Image Analysis

Lecture 2: Image acquisition, compression, storage and change detection in videos

Rasmus R. Paulsen
Tim B. Dyrby

DTU Compute

http://compute.dtu.dk/courses/02515

Lecture 2

- Image acquisition, compression, storage and change detection in videos

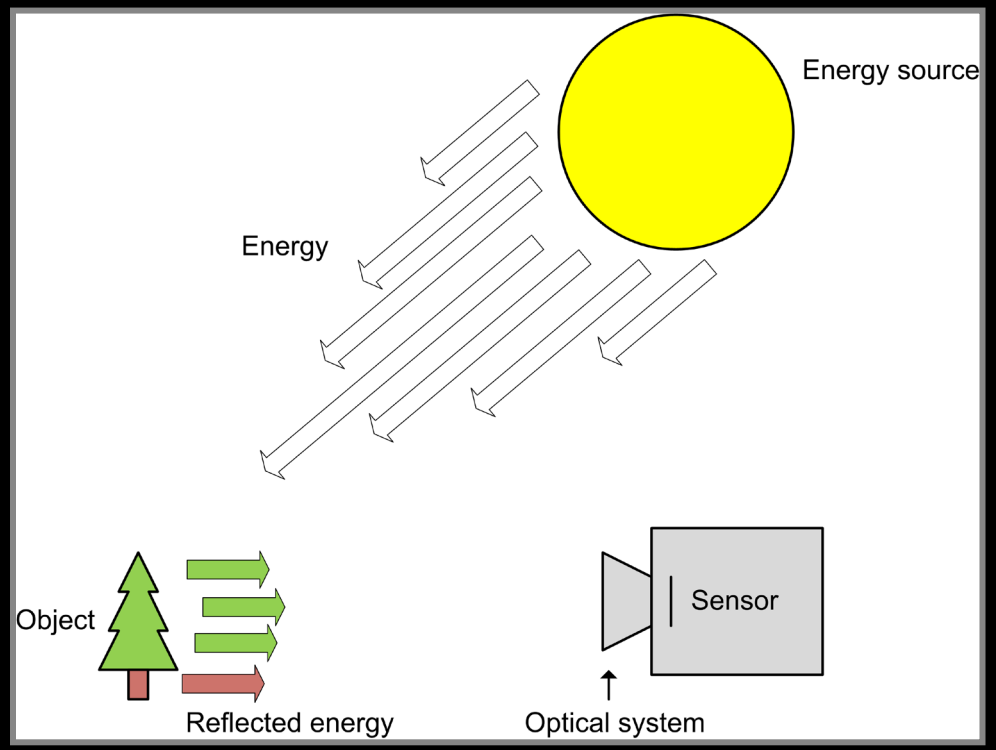

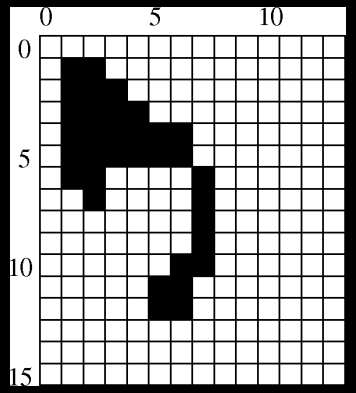

Learning objectives – cameras and lenses

- Explain where visible light is in the electromagnetic spectrum
- Describe the pin hole camera
- Describe the properties of a thin-lens including focal-length, the optical center, and the focal point
- Estimate the focal length of a thin lens
- Compute the optimal placement of a CCD chip using the thin lens equation
- Describe depth-of-field
- Compute the field-of-view of a camera
- Explain the simple CCD model

How is an image created?

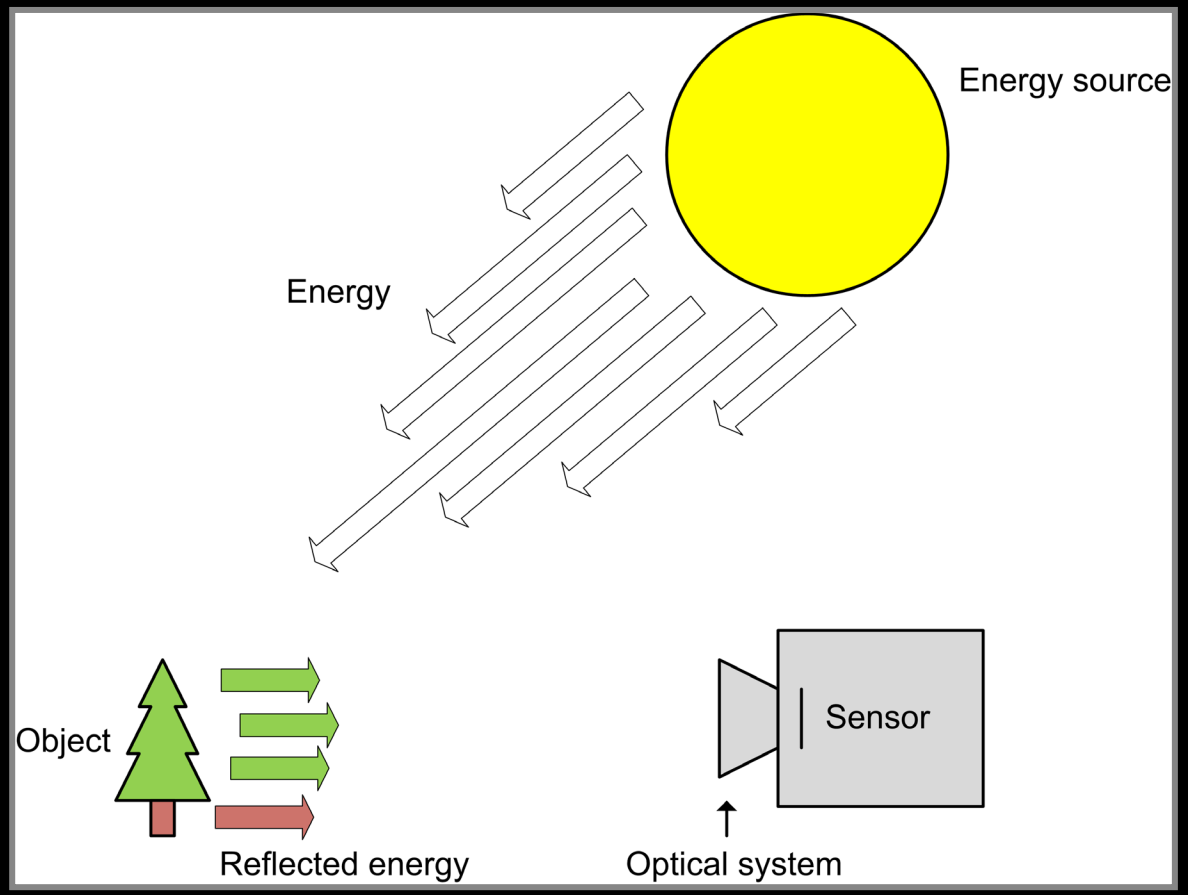

This is just one way! Other methods will be described later in the course.

What is light?

