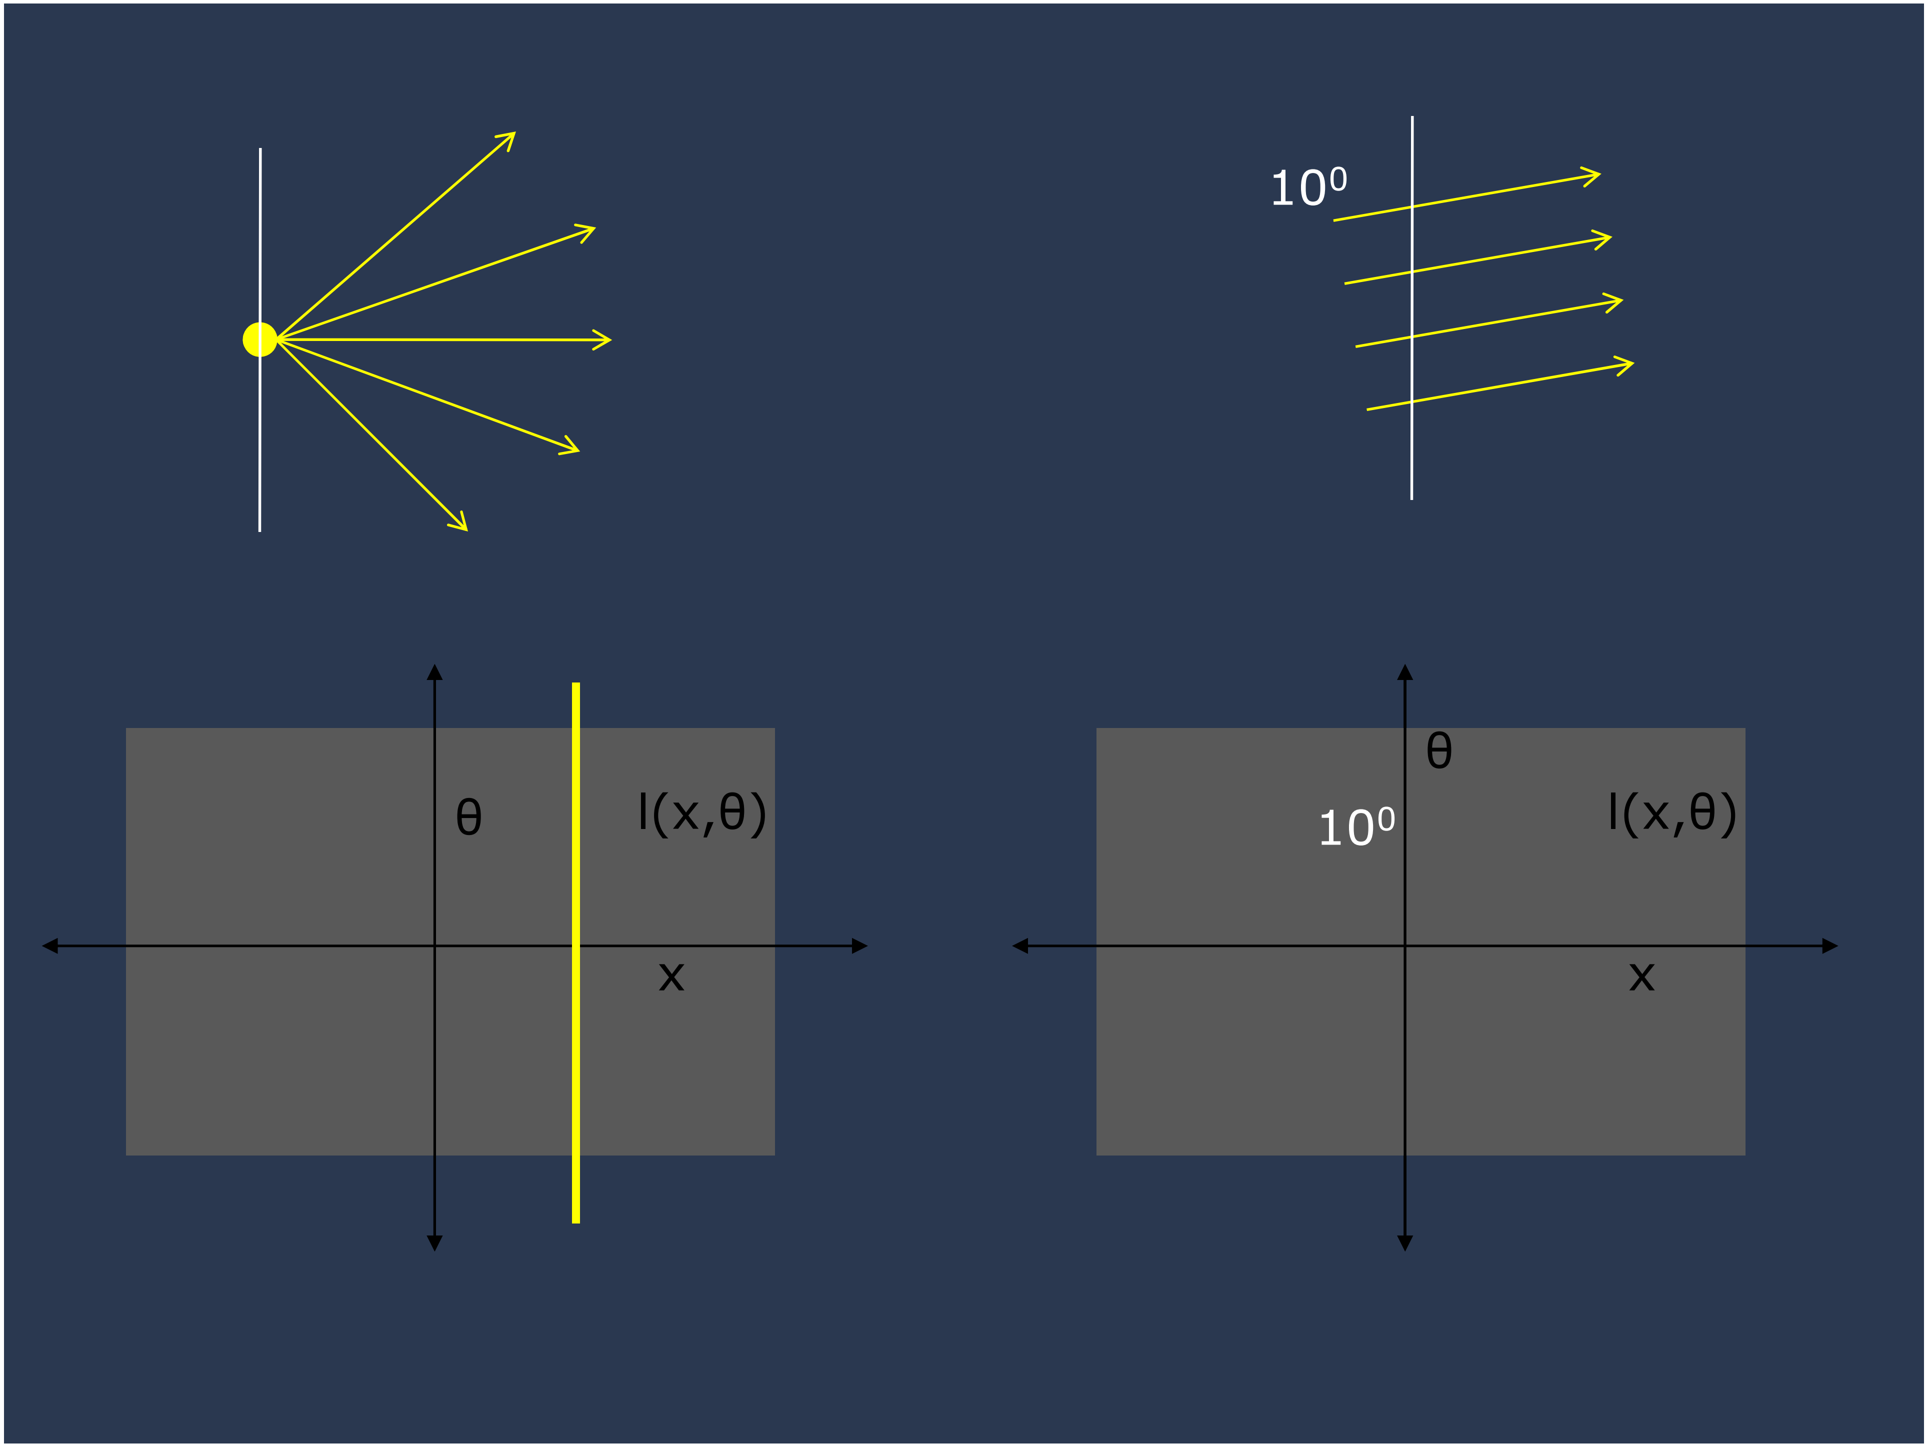

10^0
θ
I(x,θ)
x
θ
10^0
I(x,θ)
x

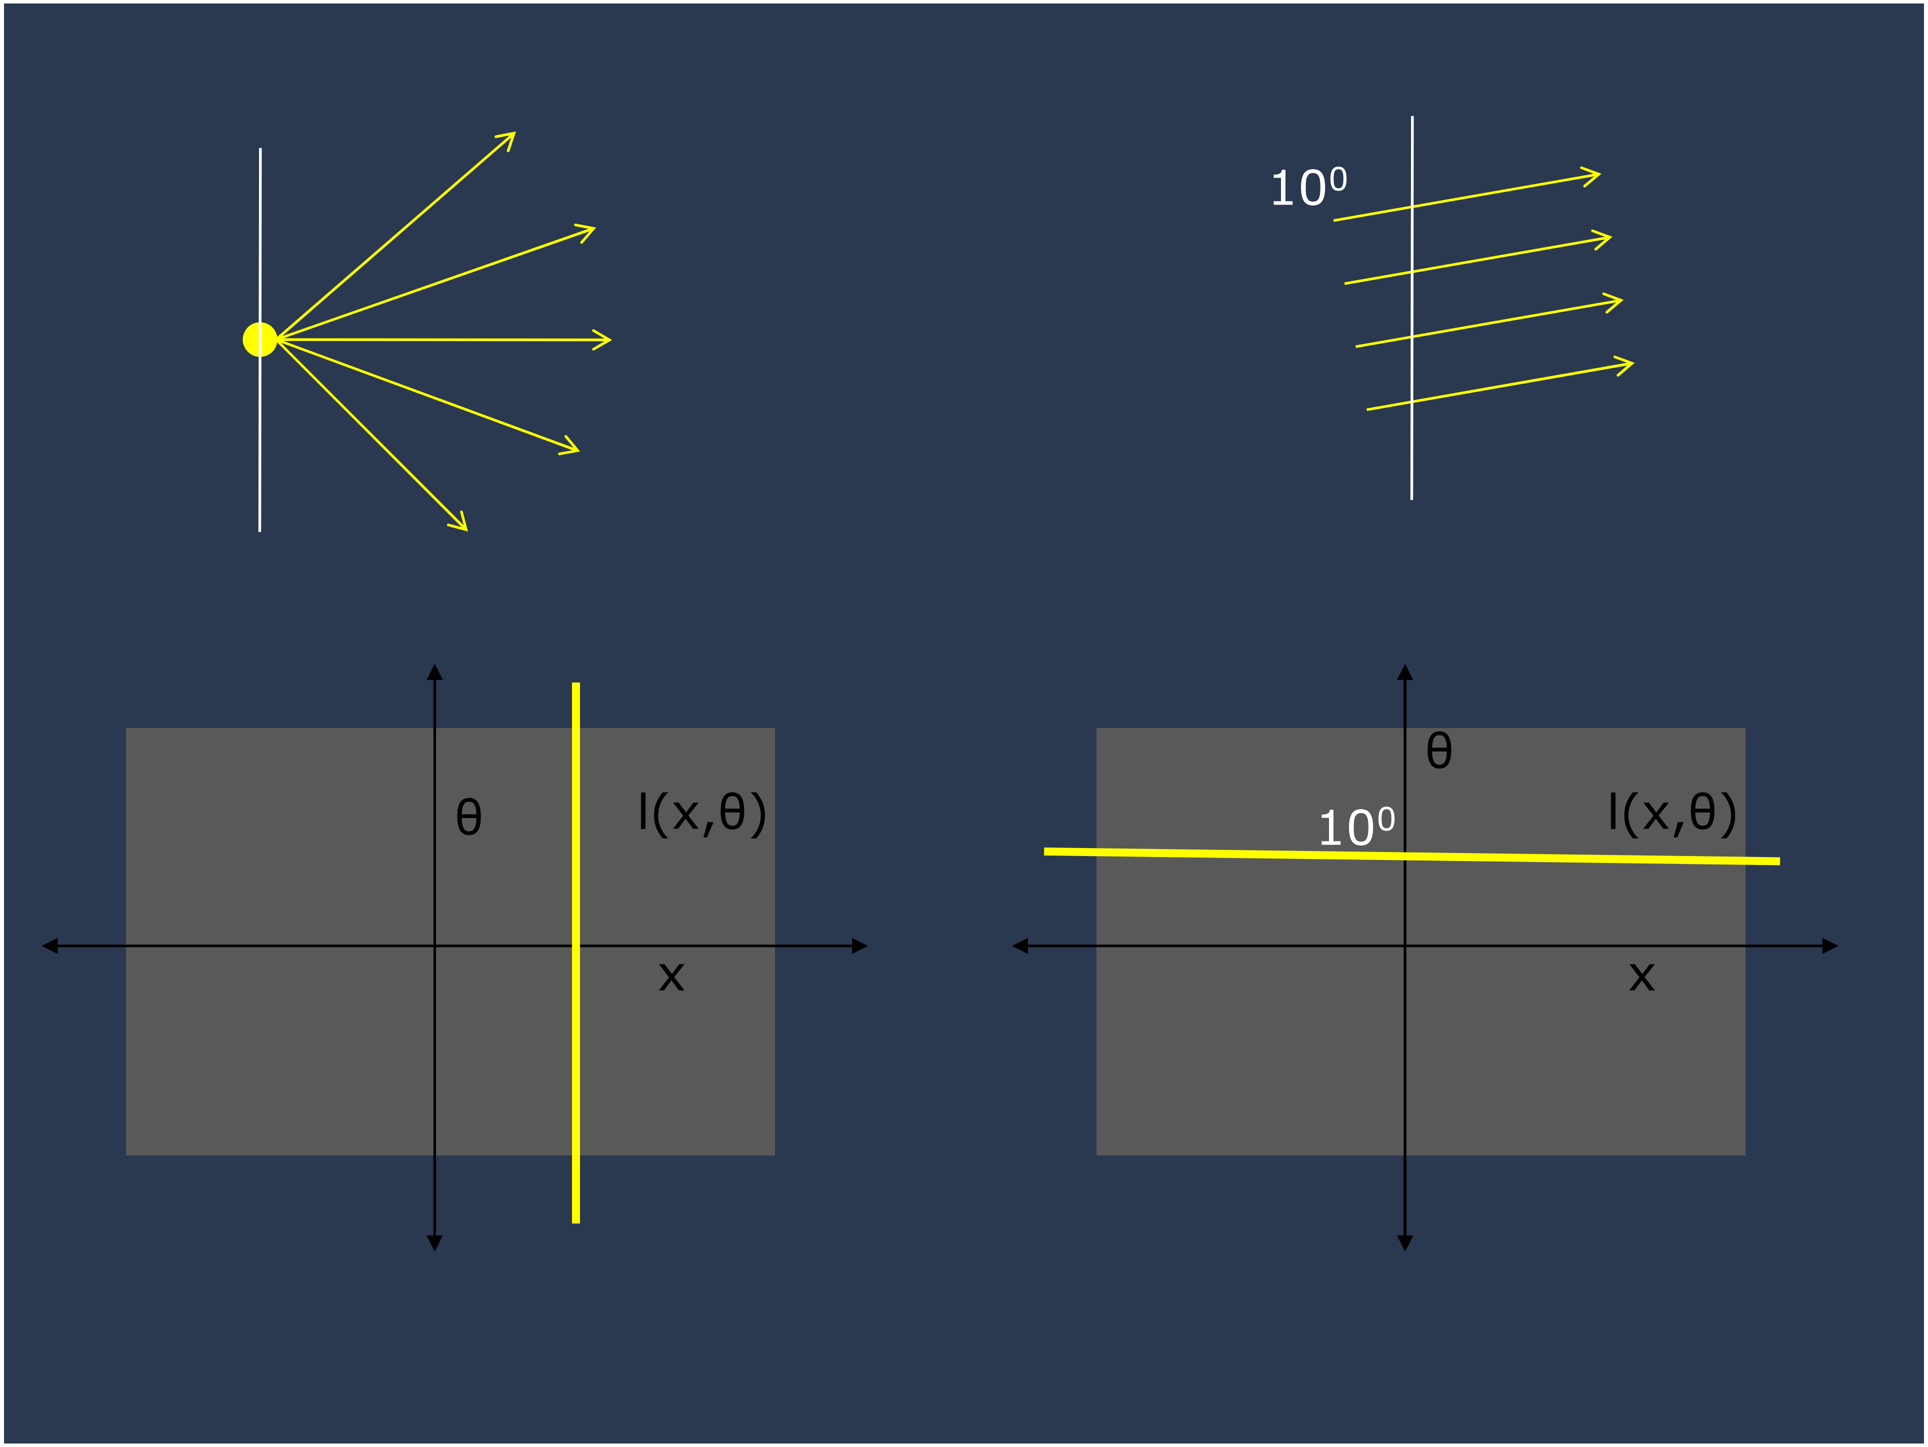
10^0
θ
I(x,θ)
x
θ
10^0
I(x,θ)
x

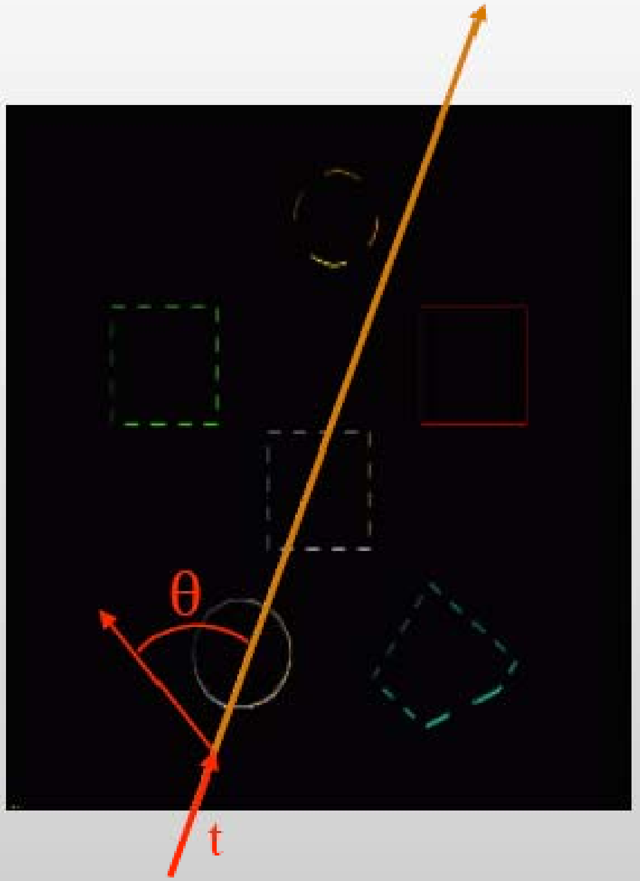

flight path through
a flatland scene

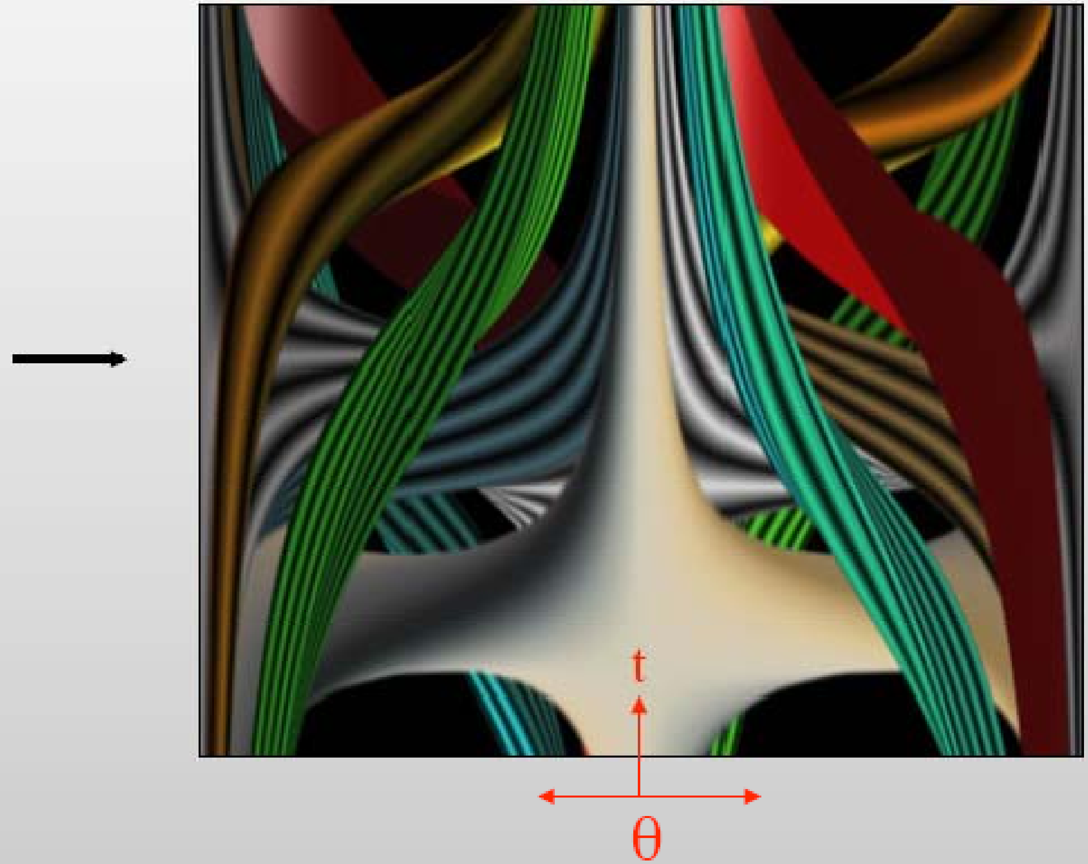

corresponding looming light field
(see also [Hasinoff 2006])

Lego gantry for capturing light fields
(built by Andrew Adams)

calibration point

plane + parallax [Vaish 2004]

Light Field = Array of (virtual) Cameras

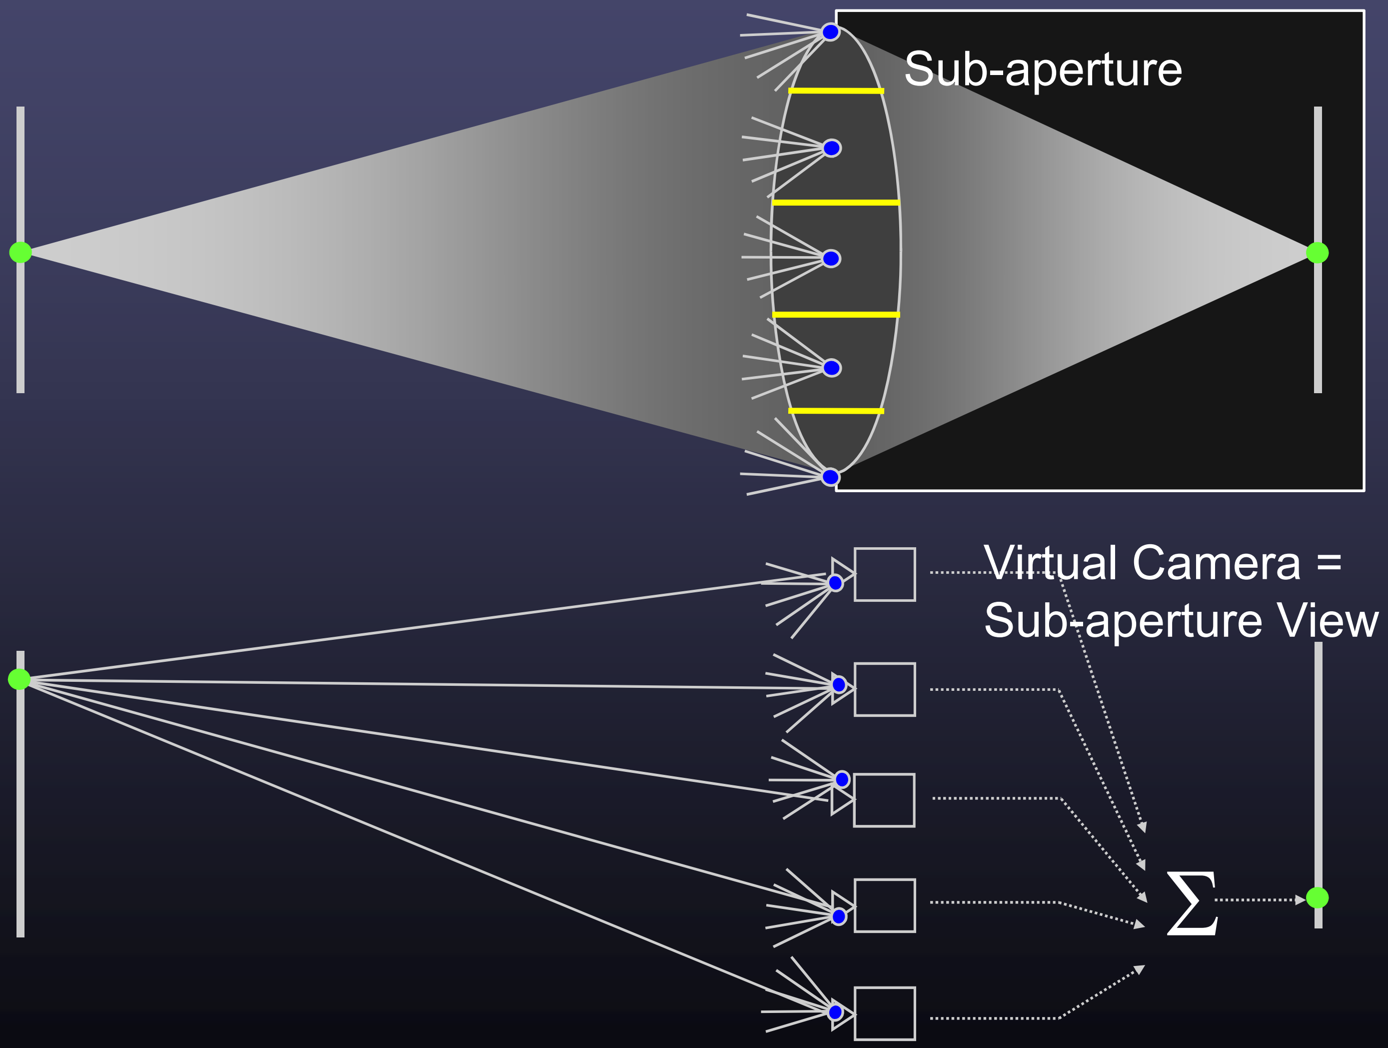

Three ways to capture LF inside a camera

- Shadows using pin-hole array
- Refraction using lenslet array
- Heterodyning using masks

