Image Sensors

Lecture 13
Specialized Cameras

Robert Forchheimer, Klas Nordberg
(small modifications by Maria Magusson)

1

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”
  - This becomes a problem if one and the same scene contains both bright and dark parts.

2

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)

3

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

4

External HDR

Simple reconstruction: add the images

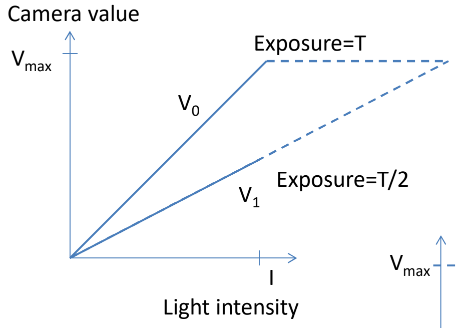

$$V_{HDR} = (V_0 + V_1) / 2$$

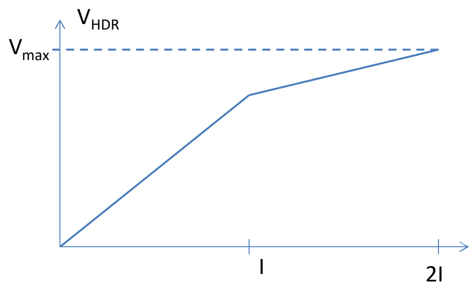

External HDR

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).

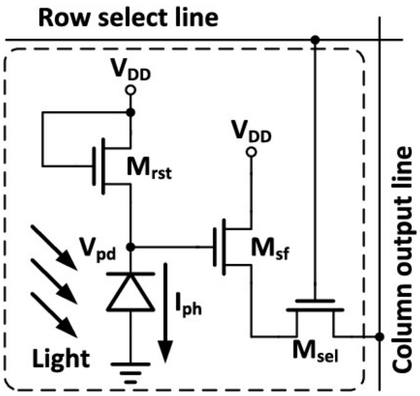

Image: Wei-Fan Chou et al. IEEE Sensors Journal 2014.

Log-intensity cameras

- Example: PhotonFocus AG Lin-log camera

Linear mode

Logarithmic mode

Dual diode

- CMOS
- Each pixel contains two photo diodes with 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

One 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.

Piecewise linear response

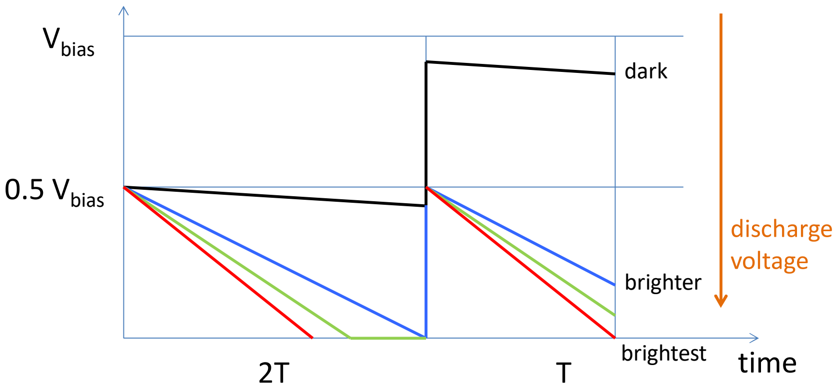

Piecewise linear response

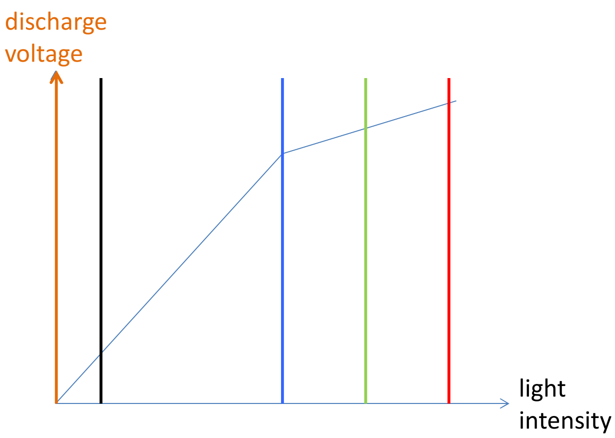

Seeing in the dark

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

Then 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
  - Further on, the electron enters a photo-multiplier …

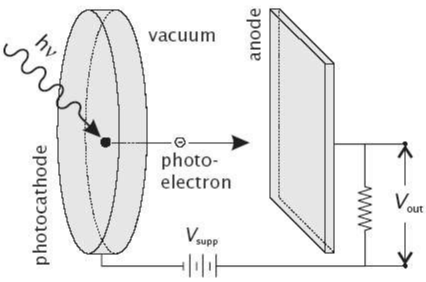

Photo-multiplier

- 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

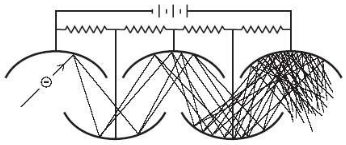

photomultiplier tube

3 different types of photo-multipliers

Type I

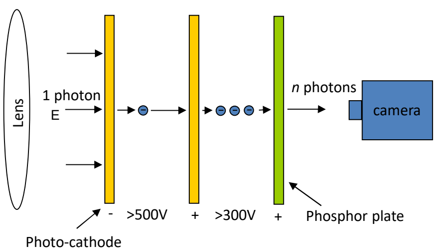

Type II

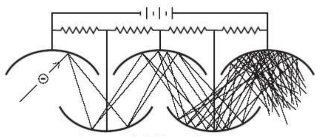

photomultiplier tube

Type III

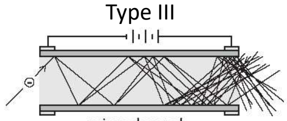

microchannel

Micro-channel plate (MCP)