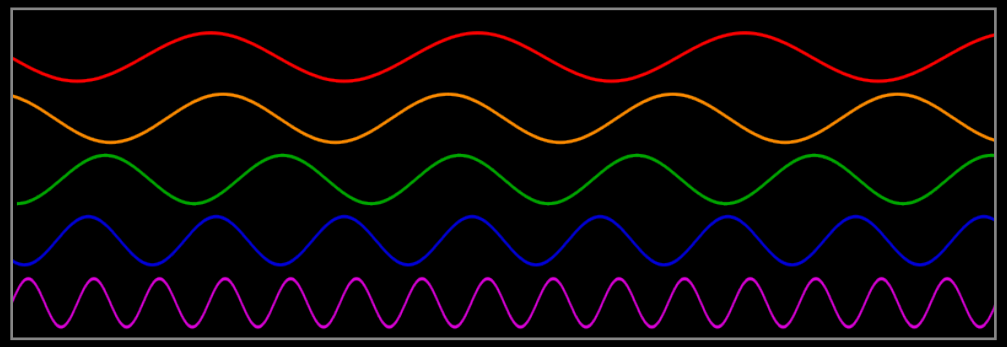

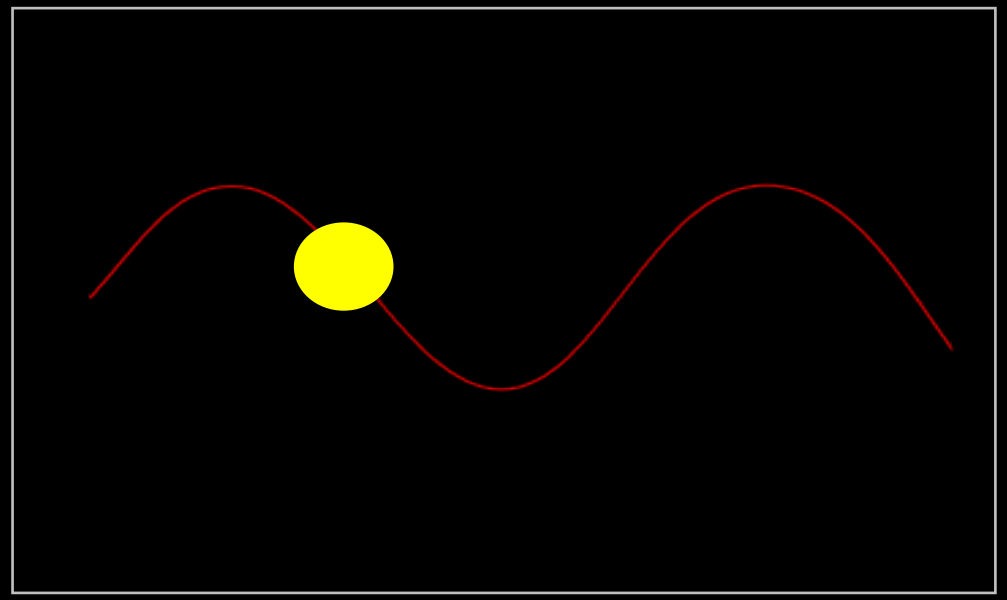

- Can be seen as electromagnetic waves
- Or as a photon (from Greek *phōtos*, "light")
  - Mass less fundamental particle

Light as a wave

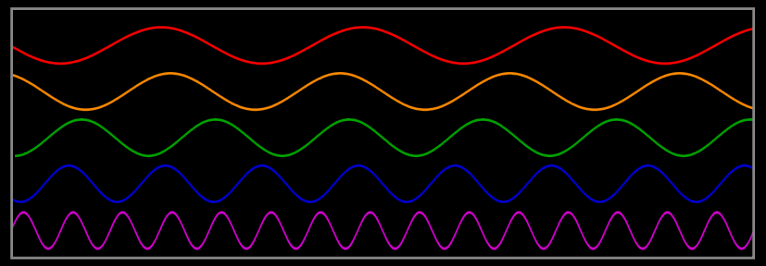

$$\lambda = \frac{c}{f}$$

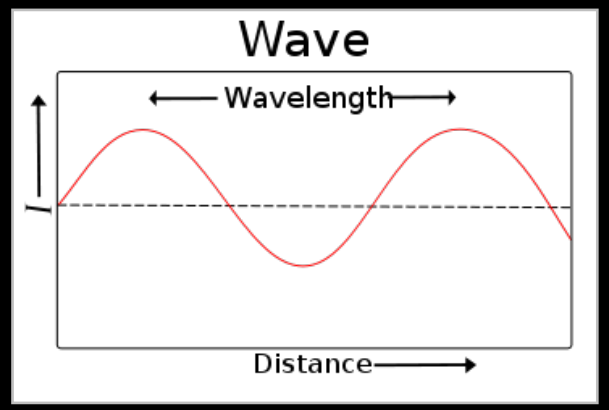

- It has a frequency f
  - Measured in Hertz [Hz]
- It has a wavelength λ (lambda)
  - Measured in meters [m]
- It has a speed
  - "The speed of light" c
  - 299.792 458 [m/s]

- High frequency -> short waves
- Low frequency -> long waves

Energy of light

$$E = h \cdot f$$

- Light has energy
  - You can feel it in the sun!
- Planck's constant h
- High frequency -> high energy
- Long waves -> low energy

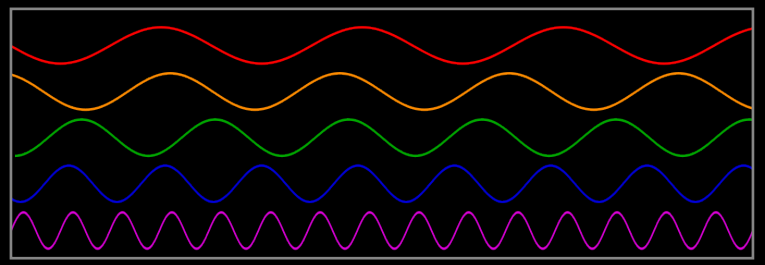

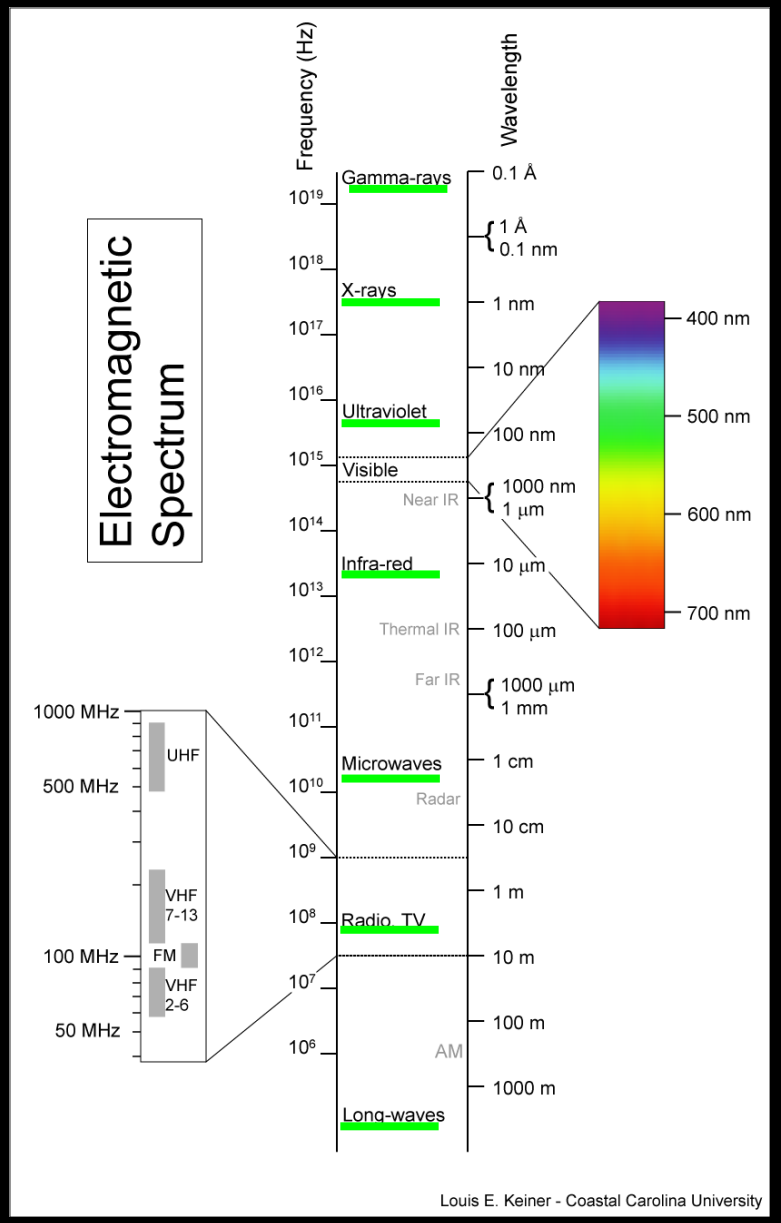

- Electromagnetic spectrum
  - Range of all frequencies
  - Divided into 7 regions
- Wavelengths
  - 1 µm = 1 micrometer = 0.001 mm
  - 1 nm = 1 nanometer = 0.0000001mm

PollEverywhere quizzes

https://pollev.com/rasmuspaulse538

What has the most energy?