Sub-Aperture = Pin-hole + Prism

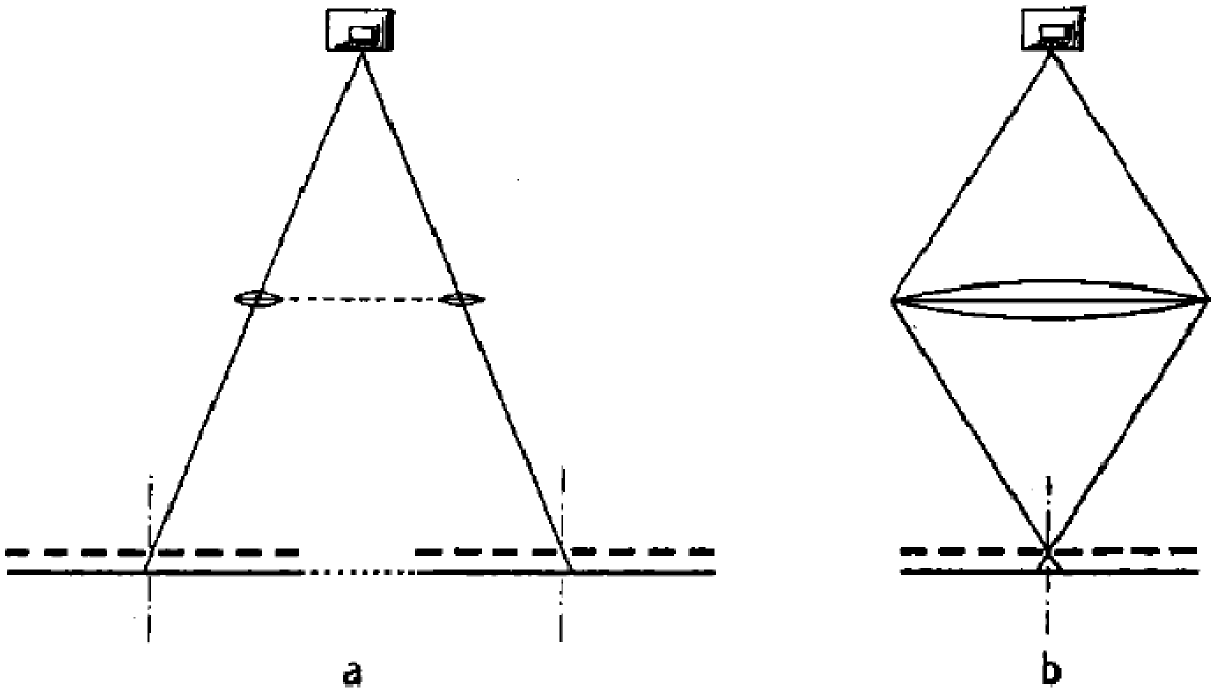

FIG. 1. *Two methods of making parallax panoramagram negatives. (a) A moving lens exposing a sensitive plate behind a grating slightly separated from it; lens, grating and plate being maintained in line during the exposure. (b) A large stationary lens, projecting an image on a stationary plate through a grating slightly separated from it.*

Ives 1933

vertical axis than is called for by the simple formula above developed. This correction, which is roughly proportional to the cosine of the angle between ww' and the sensitive surface, and so is of importance only for

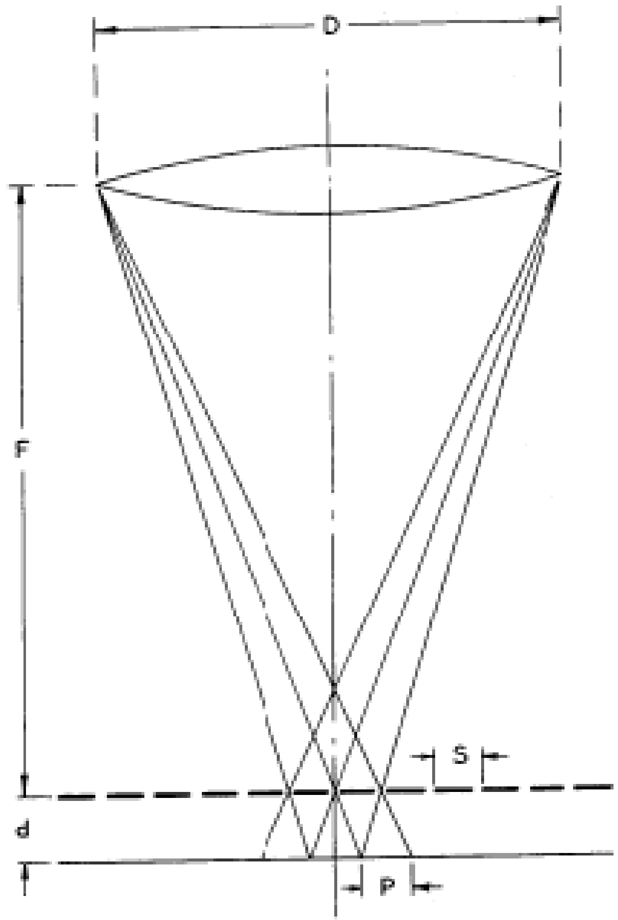

FIG. 4. *Determination of separation of grating and plate as function of grating spacing, lens diameter and focal distance.*

large angles, also varies with the angle of observation. A diameter of taking lens and size of picture can theoretically be attained such that this second order correction will fail. The slightly greater magnification of the viewing grating called for over the amount given by the

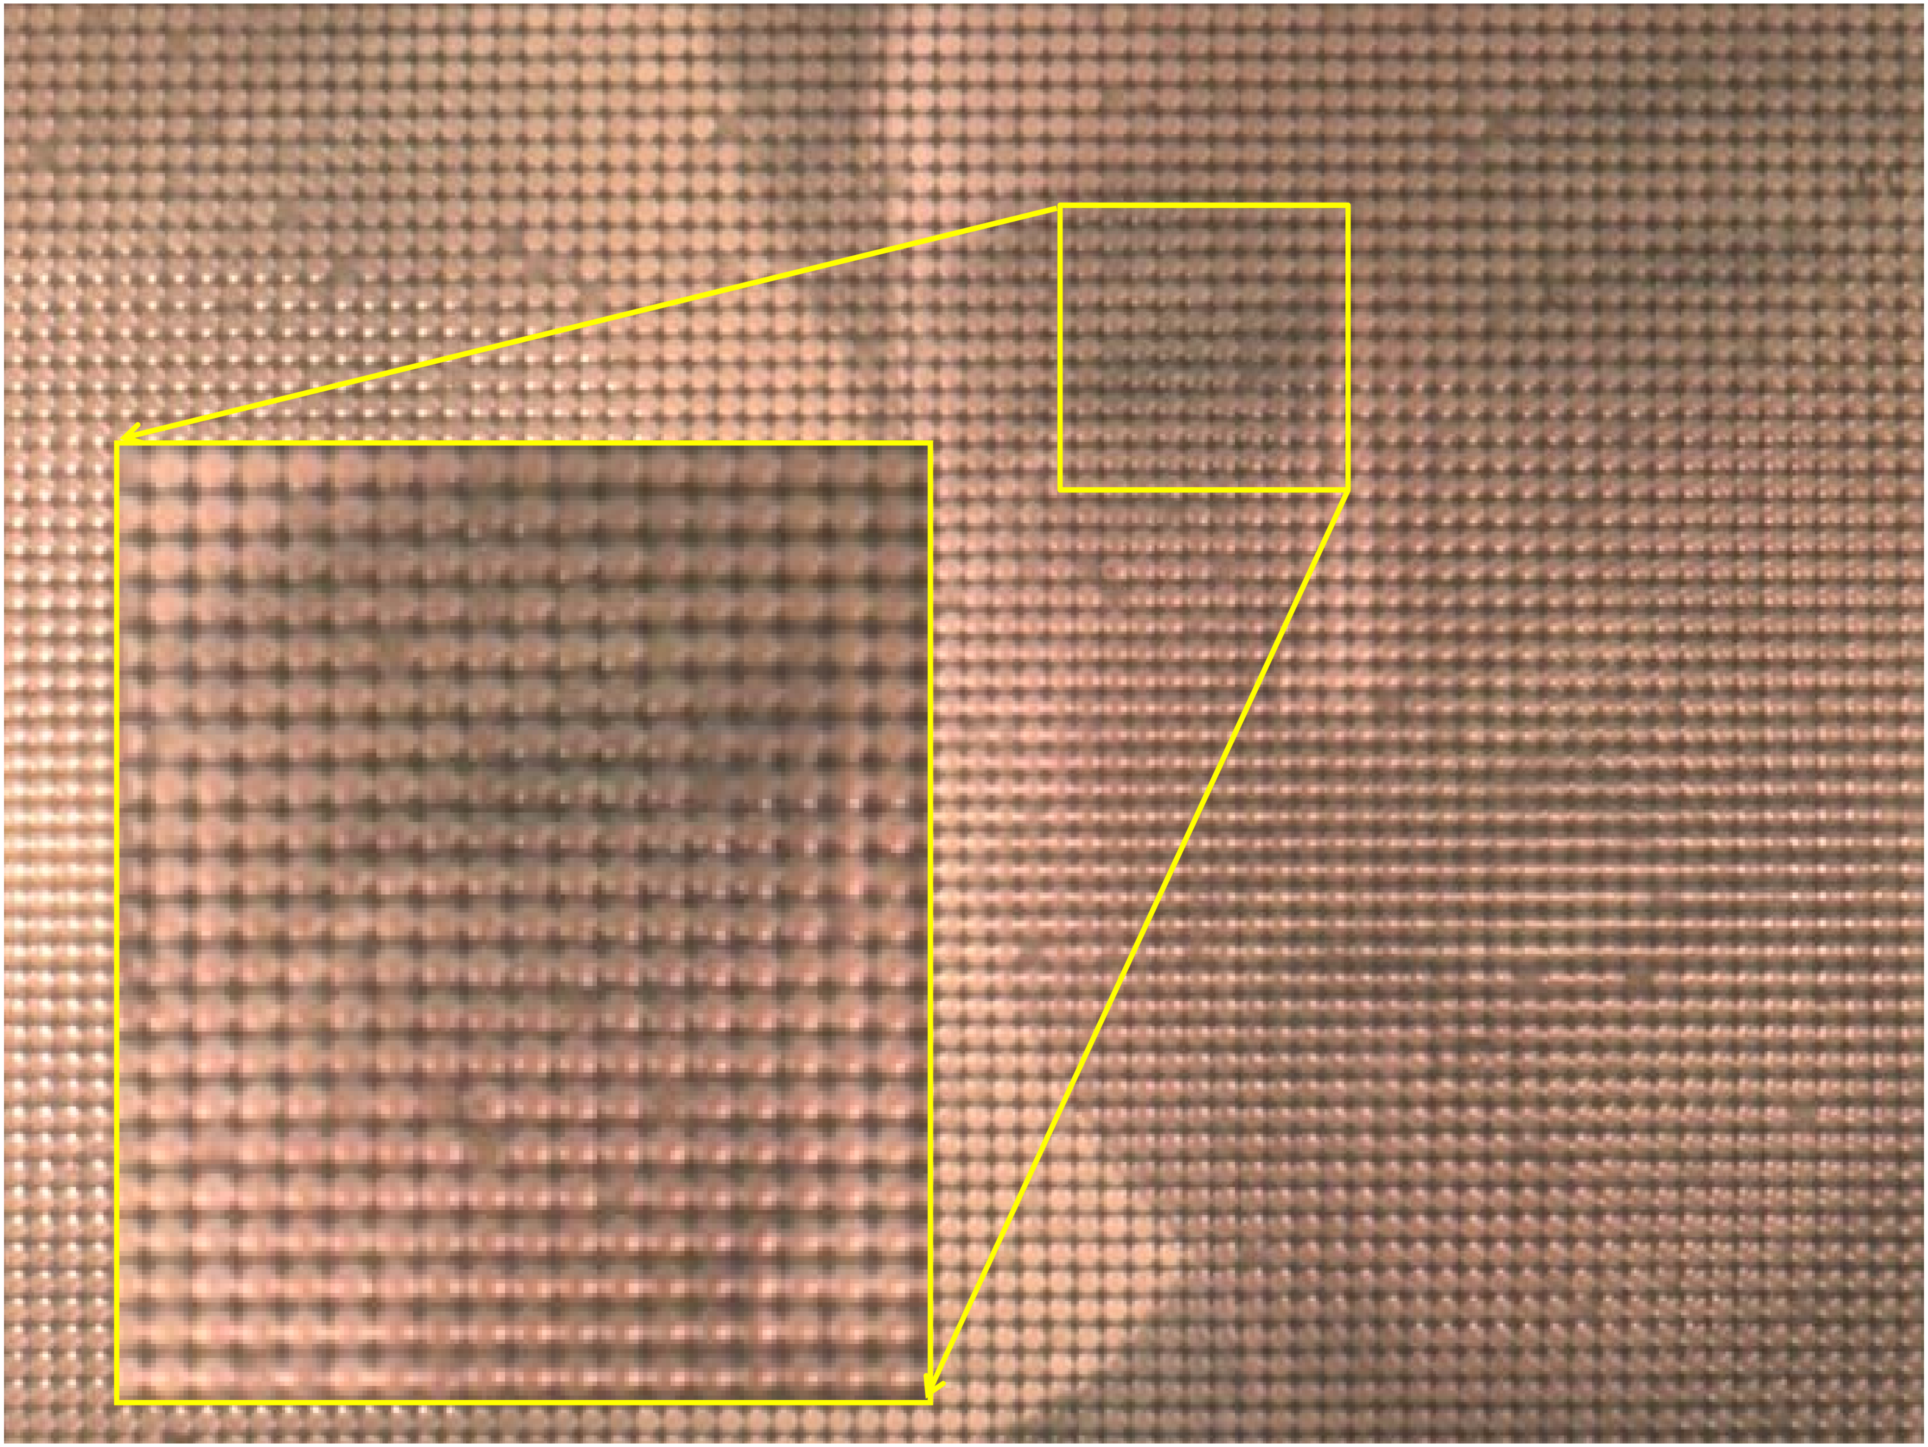

Lens Glare Reduction

[Raskar, Agrawal, Wilson, Veeraraghavan SIGGRAPH 2008]

Glare/Flare due to camera lenses reduces contrast

Reducing Glare

Conventional Photo

After removing outliers

Glare Reduced Image

Glare = low frequency noise in 2D

- But is high frequency noise in 4D
- Remove via simple outlier rejection

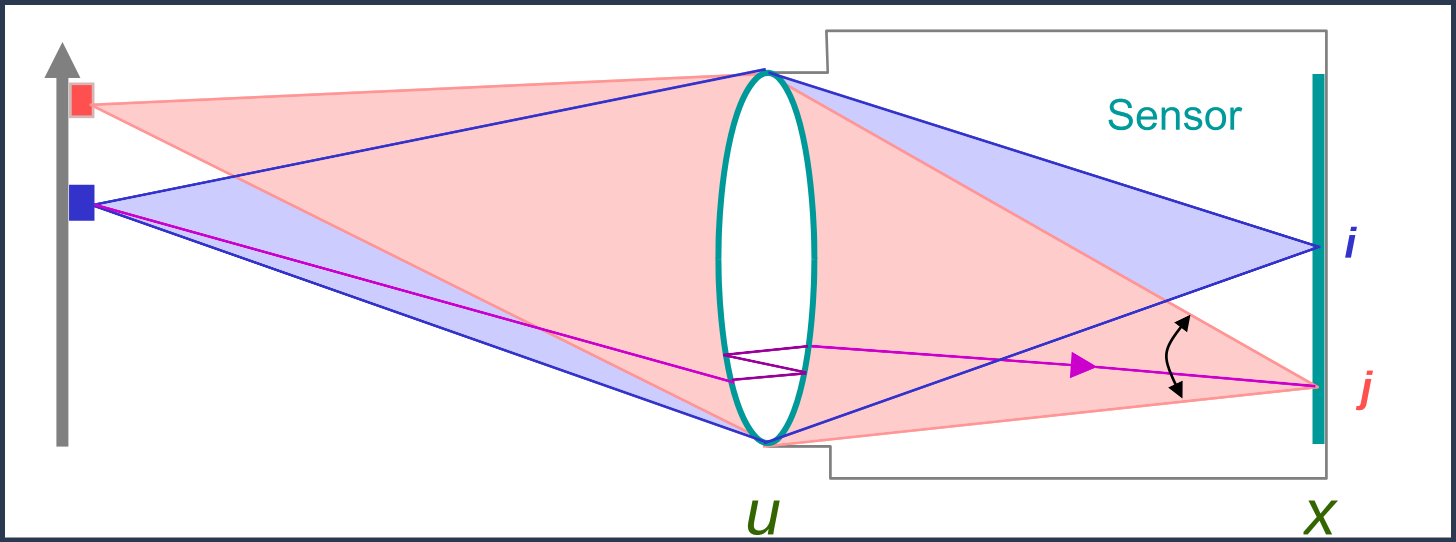

Enhancing Glare

Conventional Photo

Glare Enhanced Image

Glare due to Lens Inter-Reflections

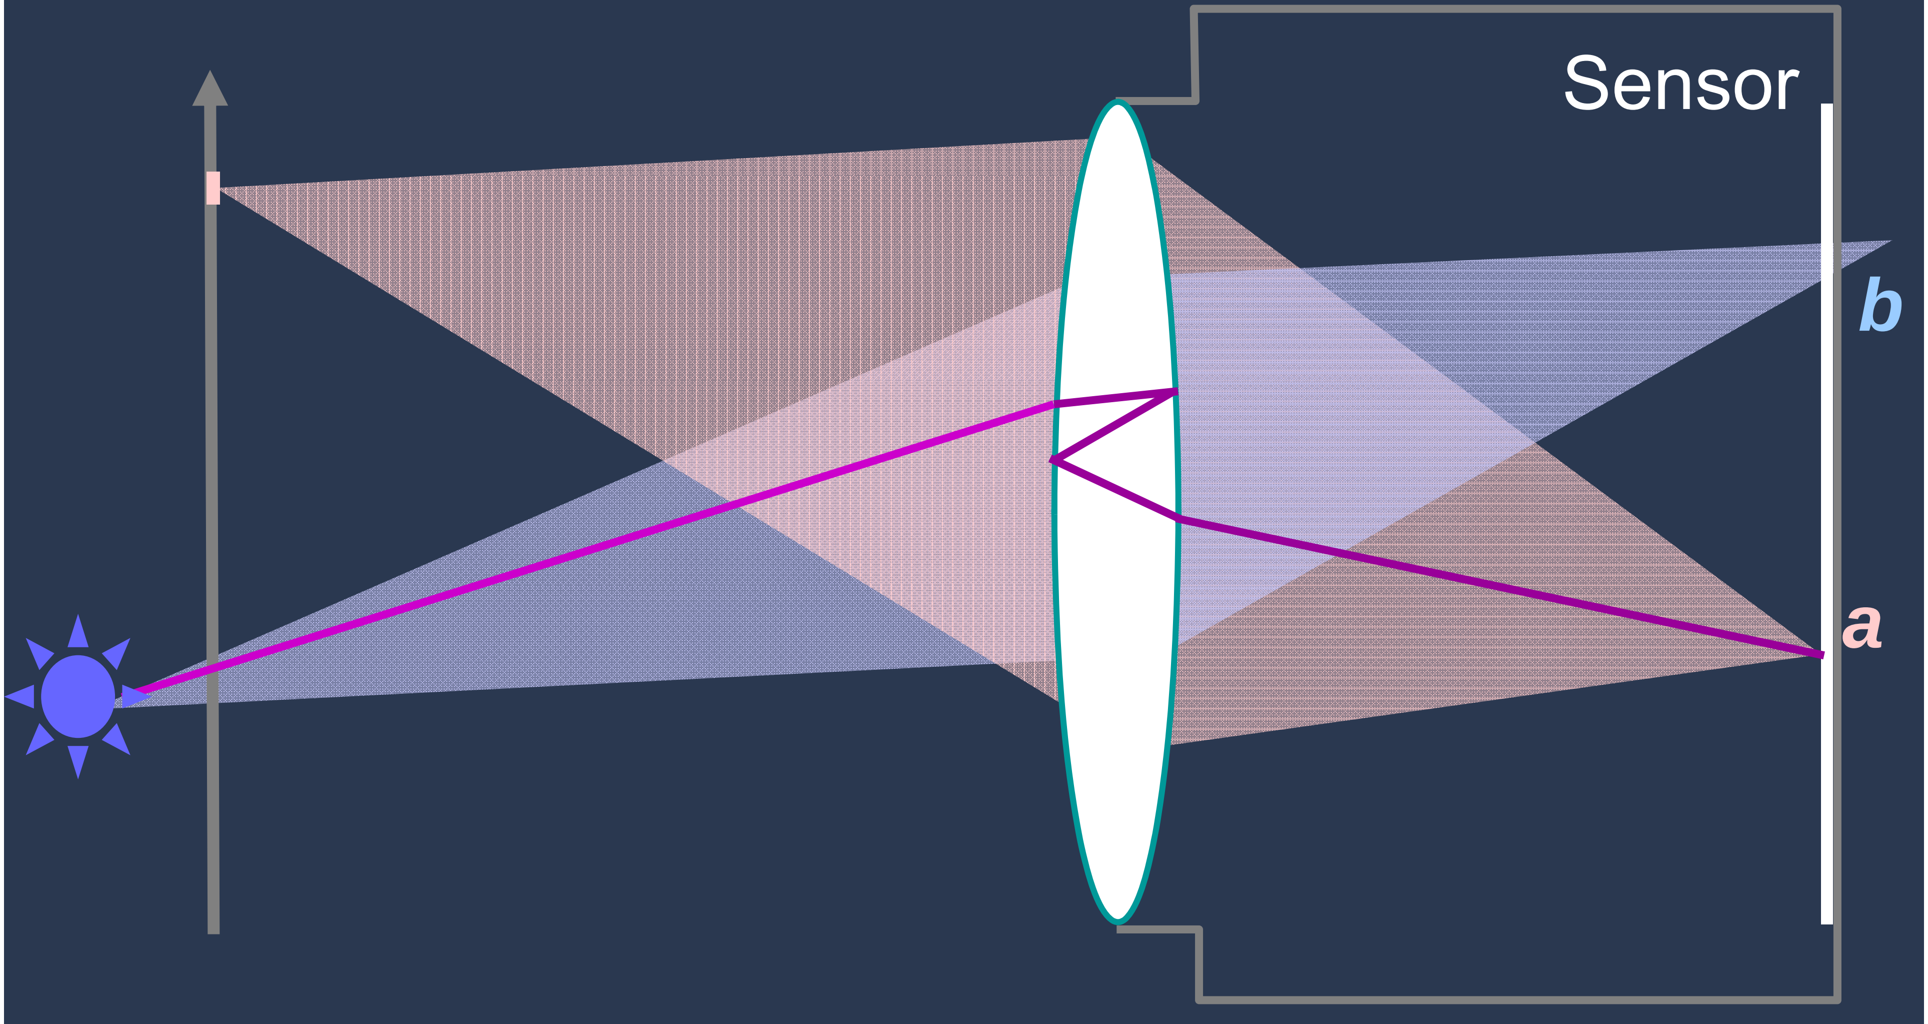

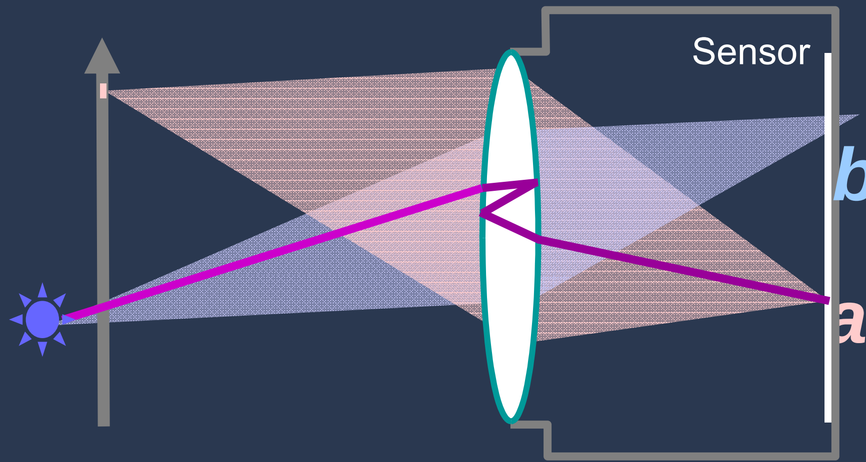
Sensor
b
a

Effects of Glare on Image

- Hard to model, Low Frequency in 2D
- But reflection glare is outlier in 4D ray-space