- 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

19

Micro-channel plate (MCP)

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

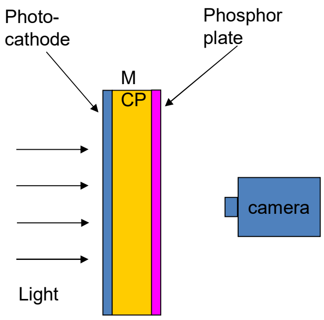

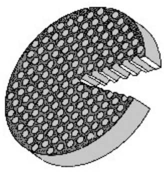

Tectra GmbH, MCP

50 mm diameter

>3000 channels across

20

Line camera

- Are used where very high 1-dimensional resolution is required, and where fast read-out is required.
- Can produce 2D images by
  - Translating the scene relative to the camera
  - Translating the camera relative to the scene
  - Rotating the camera relative to 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

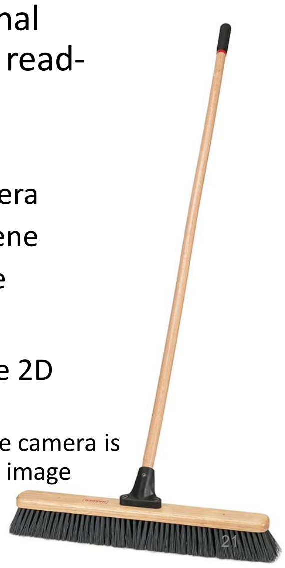

21

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)

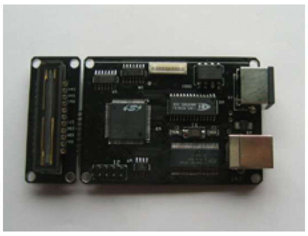

Mightex 3648-pixel line camera

22

Line camera, applications

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

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

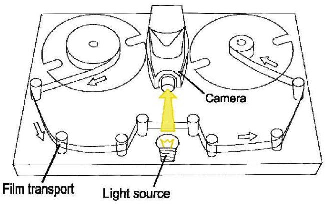

ImageSystems GoldenEye film scanner

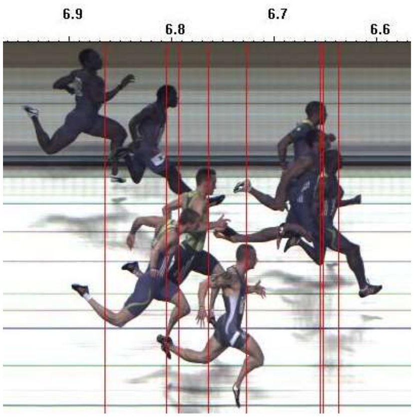

Photo finish

23

Push-broom camera, applications

ground mapping from air

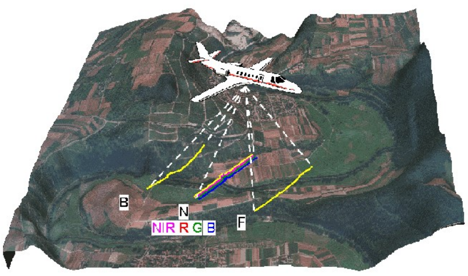

ground mapping from satellite

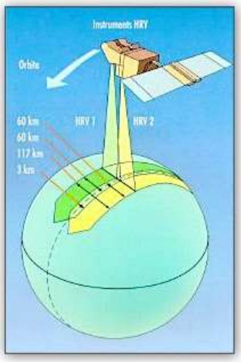

24

Smart cameras

for machine vision applications

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

PU and sensor on separate chips but in the same housing	Near sensor processing (PU + sensor on same chip) CMOS
Dedicated function PU	General purpose PU
One PU per chip	One PU per pixel

25

Smart cameras

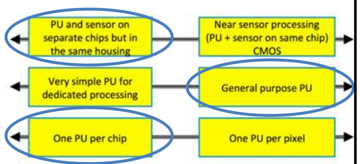

- Some cameras come with a completely integrated development environment (IDE)
  - Processing is defined in a GUI on a standard PC
  - Processing code is downloaded into the smart camera and executed
- Some cameras are integrated with illumination, e.g.
  - IR
  - visual
- Most cameras have a simple interface to a network, e.g.
  - TCP/IP protocol over Ethernet
  - Integrated web server

Cognex Checker 252

26

Smart Camera: SICK Ranger

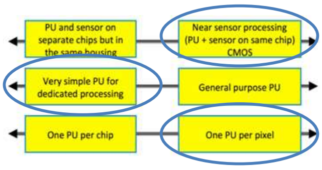

- The smart cameras from the company SICK-IVP are quite different.
  - The sensor and image processor are integrated into the same chip.
  - The processor is programmable but not general purpose.
  - There is a dedicated processing unit per pixel.
- These cameras are aimed at very fast industrial and robotic applications.

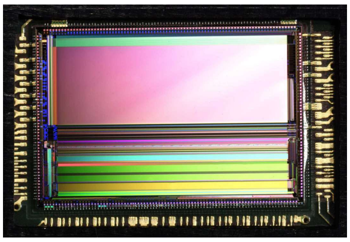

sensor

27

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
  - …

28

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*)

29

Event cameras

Example pixel circuit

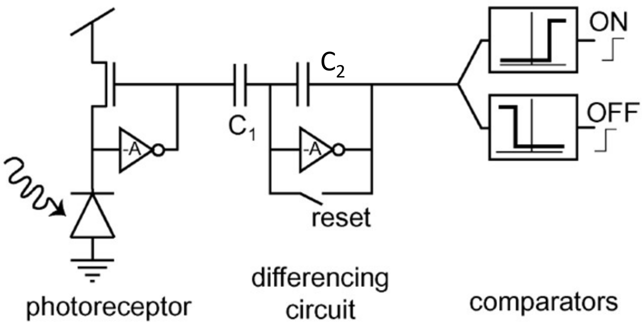

1) Photoreceptor measures continuously (no integration)
2) C_1 stores current light intensity during reset as a charge
3) After reset, if the charge in C_1 changes, a corresponding change will happen in C2 and an ON or OFF event is generated (with some threshold)
4) The event is read-out and a new reset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