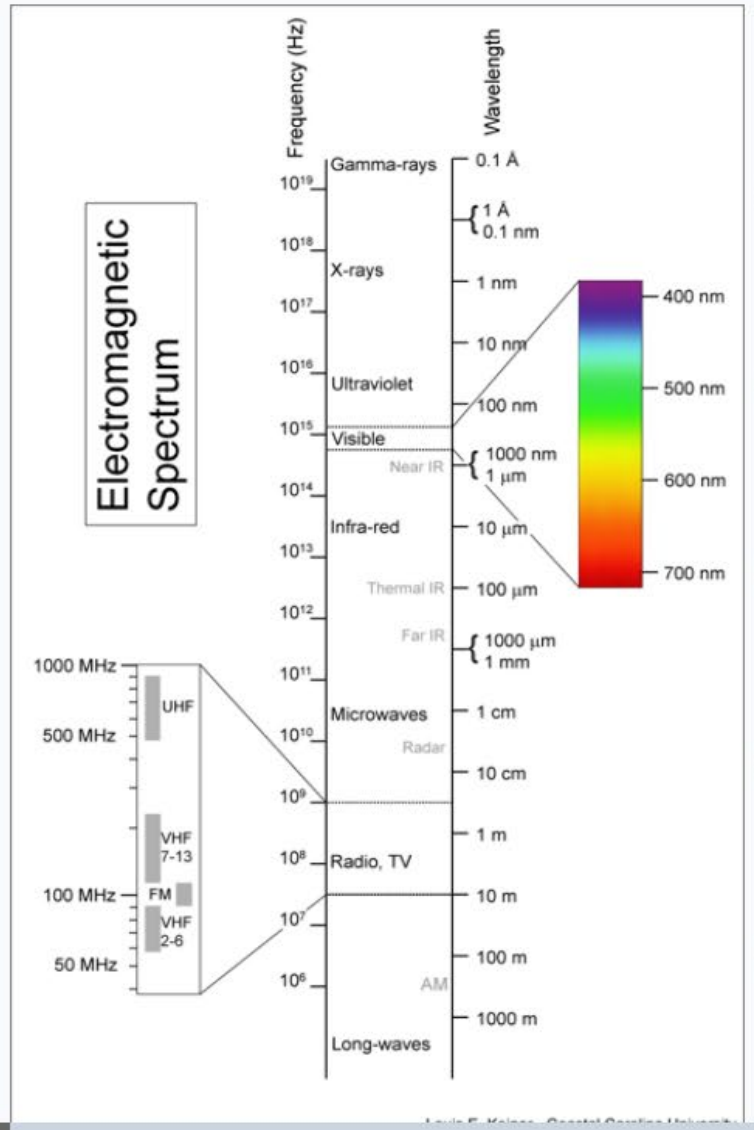

- Radiowaves
- X-rays
- Red light
- Microcwaves
- Ultraviolet light
- I do not know

Start the presentation to see live content. For screen share software, share the entire screen. Get help at **pollev.com/app**

What has the most energy?

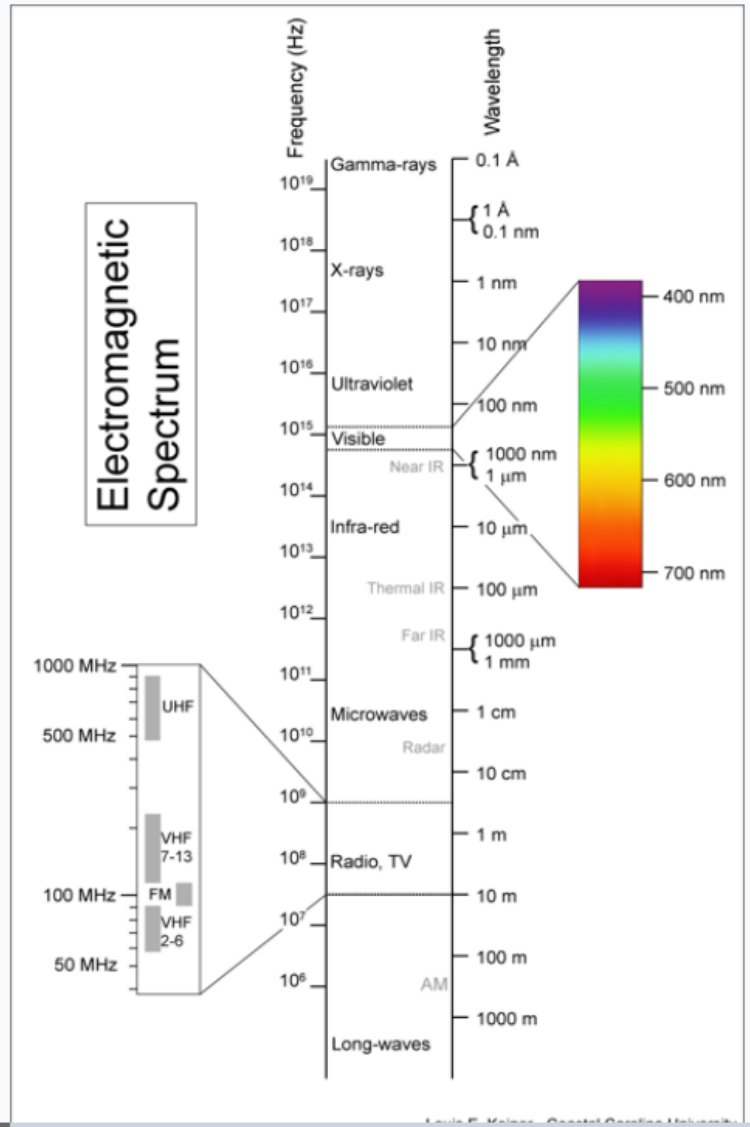

Radiowaves	0%
X-rays ✓	80%
Red light	0%
Microcwaves	0%
Ultraviolet light	20%
I do not know	0%

Start the presentation to see live content. For screen share software, share the entire screen. Get help at **pollev.com/app**

What has the most energy?

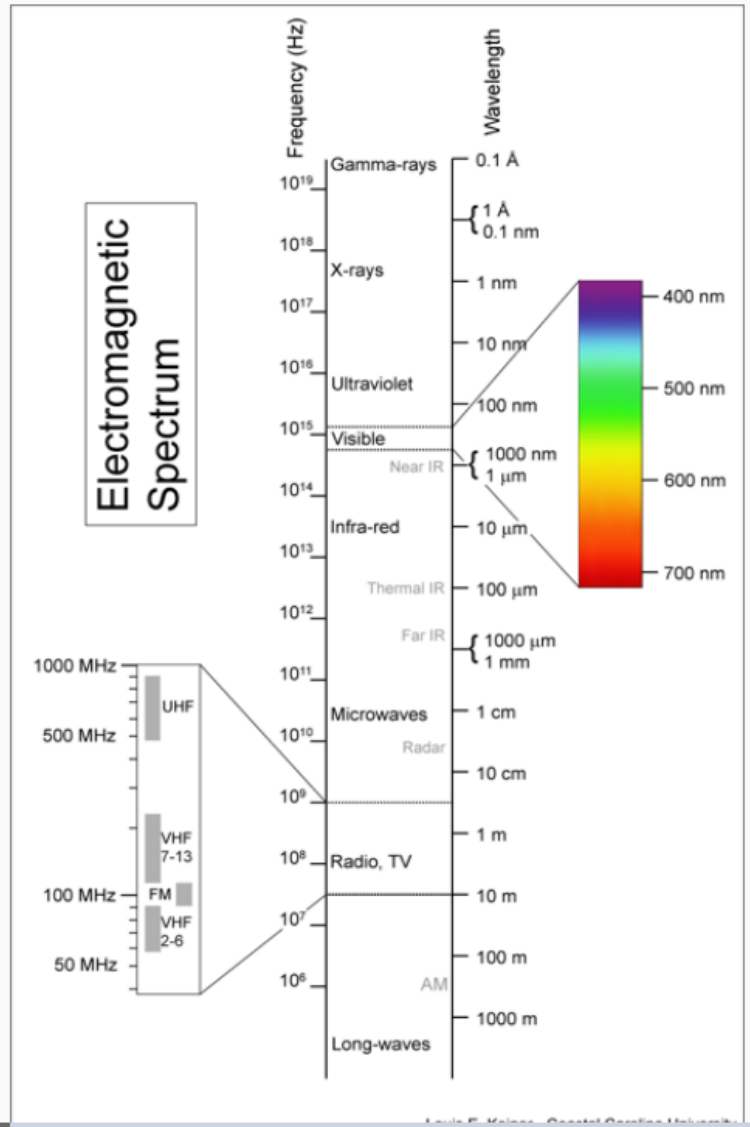

Radiowaves	0%
X-rays ✓	80%
Red light	0%
Microcwaves	0%
Ultraviolet light	20%
I do not know	0%

Start the presentation to see live content. For screen share software, share the entire screen. Get help at **pollev.com/app**

How do light become a digital image?