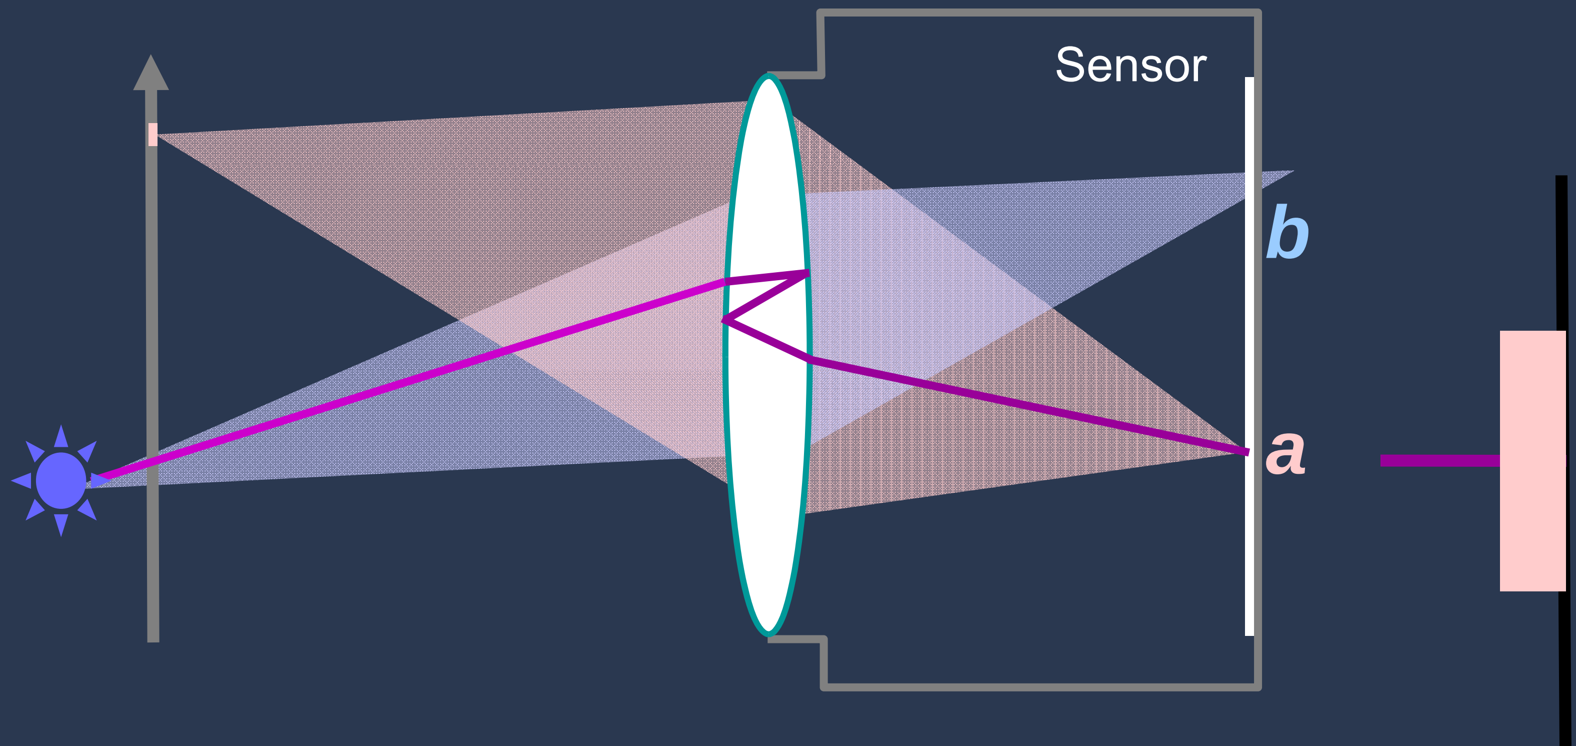

Lens Inter-reflections

Key Idea

- Lens Glare manifests as low frequency in 2D Image
- But Glare is highly view dependent
  - manifests as outliers in 4D ray-space
- Reducing Glare == Remove outliers among rays

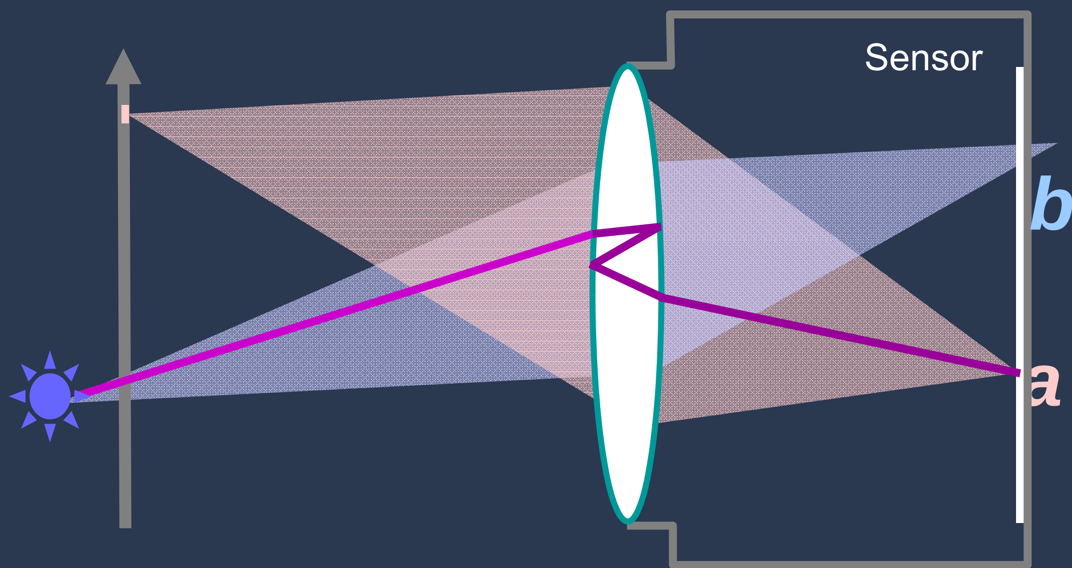

Light Field Inside a Camera

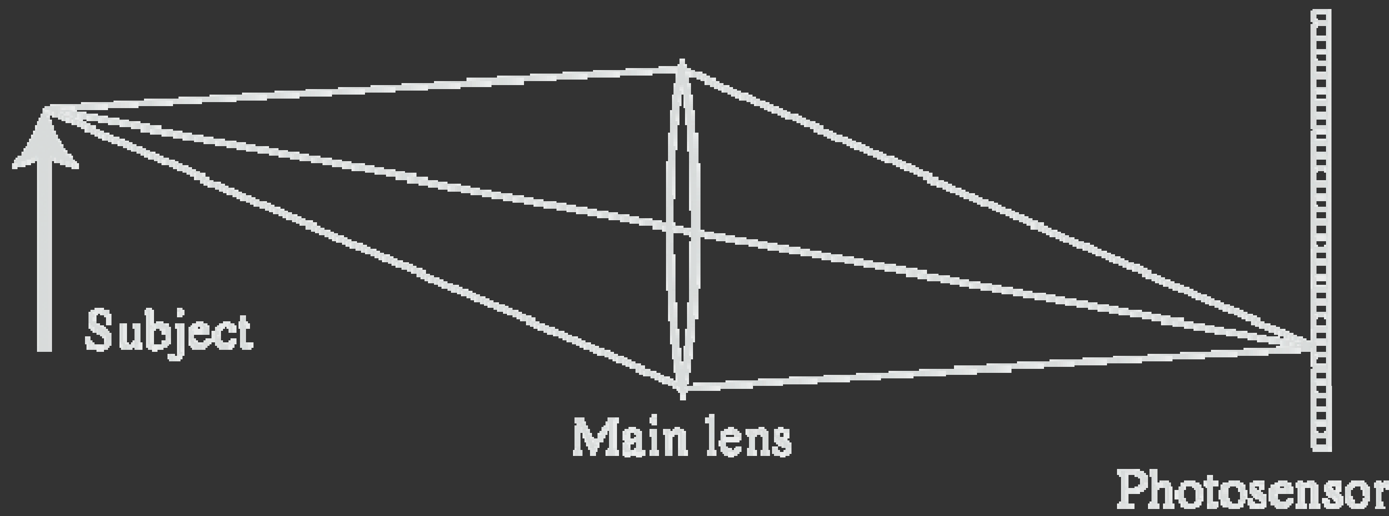

Lenslet-based Light Field camera

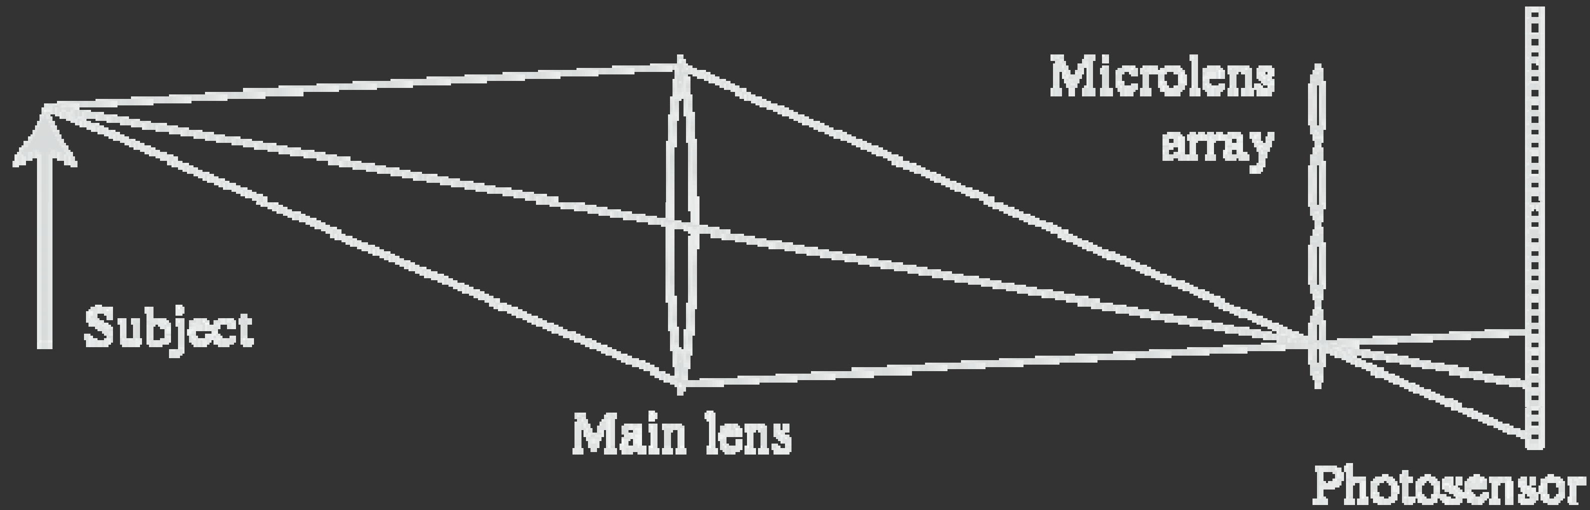

[Adelson and Wang, 1992, Ng et al. 2005]

Prototype camera

Contax medium format camera

Kodak 16-megapixel sensor

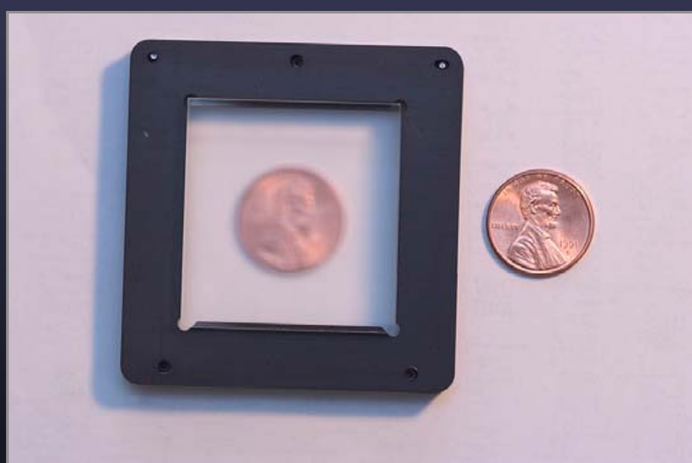

Adaptive Optics microlens array

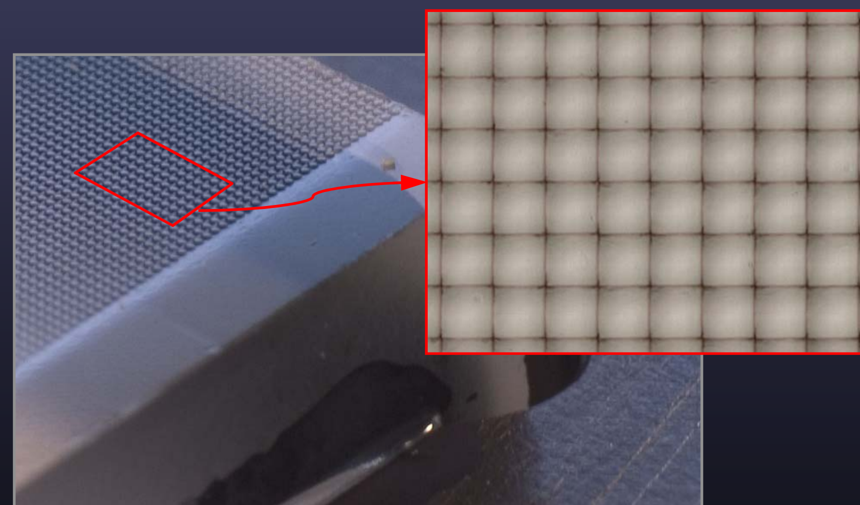

125µ square-sided microlenses

4000 × 4000 pixels ÷ 292 × 292 lenses = 14 × 14 pixels per lens

[Fife 2008]

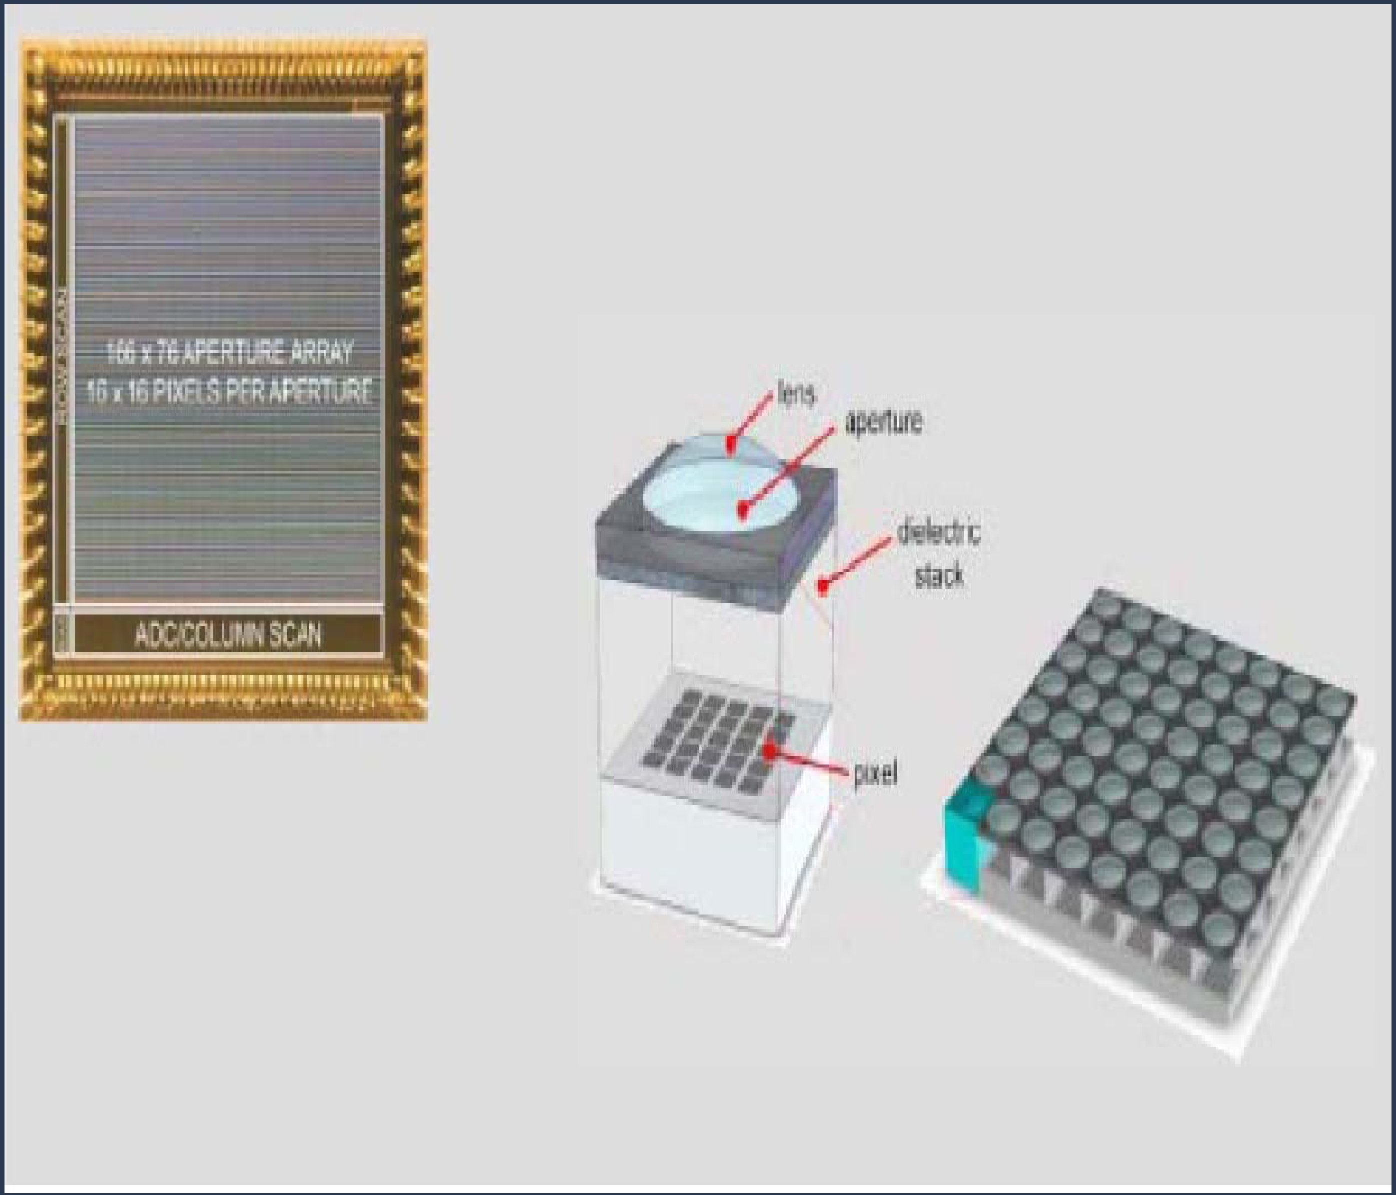

166 x 76 APERTURE ARRAY
16 x 16 PIXELS PER APERTURE
ADC/COLUMN SCAN
ROW SCAN
lens
aperture
dielectric
stack
pixel

Zooming into the raw photo

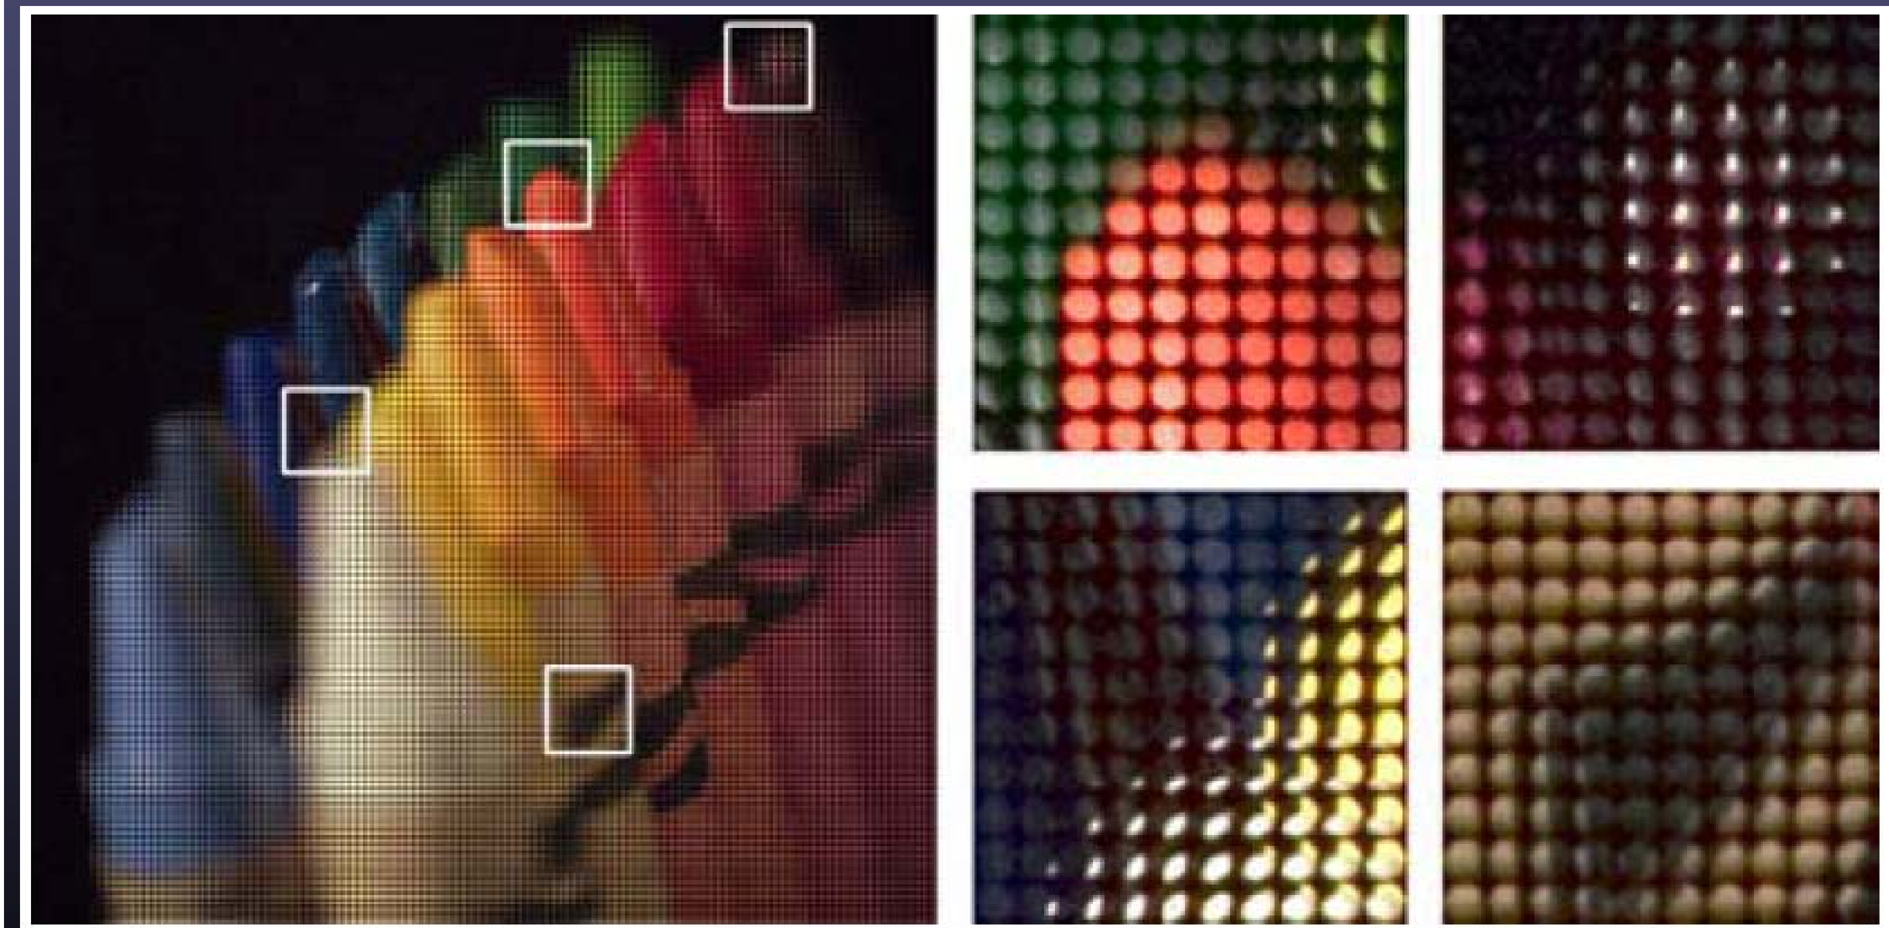

Digital Refocusing

[Ng et al 2005]

Can we achieve this with a Mask alone?

Mask based Light Field Camera

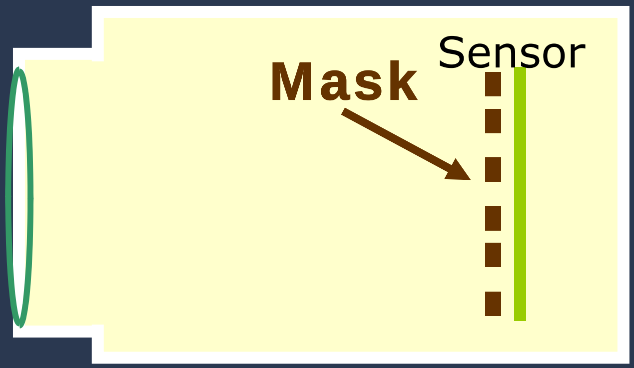

[Veeraraghavan, Raskar, Agrawal, Tumblin, Mohan, Siggraph 2007]

How to Capture 4D Light Field with 2D Sensor ?

What should be the pattern of the mask ?

