Image Sensors

Lecture J
Specialized Cameras

Robert Forchheimer

Dynamic range extension (HDR)

- A standard camera takes images that can be
  - Overexposed if we want to see details in the dark areas
    - Bright areas become “too bright”
  - Underexposed if we want to see details in the bright areas
    - Dark areas become “too dark”

Dynamic range extension (HDR)

- Several techniques can be used to produce images of higher dynamic range (HDR) than provided by the basic sensor technology
  - External HDR (outside the sensor chip)
  - Internal HDR (inside the sensor chip)

External HDR

- Take two or more images, one immediately after the other, with different exposure time
- Read out each image as normal
- Produce an HDR image outside the sensor chip by, for example, at each pixel:
  - Combine the measured intensity values from the different exposures, after suitable normalization (how?)
- Assumes fast exposure of all images
  - Works fine with CMOS
  - Or static scene

External HDR

Simple reconstruction: add the two images and renormalize

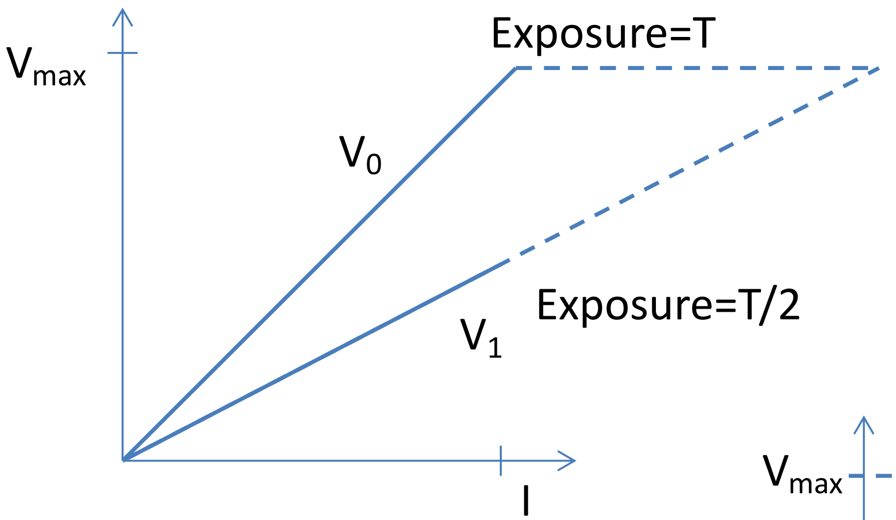

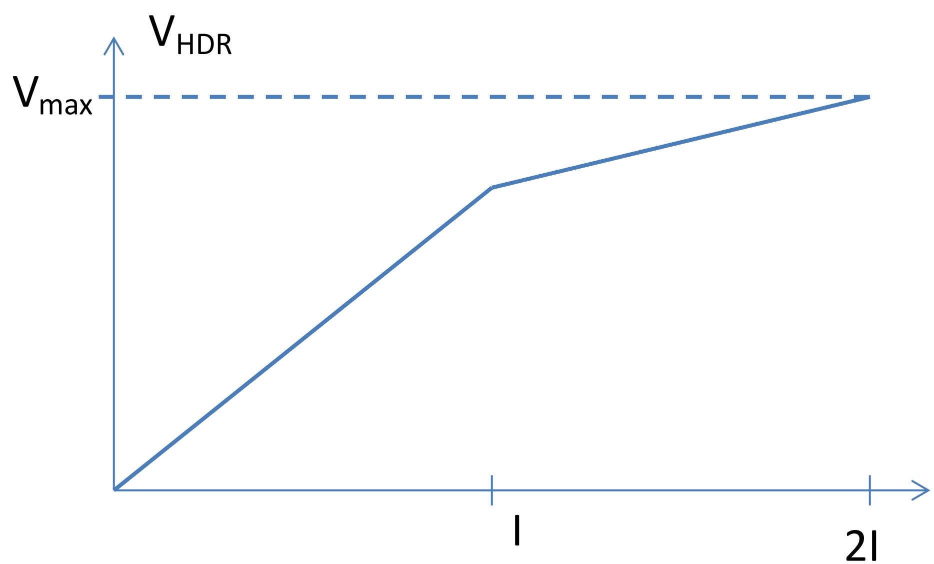

$$V_{HDR} = (V_0 + V_1) / 2$$

External HDR

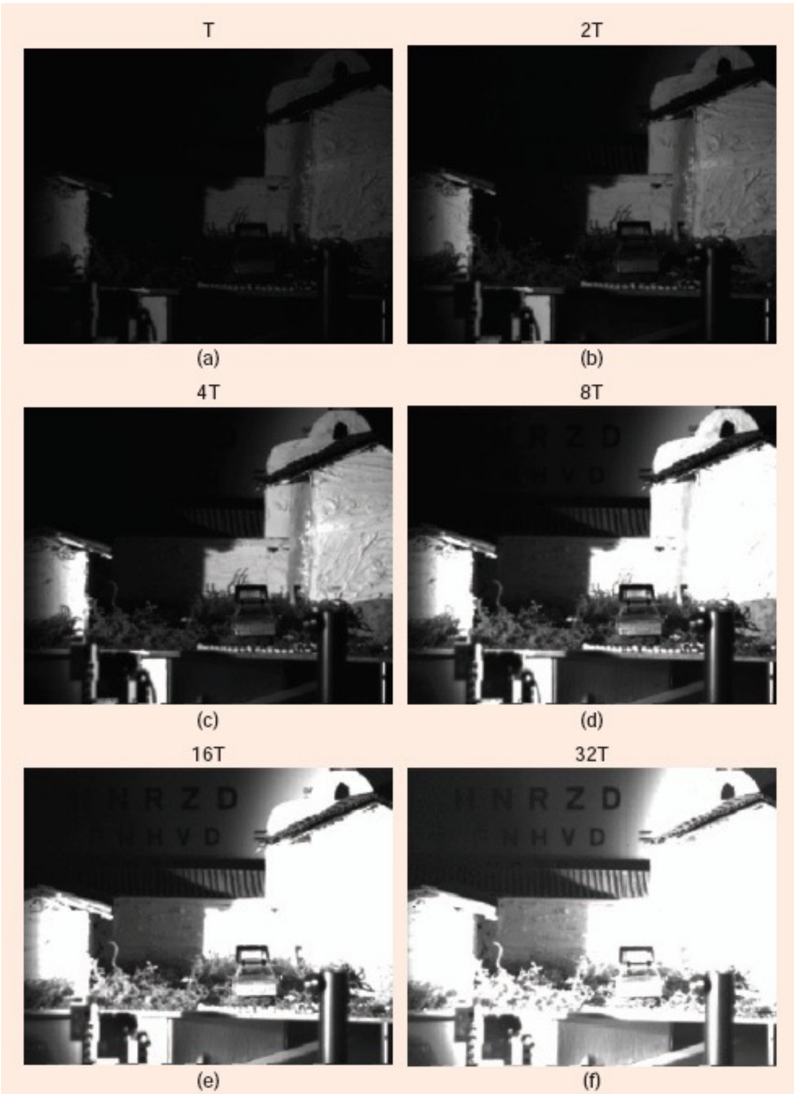

The same scene is viewed with 6 different exposure times

Internal HDR

- Log-intensity
- Dual diode
- Piecewise linear response

Log-intensity cameras

- CMOS technology can be used to achieve a logarithmic dependence between the absorbed light and the resulting photo-voltage
- Enables a camera with very high dynamic range (>120 dB = 6 decades)
- Linear and logarithmic mode can be combined in the same camera

Logarithmic pixel

- 3T logarithmic pixel
- Transistor M_{rst} acts as a non-linear resistor (”diode”)
- V_{pd} will be a logarithmic function of the light intensity (continuous sensing, no integration).

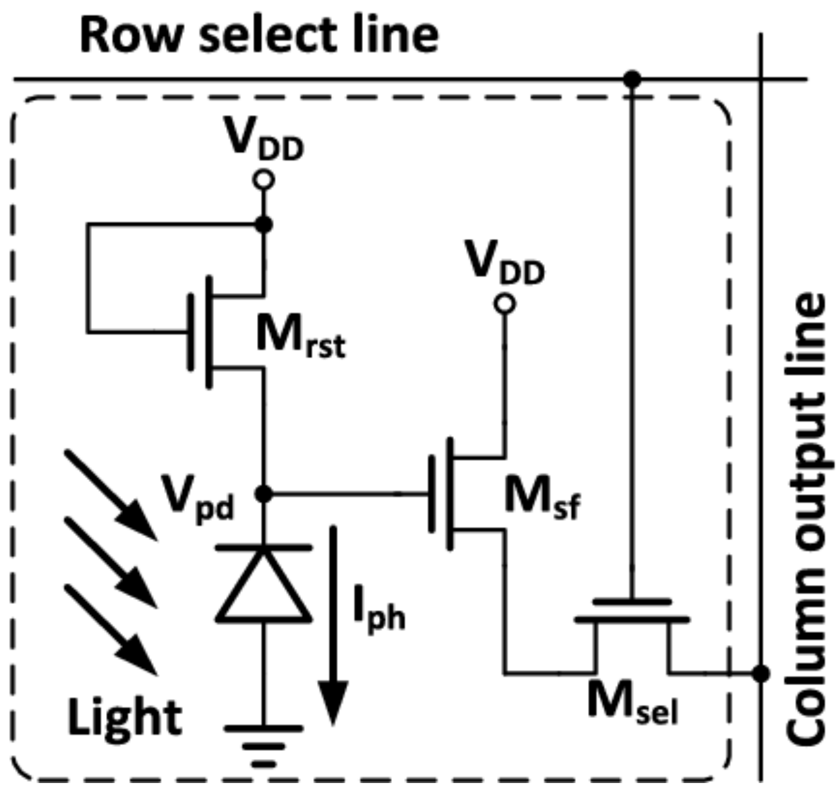

Image: Wei-Fan Chou et al.
IEEE Sensors Journal 2014.

Log-intensity cameras

- Example: PhotonFocus AG Lin-log camera

Linear mode

Logarithmic mode

Dual diode

- CMOS
- Each pixel contains two photo diodes, each with different size and/or other characteristics that give them different sensitivities to light
- A high-sensitivity diode
  - Operates well in dark areas
- A low-sensitivity diode
  - Operates well in bright areas

Piecewise linear response

For example:

- The initial precharge is only applied to 50% *)
- Discharge the diode during an initial and longer exposure time
- Dark areas: The diode never becomes fully discharged
- Bright areas: The diode becomes fully discharged (saturated)
- Add the missing 50% precharge (keeping any remaining charge from the first exposure) *)
- Continue discharging during a second and shorter exposure time

*) Requires a modified pixel circuit compared to those shown in lecture 3.