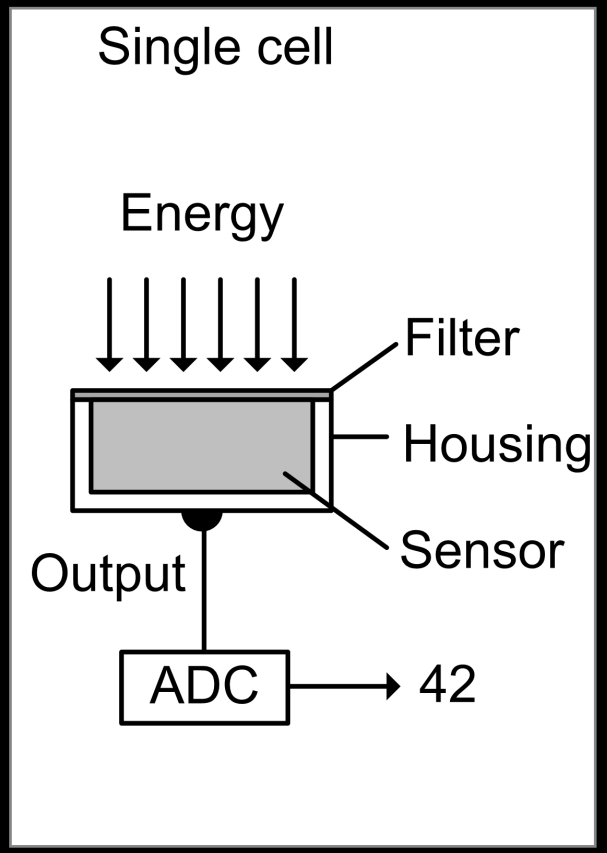

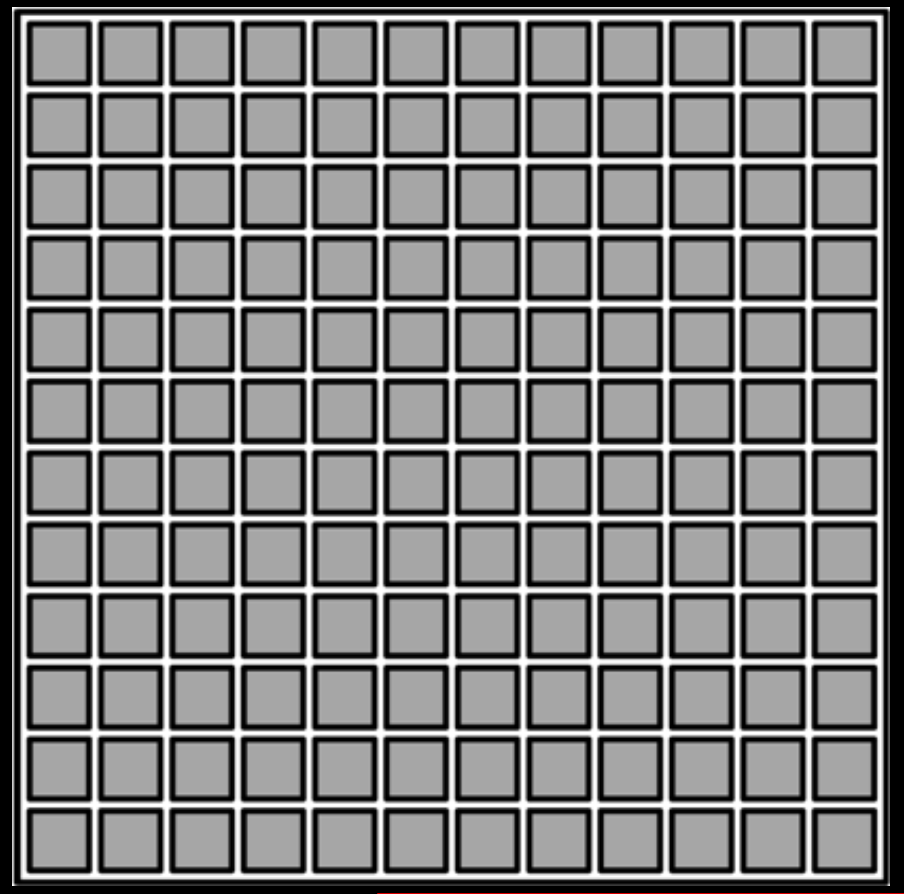

Charged coupled device (CCD-chip)

The digital film!

The CCD cell

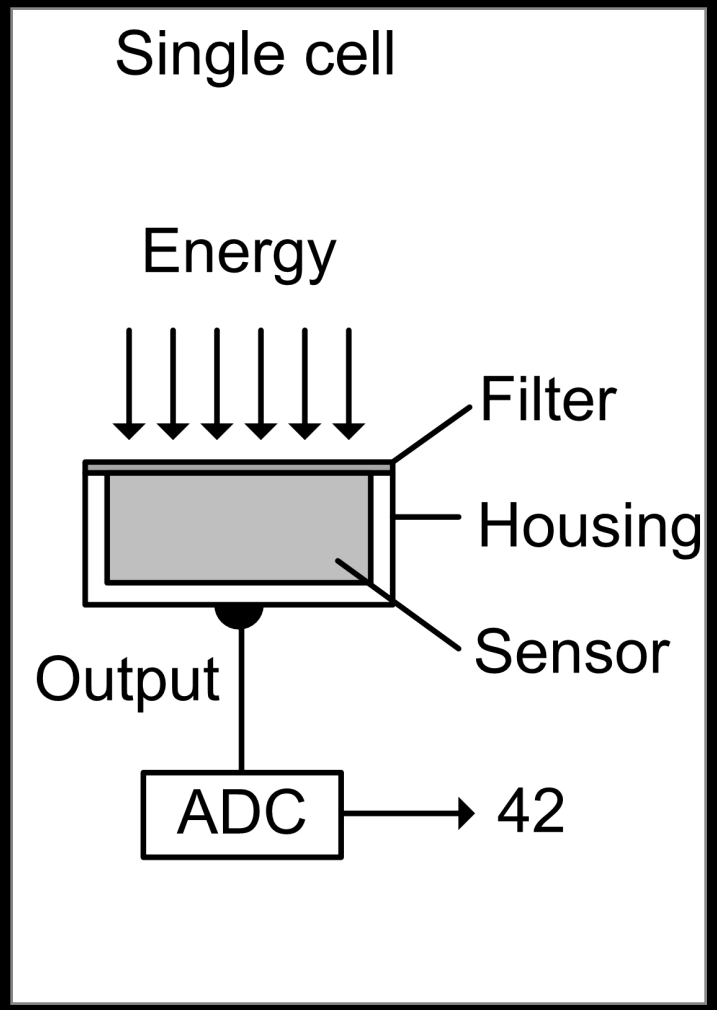

- The cell can be seen as a well that collects energy
- It collect energy for a limited time (*to be charged*)
  - Exposure time
  - Integration time
  - Shutter