Lens Copies the Lightfield of Conjugate Plane

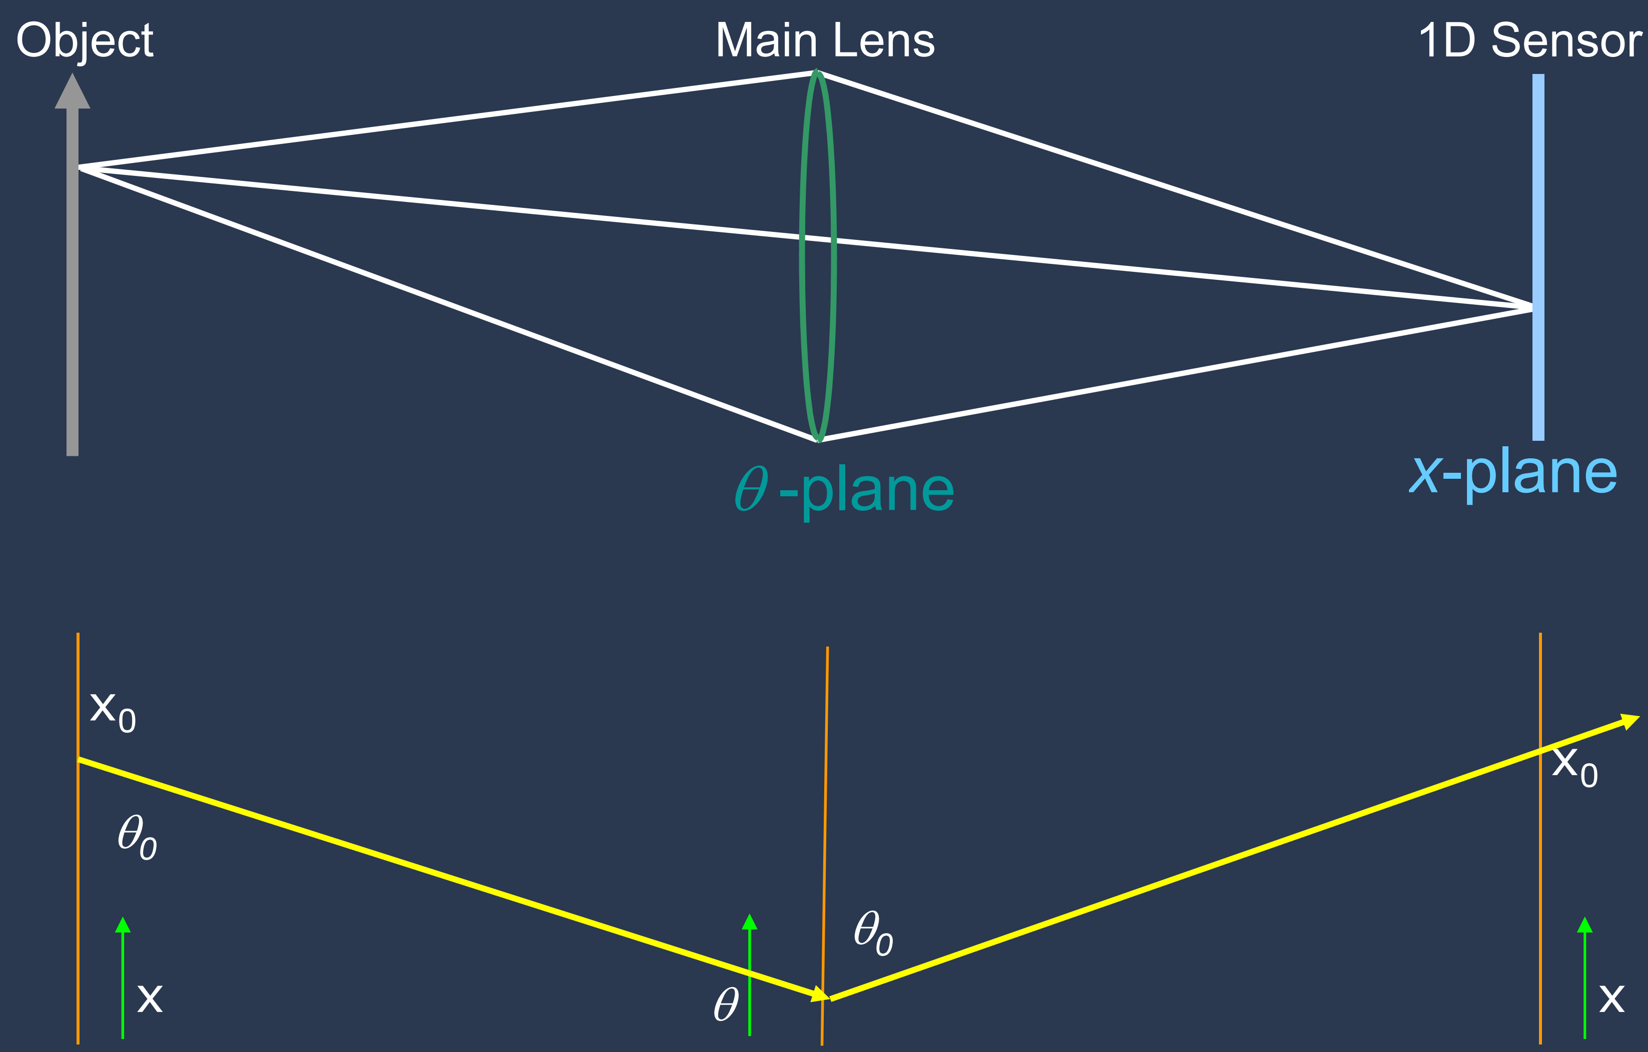

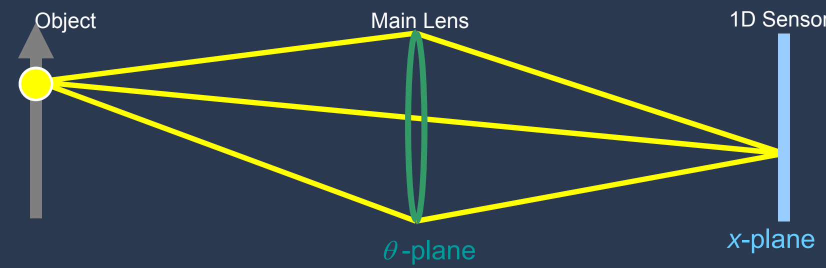

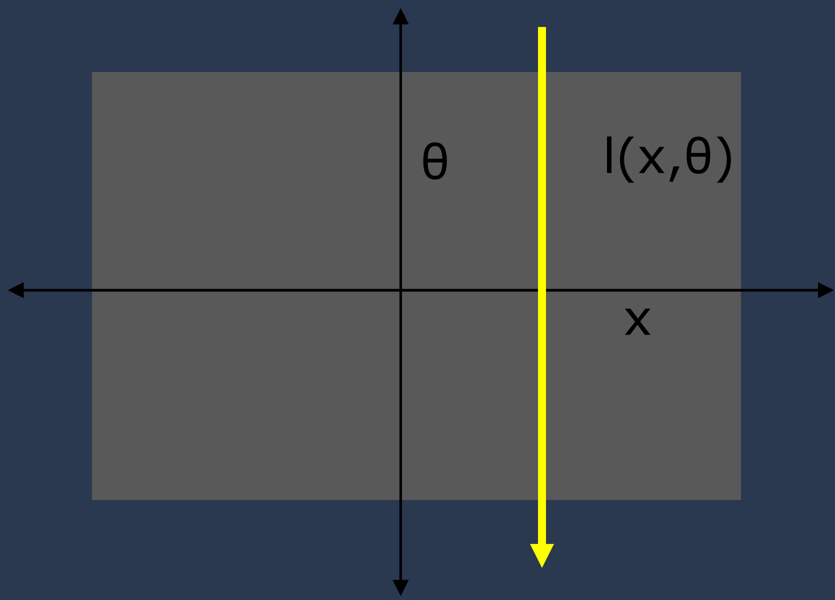

Line Integral

Captured
Photo

Fourier Slice Theorem

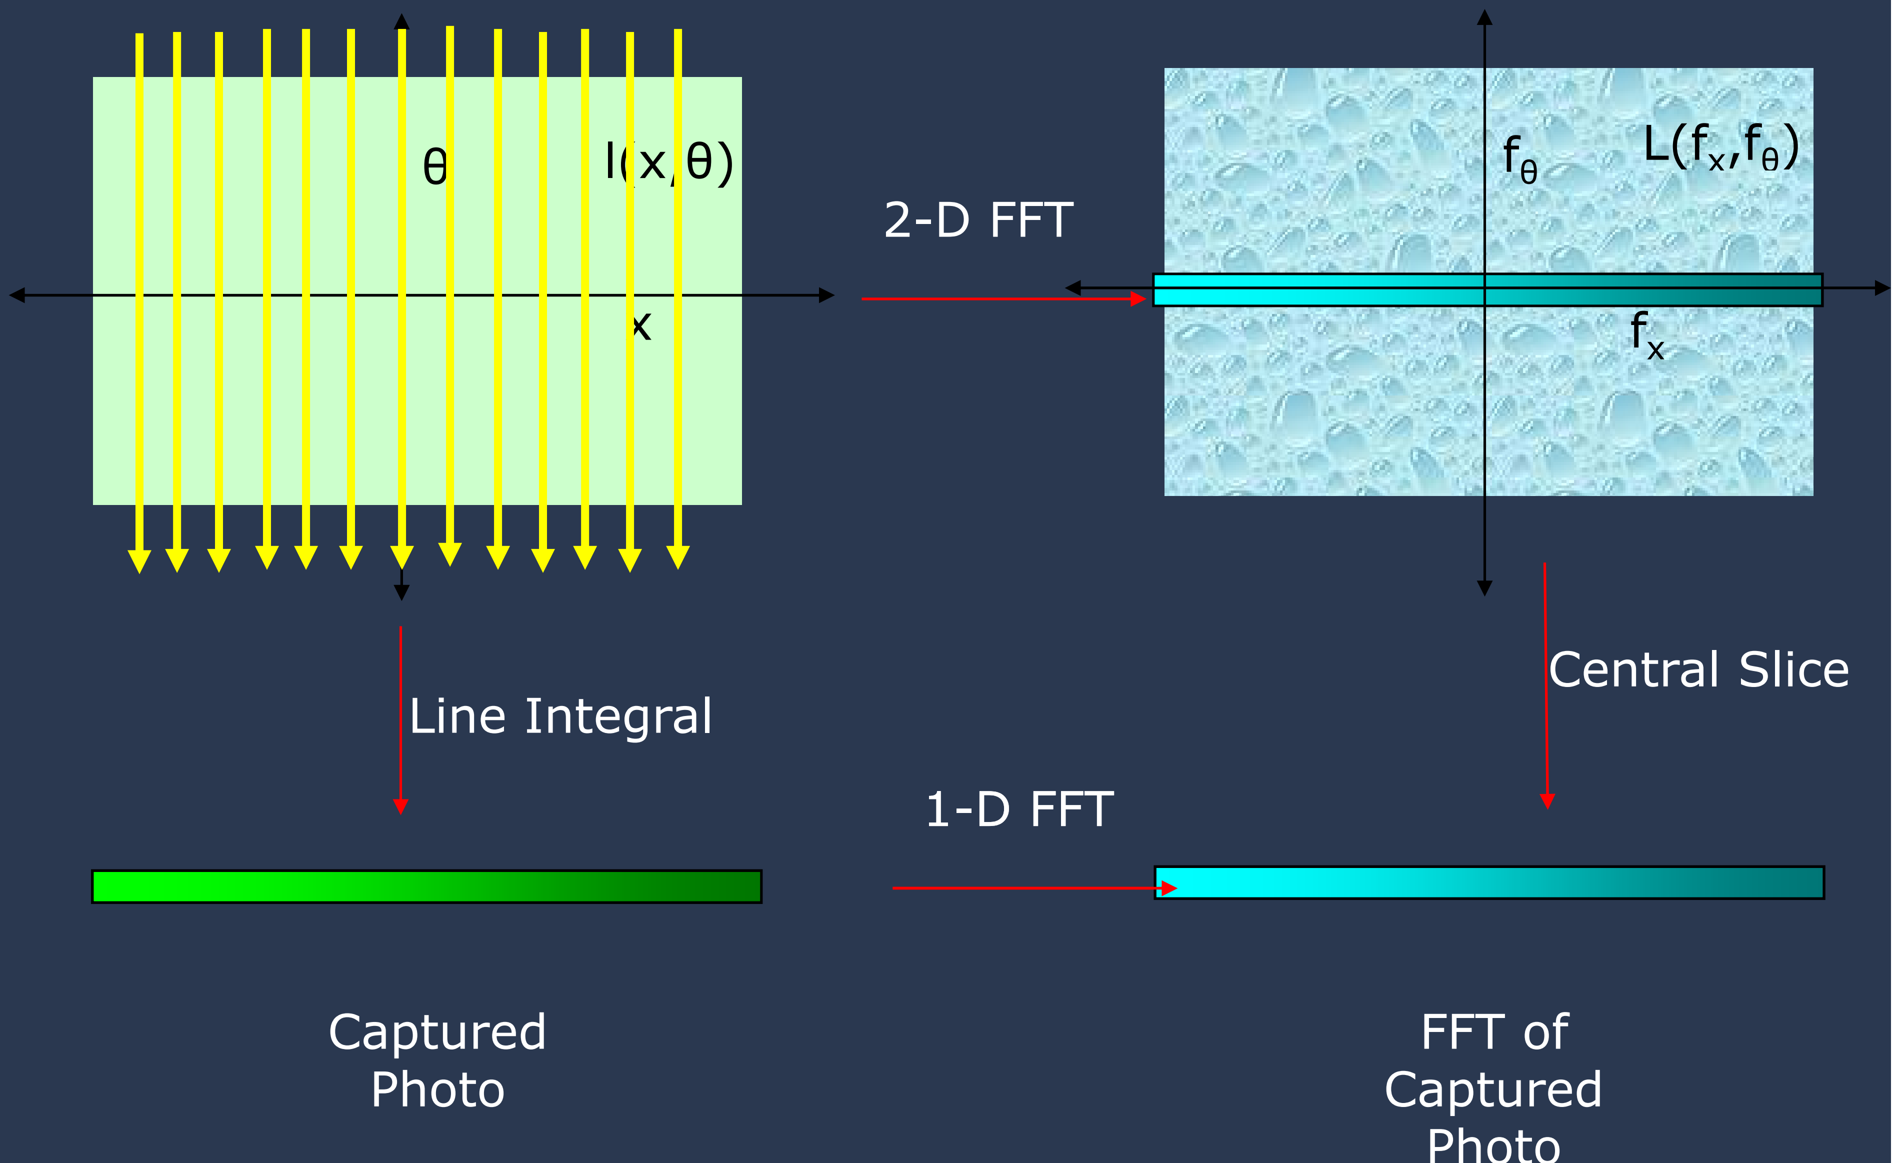

Light Propagation (Defocus Blur)