Piecewise linear response

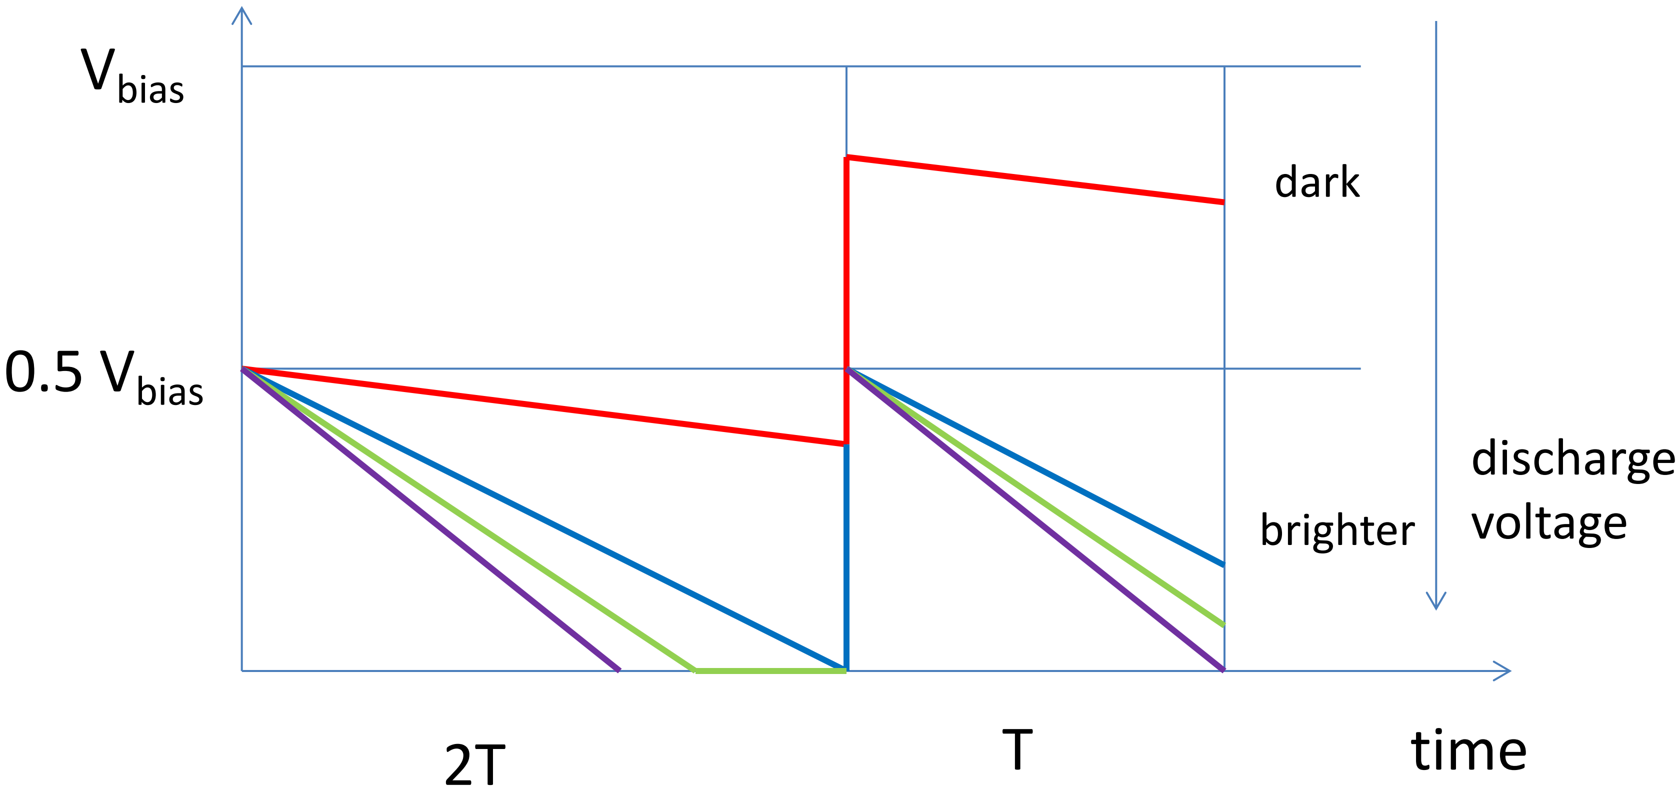

Piecewise linear response

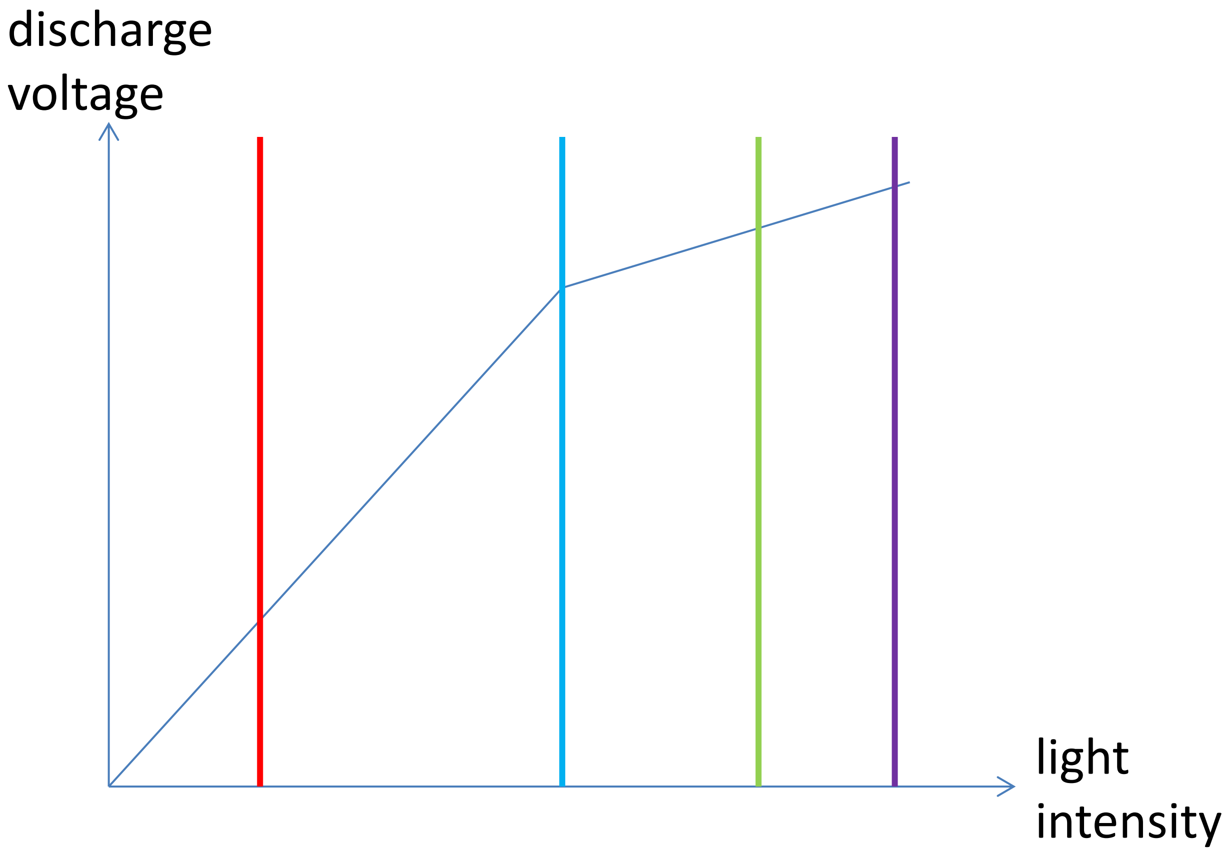

Seeing in the dark

However, if it is very dark, long exposure time or logarithmic pixels are not enough.

It is necessary to amplify the image itself.

Night goggles

Photo: DSA

Photo-electric detectors

- Basic idea
  - Each electron that has been excited by a photon is made to leave the material and is accelerated by means of an electric field

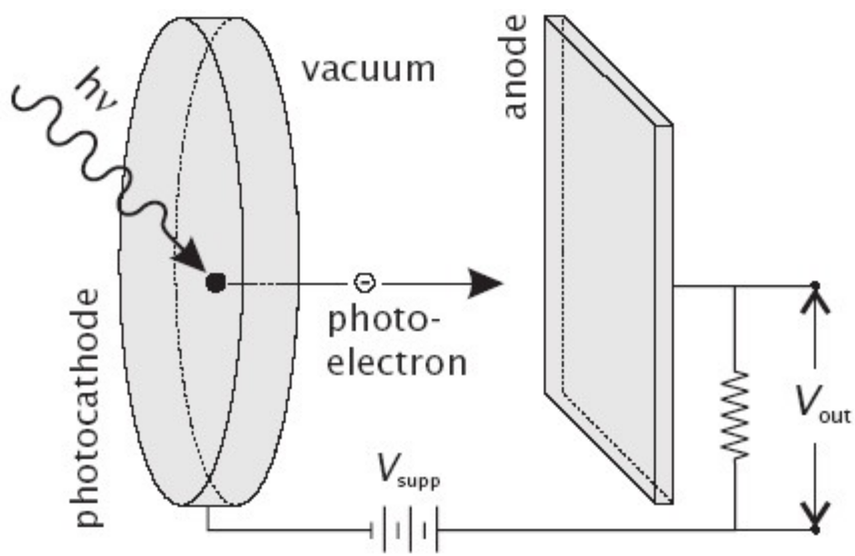

Photo-multiplier

- In a *photo-multiplier*, the field is **strong** enough to make the electron, on impact, knock out two or more electrons
- These, in turn, knock out **several electrons** and in the end an amplification of > 10^6 can be accomplished

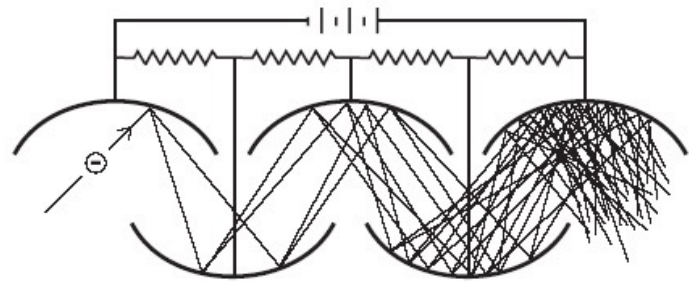

photomultiplier tube

Different types of photo-multiplier

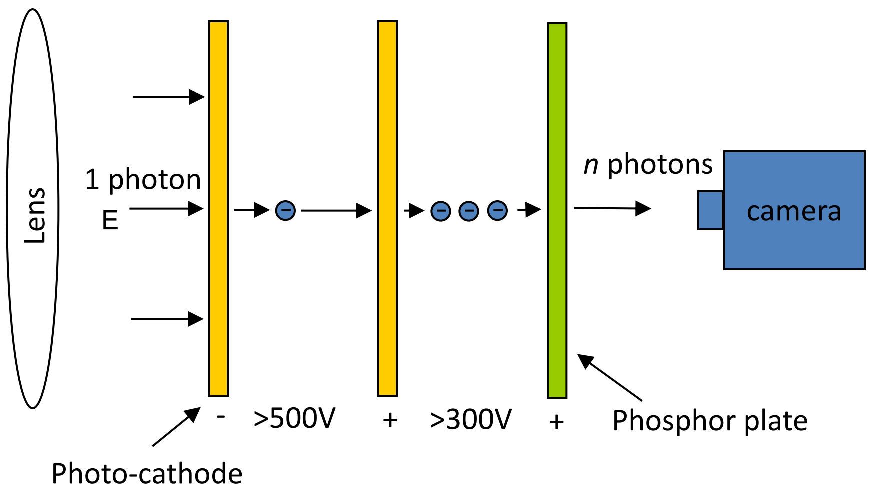

Type I

Type II

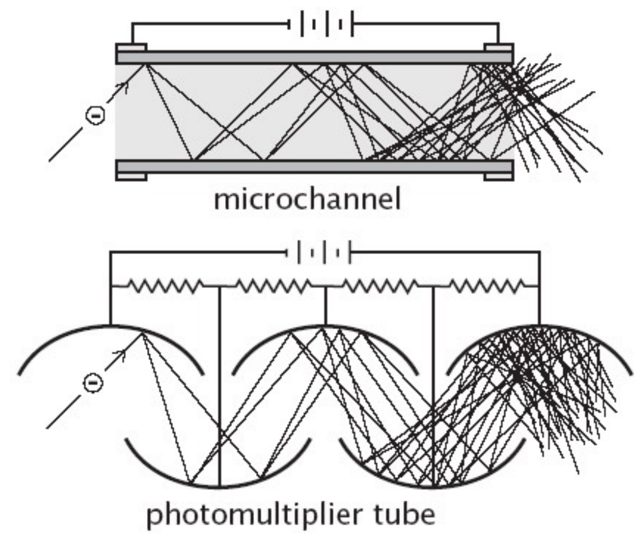

Type III

Micro-channel plate

- A micro-channel is a photomultiplier in the form of a tube
  - Can be as small as 10 μm in diameter and a few mm long
  - Electron gain $>10^4$
- A micro-channel plate consists of an array of such micro-channels stacked side by side