30

Event cameras

Event address representation (AER)

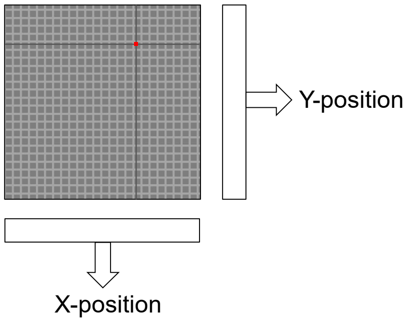

- An event (ON or OFF) triggers a readout of the pixel address.
- For every event, its X and Y position is delivered from the sensor.
- A modern event camera can deliver millions of events per second.

31

Event cameras

- Illustration of the difference between an event camera (Samsung DVS) and a standard camera.

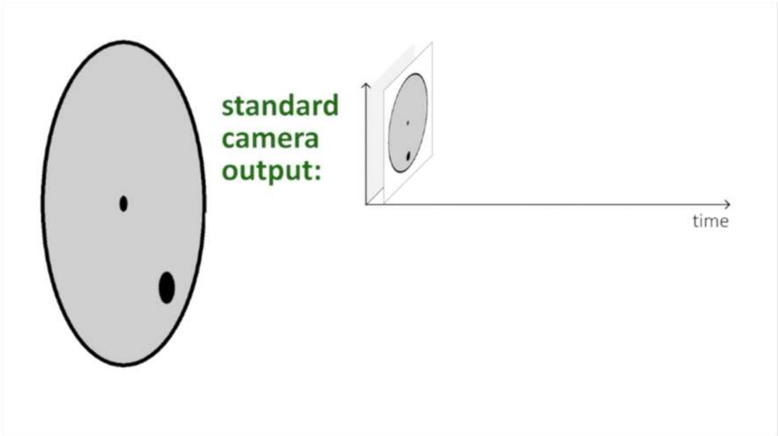

For the full video, see https://www.youtube.com/watch?v=LauQ6LWTkxM

32

Event cameras

- Motion in a natural scene.
- The white screen shows the output from the event camera.

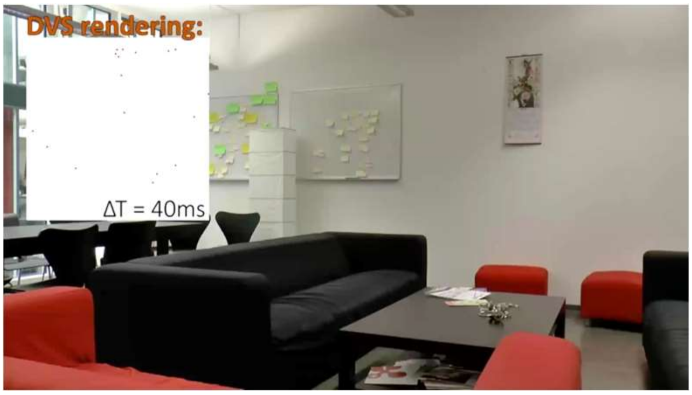

For the full video, see https://www.youtube.com/watch?v=LauQ6LWTkxM

33

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

34

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…

35

Light field

Repetition from lecture 1:

- At each point **x** in 3D space,
  in each direction **n**,
  there is an amount of light passing through **x**
- The *plenoptic function* I(**x**, **n**)
  – also known as *Light field*

36

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

37

Light field camera

A more practical implementation:

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

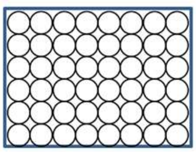

lenslet array

38

Light field camera

Examples

Lytro

Raytrix

Stanford plenoptic camera

Light field 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-directional cameras represent another class of light-field cameras. Here, it is the ability to capture many different directions that is of importance.
- Omni: Latin for *every, all*
  - In theory: a camera that sees in **all directions**
  - Often 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, but we will look into some approximations.

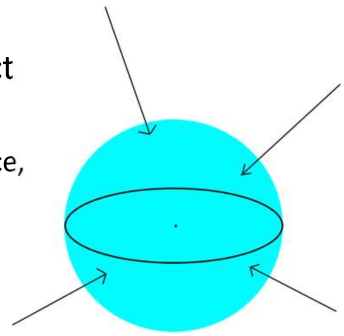

Multiple pinhole cameras

- 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.

43

Multiple pinhole cameras

- The multiple pinhole camera can produce a wide field-of-view by stiching several images together.

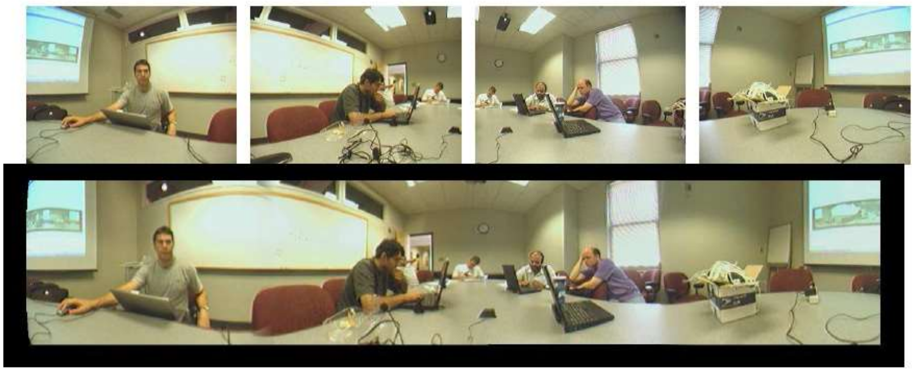

From: de la Torres, et al,
Learning to Track Multiple People in Omnidirectional Video