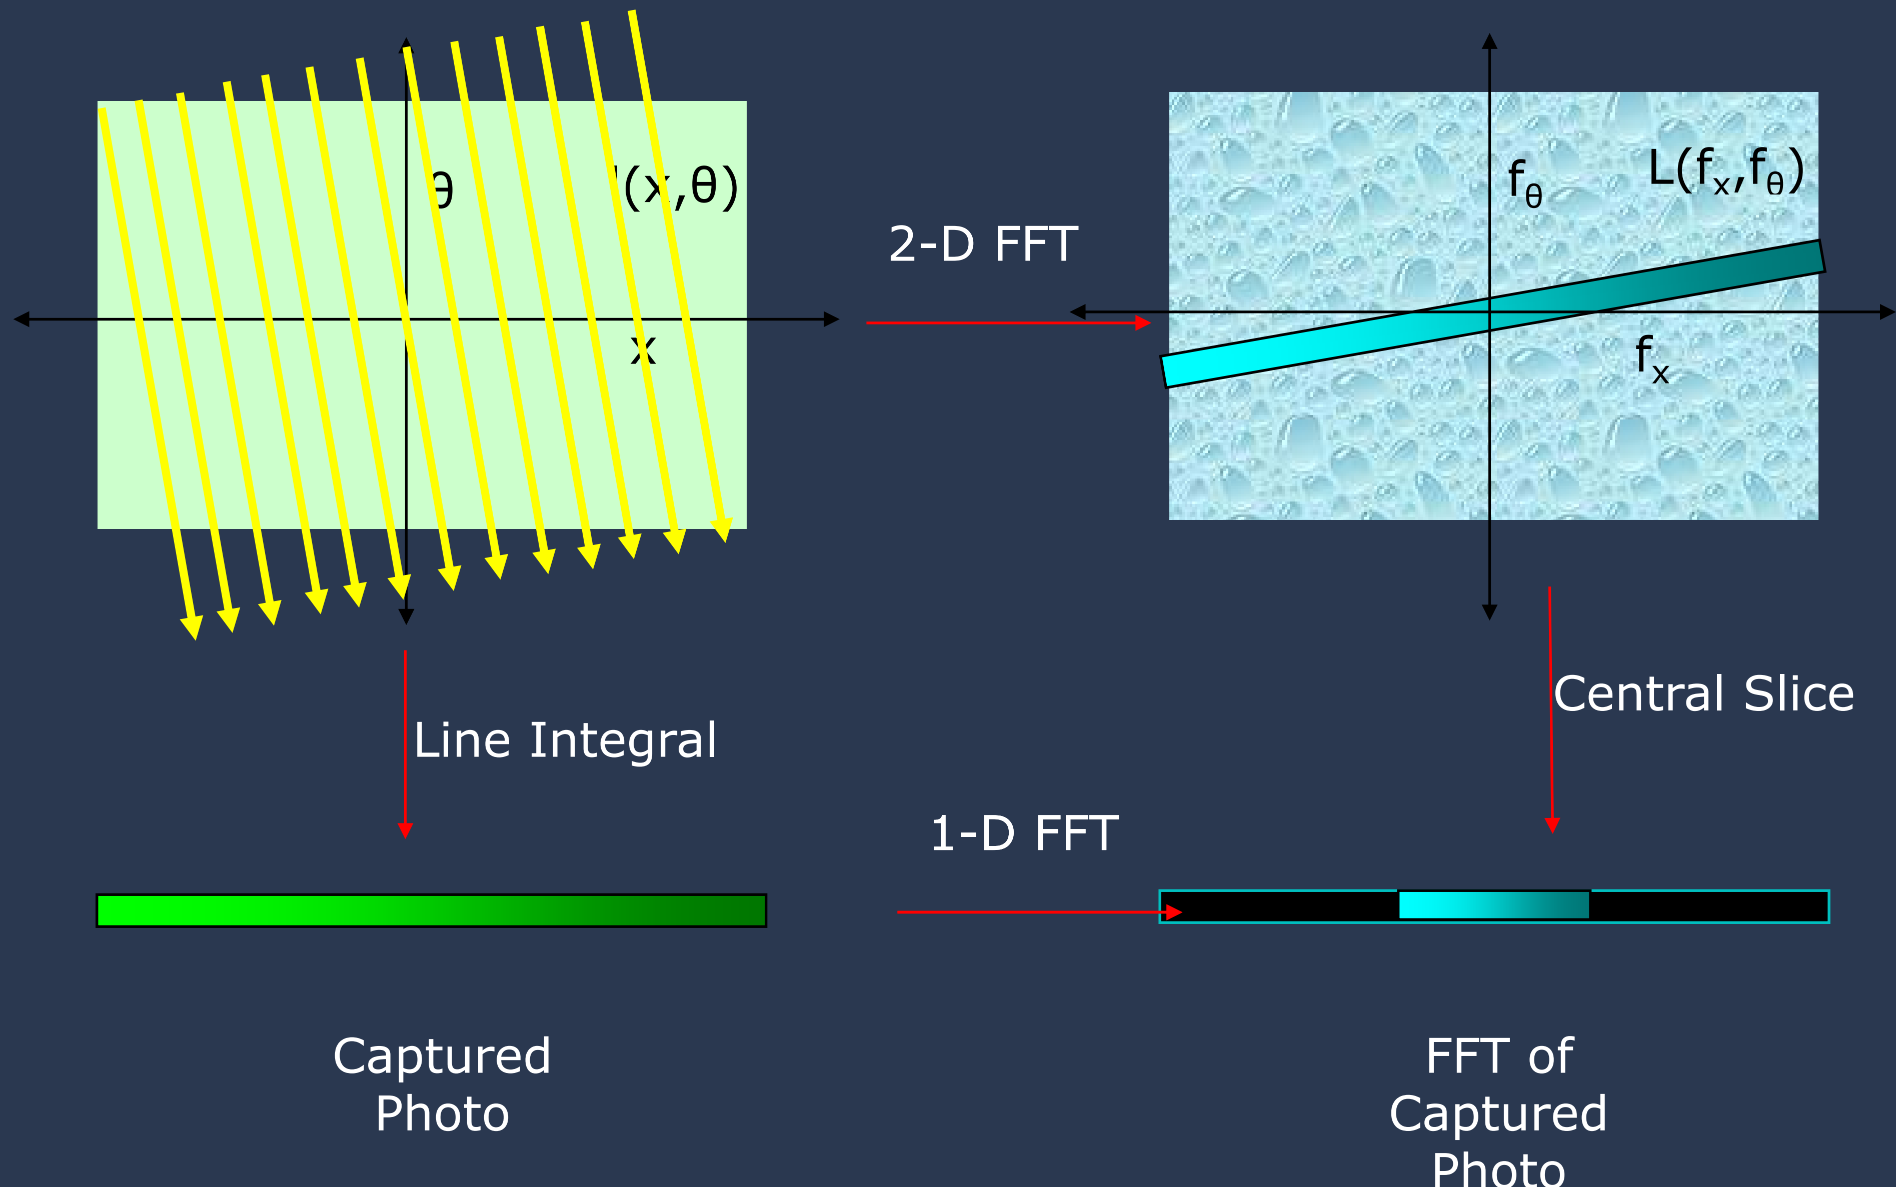

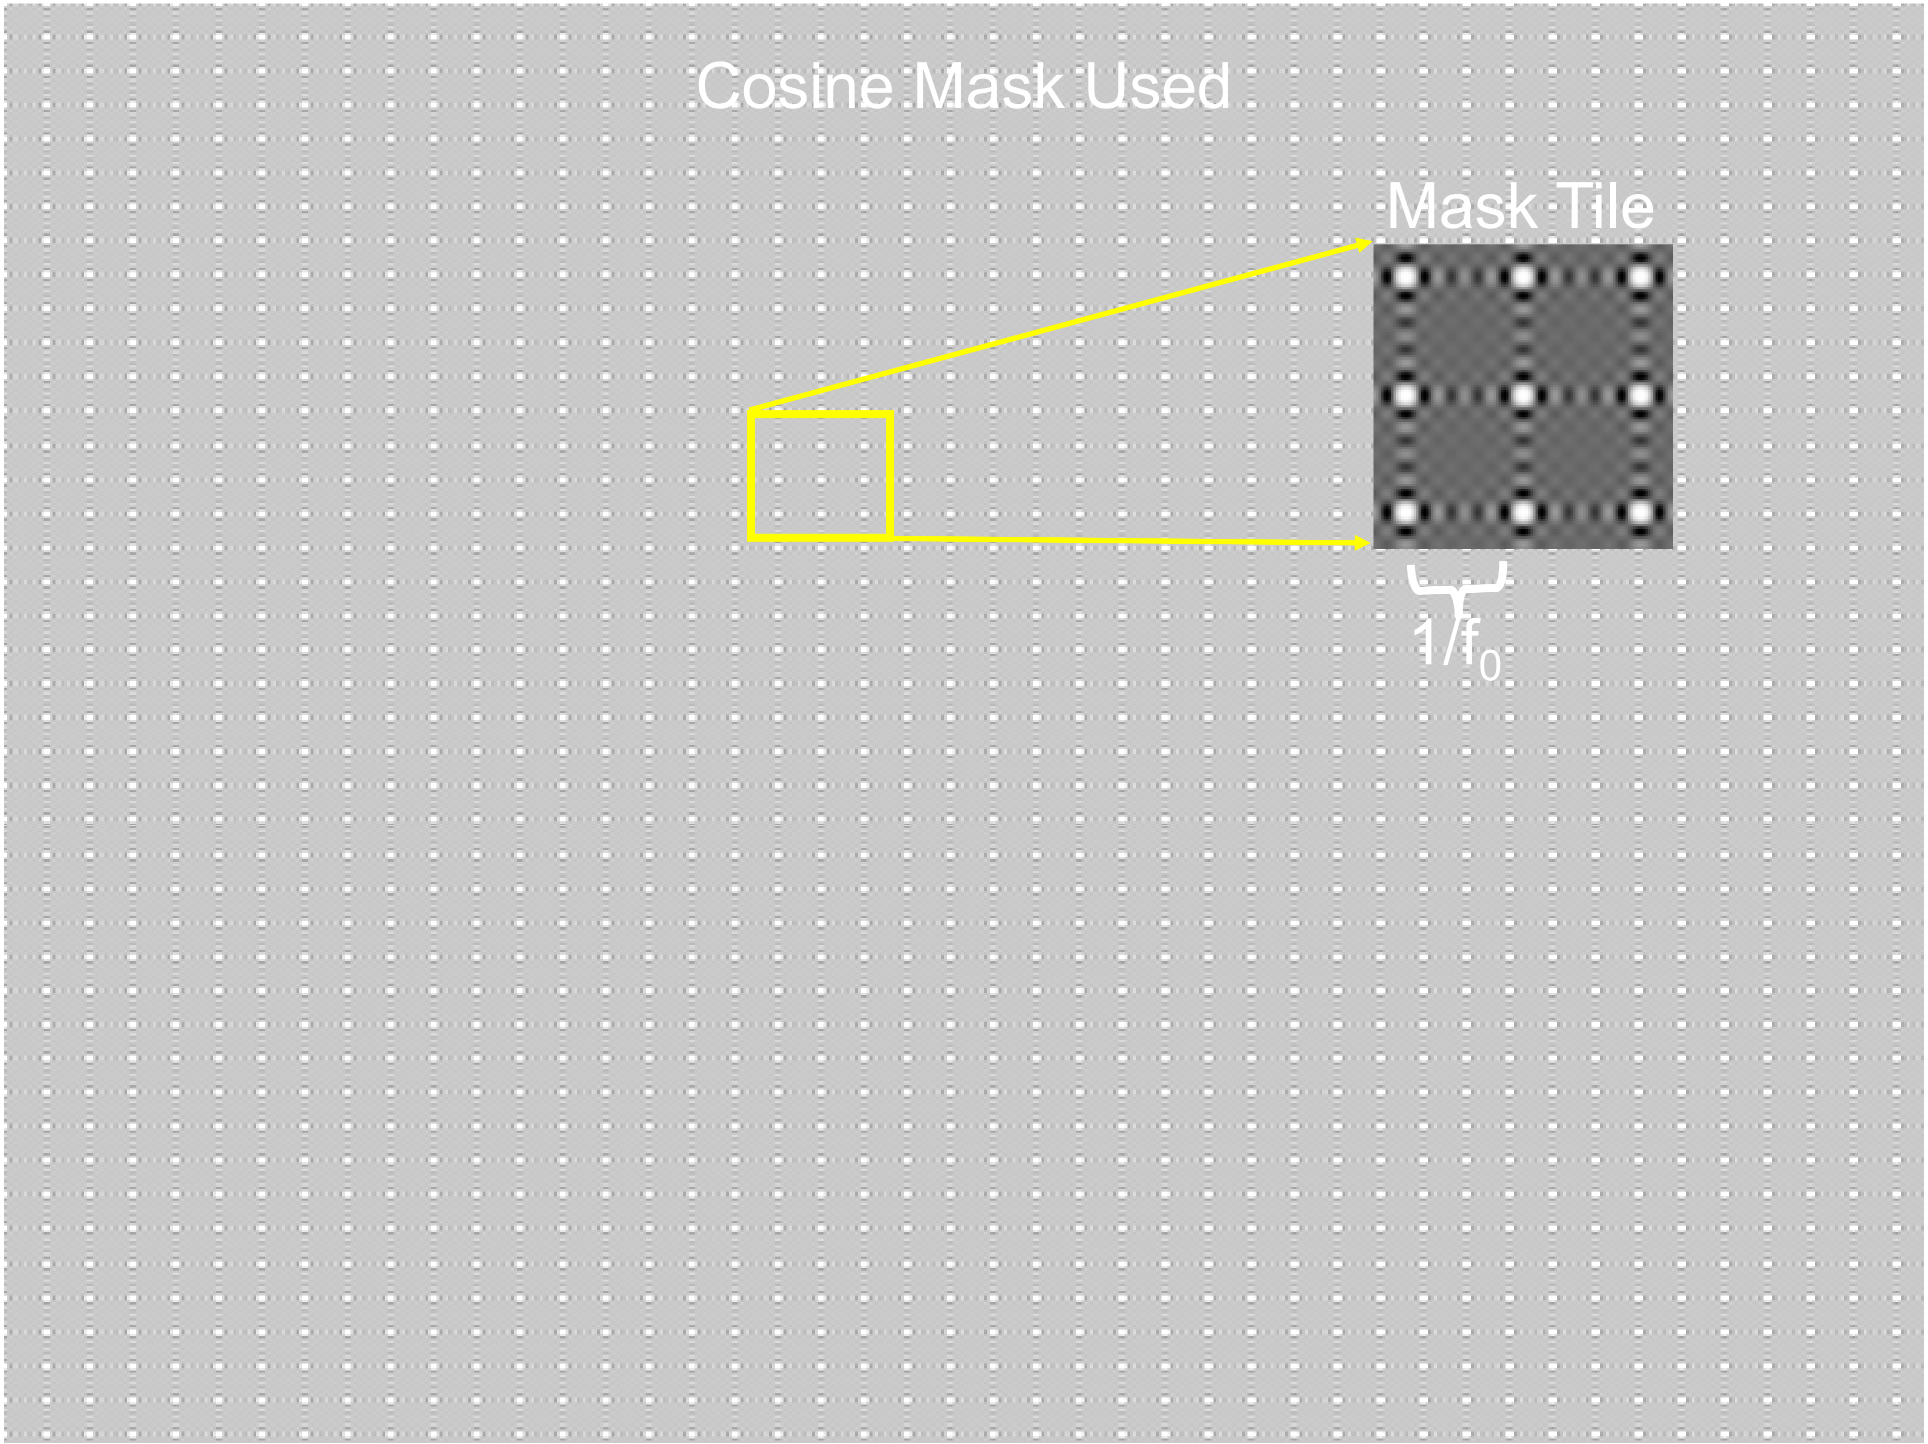
Cosine Mask Used
Mask Tile
1/f0

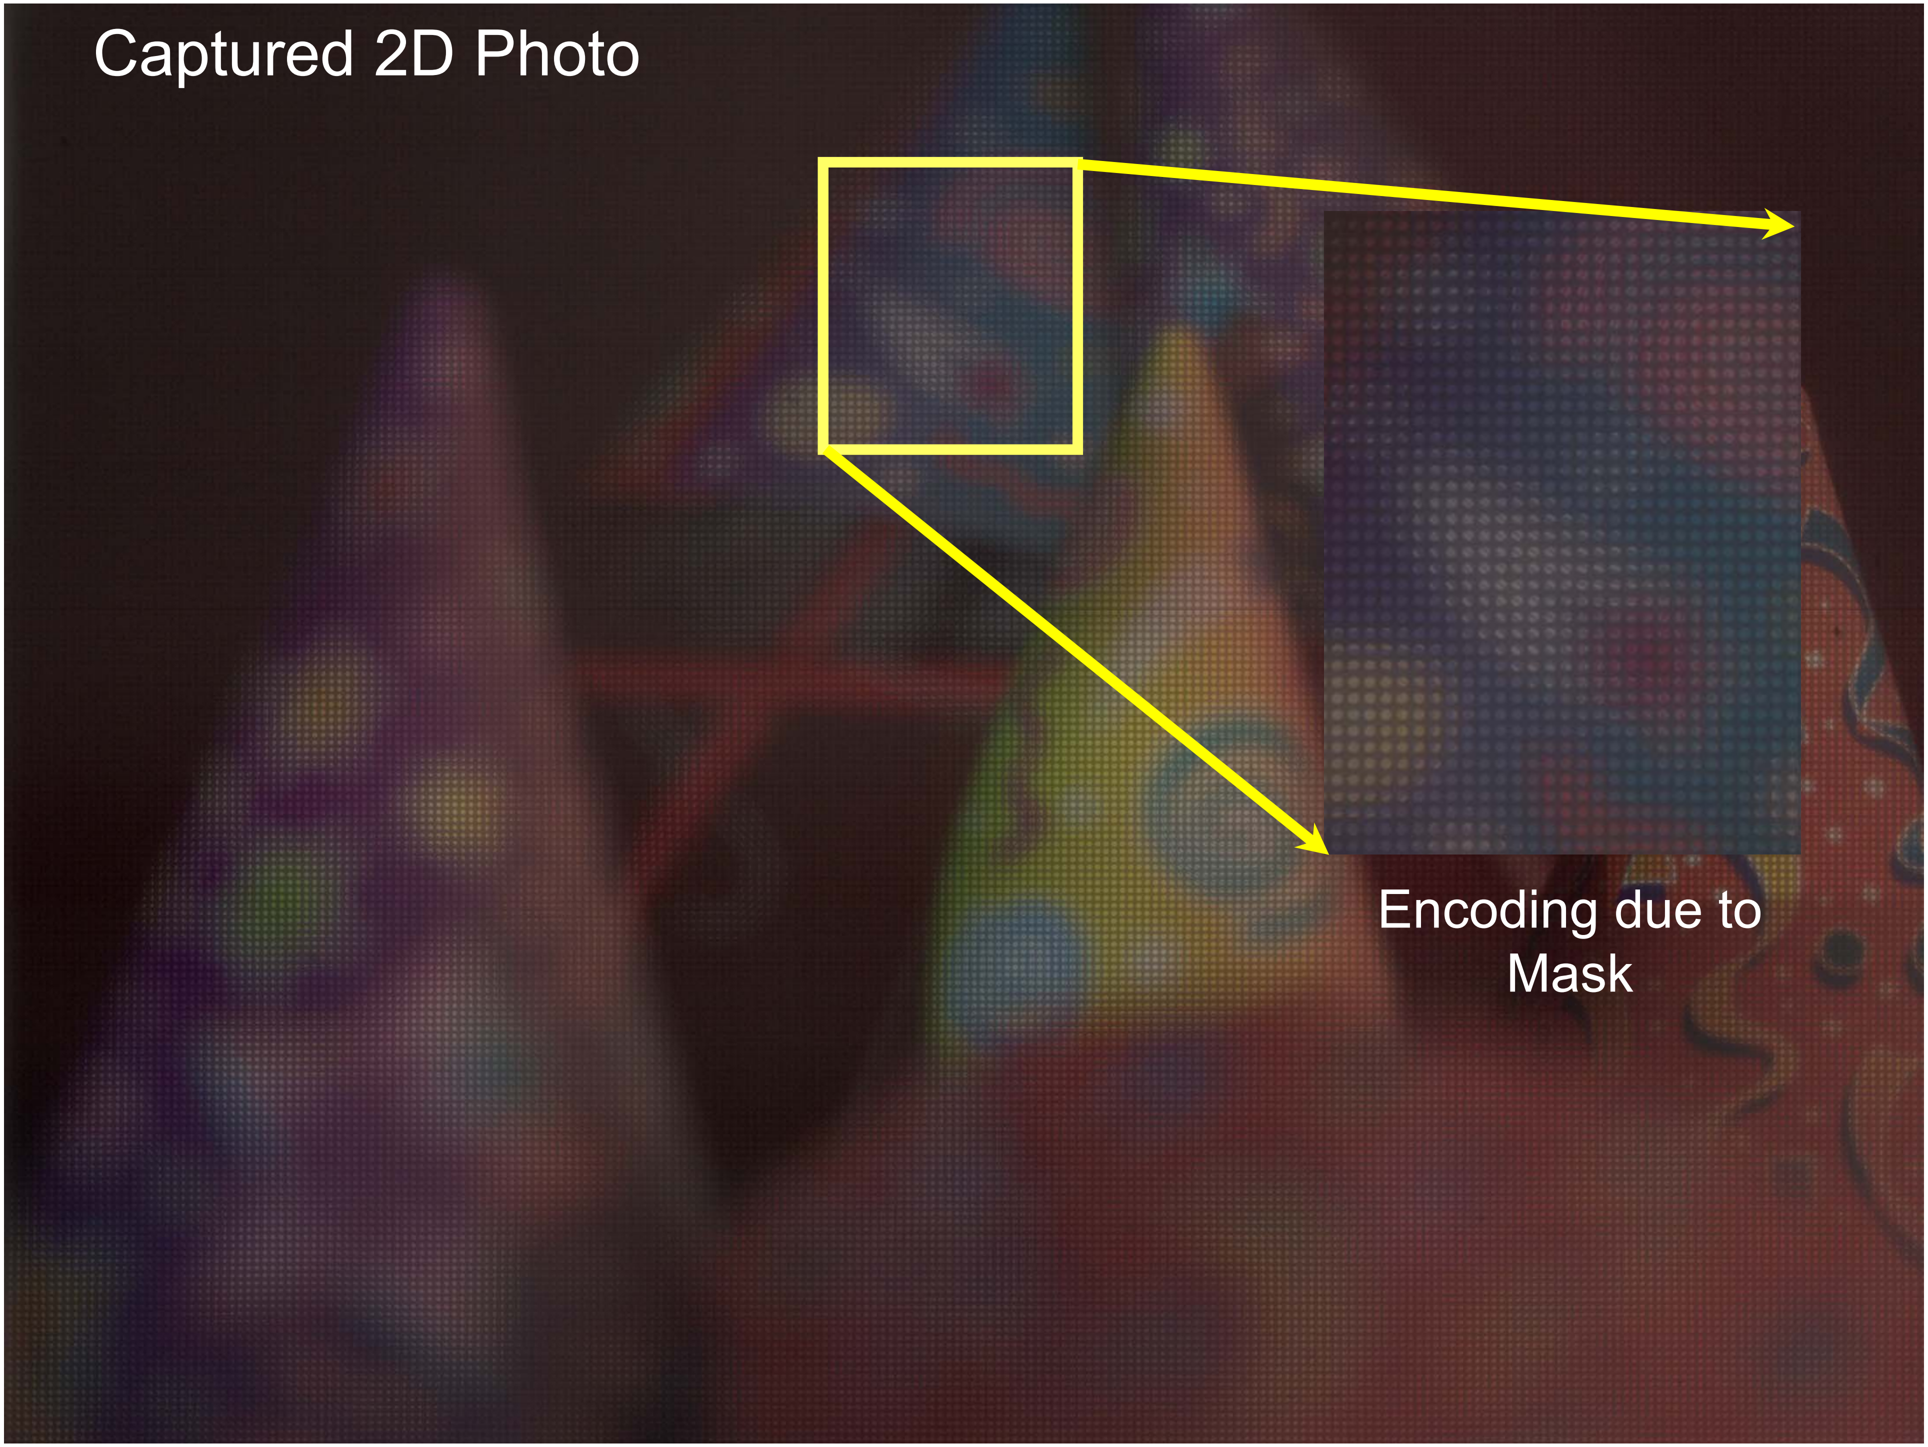
Captured 2D Photo
Encoding due to Mask

2D FFT

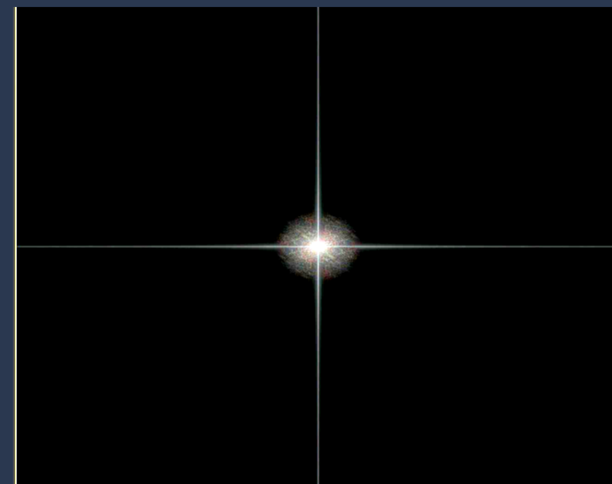

Traditional Camera Photo

Magnitude of 2D FFT

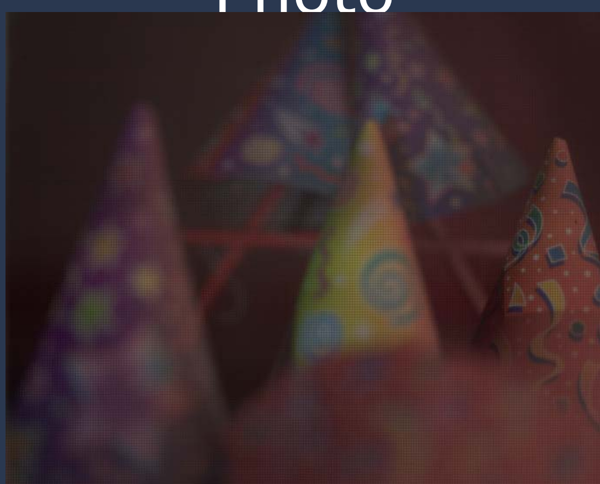

2D FFT

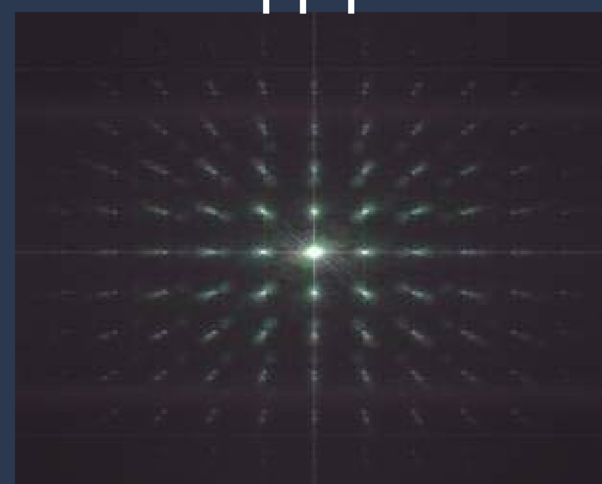

Heterodyne Camera

Magnitude of 2D FFT

Extra sensor bandwidth cannot capture extra *angular dimension* of the light field

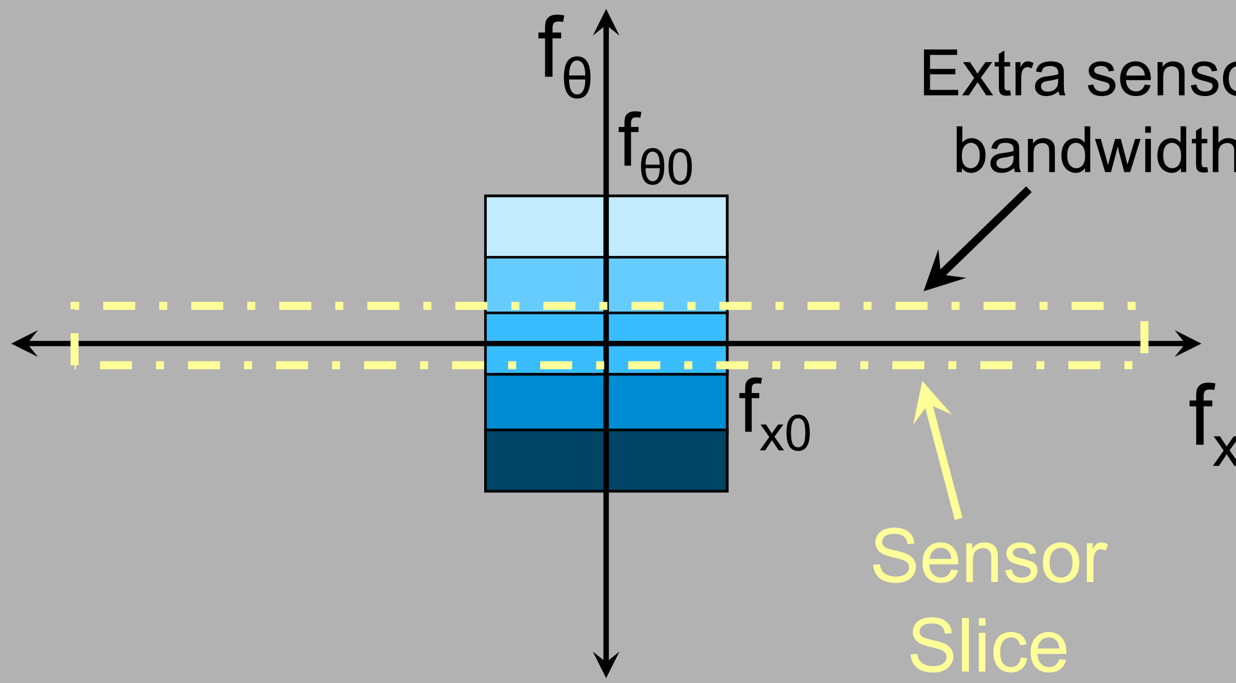

Fourier Light Field Space (Wigner Transform)

Sensor Slice captures entire Light Field

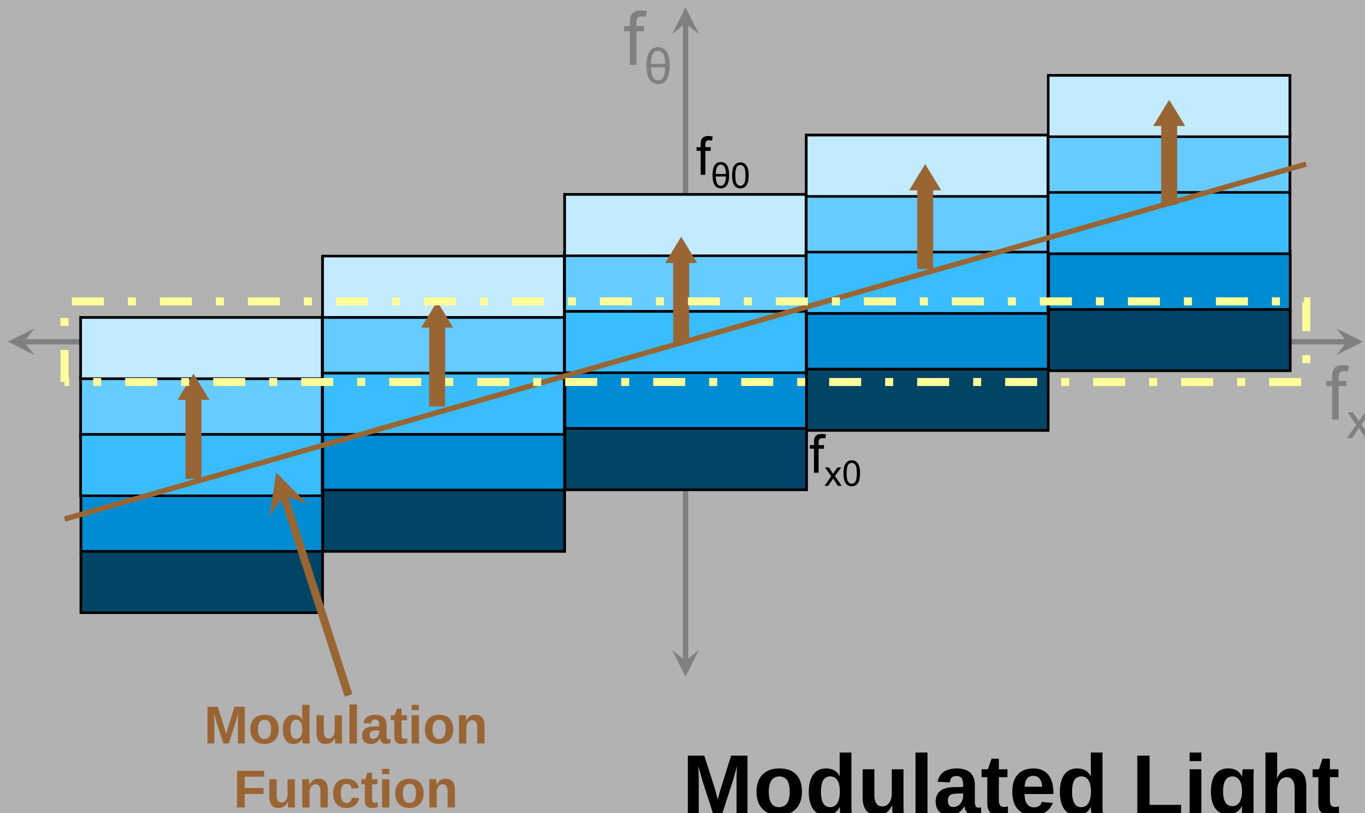

Modulated Light Field

Where to place the Mask?