The CCD cell - conversion

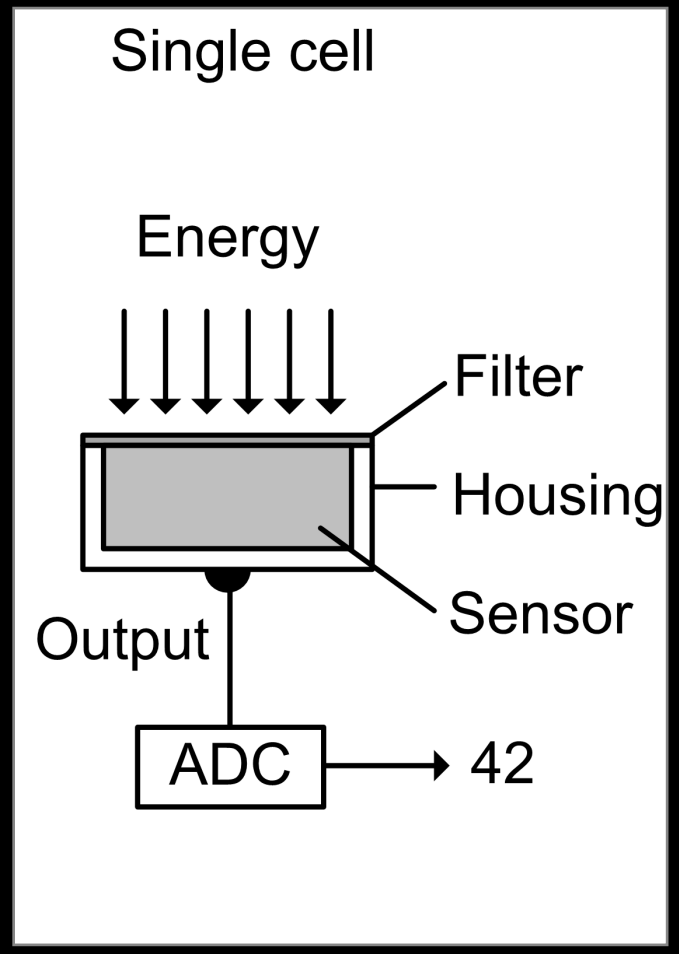

- Energy transformed to a digital number
  - Analog-to-Digital converter (ADC)
- Takes a an "analogue signal" and converts it to a digital signal

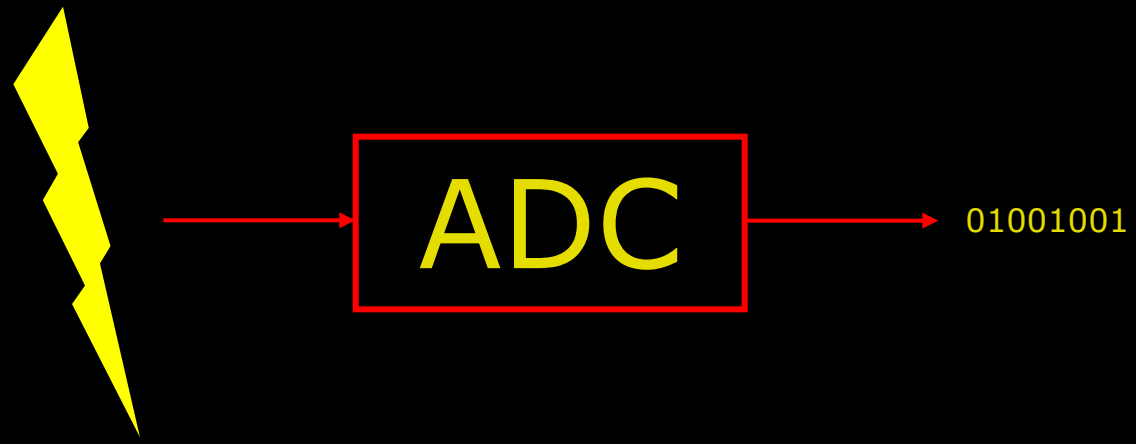

$(01001001)_2 = (0 \times 2^7) + (1 \times 2^6) + (0 \times 2^5) + (0 \times 2^4) + (1 \times 2^3) + (0 \times 2^2) + (0 \times 2^1) + (1 \times 2^0) = (73)_{10}$

CCD and images

- Surprise! 1 CCD cell = 1 pixel
  - Only for grayscale images
  - More complex for RGB images
- 10 MPixel camera
  - 10 millions analog to digital conversions for one image!

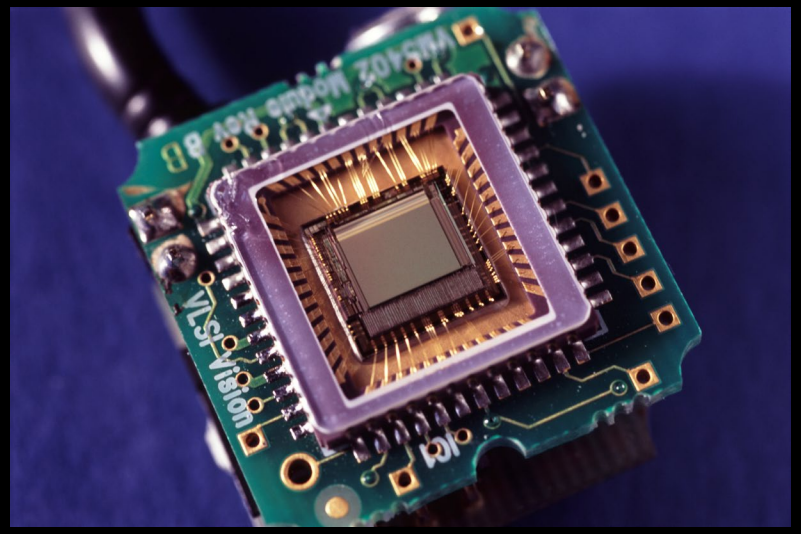

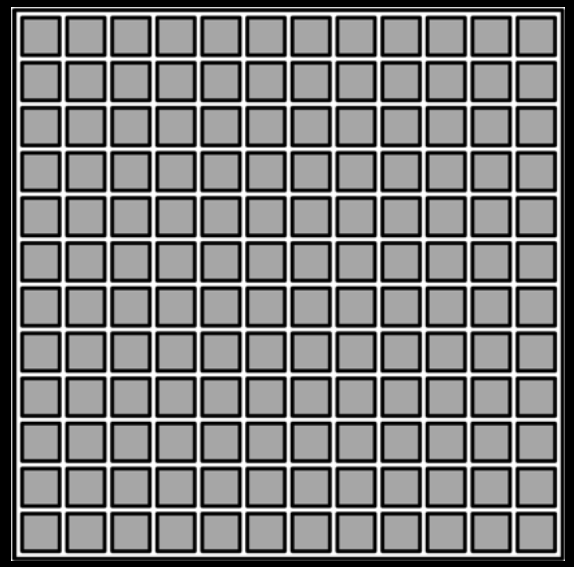

What happens when you press the button?