44

Fish-eye lens

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

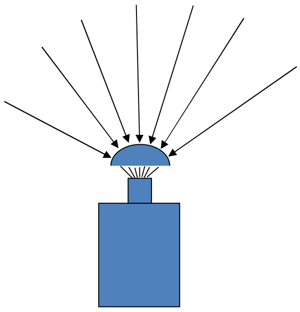

Image: Dan Slater

Catadioptric cameras

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

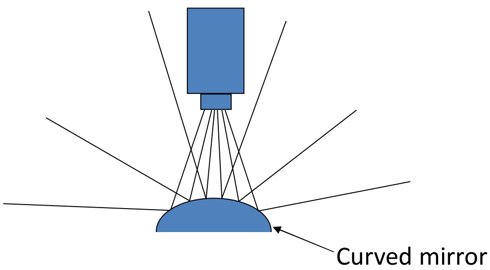

The mirror can be

- 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 also capture a larger field of view, nearly 360 degrees, with camera occlusion.

Omni-directional cameras

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

49

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.

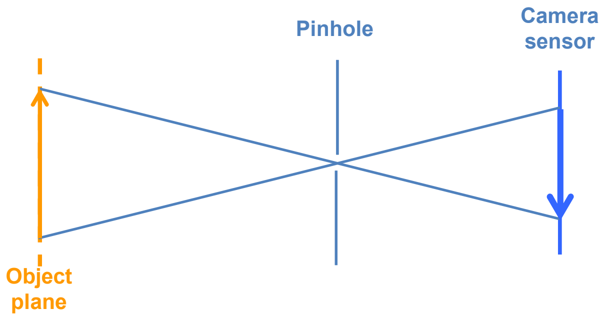

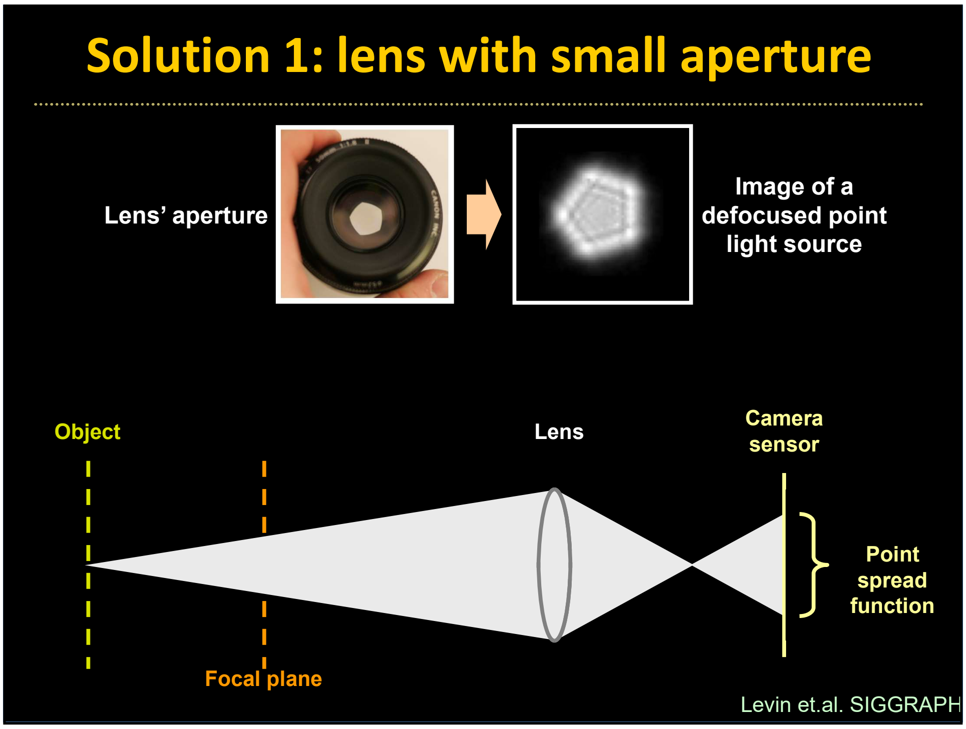
Solution 1: lens with small aperture
Lens' aperture
Image of a defocused point light source
Object
Lens
Camera sensor
Point spread function
Focal plane
Levin et.al. SIGGRAPH

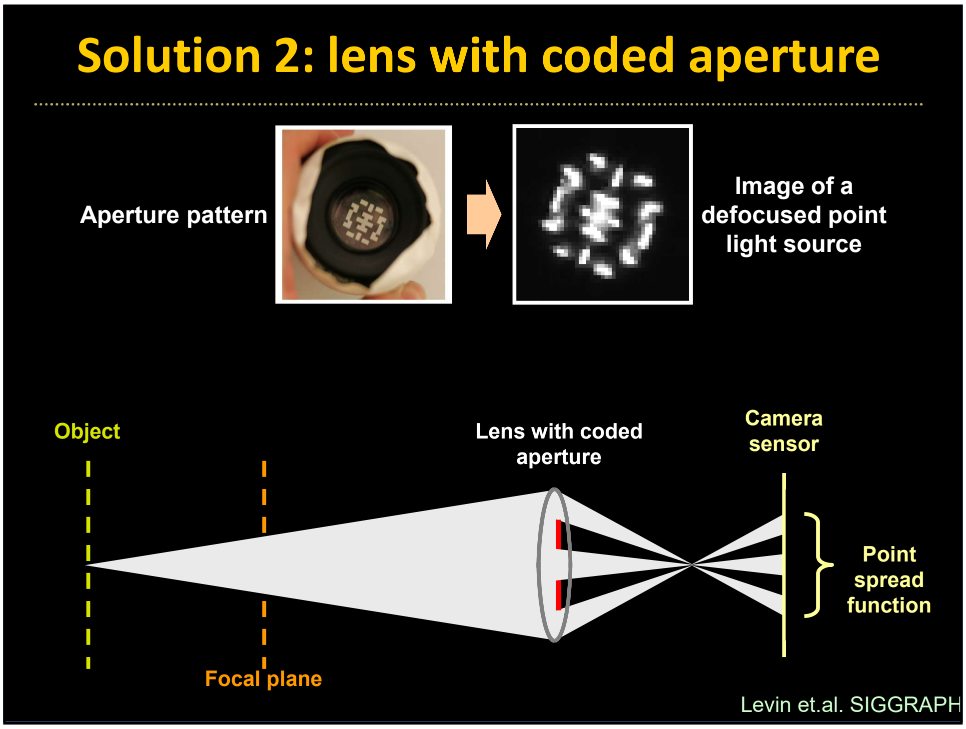
Solution 2: lens with coded aperture
Aperture pattern
Image of a defocused point light source
Object
Lens with coded aperture
Camera sensor
Point spread function
Focal plane
Levin et.al. SIGGRAPH