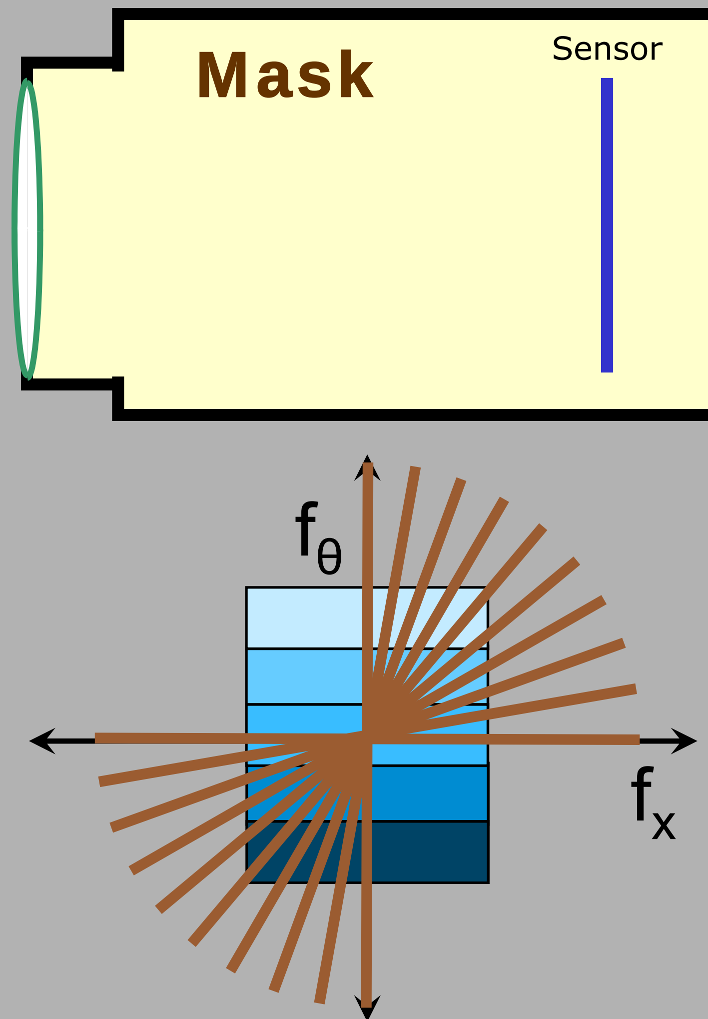

Mask Modulation Function

Computing 4D Light Field

2D Sensor Photo, 1800*1800

2D FFT

2D Fourier Transform, 1800*1800

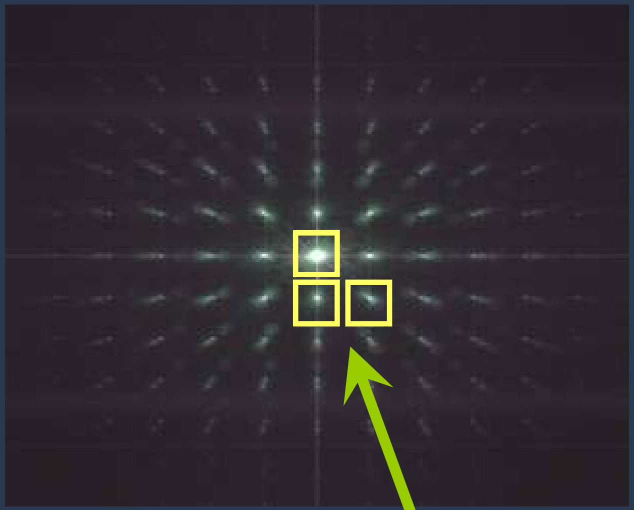

9*9=81 spectral copies

Rearrange 2D tiles into 4D planes

200*200*9*9

4D IFFT

4D Light Field

200*200*9*9

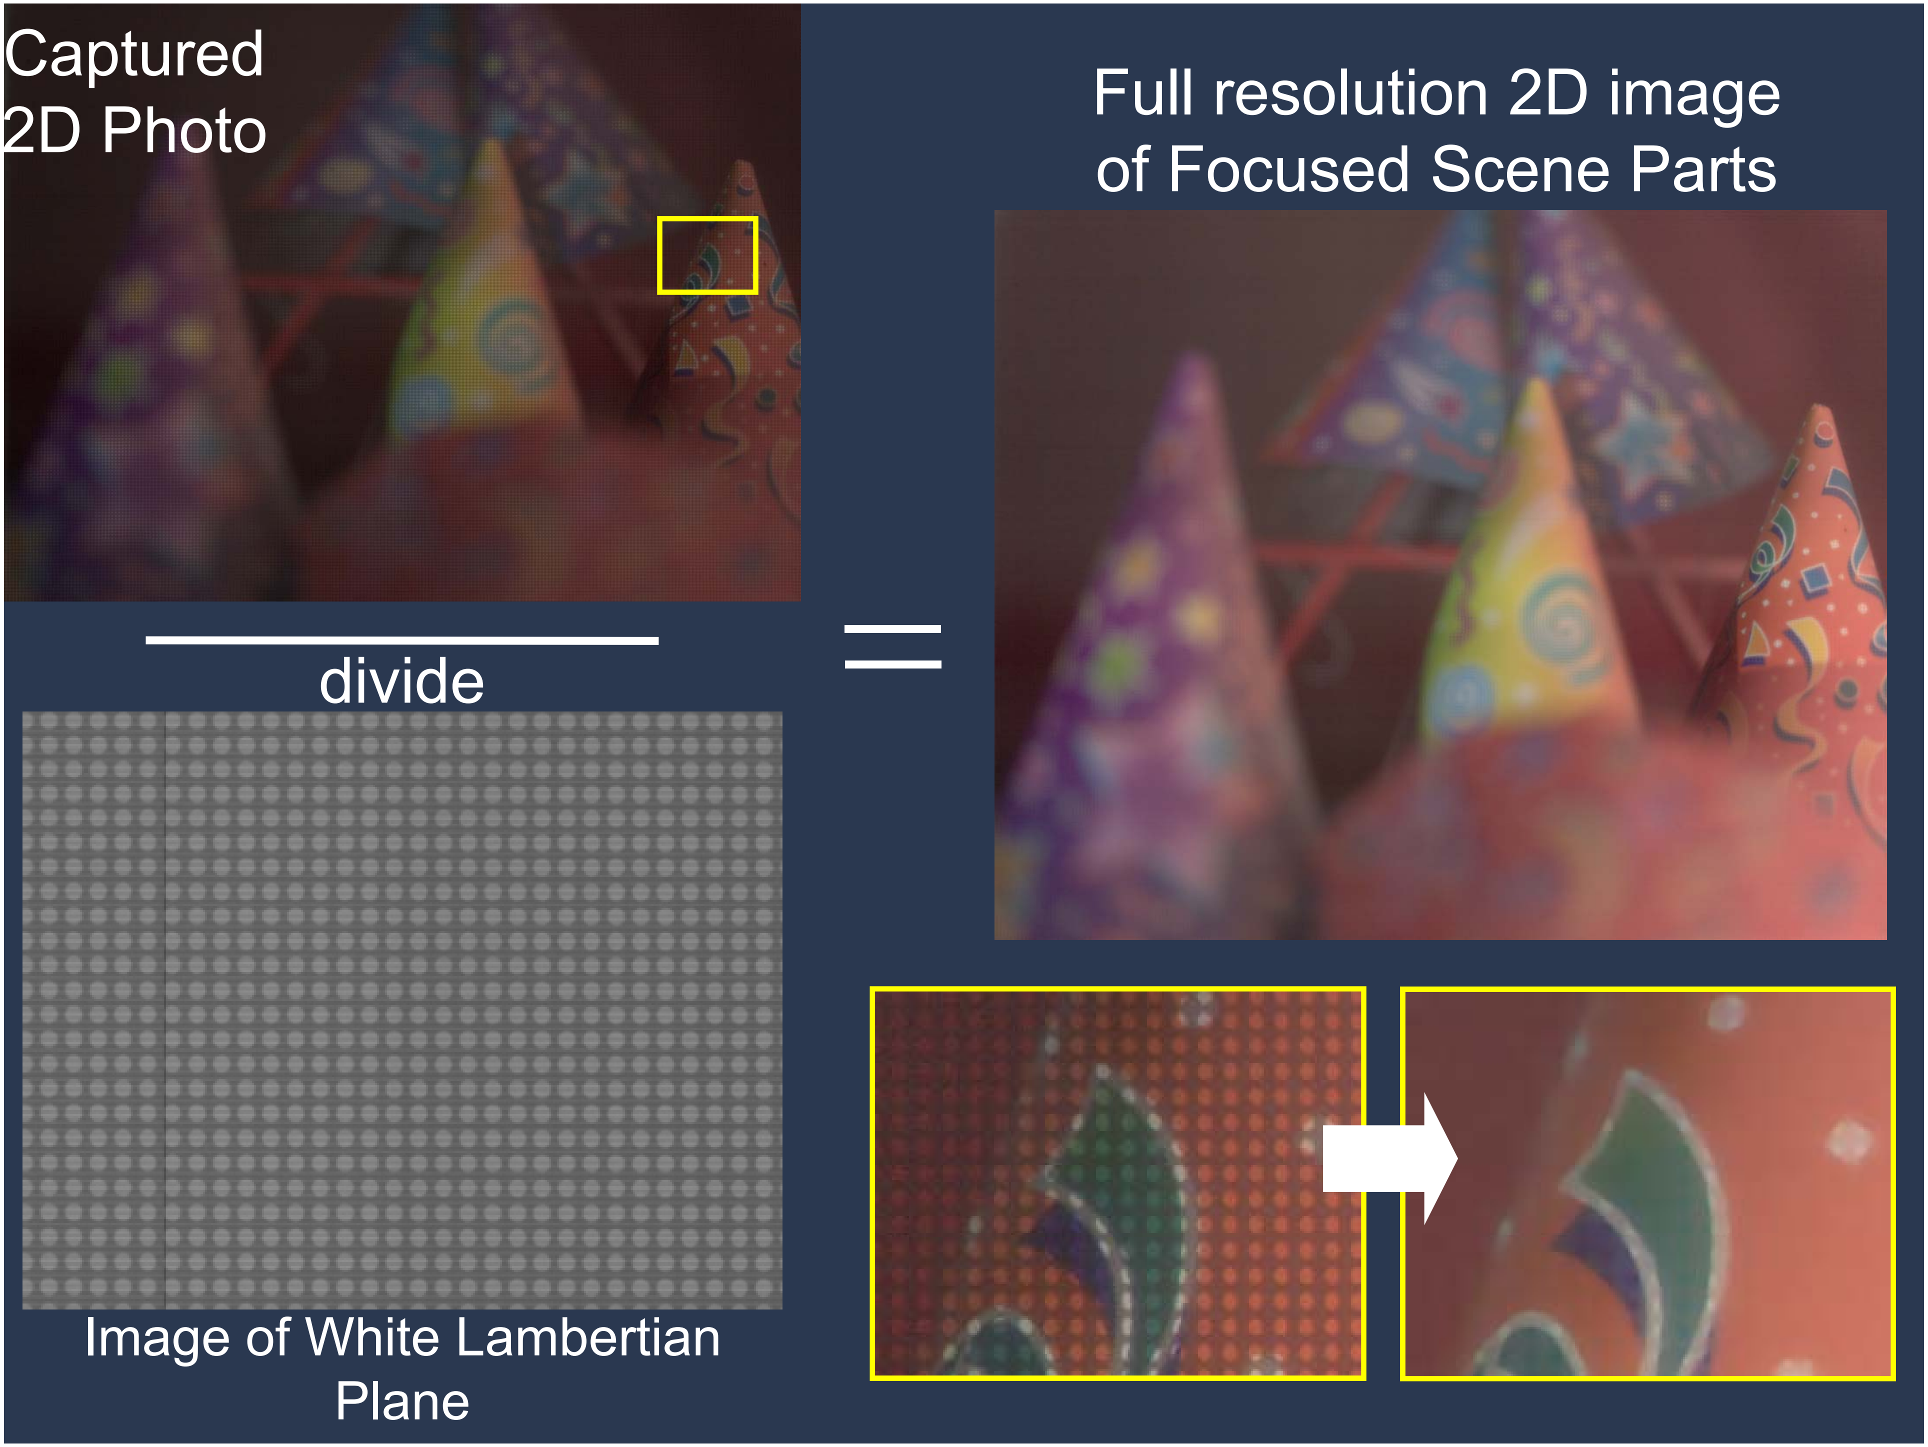
Captured
2D Photo
divide
Image of White Lambertian
Plane
=
Full resolution 2D image
of Focused Scene Parts

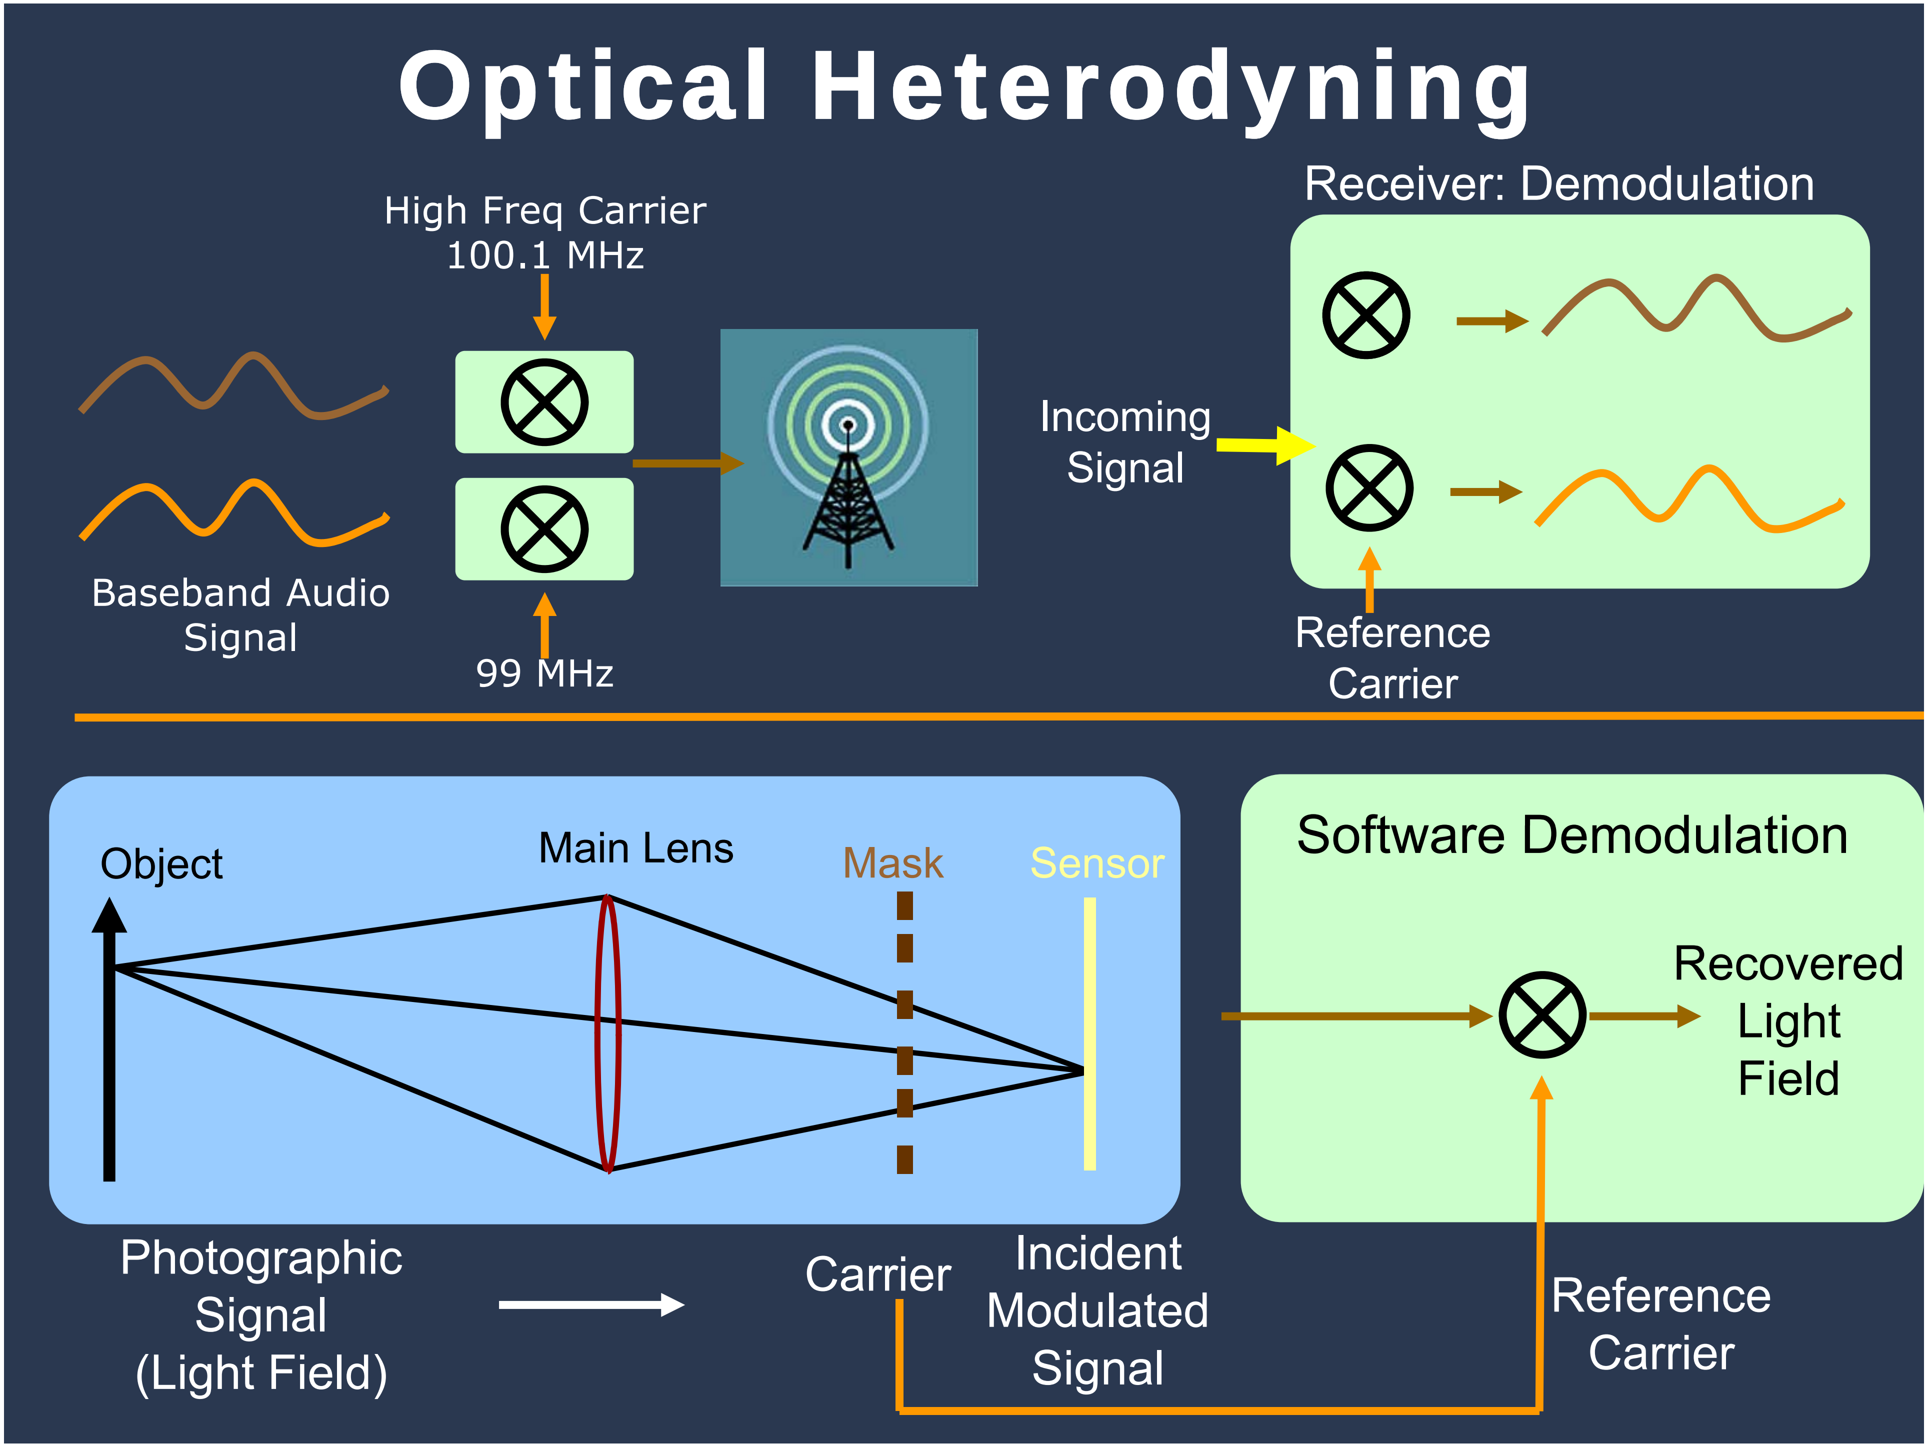
Optical Heterodyning
High Freq Carrier
100.1 MHz
Receiver: Demodulation
Incoming
Signal
Baseband Audio
Signal
99 MHz
Reference
Carrier
Object
Main Lens
Mask
Sensor
Software Demodulation
Recovered
Light
Field
Photographic
Signal
(Light Field)
Carrier
Incident
Modulated
Signal
Reference
Carrier

Light Fields

- What are they?
- What are the properties?
- How to capture?
- What are the applications?

Light Field Applications

- Lens effects
  - Refocussing
  - New aperture setting
  - All in focus image
- Geometric
  - Estimate depth
  - (Create new views)
  - Synthetic aperture (Foreground/background)
  - (Insert objects)
- Statistical
  - Lens glare
  - Specular-diffuse
- Note:
  - LF not required, 4D sampling sufficient
  - Similar HD analysis also works for motion, wavelength, displays

Compound Lens of Dragonfly

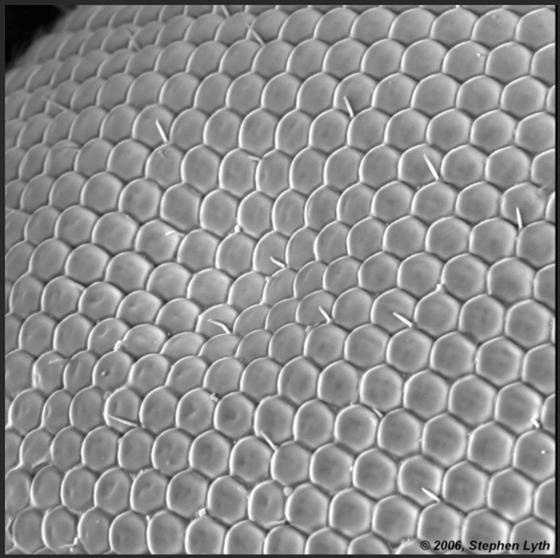

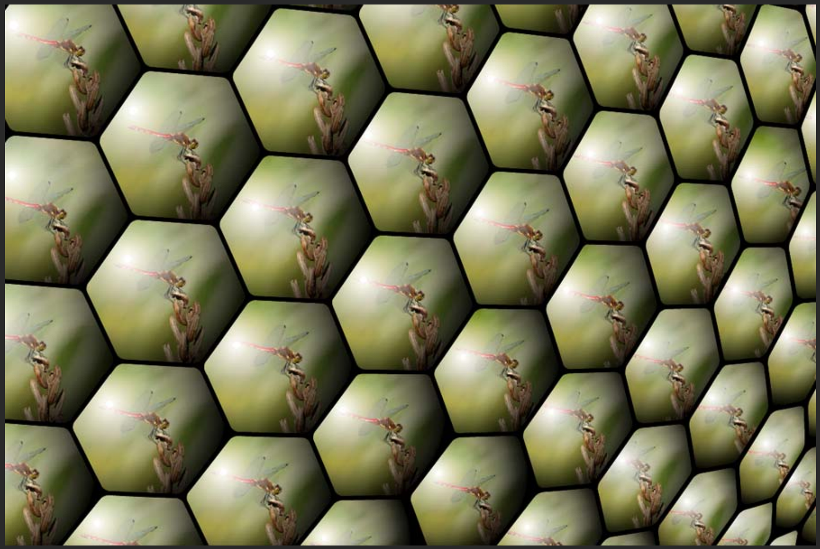

TOMBO: Thin Camera (2001)

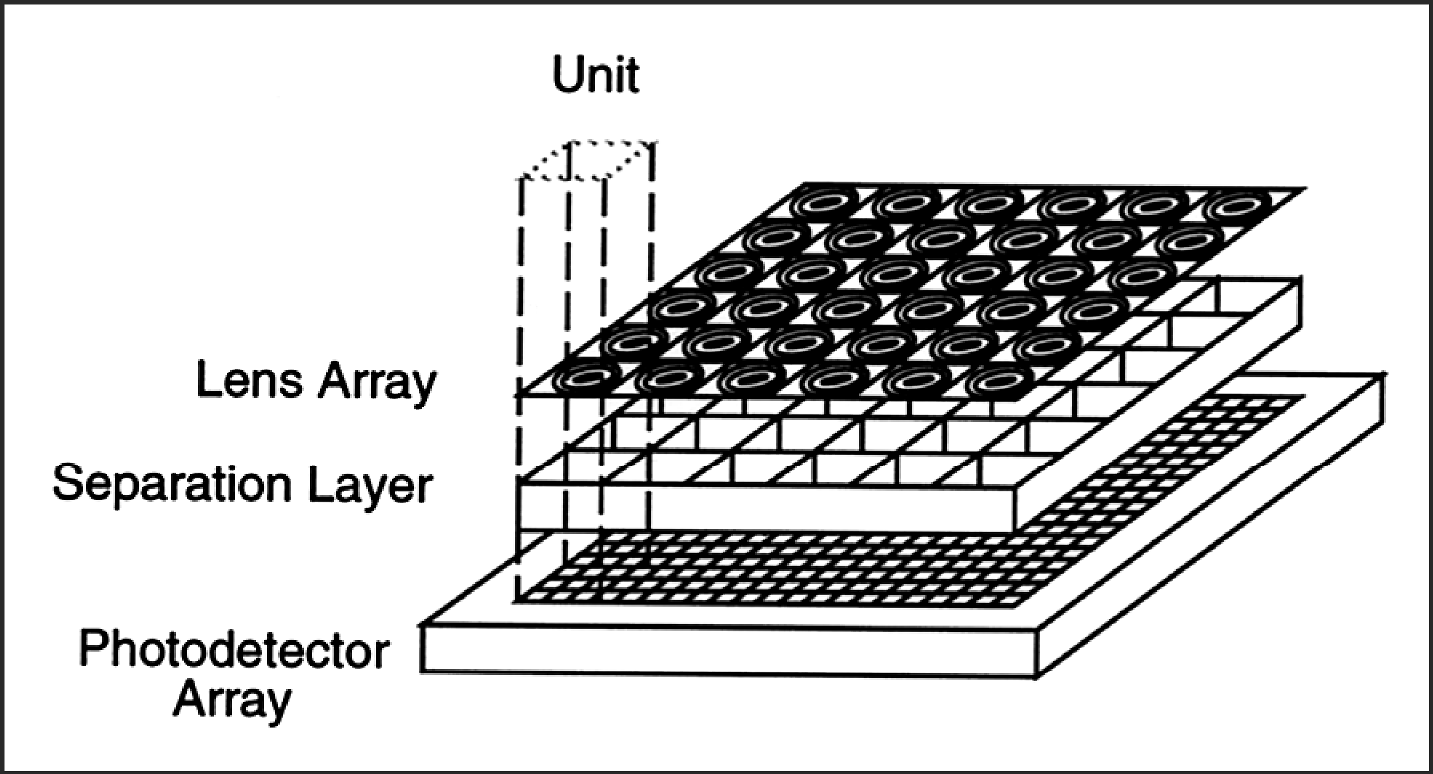

"Thin observation module by bound optics (TOMBO),"
J. Tanida, T. Kumagai, K. Yamada, S. Miyatake
Applied Optics, 2001