Micro-channel plate

- Applications
  - Photo-multiplication of visual light
  - X-ray detectors

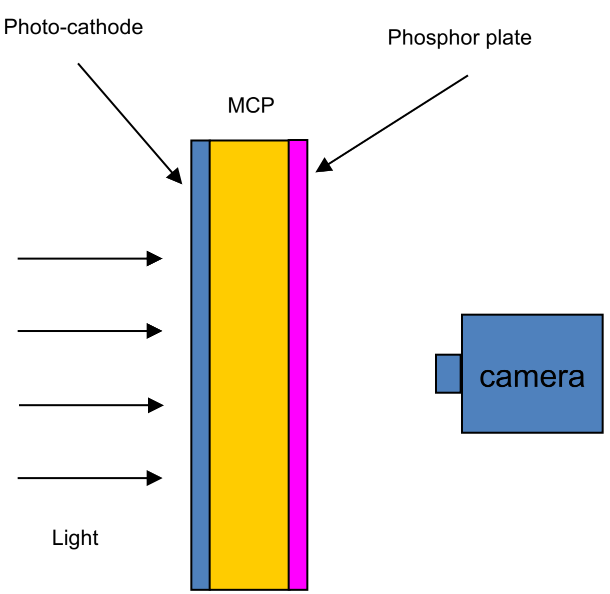

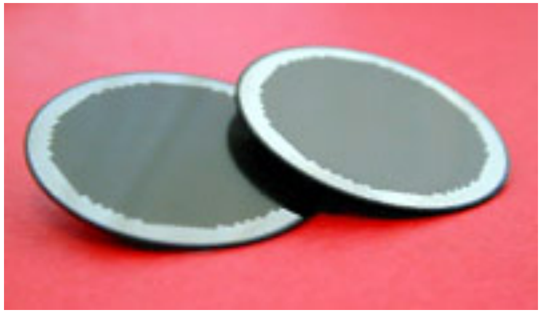

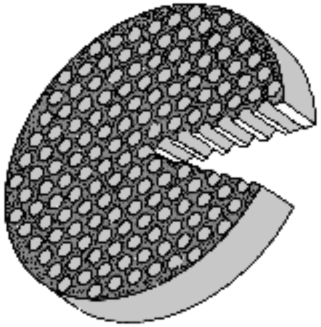

Tectra GmbH, MCP

50 mm diameter

>3000 channels across

Line camera

- Can produce 2D images by
  – Translating the scene relative the camera
  – Translating the camera relative the scene
  – Rotating the camera relative the scene
- Also known as
  – *Push-broom camera (PBC)*: it “paints” the 2D image by moving the camera.
    - Push-broom camera normally implies that the camera is moving relative to the scene to produce a 2D image

Line camera

- Why a line camera?
  - High resolution along one axis (>10 kpixels)
  - High resolution in bits/pixel (>16 bits)
  - High scanning rate (>10.000 scans/sec)
  - Allows integration on the chip of a processing unit per pixel (smart camera)

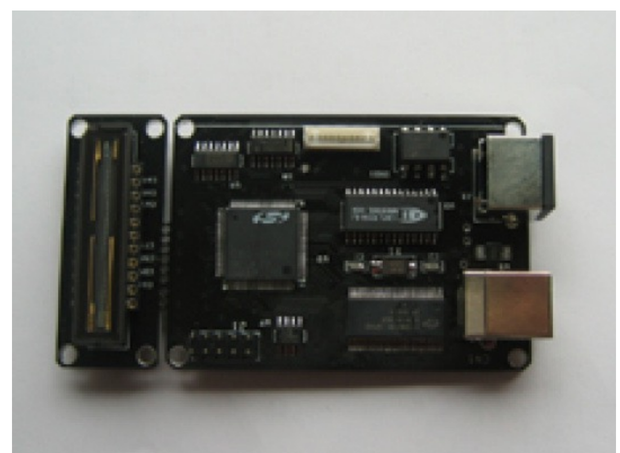

Mightex 3648-pixel line camera

Line camera, applications

- High resolution inspection
- Photo finish
- Fax, copy machines
- Film scanners
- …

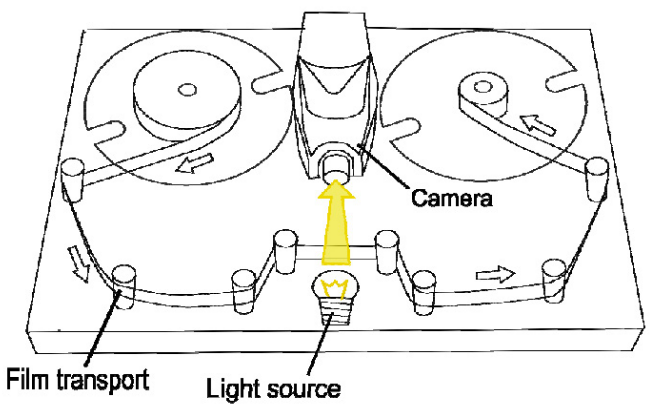

ImageSystems GoldenEye film scanner

This is NOT an image of a 3D scene at some point in time.
It is an image of a LINE at different points in time

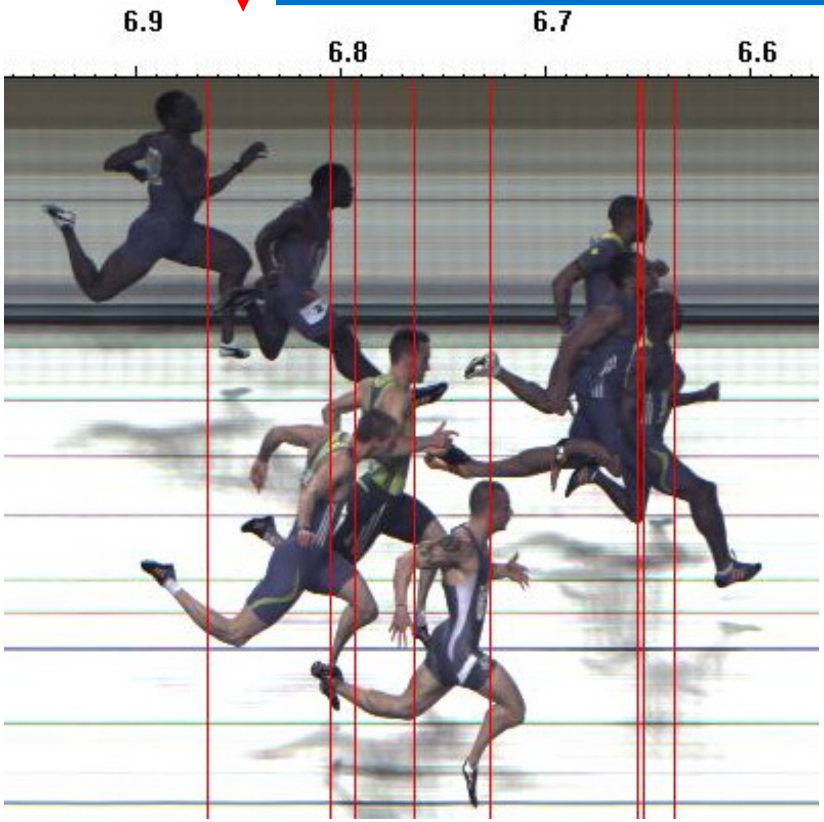

Photo finish

Push-broom camera, applications

ground mapping from air

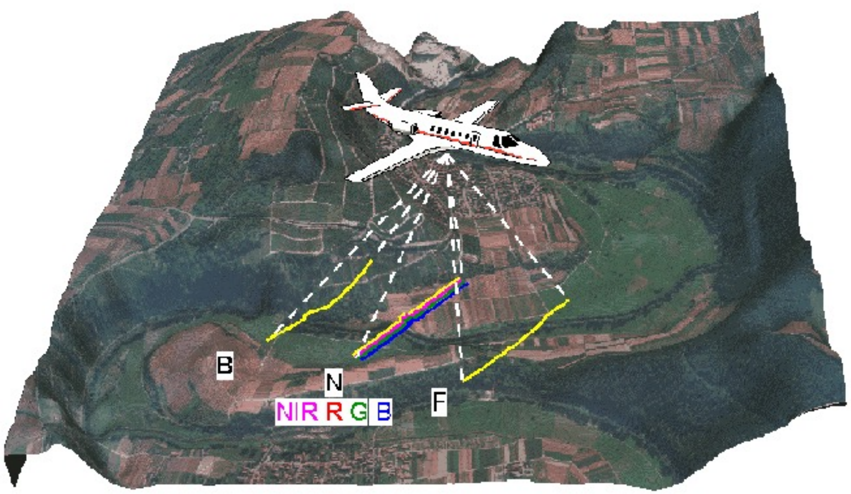

ground mapping from space

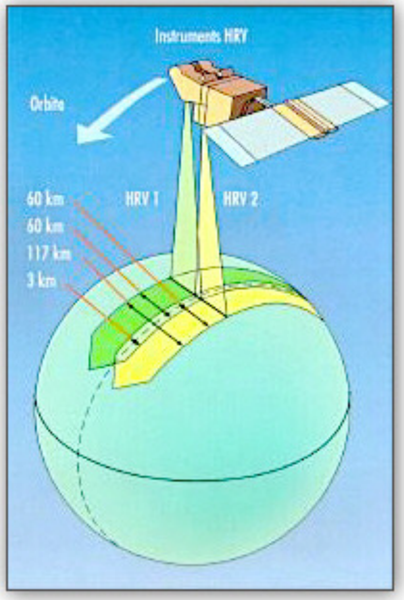

Smart cameras

for machine vision applications

- Integrates processing unit(s) (PU) with a camera chip. Several different solutions exists:

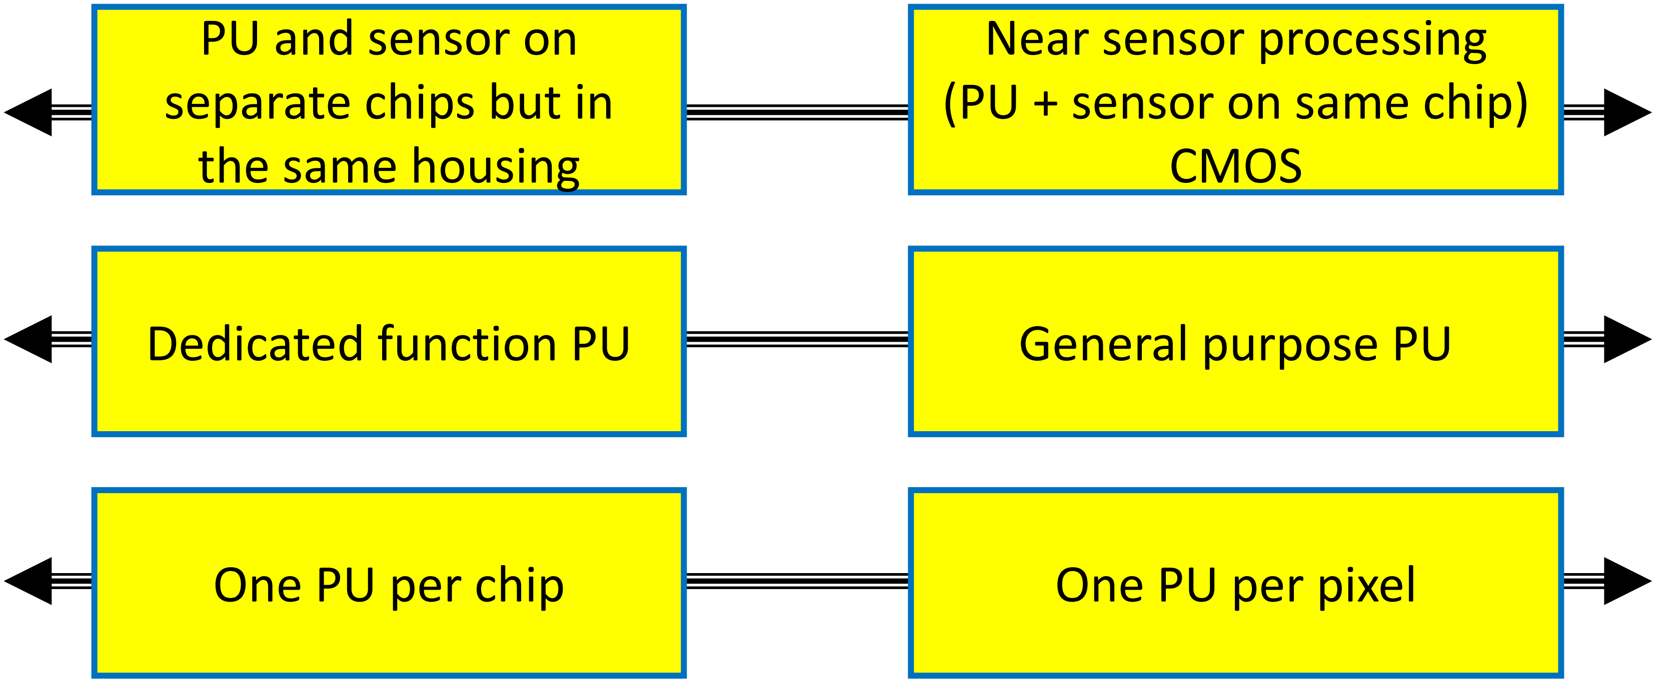

Smart cameras

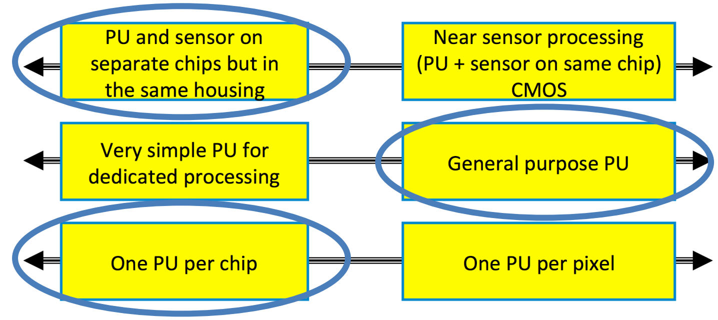

- Some cameras come with a complete IDE
  - Processing is defined in a GUI on a standard PC
  - Processing code is downloaded into the smart camera and executed
- Some cameras are integrated with illumination
  - IR
  - visual
- Most cameras have simple interfaces
  - TCP/IP protocol over Ethernet
  - Integrated web server