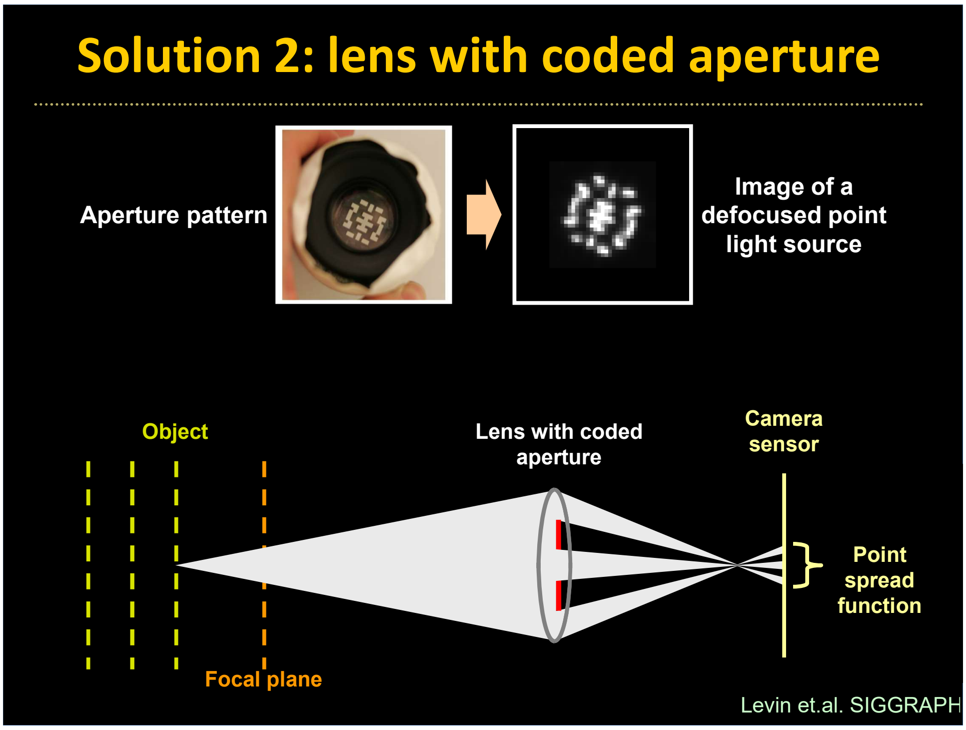
Solution 2: lens with coded aperture
Aperture pattern
Image of a defocused point light source
Object
Lens with coded aperture
Camera sensor
Point spread function
Focal plane
Levin et.al. SIGGRAPH

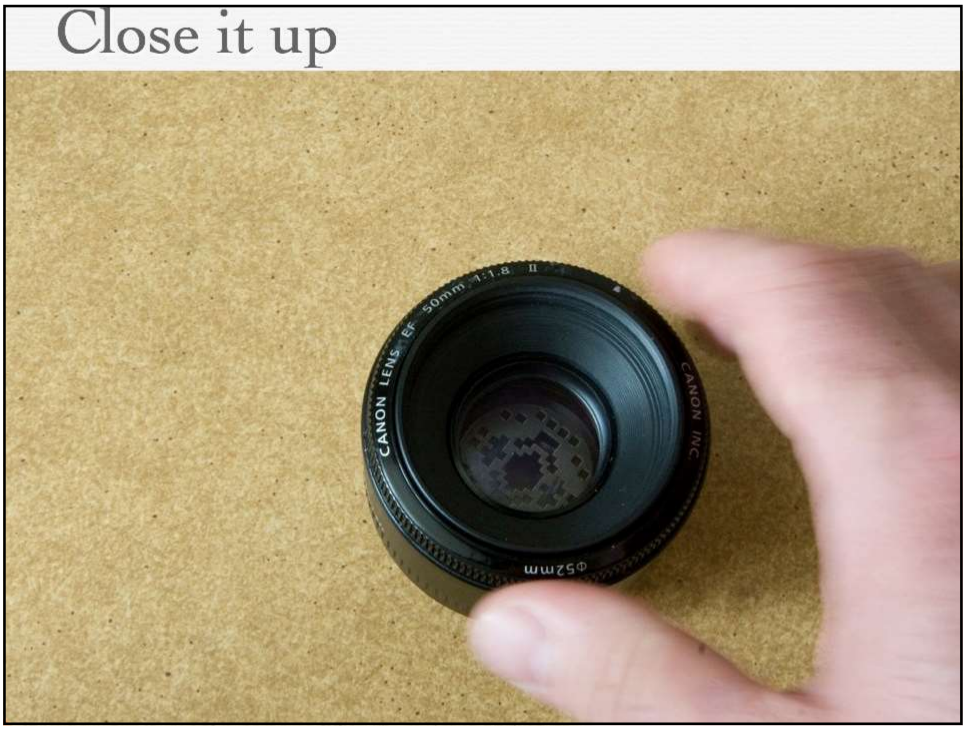
Close it up
CANON LENS EF 50mm 1:1.8 II
CANON INC.
Φ52mm