The shutter opens and the CCD is hit by light

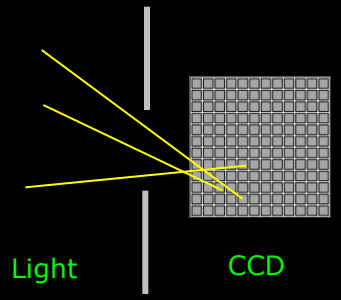

CCD is integrating

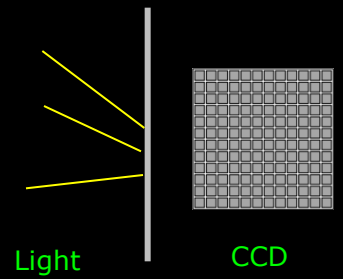

Shutter closes

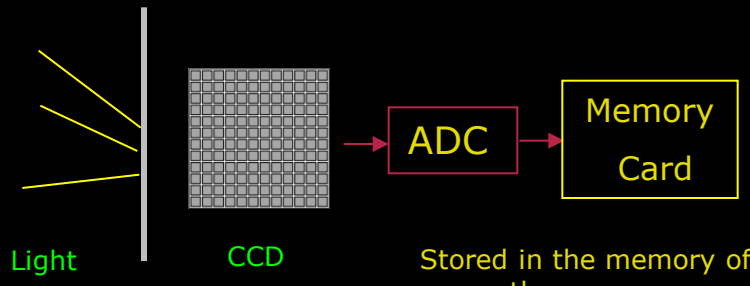

The collected energies are transferred and converted to digital values

Stored in the memory of the camera

Question: Integration time

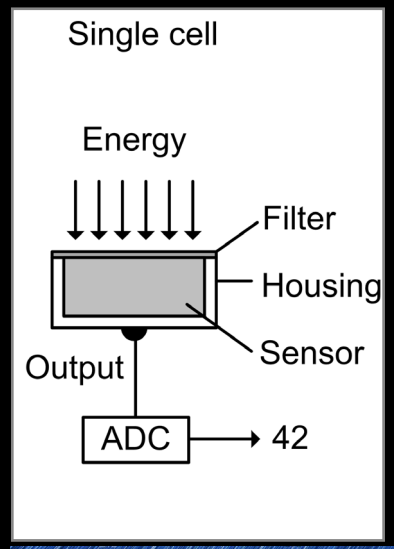

- What happens if we integrate over long time?
  - Motion blur
  - Over-exposure (the well is overrunning)
  - Blooming
- Short integration time
  - Noise
  - Lack of contrast

Motion blur

- Causes blurring of the moving object

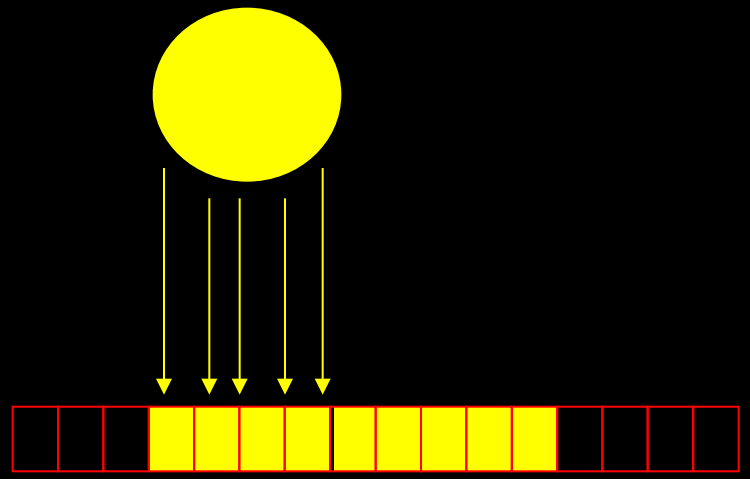

The bigger picture

- A camera is more than a CCD!
- The CCD is the sensor!
- There is also “an optical system”

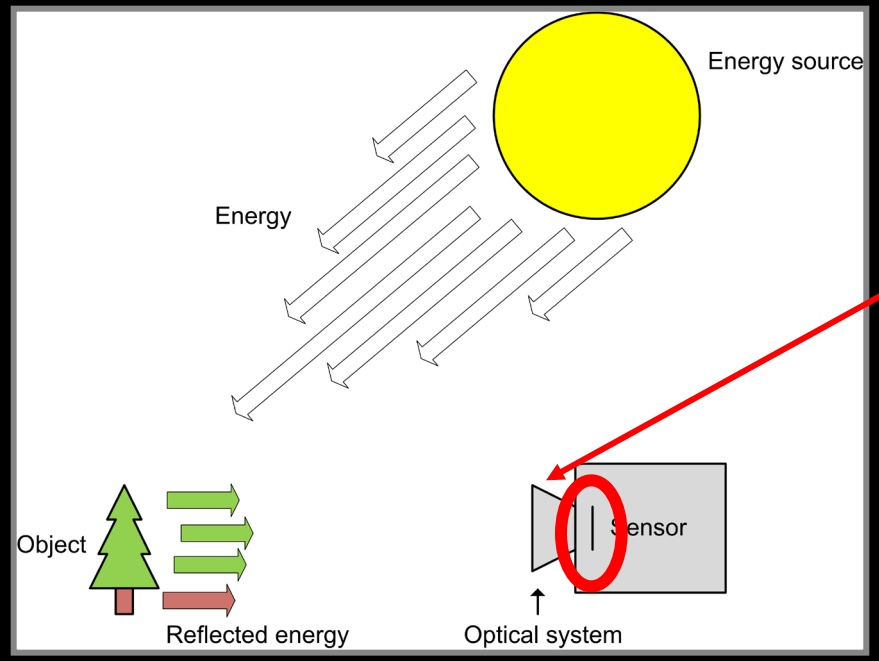

Optical system

(lenses)

Optical system

- How do we get an image on the CCD?
- Light follows a straight line
- Light that hit one spot reflects in many directions

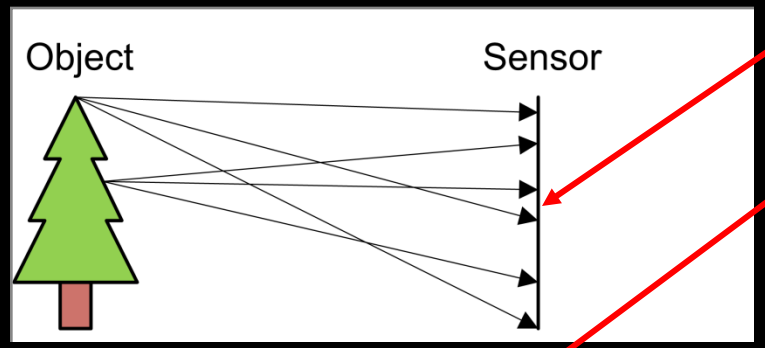

Same point hit by rays from all over the object

Barrier with tiny hole

mera

Pinhole camera

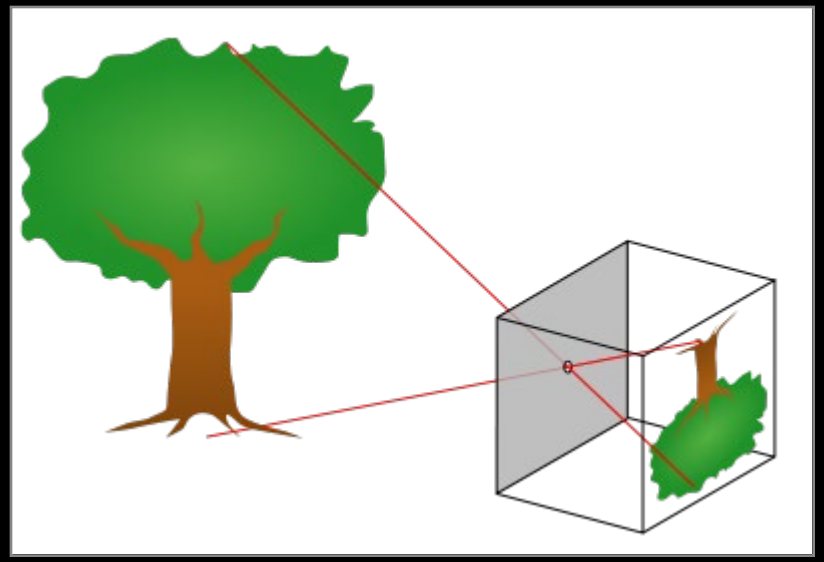

- Light coming through the tiny hole – any problems?
  - Very little light!
- How do we get more light inside the camera?
  - While keeping the focus?

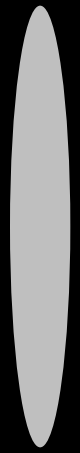

A lens!