Cognex Checker 252

M12 Smart Camera

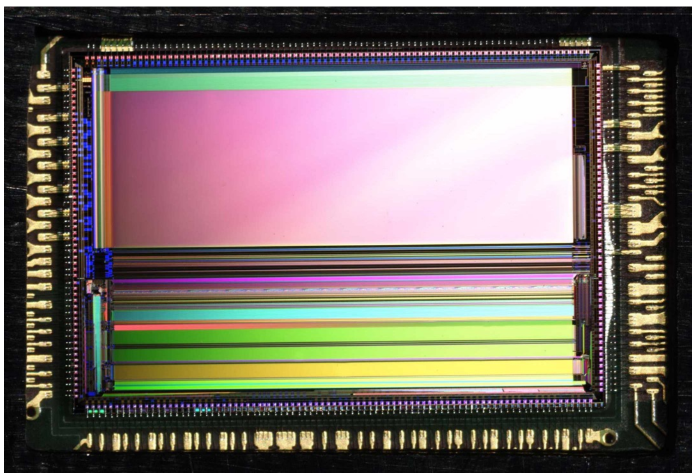

PU and sensor on separate chips but in the same housing — Near sensor processing (PU + sensor on same chip) CMOS

Very simple PU for dedicated processing — General purpose PU

One PU per chip — One PU per pixel

Smart cameras

- Used where a low-cost solution is required, or alternatively, extremely high performance.

- Applications
  - Range imaging
  - Bar code reading / data matrix / OCR
  - Event/motion detection
  - Counting objects/people passing by
  - Surveillance
  - Gaze measurement
  - …

Event cameras

- **Short definition**: An event camera only senses motion in the scene.

- **Long definition**: An event camera, also known as a *neuromorphic camera*, *silicon retina* or *dynamic vision sensor*, is an imaging sensor that responds to local changes in brightness. Event cameras do not capture images using a shutter as conventional cameras do. Instead, each pixel inside an event camera operates independently and asynchronously, reporting changes in brightness as they occur, and staying silent otherwise. (*Wikipedia*)

Event cameras
Example pixel circuit

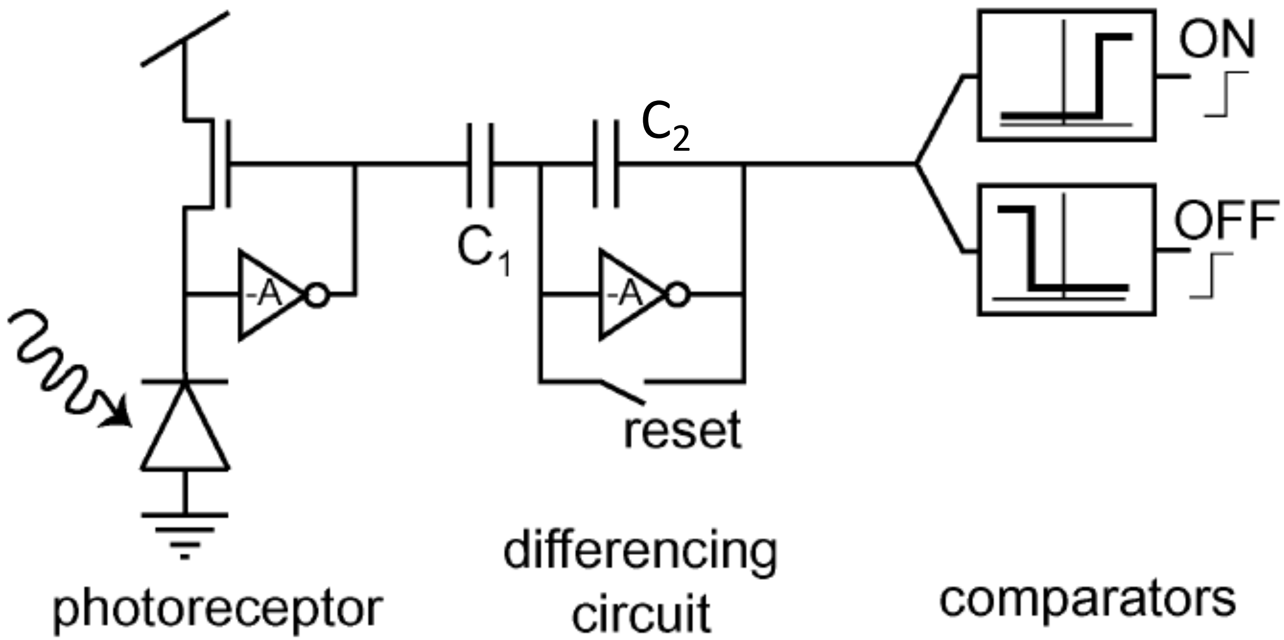

1) Photoreceptor measures continuously (no integration)
2) C_1 stores current light intensity during reset as a charge
3) After reset, if the charge in C_1 changes, a corresponding change will happen in C2 and an ON or OFF event is generated (with some threshold)
4) The event is read-out and a new reset is done (typically within μs)

Event cameras

Event address representation (AER)

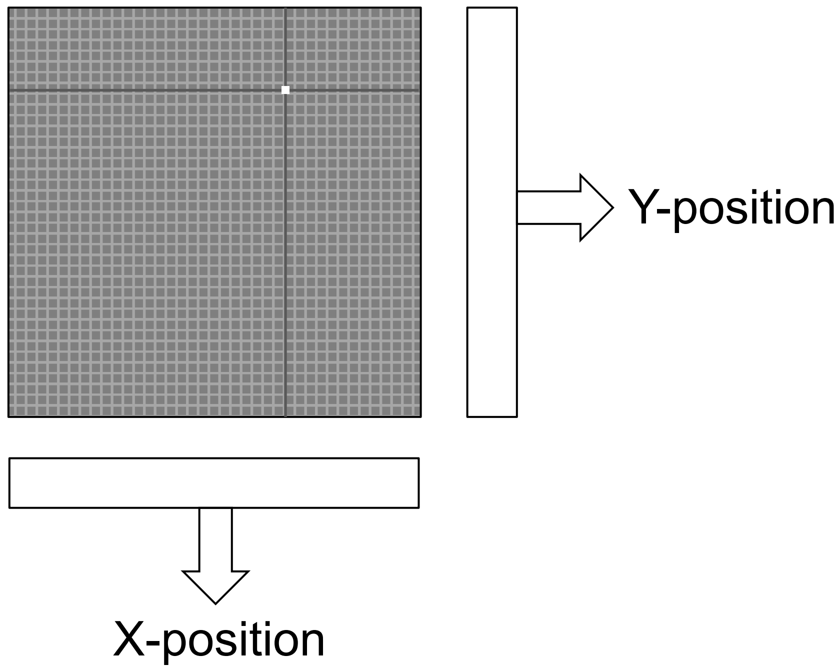

An event (ON or OFF) triggers a redout of the pixel address

Event cameras

Illustration

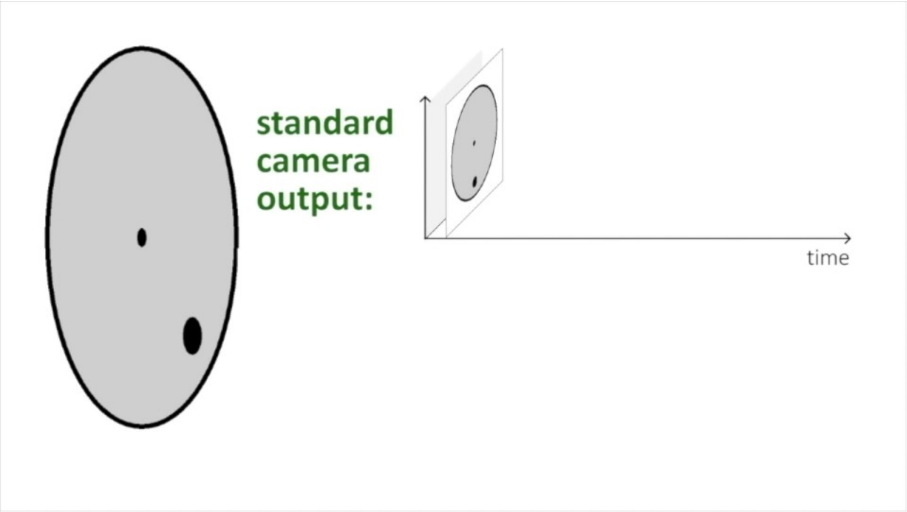

Motion

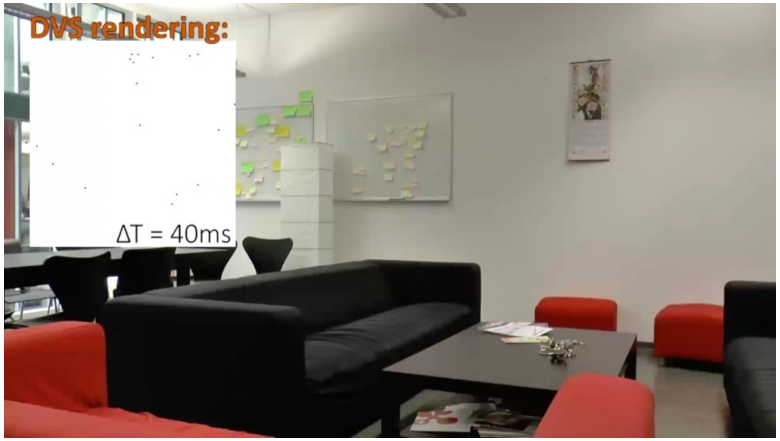

For the full video, see https://www.youtube.com/watch?v=LauQ6LWTkxM

Event cameras

Modern event cameras have microsecond temporal resolution, 120 dB dynamic range, and less under/overexposure and motion blur than frame cameras.

Commercially available event cameras

Inivation:
DAVIS sensor: frames, events, IMU.
Resolution: ~QVGA (346x260 pixels) Cost: 6,000 USD

Insightness:
RINO sensor: frames, events, IMU.
Resolution: ~QVGA (320x262 pixels) Cost: 6,000 USD

Prophesee:
ATIS sensor: events, IMU, absolute intensity at the event pixel. Resolution: 1M pixels
Cost: 4,000 USD.

CelexPixel Technology:
Celex One: events, IMU, absolute intensity at the event pixel. Resolution: 1M pixels
Cost: 1,000 USD.