TOMBO: Thin Camera

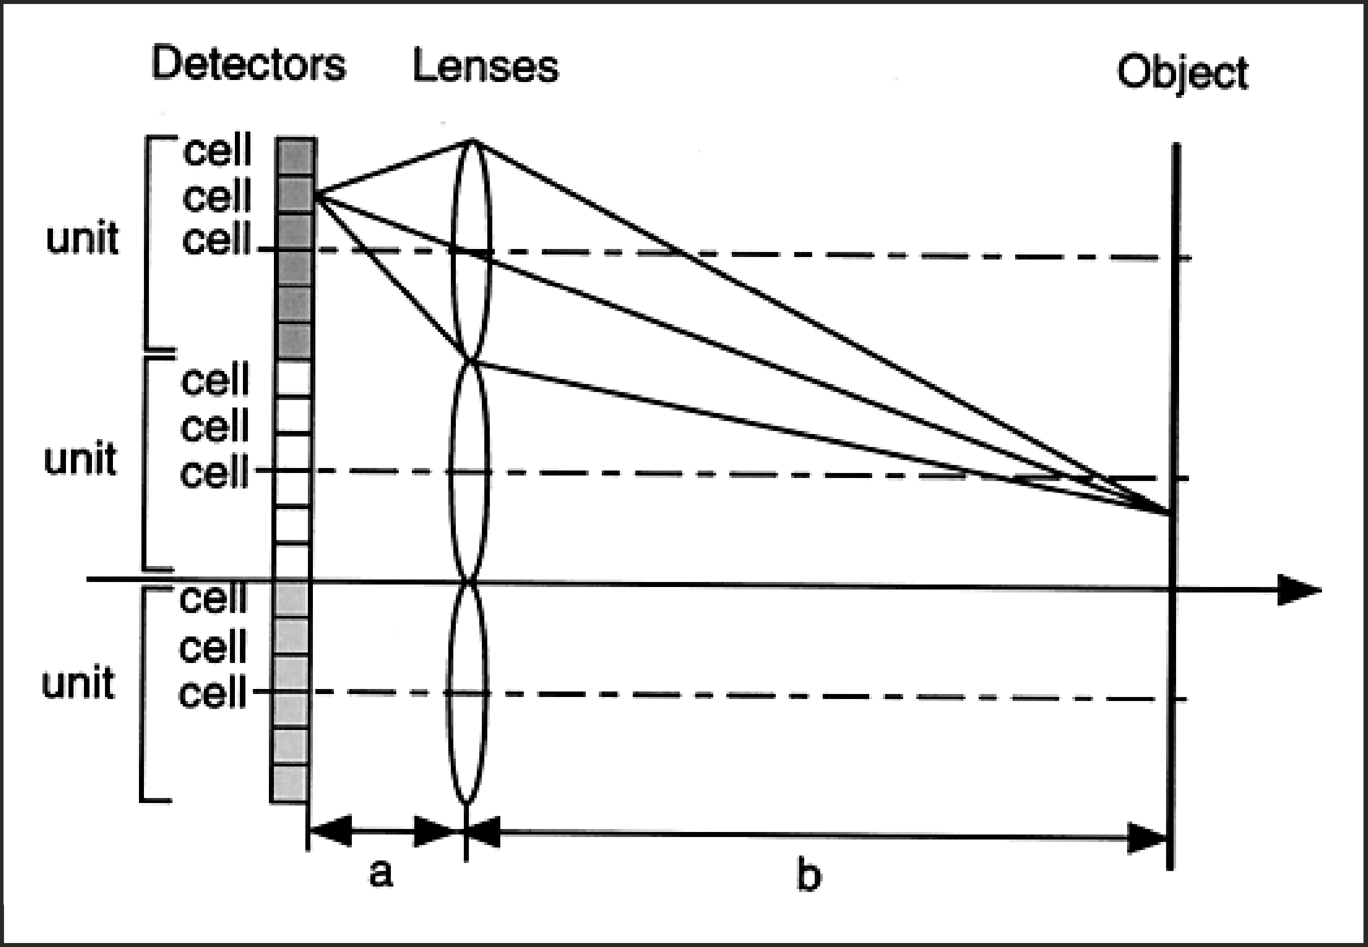

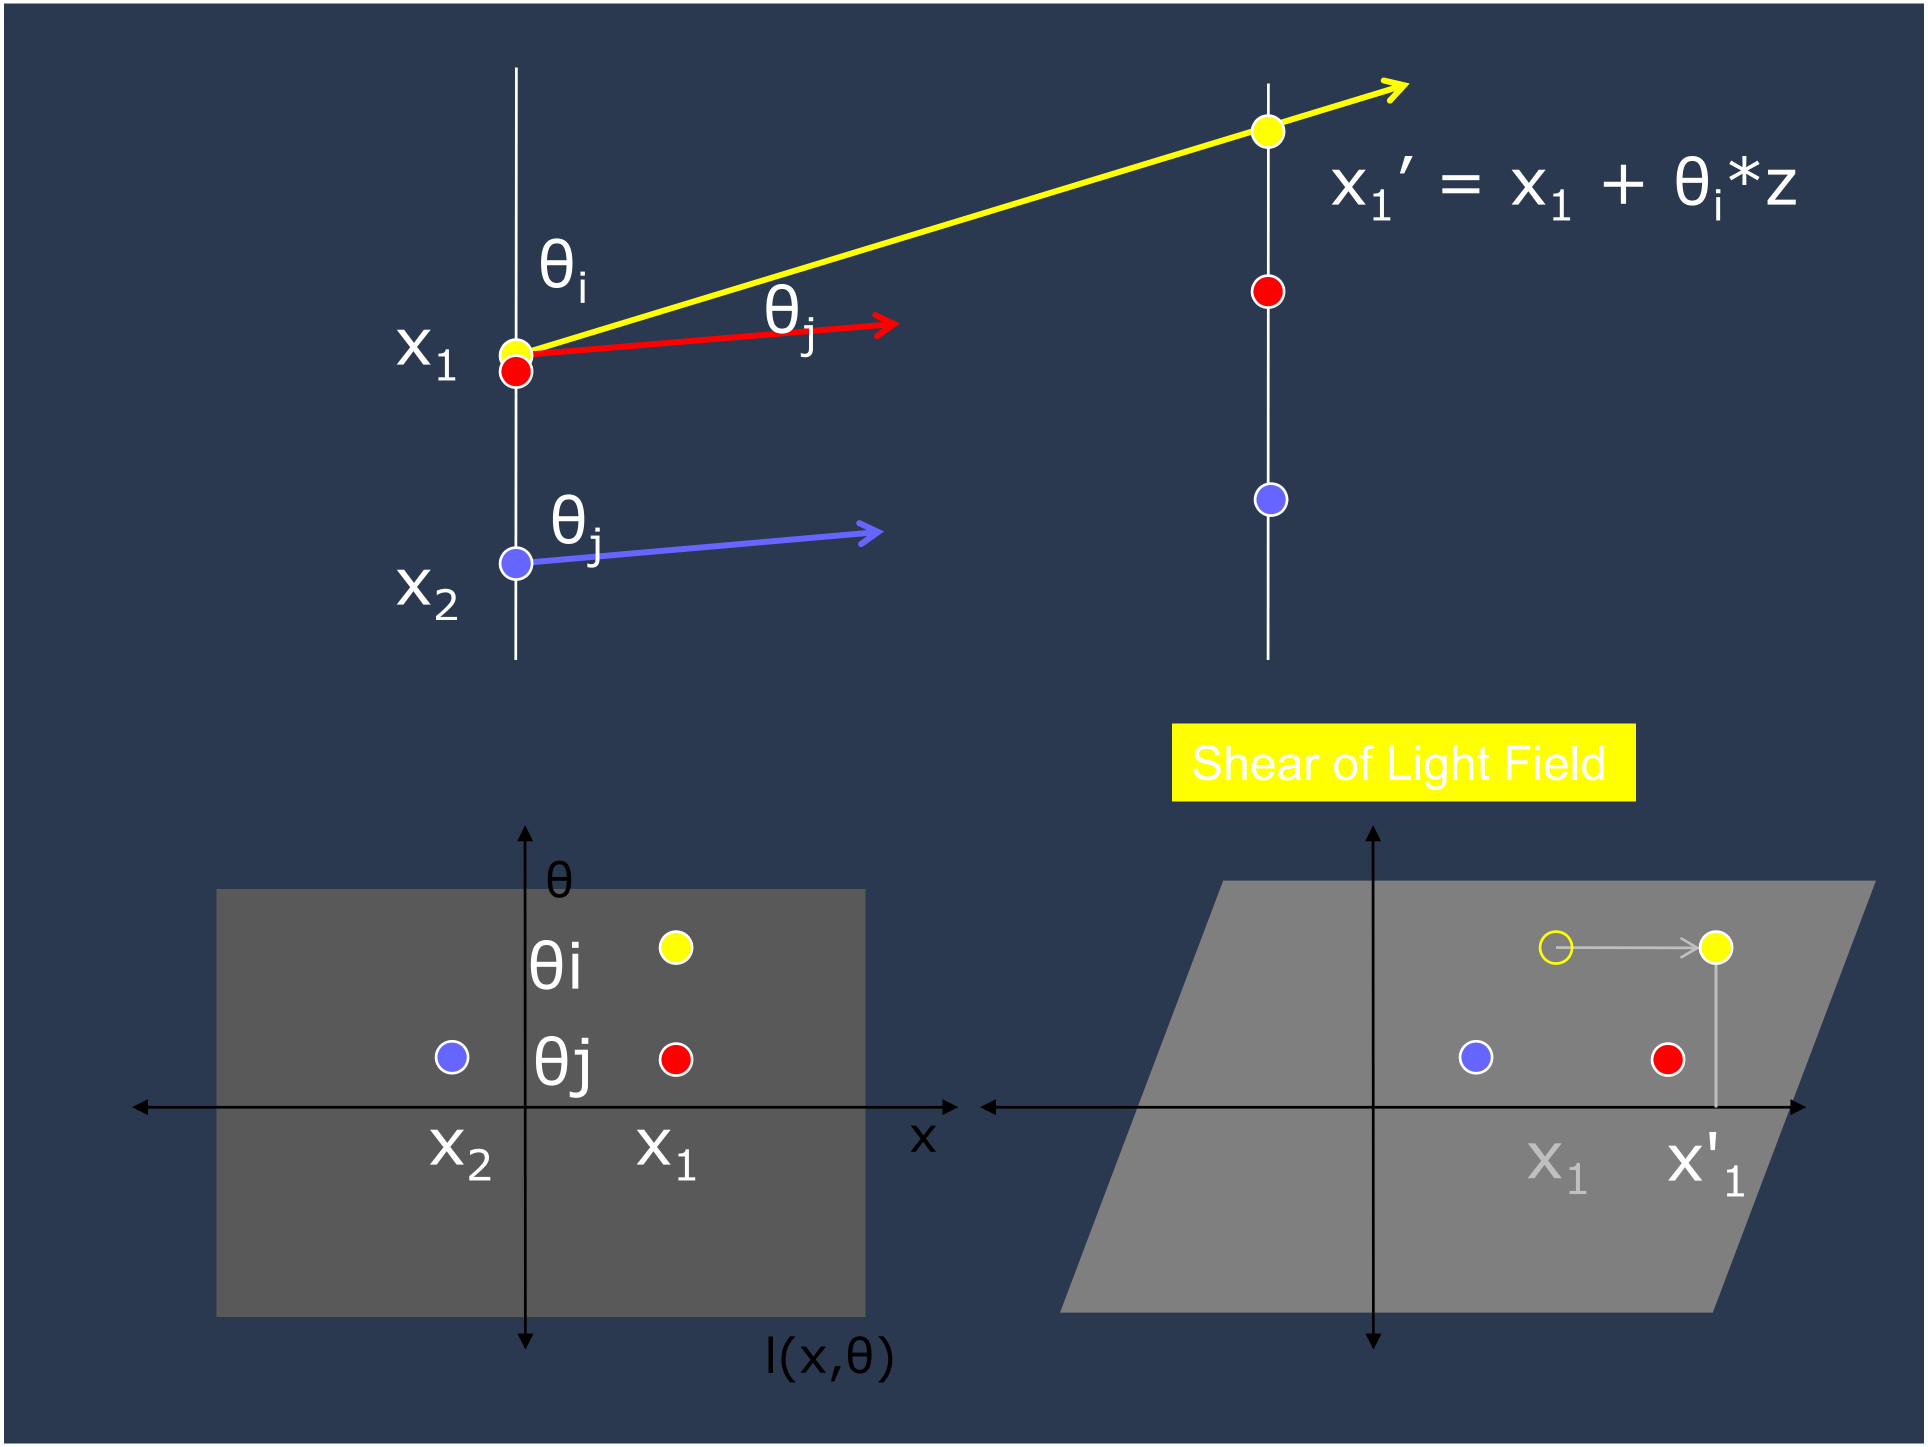

θi
θj
x1
x2
θj
x1' = x1 + θi*z
Shear of Light Field
θ
θi
θj
x2
x1
x
I(x,θ)
x1
x'1

Light Propagation (Defocus Blur)

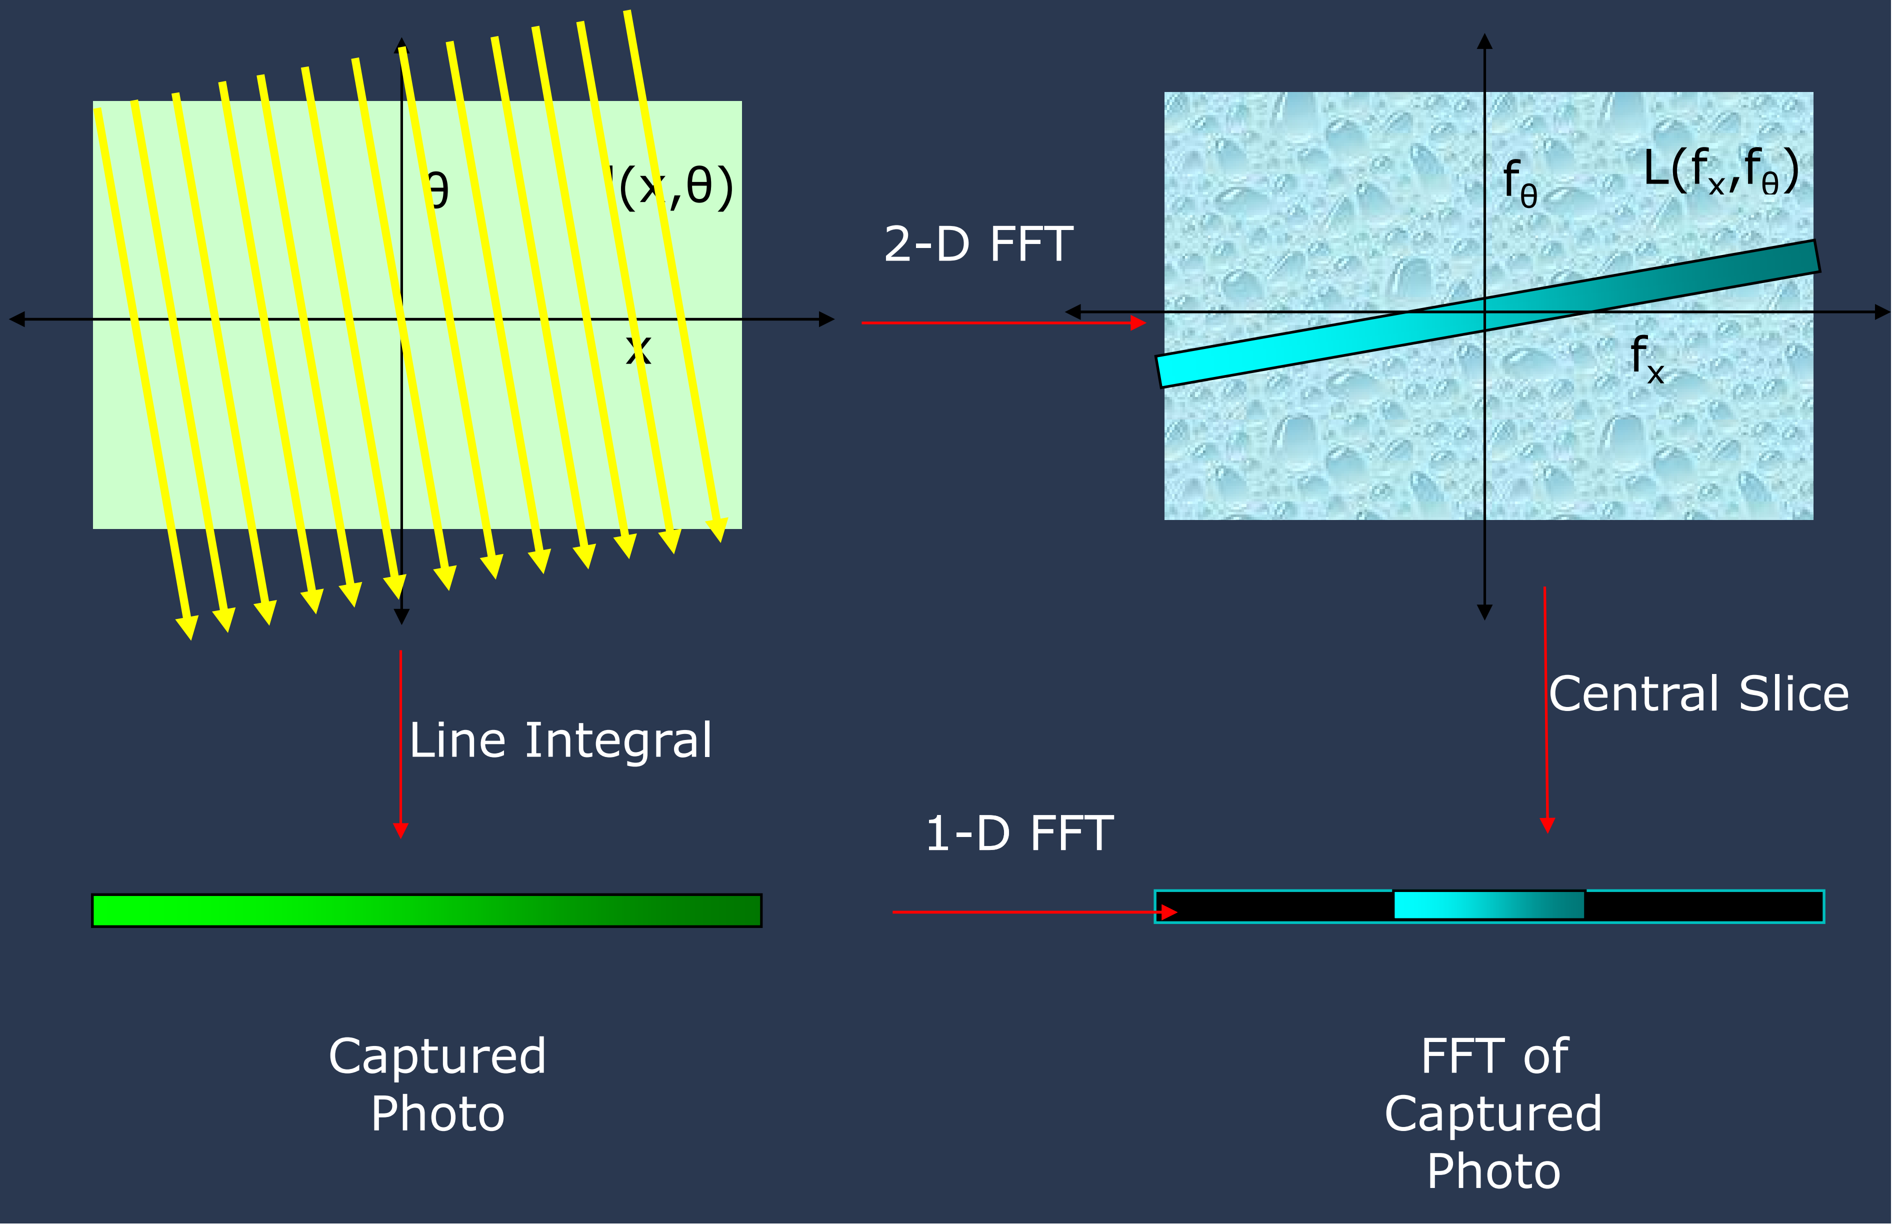

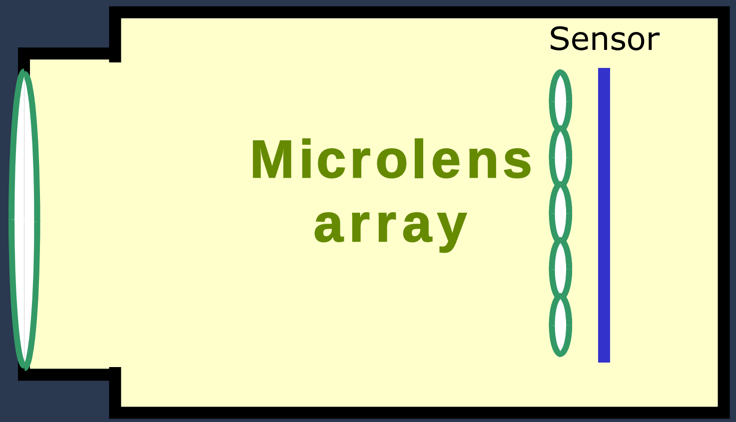

Plenoptic Camera

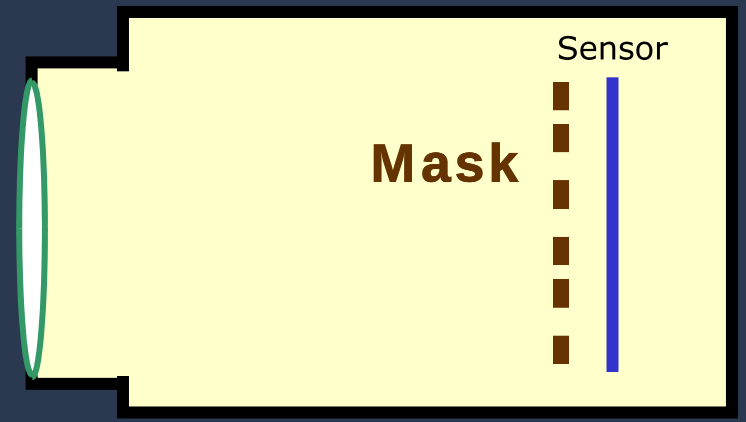

Heterodyne Camera

- Samples individual rays
- Predefined spectrum for lenses
- Chromatic abberration
- High alignment precision
- Peripheral pixels wasted pixels
- Negligible Light Loss

- Samples coded combination of rays
- Supports any wavelength
- Reconfigurable f/#, Easier alignment
- No wastage
- High resolution image for parts of scene in focus
- 50 % Light Loss due to mask

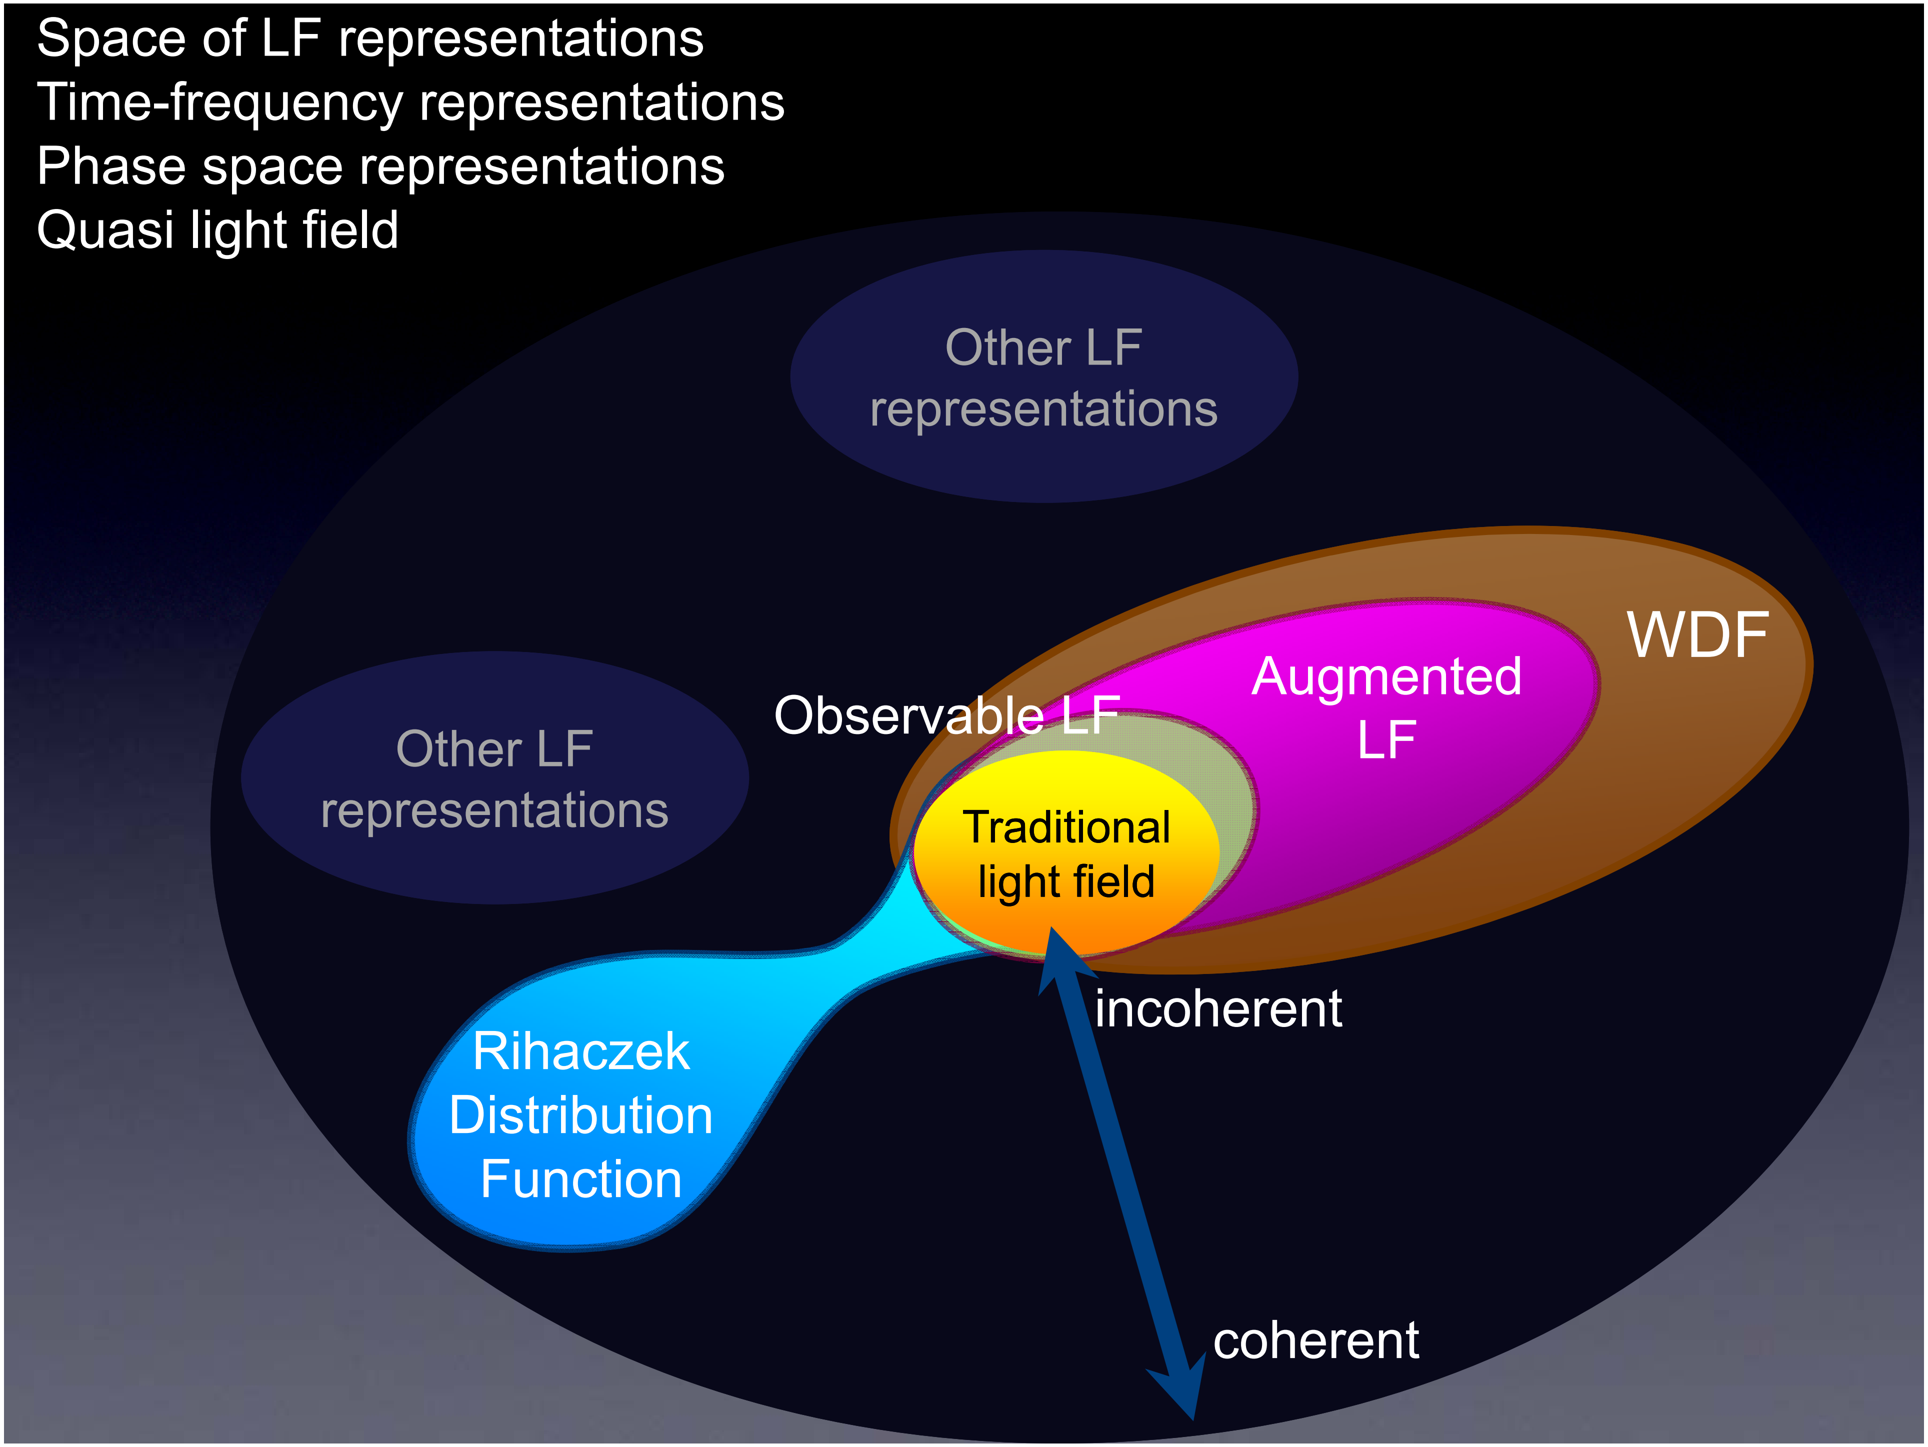
Space of LF representations
Time-frequency representations
Phase space representations
Quasi light field
Other LF representations
Other LF representations
WDF
Augmented LF
Observable LF
Traditional light field
Rihaczek Distribution Function
incoherent
coherent