The lens

- A lens focuses a bundle of rays to one point
- Parallel rays pass through a focal point **F** at a distance f beyond the plane of the lens. f is the focal length
- O is the optical centre. F and O span the optical axis

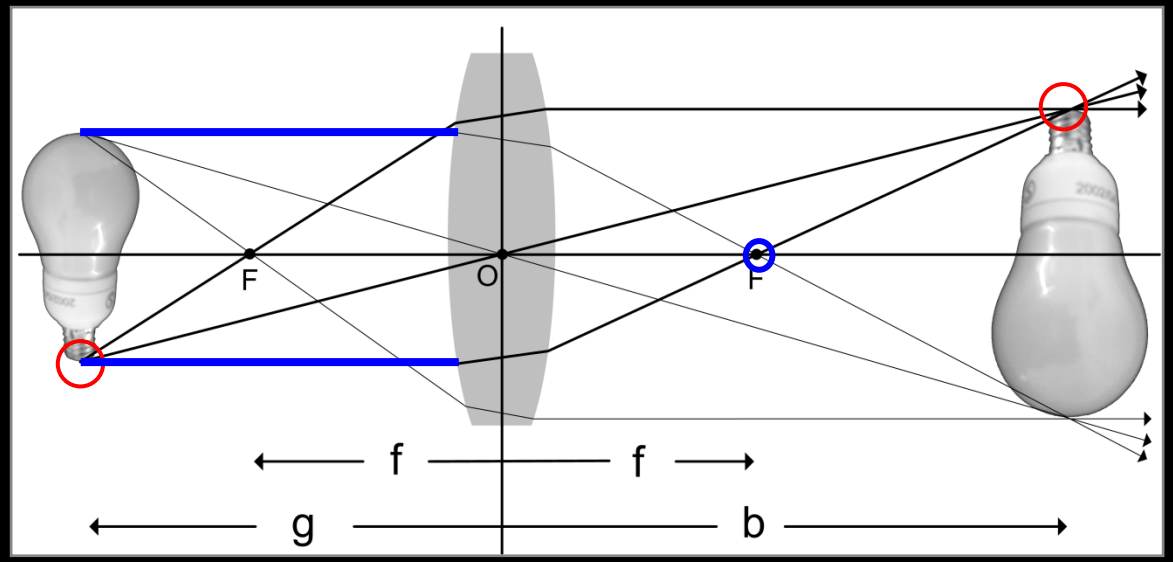

World

Inside camera

Focal point – focal length

- Light coming from "really far away" can be seen as parallel rays
- Rays intersect at the focal point
- Distance from optical centre O to focal point F is called *focal length f*

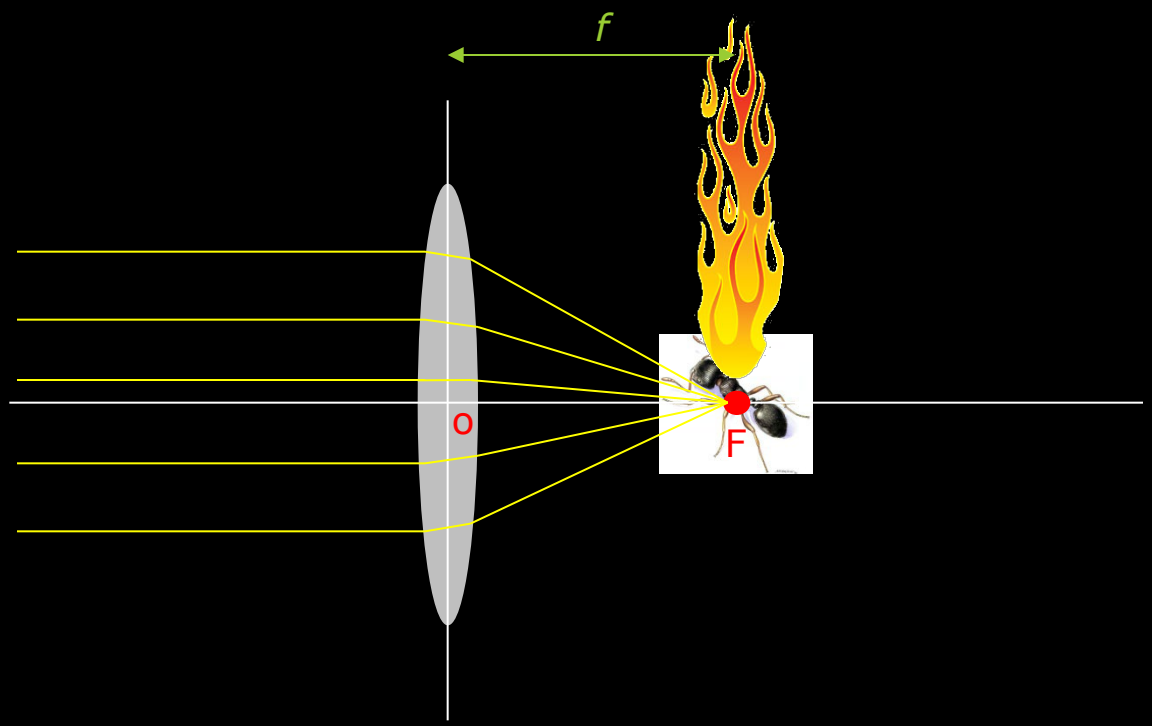

Where do non-parallel rays meet?

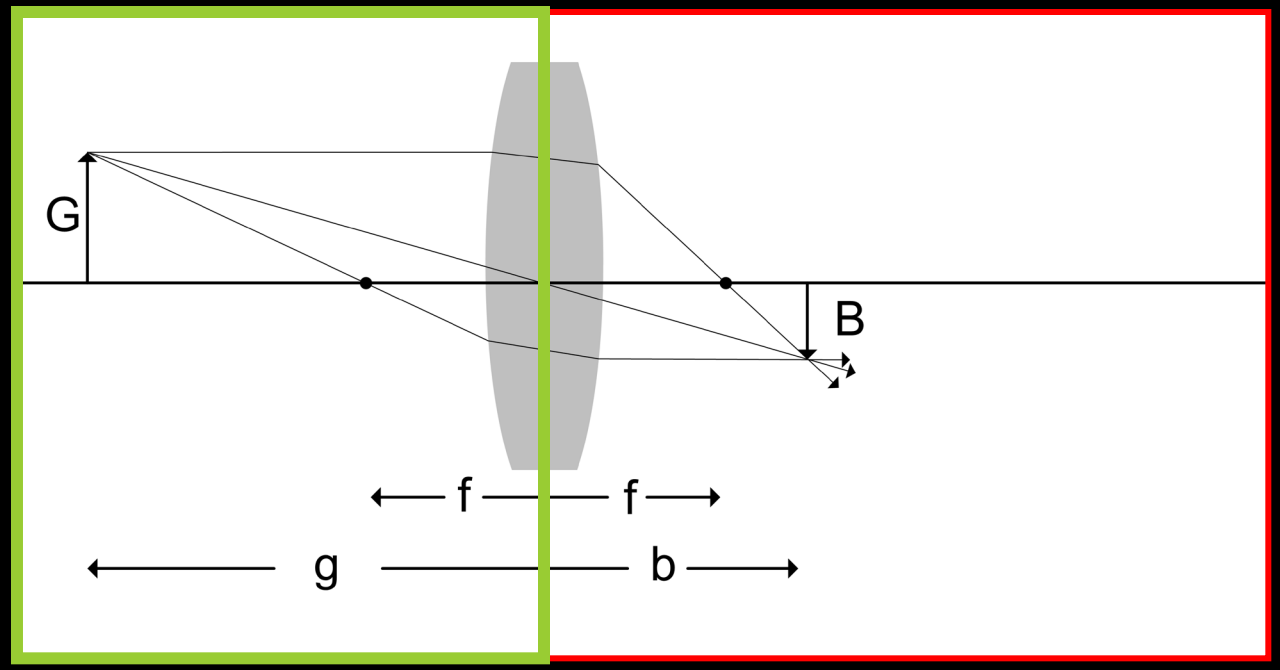

$$\frac{1}{g} + \frac{1}{b} = \frac{1}{f}$$

Thin lens equation

or

Gauss' lens equation

g – distance to object

b – distance to intersection

Where do the rays meet

b=1 mm

b=4 mm

b = 5 mm

b=6 mm

b=7 mm

Do not know

Start the presentation to see live content. For screen share software, share the entire screen. Get help at **pollev.com/app**