Samsung Electronics
Samsung DVS: events, IMU ▪ Resolution: up to 1M pixels. Cost: not listed

Computational cameras

- Cameras that capture optically coded images and require computations to produce “real images”

- Examples

  Light-field cameras, Coded aperture, Catadioptric imaging, holographic imaging, flexible depth-of-field…

Light field

Repetition from lecture A:

- At each point **x** in 3D space,
  in each direction **n**,
  there is an amount of light passing through **x**
- The *plenoptic function* $I(\mathbf{x}, \mathbf{n})$
  - A.k.a. *Light field*

Light field camera

- A *light field camera* makes a denser sampling of $I(\mathbf{x}, \mathbf{n})$ than a standard camera
  - Ideally all **x** and all **n** (not practically possible)
- Practical implementation:
  an array of pin-hole cameras *)

*) Approximated by lens cameras

Light field camera

A more compact implementation:

- Use a large sensor chip: lots of pixels
- Divide the chip into several small “cameras”
- Sophisticated lens system,“lenslet array”, to emulate an array of pin-hole camera projections onto the chip

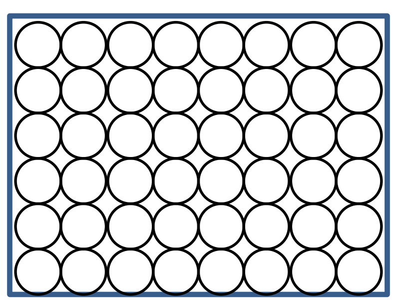

lenslet array

Light field camera

Examples

Lytro

Raytrix

Stanford
plenoptic
camera

Applications

- A point in the scene is viewed from multiple directions:
  - 3D reconstruction possible from "single image"
  - Extended depth of field
  - Adjustable object plane (and thus focus) *after* exposure

Images: Lytro

Omni-directional cameras

- Omni: Latin for *every, all*
  - In theory: a camera that sees in *all directions*
  - In practice: not all directions, but a much larger field of view compared to a standard camera
- There are several design approaches to omni-directional cameras
  - Multiple pinhole cameras
  - Fish-eye lens
  - Catadioptric camera

Omni-directional cameras

- The image is best represented as a sphere instead of a plane
  - Can be mapped to a plane image but with severe distortion
  - Ideally: all light rays intersect at a single point
    - Cannot be achieved in practice!

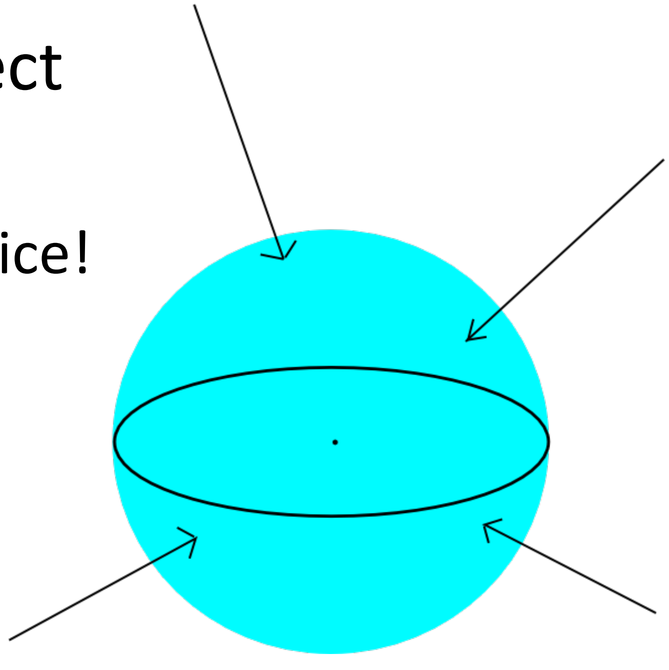

Multiple pinhole cameras

- Set up multiple pinhole cameras to cover the desired set of directions
- Use image stitching to produce a representation of the image sphere (or parts thereof)

Ladybug 2 and 3 from
Point Gray Research Inc.

Multiple pinhole cameras

From: de la Torres, et al,
Learning to Track Multiple People in Omnidirectional Video

Fish-eye lens

- A single camera with a fish-eye lens can cover approximately a hemi-sphere

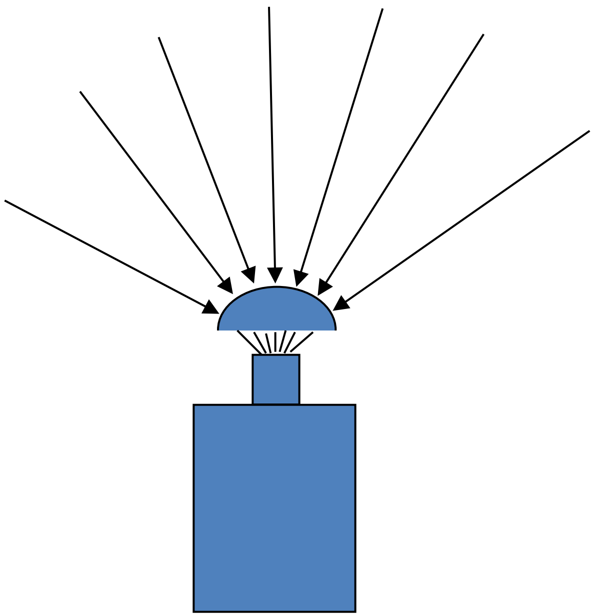

Image: Dan Slater

Catadioptric cameras

- Mirrors and/or lenses re-project the light rays into a single camera lens
  - Special case: fish-eye lens

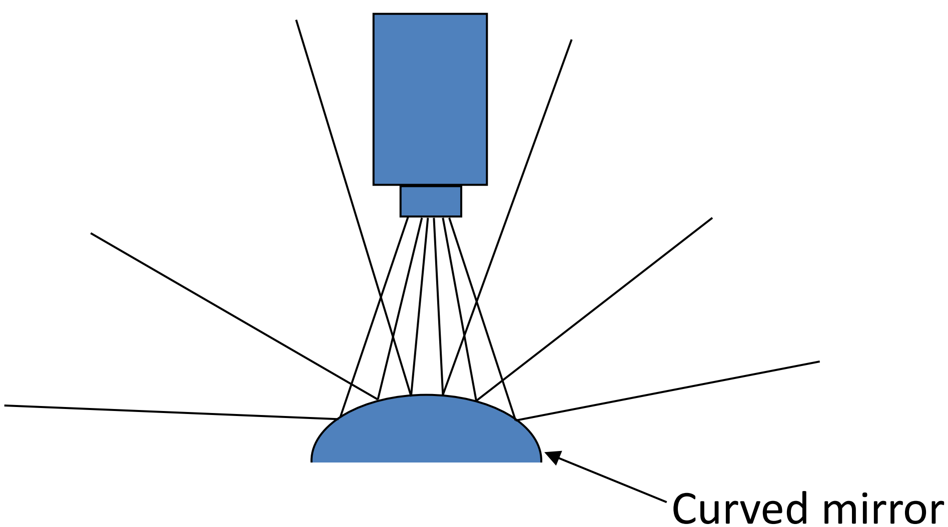

The mirror can be

- Spherical
- Hyperbolic
- Conic
- …

Catadioptric cameras

Kaidan Inc. catadioptric objective

Image: Tomas Pajdla

Catadioptric vs. fish-eye lens

- Fish-eye lens: cheap and simple
  - For example, a door peep-hole: 50 SEK
  - Can give approx 180^o field of view
- Catadioptric camera system
  - Exact control of how the plenoptic function is sampled by choosing the curvature of the mirror
    - User specified curved mirrors: expensive
  - Single optic center can be accomplished
  - Can give > 180^o FOV, with camera occlusion

Omni-directional cameras

- Applications
  - Video conferences
  - Surveillance
  - Environment mapping
  - …

Coded aperture

- Aimed at extending depth of field
- Originally developed for X-ray and gamma ray imaging.
- Simple coded aperture: pinhole camera
- More complex apertures can be used if combined with computational algorithms to reconstruct the image.