Levin et.al. SIGGRAPH

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

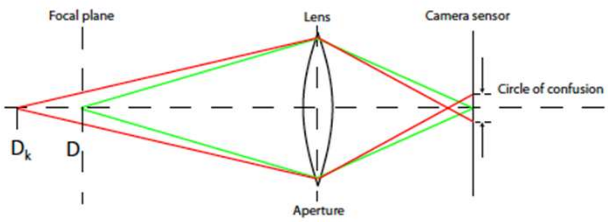

Reconstruction

- 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.

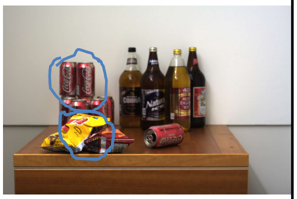

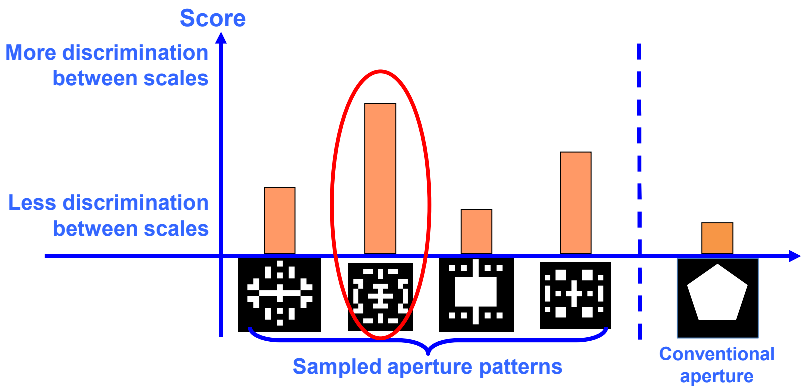

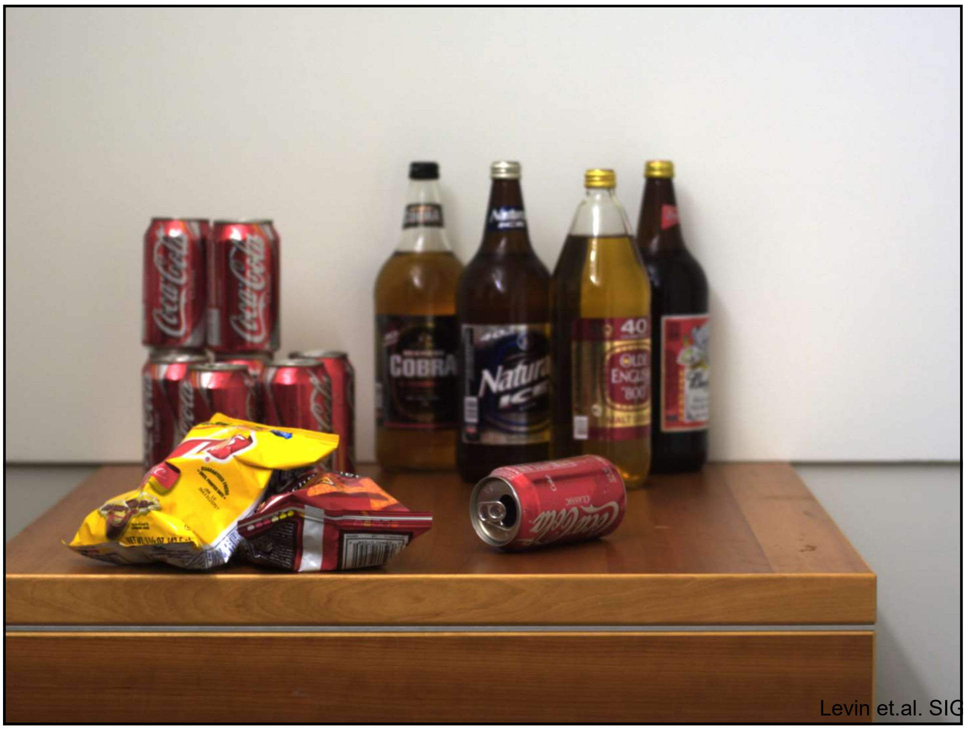

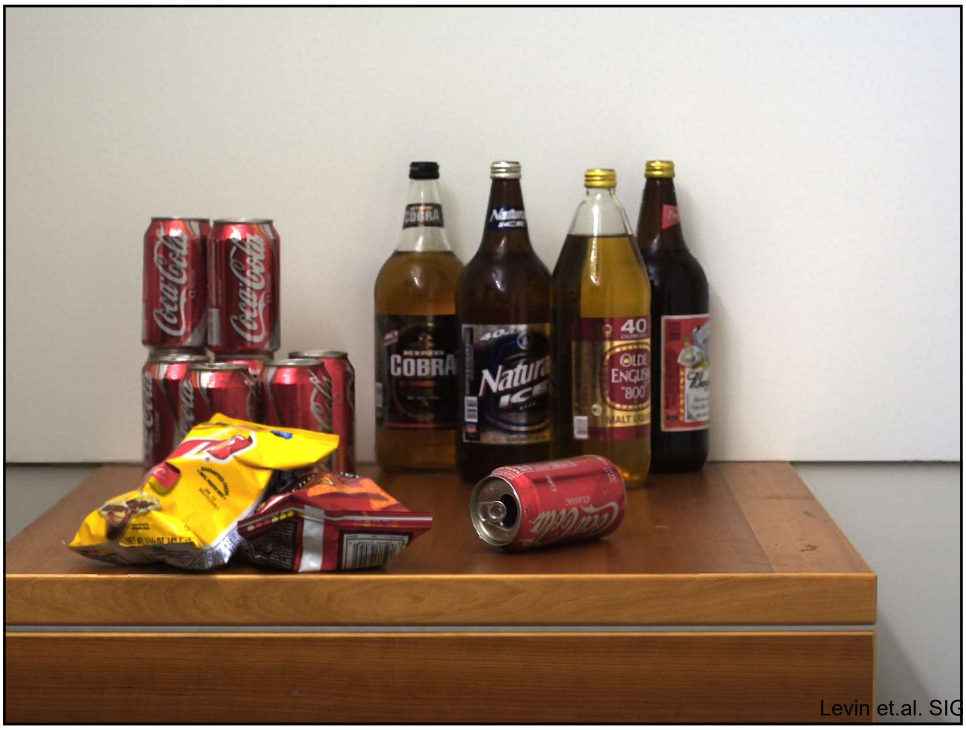

Close-up

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

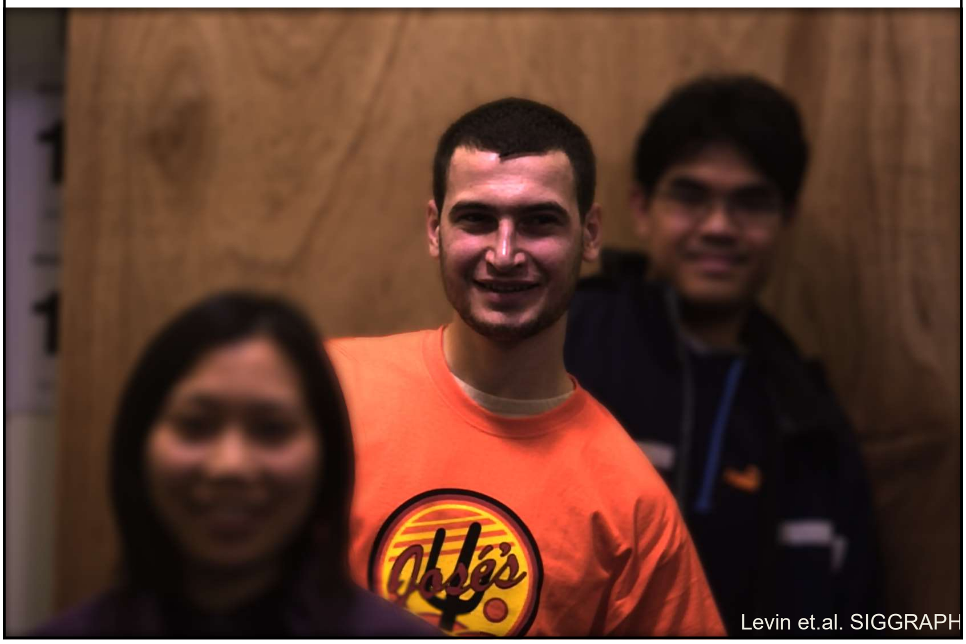

Levin et.al. SIGGRAPH

Application: Digital refocusing from a single image