Quasi light fields

the utility of light fields, the versatility of Maxwell

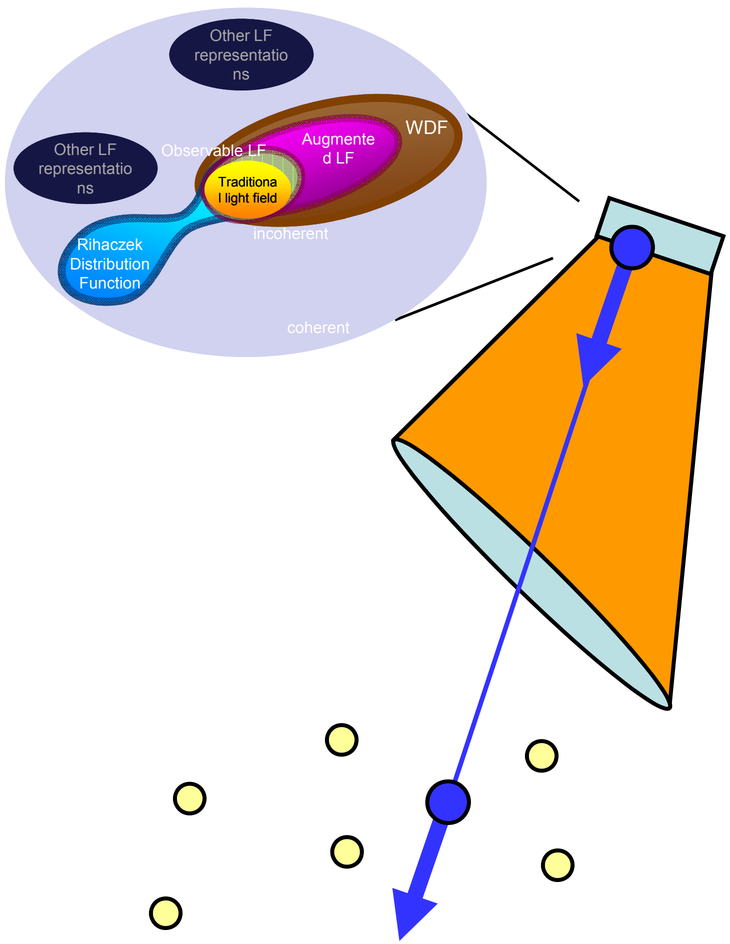

We form coherent images by

formulating,

capturing,

and integrating

quasi light fields.

(i) Observable Light Field

- move aperture across plane
- look at directional spread
- continuous form of plenoptic camera

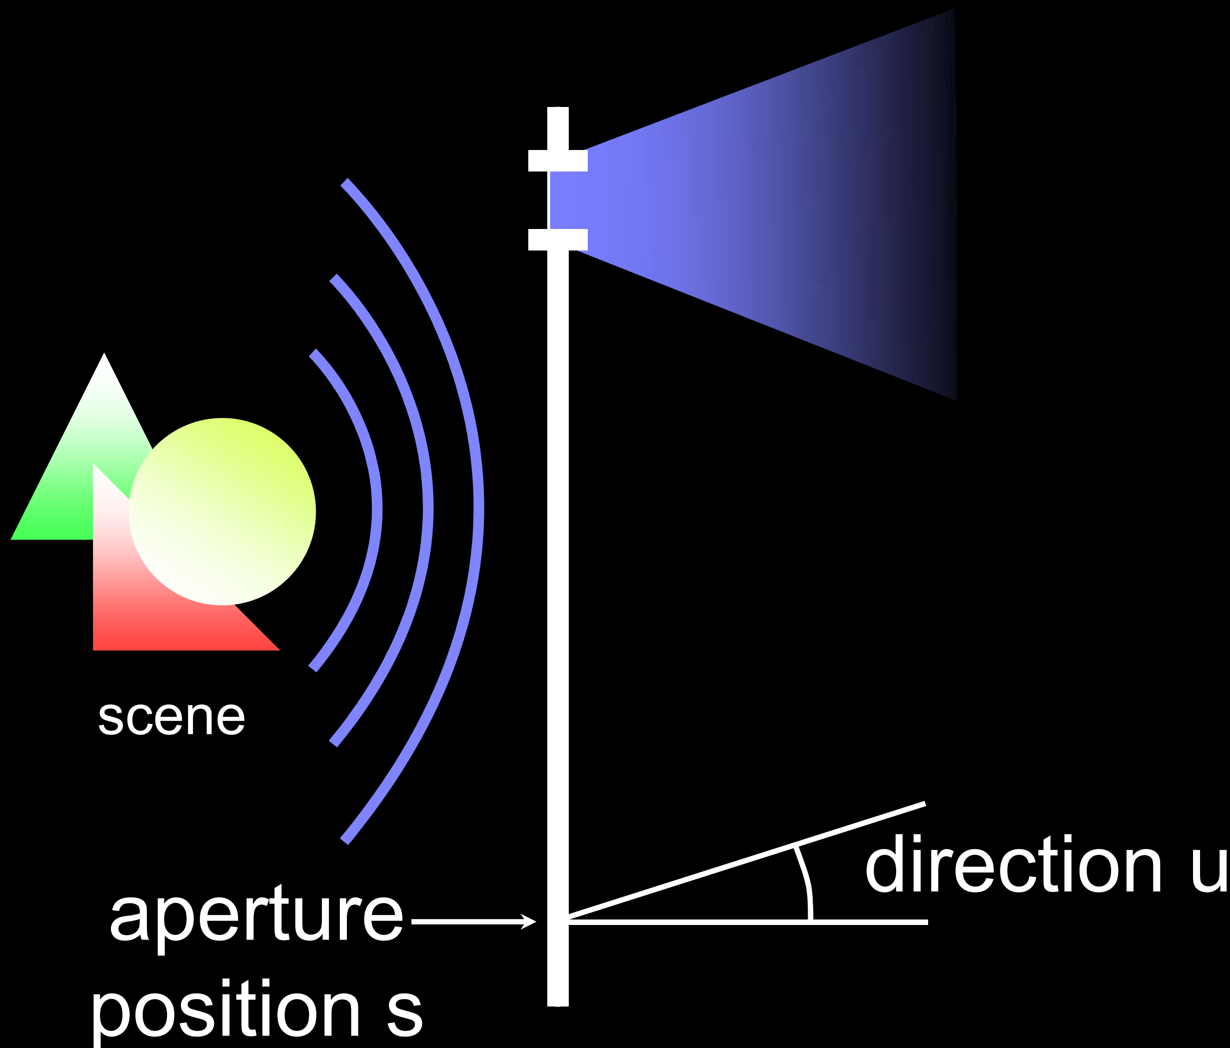

(ii) Augmented Light Field with LF Transformer

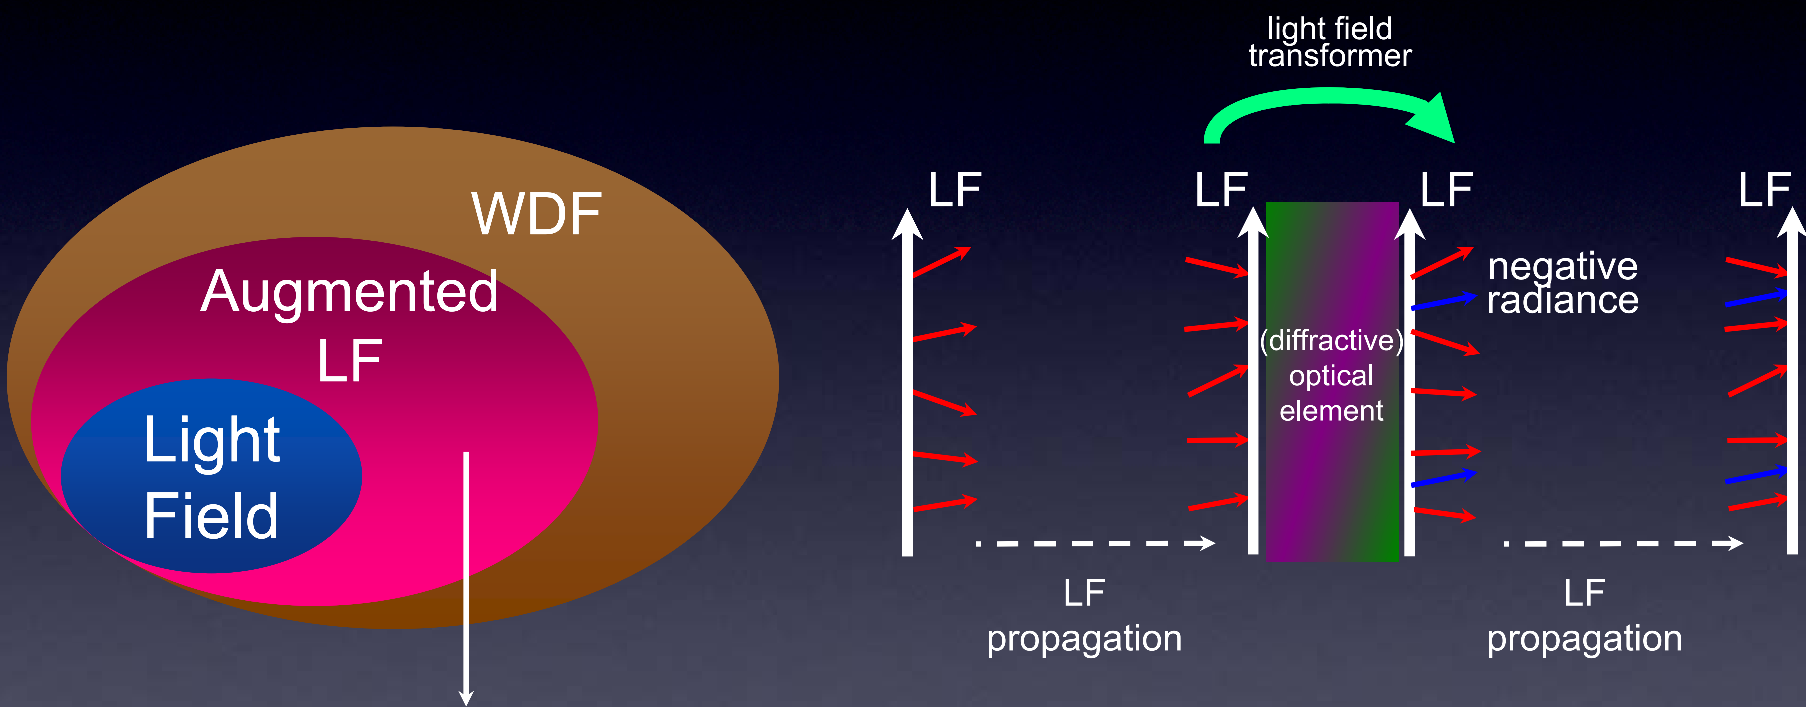

Interaction at the optical elements

Virtual light projector with real valued (possibly negative radiance) along a ray

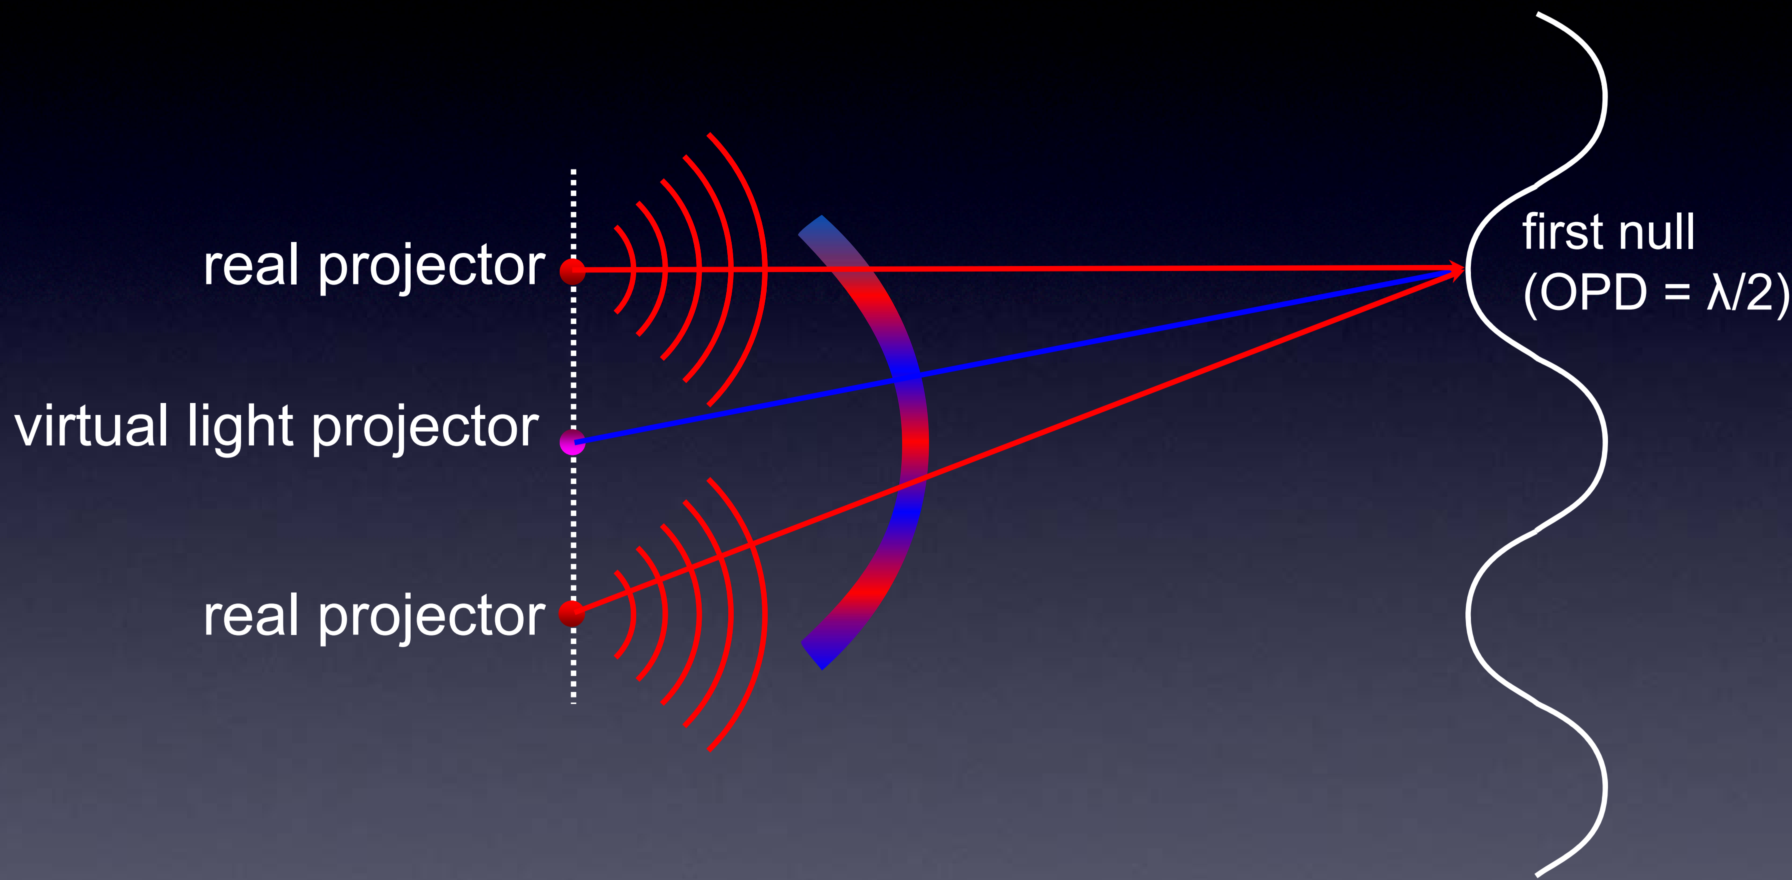

(ii) ALF with LF Transformer

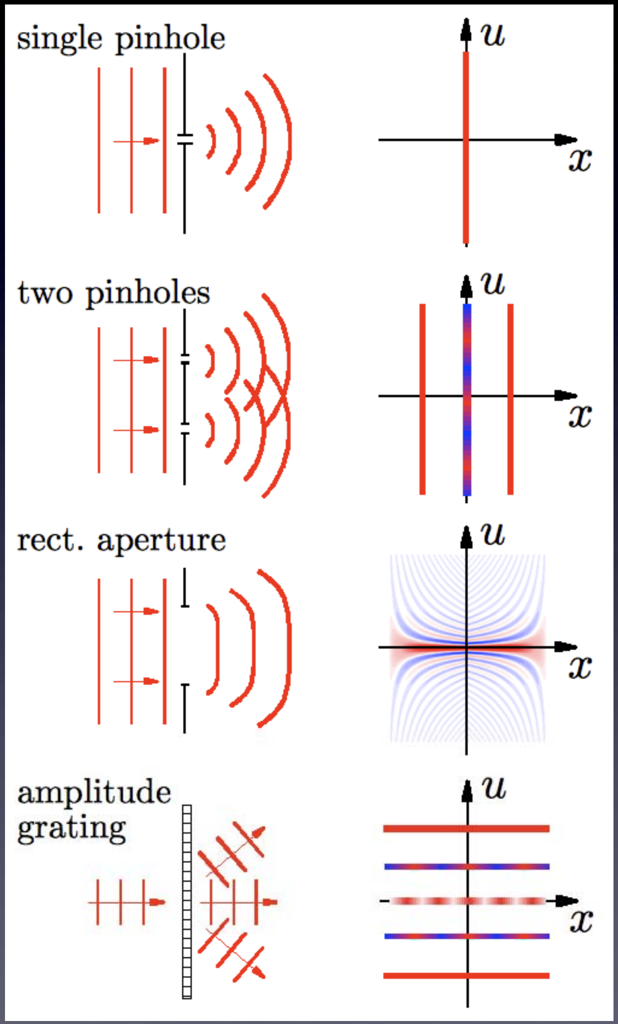

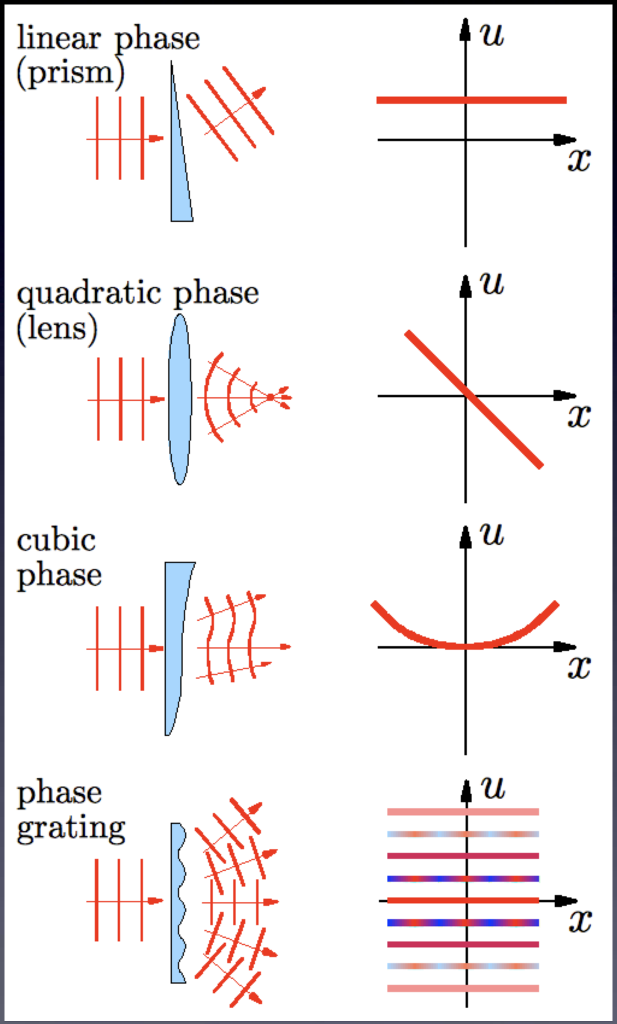

Tradeoff between cross-interference terms and localization

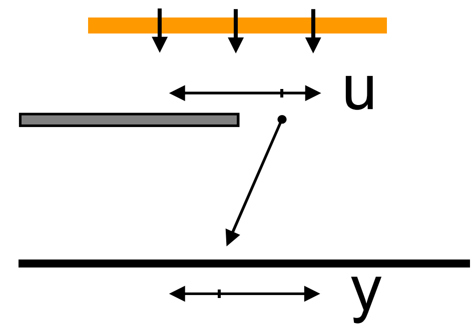

(i) Spectrogram
non-negative
localization

(ii) Wigner
localization
cross terms

(iii) Rihaczek
localization
complex

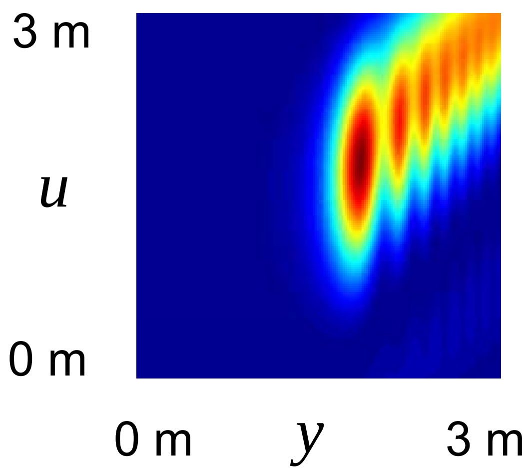

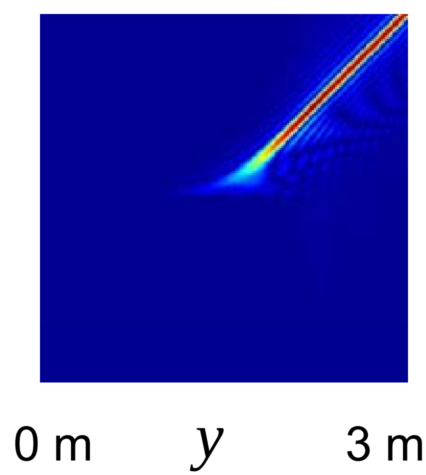

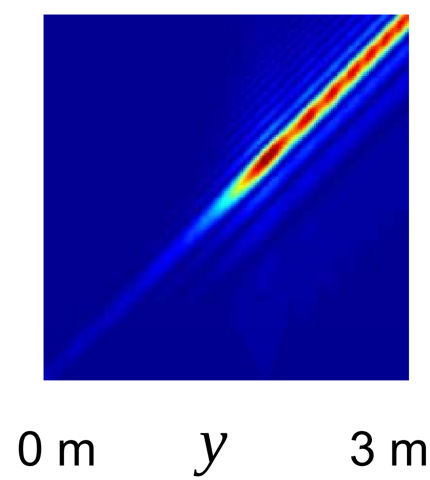

Property of the Representation

	Constant along rays	Non-negativity	Coherence	Wavelength	Interference Cross term
Traditional LF	always constant	always positive	only incoherent	zero	no
Observable LF	nearly constant	always positive	any coherence state	any	yes
Augmented LF	only in the paraxial region	positive and negative	any	any	yes
WDF	only in the paraxial region	positive and negative	any	any	yes
Rihaczek DF	no; linear drift	complex	any	any	reduced

Benefits & Limitations of the Representation

	Ability to propagate	Modeling wave optics	Simplicity of computatio n	Adaptability to current pipe line	Near Field	Far Field
Traditional LF	x-shear	no	very simple	high	no	yes
Observable LF	not x-shear	yes	modest	low	yes	yes
Augmented LF	x-shear	yes	modest	high	no	yes
WDF	x-shear	yes	modest	low	yes	yes
Rihaczek DF	x-shear	yes	better than WDF, not as simple as LF	low	no	yes

Motivation

- What is the difference between a hologram and a lenticular screen?
- How they capture ‘phase’ of a wavefront for telescope applications?
- What is ‘wavefront coding’ lens for extended depth of field imaging?

Acknowledgements

- Darthmuth
  - Marcus Testorf,
- MIT
  - Ankit Mohan, Ahmed Kirmani, Jaewon Kim
  - George Barbastathis
- Stanford
  - Marc Levoy, Ren Ng, Andrew Adams
- Adobe
  - Todor Georgiev,
- MERL
  - Ashok Veeraraghavan, Amit Agrawal

Light Fields

Ramesh Raskar

MIT Media Lab

http:// CameraCulture . info/

Light Fields in Ray and Wave Optics

Introduction to Light Fields:	Ramesh Raskar
Wigner Distribution Function to explain Light Fields:	Zhengyun Zhang
Augmenting LF to explain Wigner Distribution Function:	Se Baek Oh
Q&A	
Break	
Light Fields with Coherent Light:	Anthony Accardi
New Opportunities and Applications:	Raskar and Oh
Q&A:	All