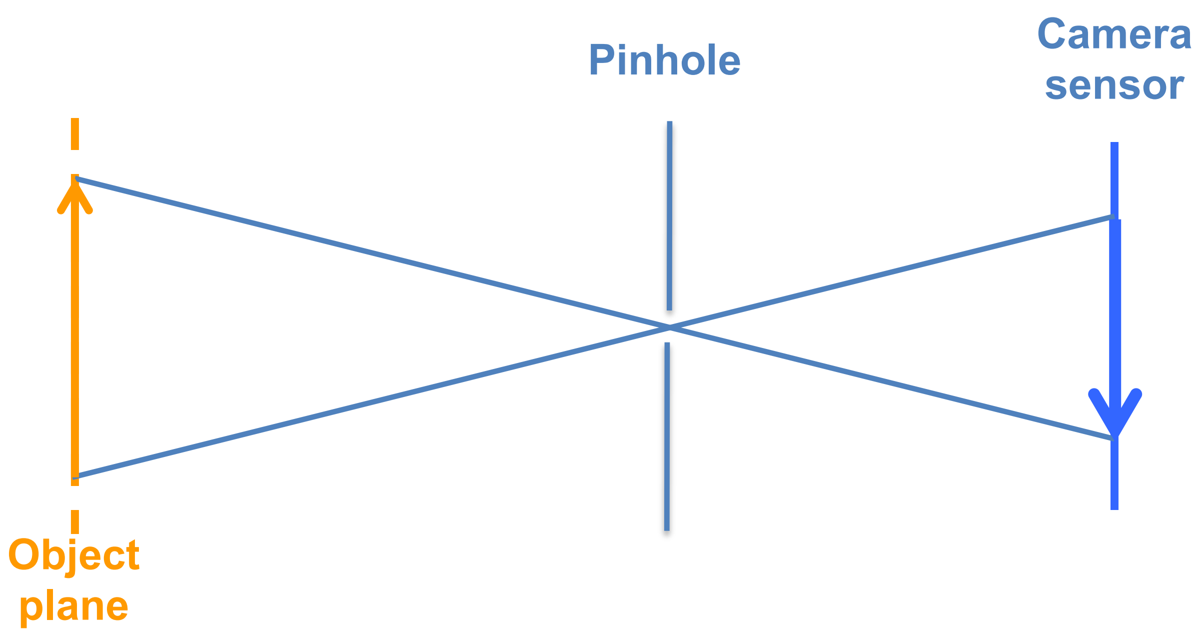

Solution 1: lens with small aperture

Lens' aperture

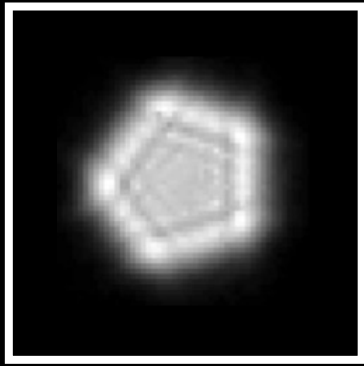

Image of a defocused point light source

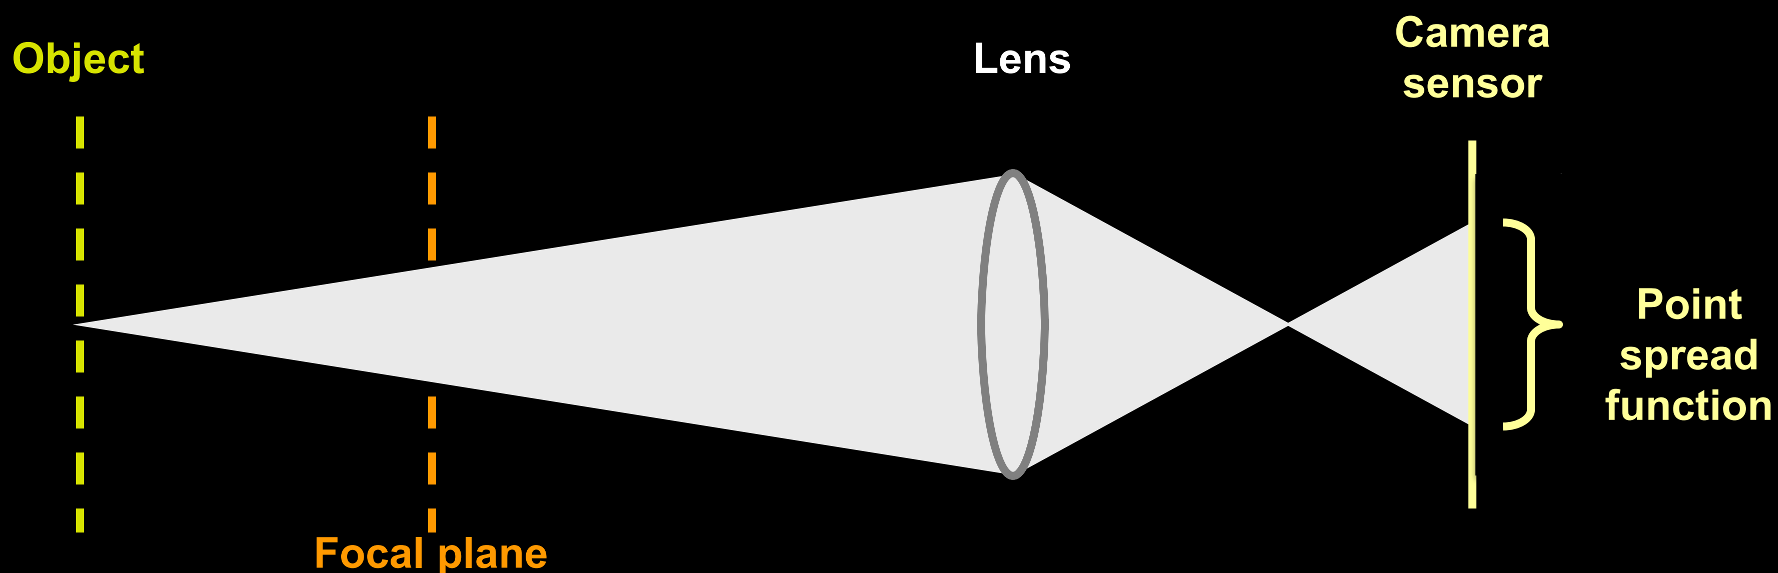

Levin et.al. SIGGRAPH

Solution 2: lens with coded aperture

Aperture pattern

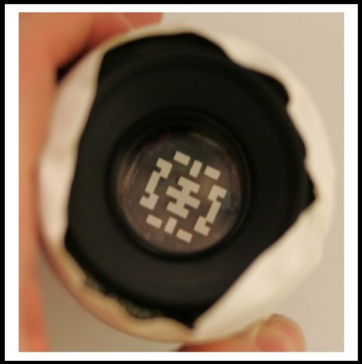

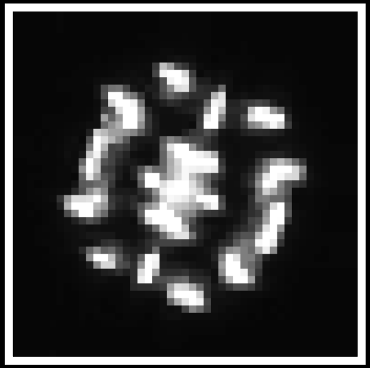

Image of a defocused point light source

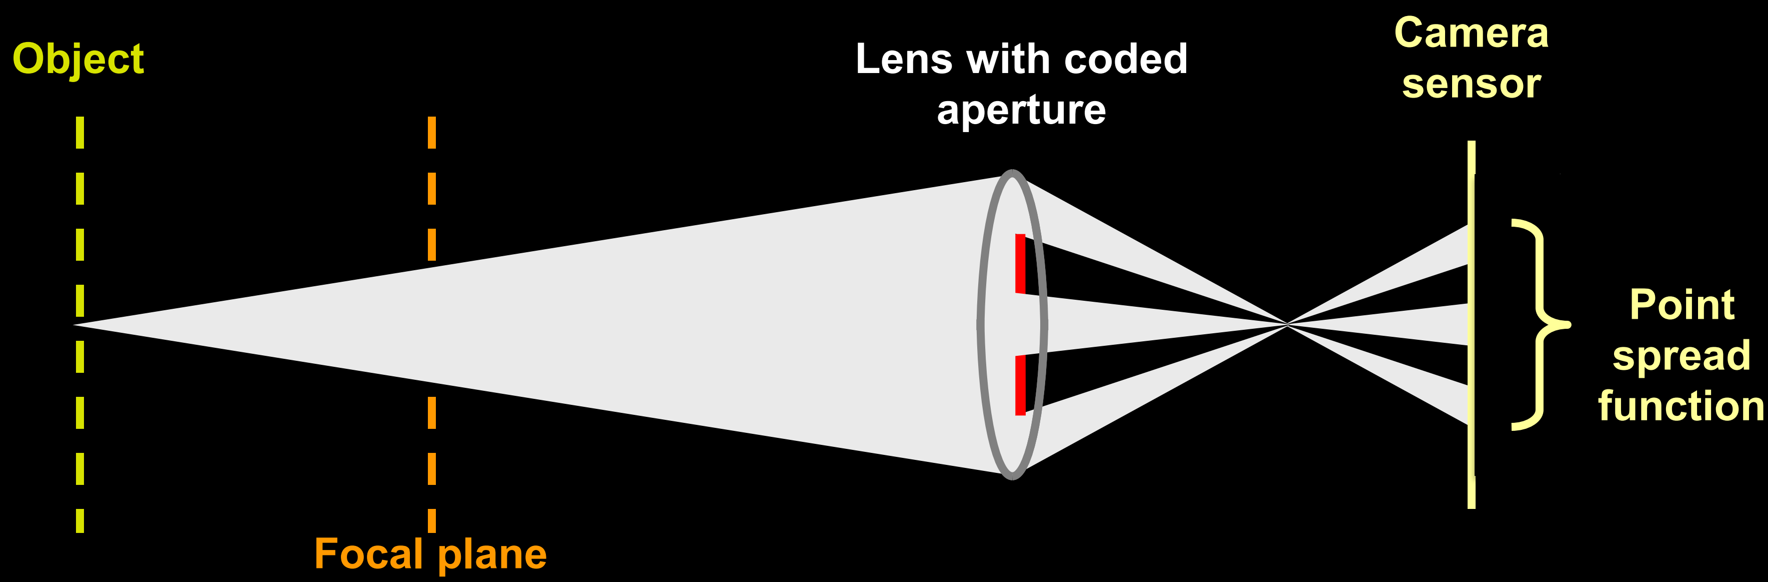

Levin et.al. SIGGRAPH

Solution 2: lens with coded aperture

Aperture pattern

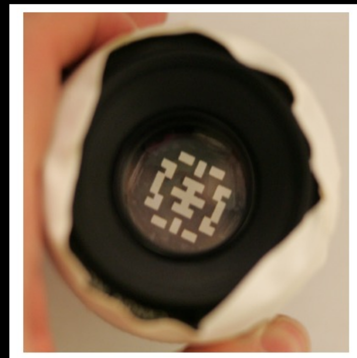

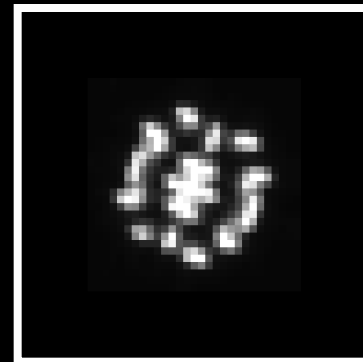

Image of a defocused point light source

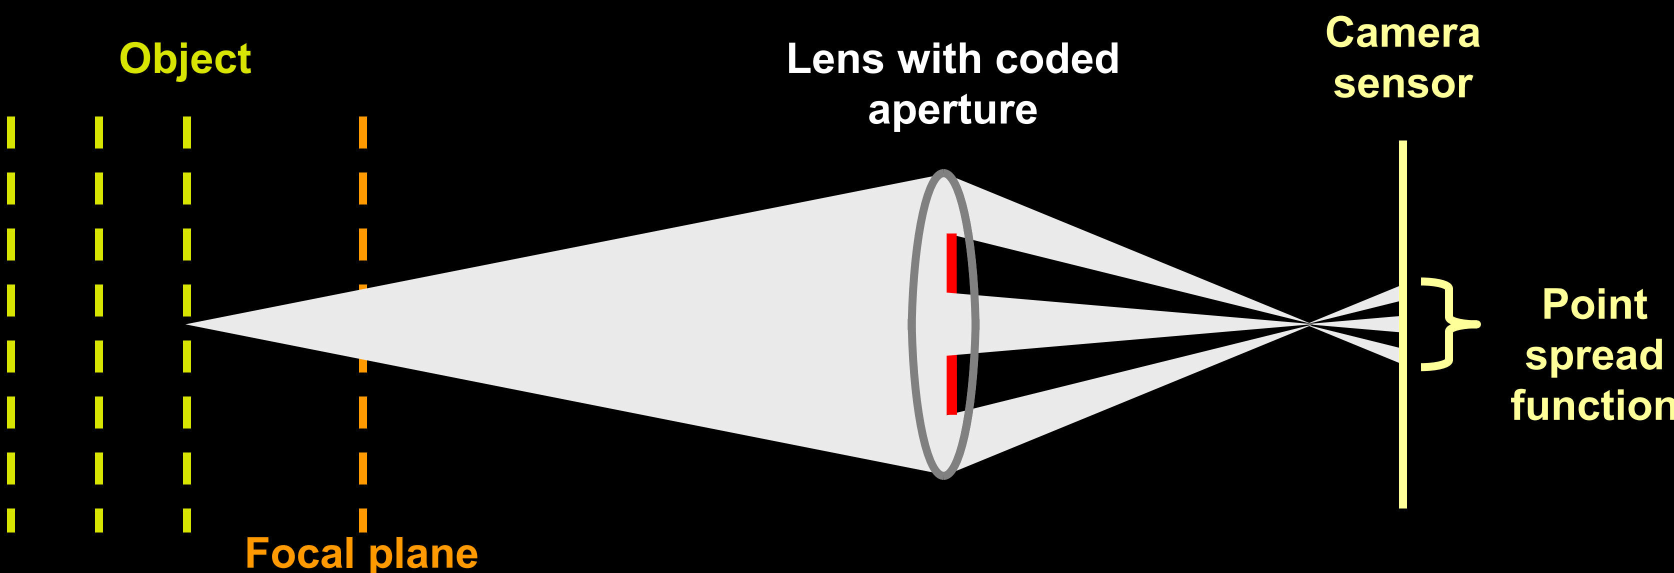

Levin et.al. SIGGRAPH

Close it up


CANON INC
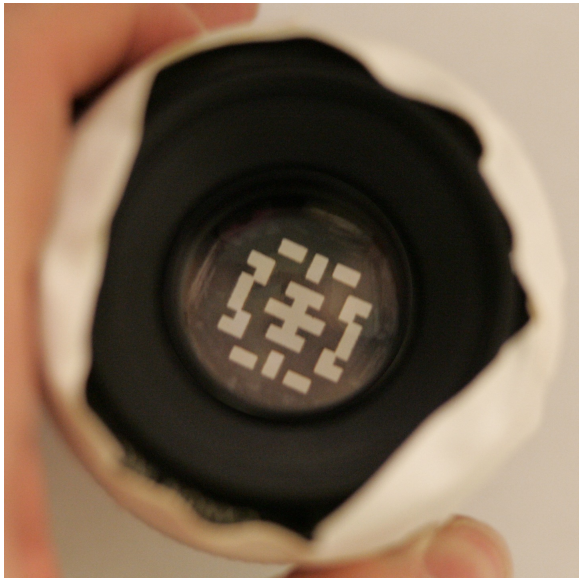

Reconstruction

- Image is formed by convolution of the object with the (scaled) aperture: $y = f_k * x$
- f_k is the blurred point spread function

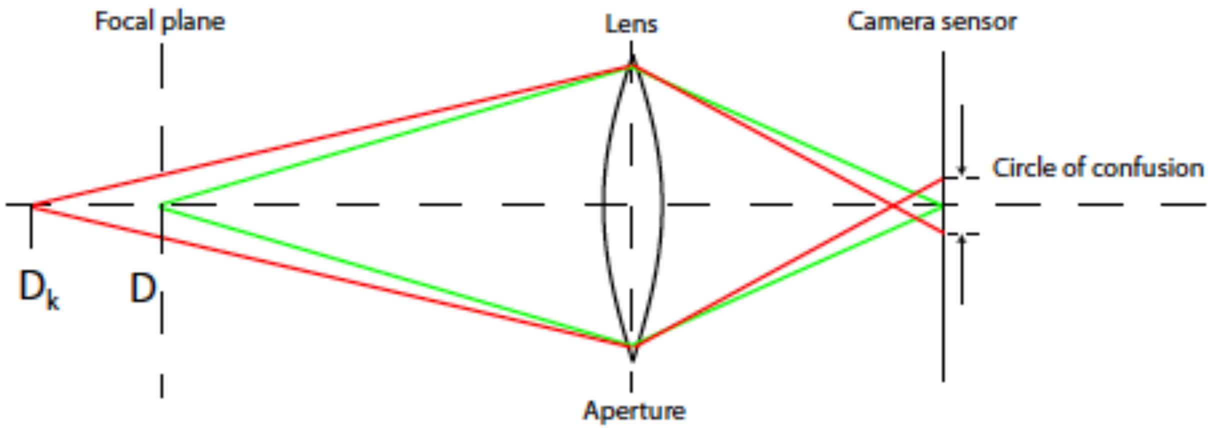

Reconstruction

- Deblur locally by minimizing $|f_k * x - y|^2$
- Search for best depths (k) in the local areas
- Reconstruct all parts of the image.
- Result depends on the particular aperture.