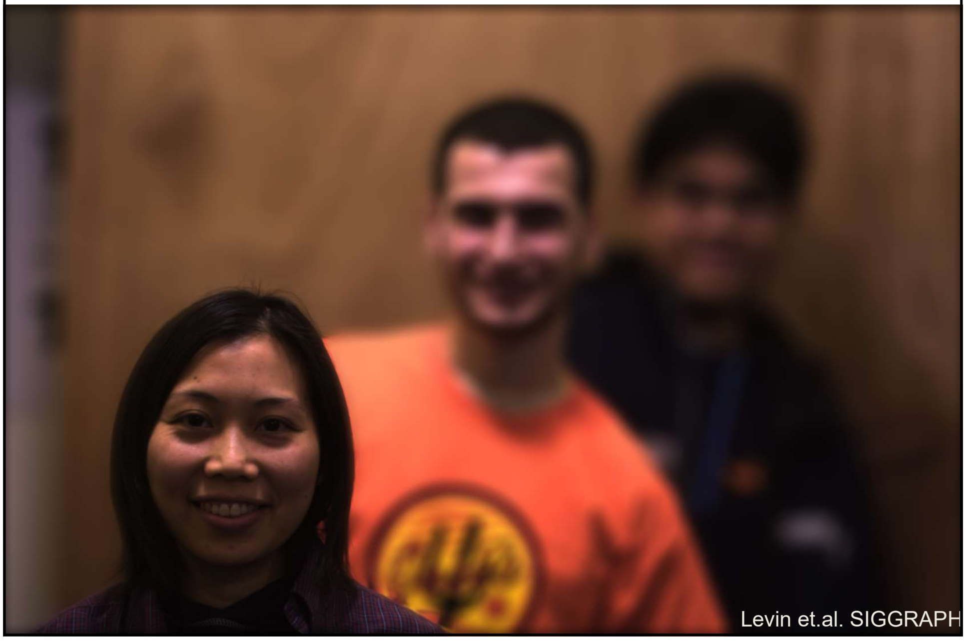

Application: Digital refocusing from a single image

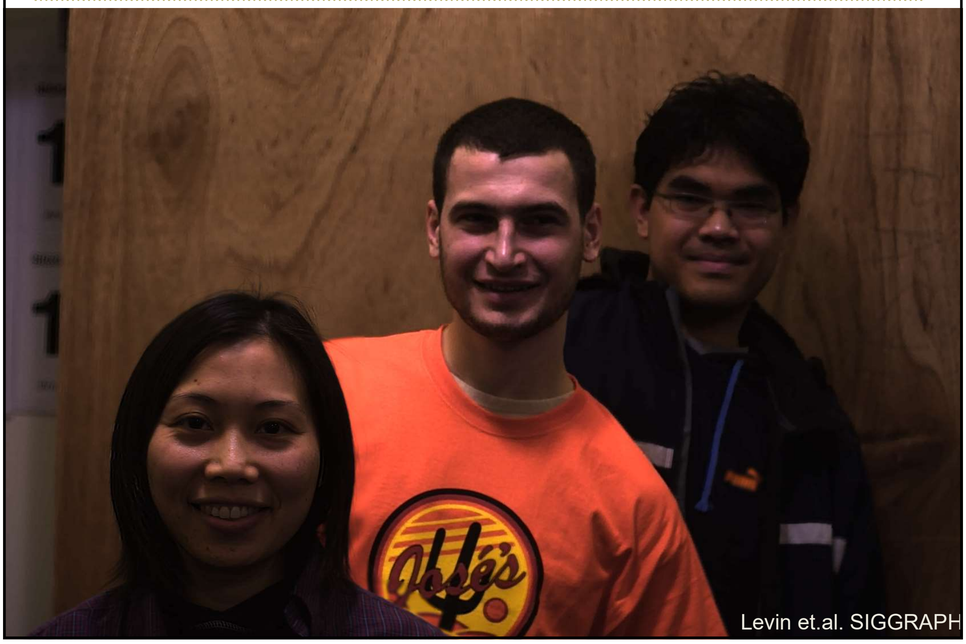

Added bonus: depth estimation

Levin et.al. SIGGRAPH

Flexible depth of field*

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

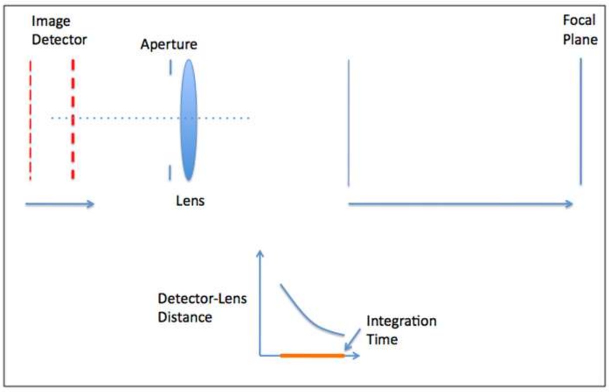

*"Flexible Depth of Field Photography," H. Nagahara, S. Kuthirummal, C. Zhou, and S.K. Nayar, European Conference on Computer Vision (ECCV), Oct, 2008.

Flexible depth of field imaging

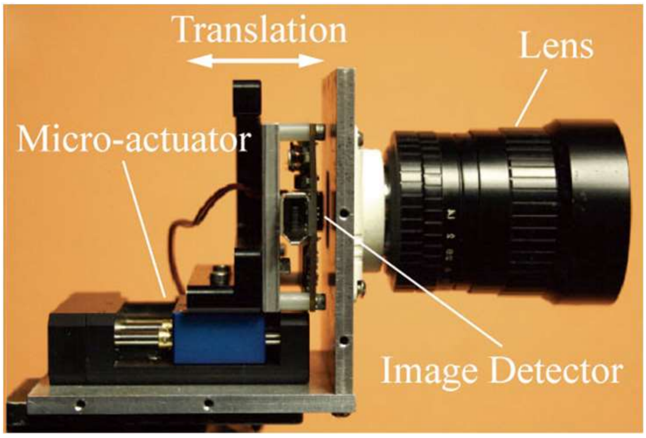