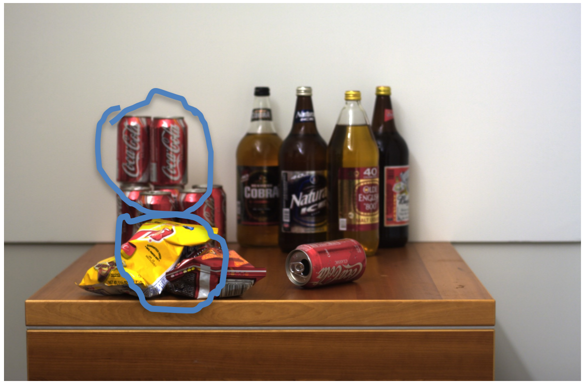

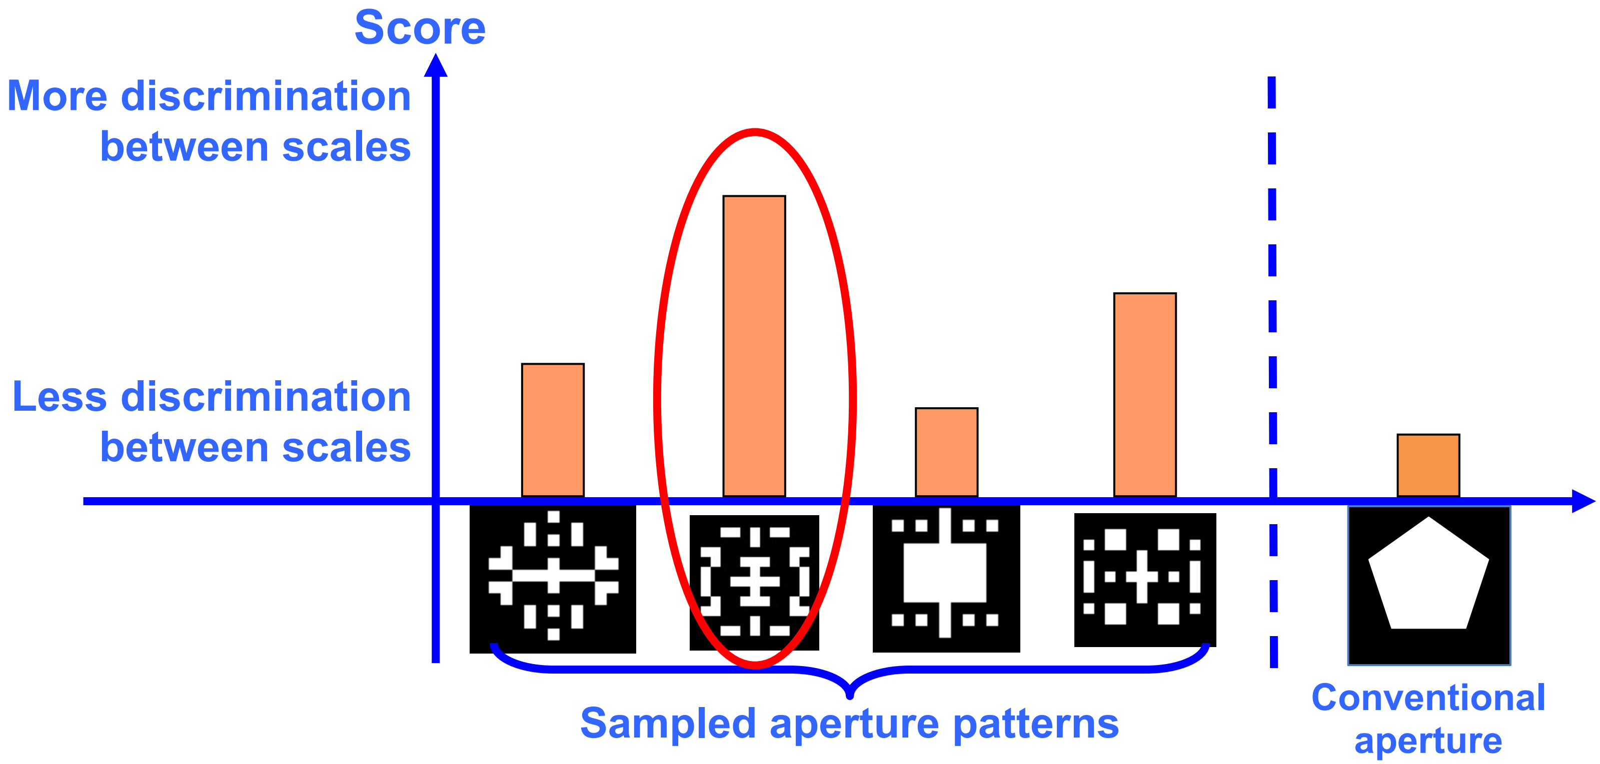

Levin et.al. SIG

Levin et.al. SIG

Close-up

Original image

All-focus image

Levin et.al. SIGGRAPH


Levin et.al. SIG


Levin et.al. SIG

Close-up

Original image

All-focus image

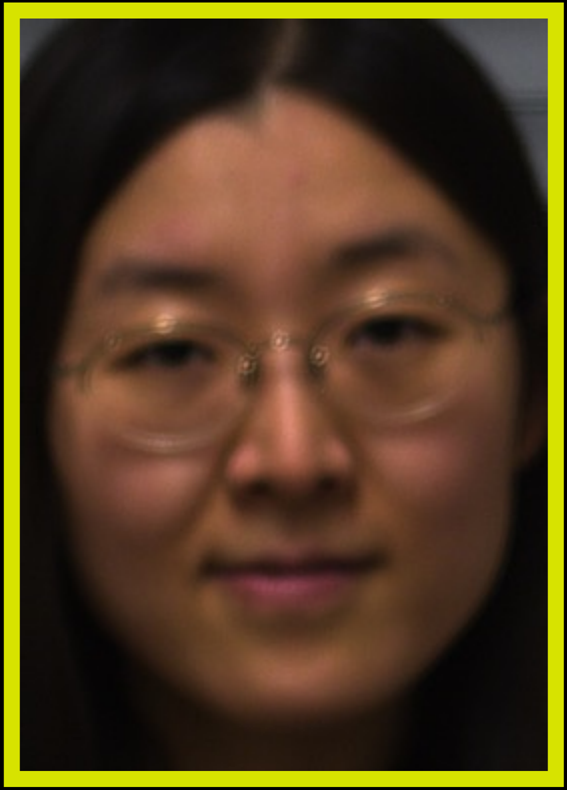

Naïve sharpening

Levin et.al. SIGGRAPH

Application: Digital refocusing from a single image

Application: Digital refocusing from a single image

Application: Digital refocusing from a single image

Levin et.al. SIGGRAPH

Added bonus: depth estimation

Levin et.al. SIGGRAPH

Flexible depth of field*

- Aimed at extending depth of field
- Based on moving the sensor during exposure

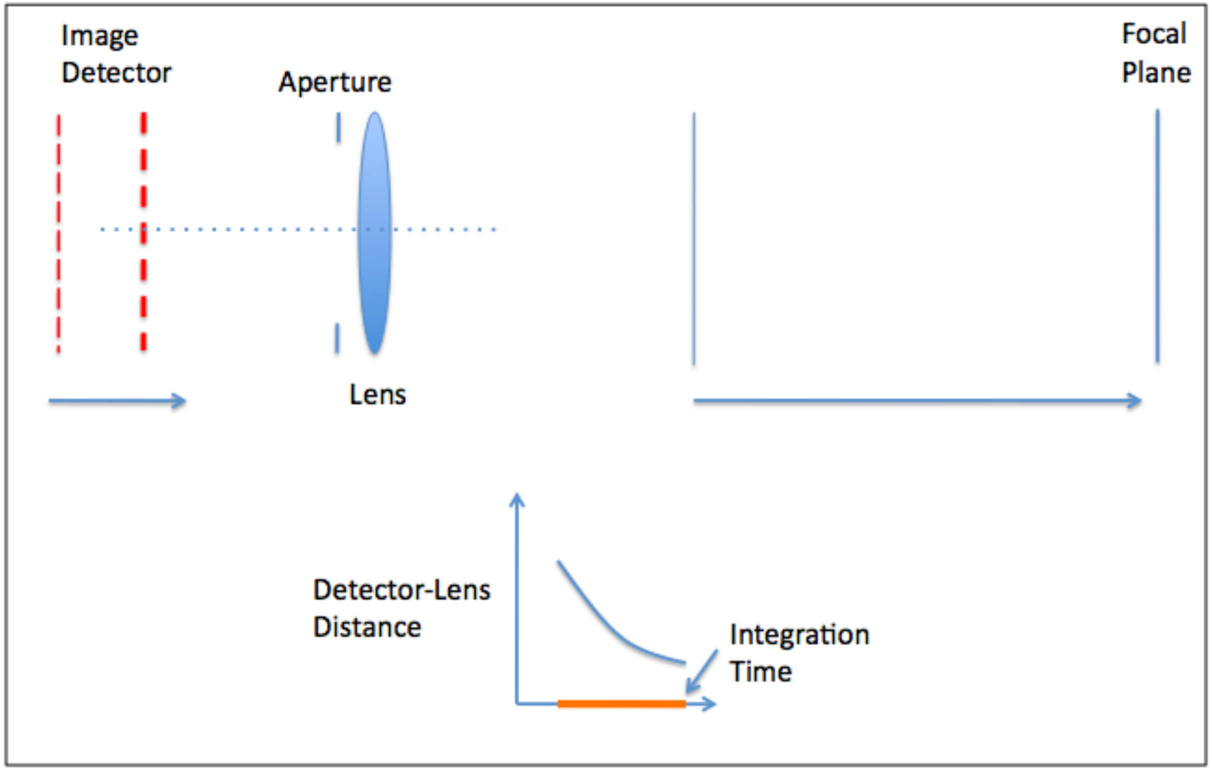

*"Flexible Depth of Field Photography," H. Nagahara, S. Kuthirummal, C. Zhou, and S.K. Nayar, European Conference on Computer Vision (ECCV), Oct, 2008.