PSF becomes independent on object distance

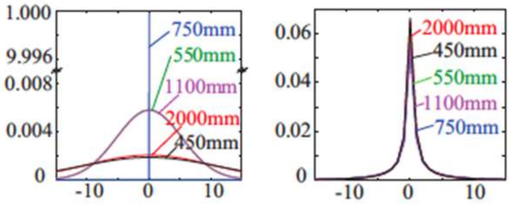

(c) Normal Camera PSF (Gaussian) (d) EDOF Camera IPSF

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

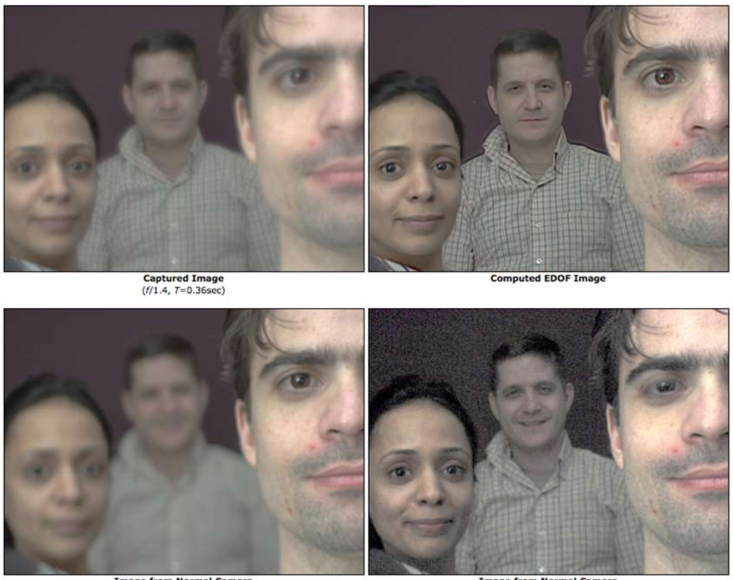

Captured Image
(f/1.4, T=0.36sec)

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 imagin

Computational cameras

- If you want to learn more about computational cameras, listen to the…
- Plenary talk by Shree Nayar (Professor at Columbia University)
- The video is available on the course home page