Flexible depth of field imaging

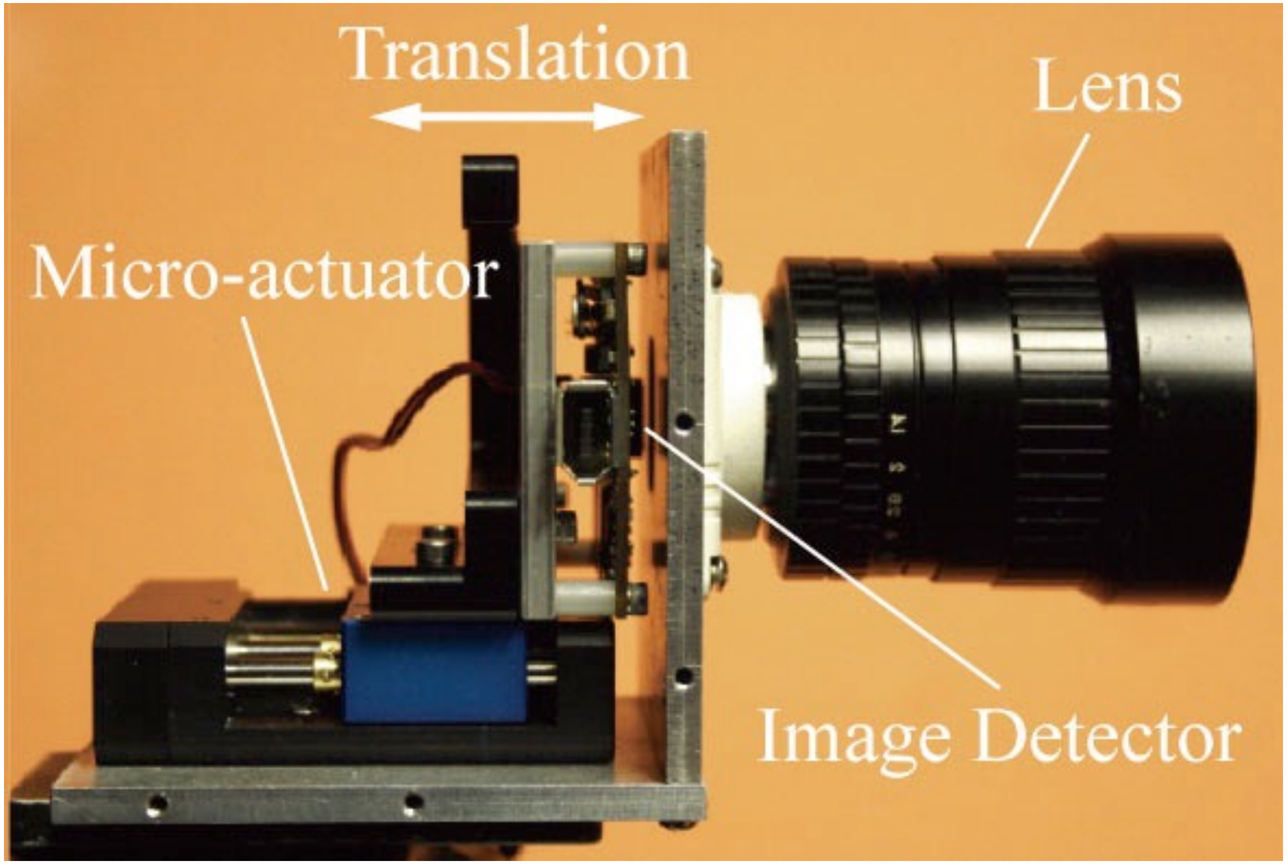

PSF becomes independent on object distance

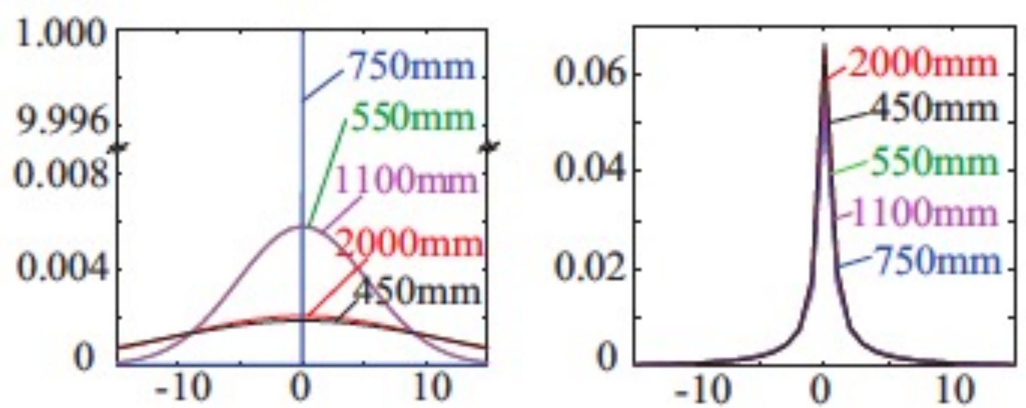

(c) Normal Camera PSF (Gaussian) (d) EDOF Camera IPSF

Thus, deblurring can be done with a fixed filter for all image parts!

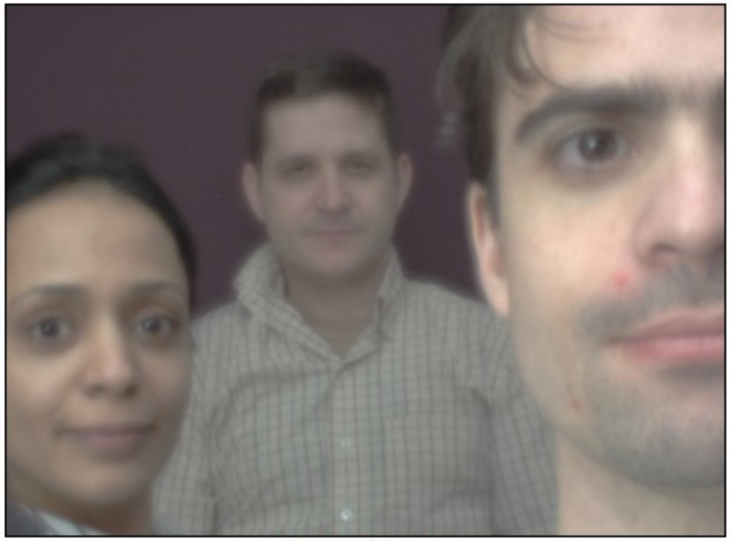

Captured Image
(f/1.4, T=0.36sec)

Computed EDOF Image

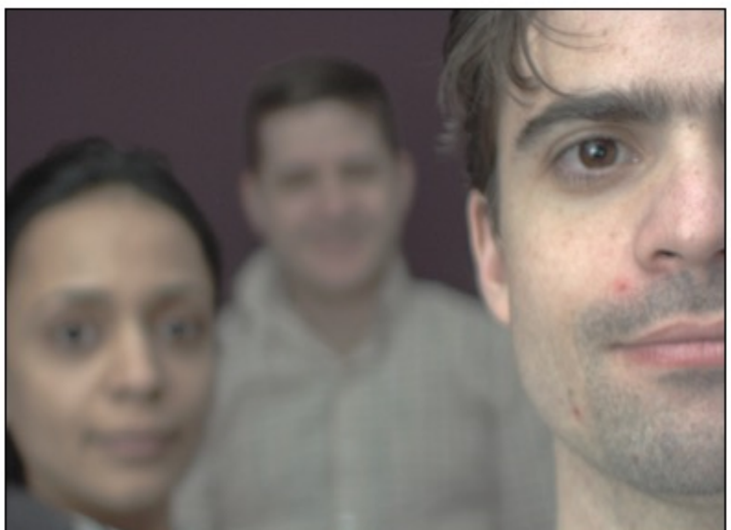

Image from Normal Camera
(f/1.4, T=0.36sec, Near Focus)

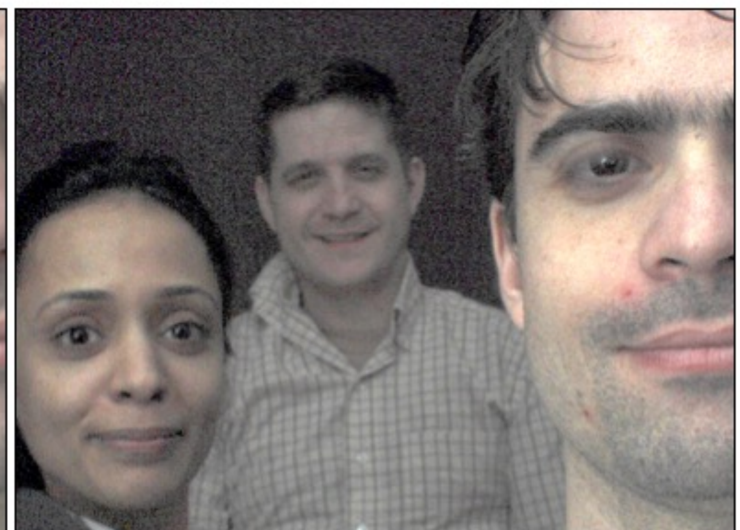

Image from Normal Camera
(f/8, T=0.36sec, Near Focus) with Scaling

Nagahara et al., Flexible depth of field imaging

Computational camera
- affected camera parts

(Inspired by W.T.Freeman, EI 2012)

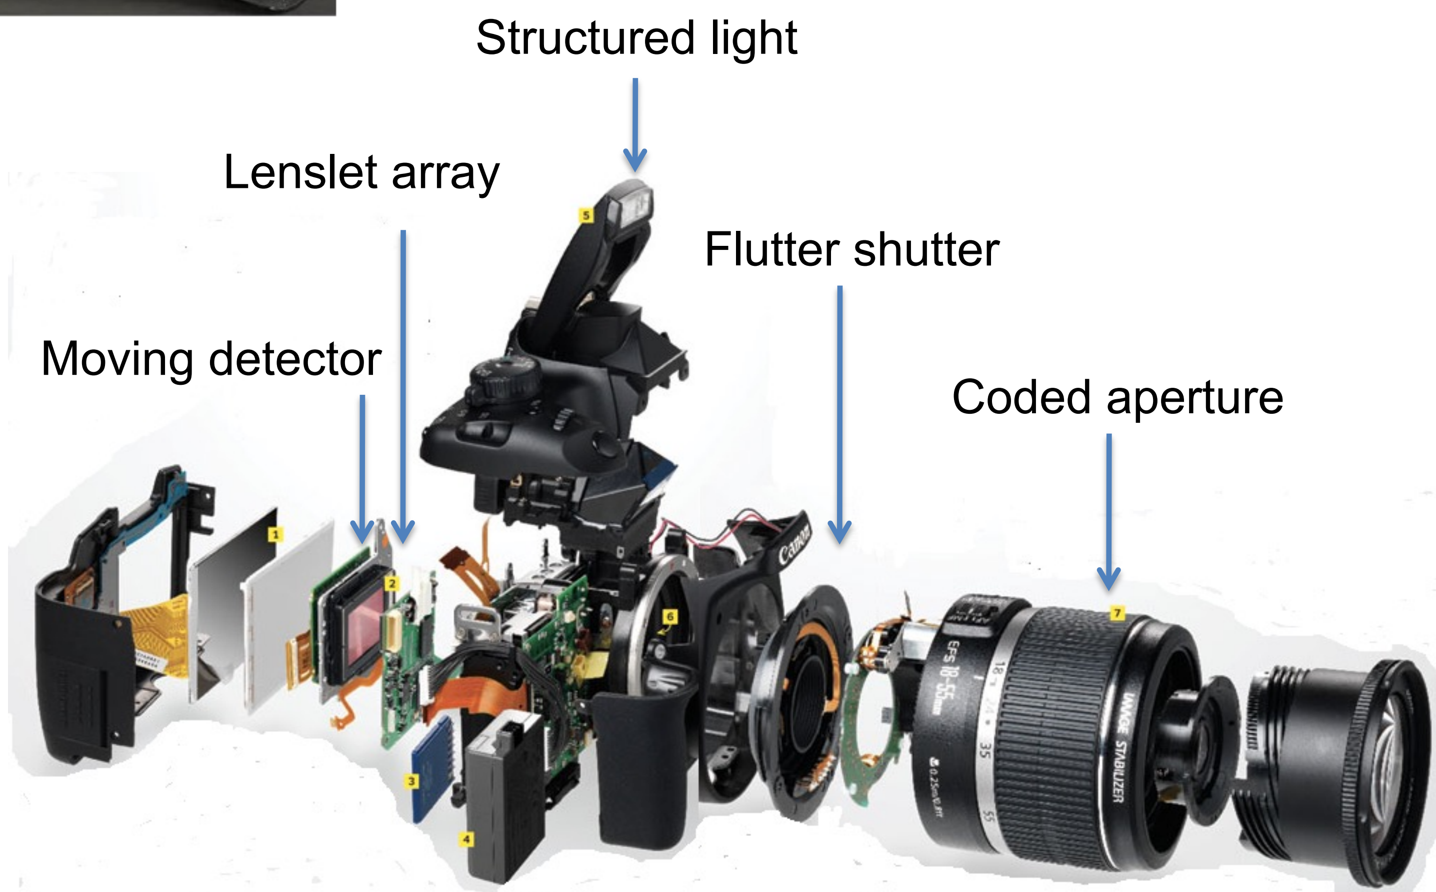

Photograph by Gregor Halenda. popularmechanics.com. Dec. 2008. http://www.spd.org/images/blog/117.jpg

Computational cameras

- Plenary talk by Shree Nayar
  – Professor at Columbia University
- Check out the video!
  – Available on the course web page