Where do the rays meet

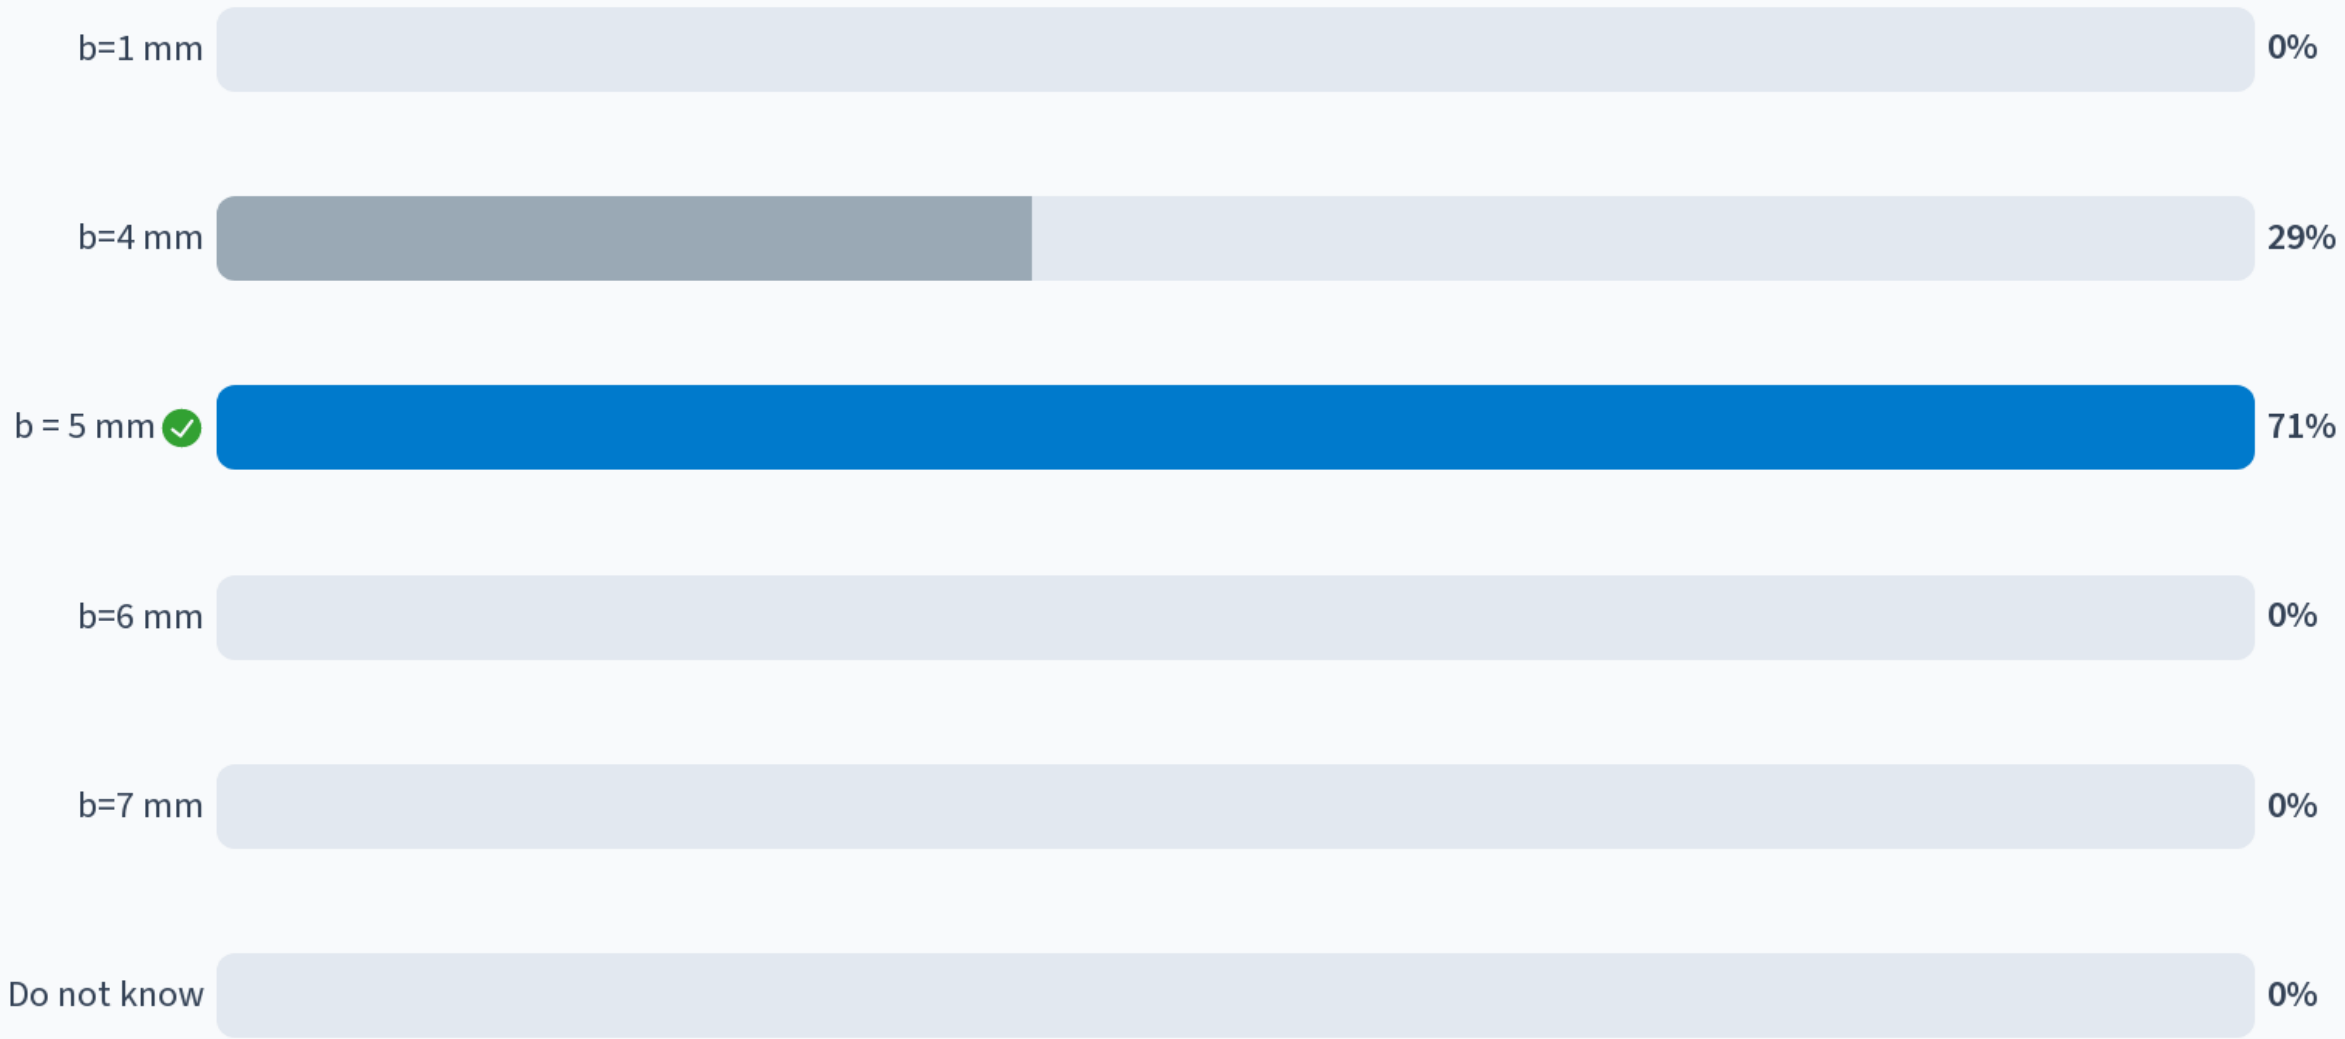

Start the presentation to see live content. For screen share software, share the entire screen. Get help at **pollev.com/app**

Where do the rays meet

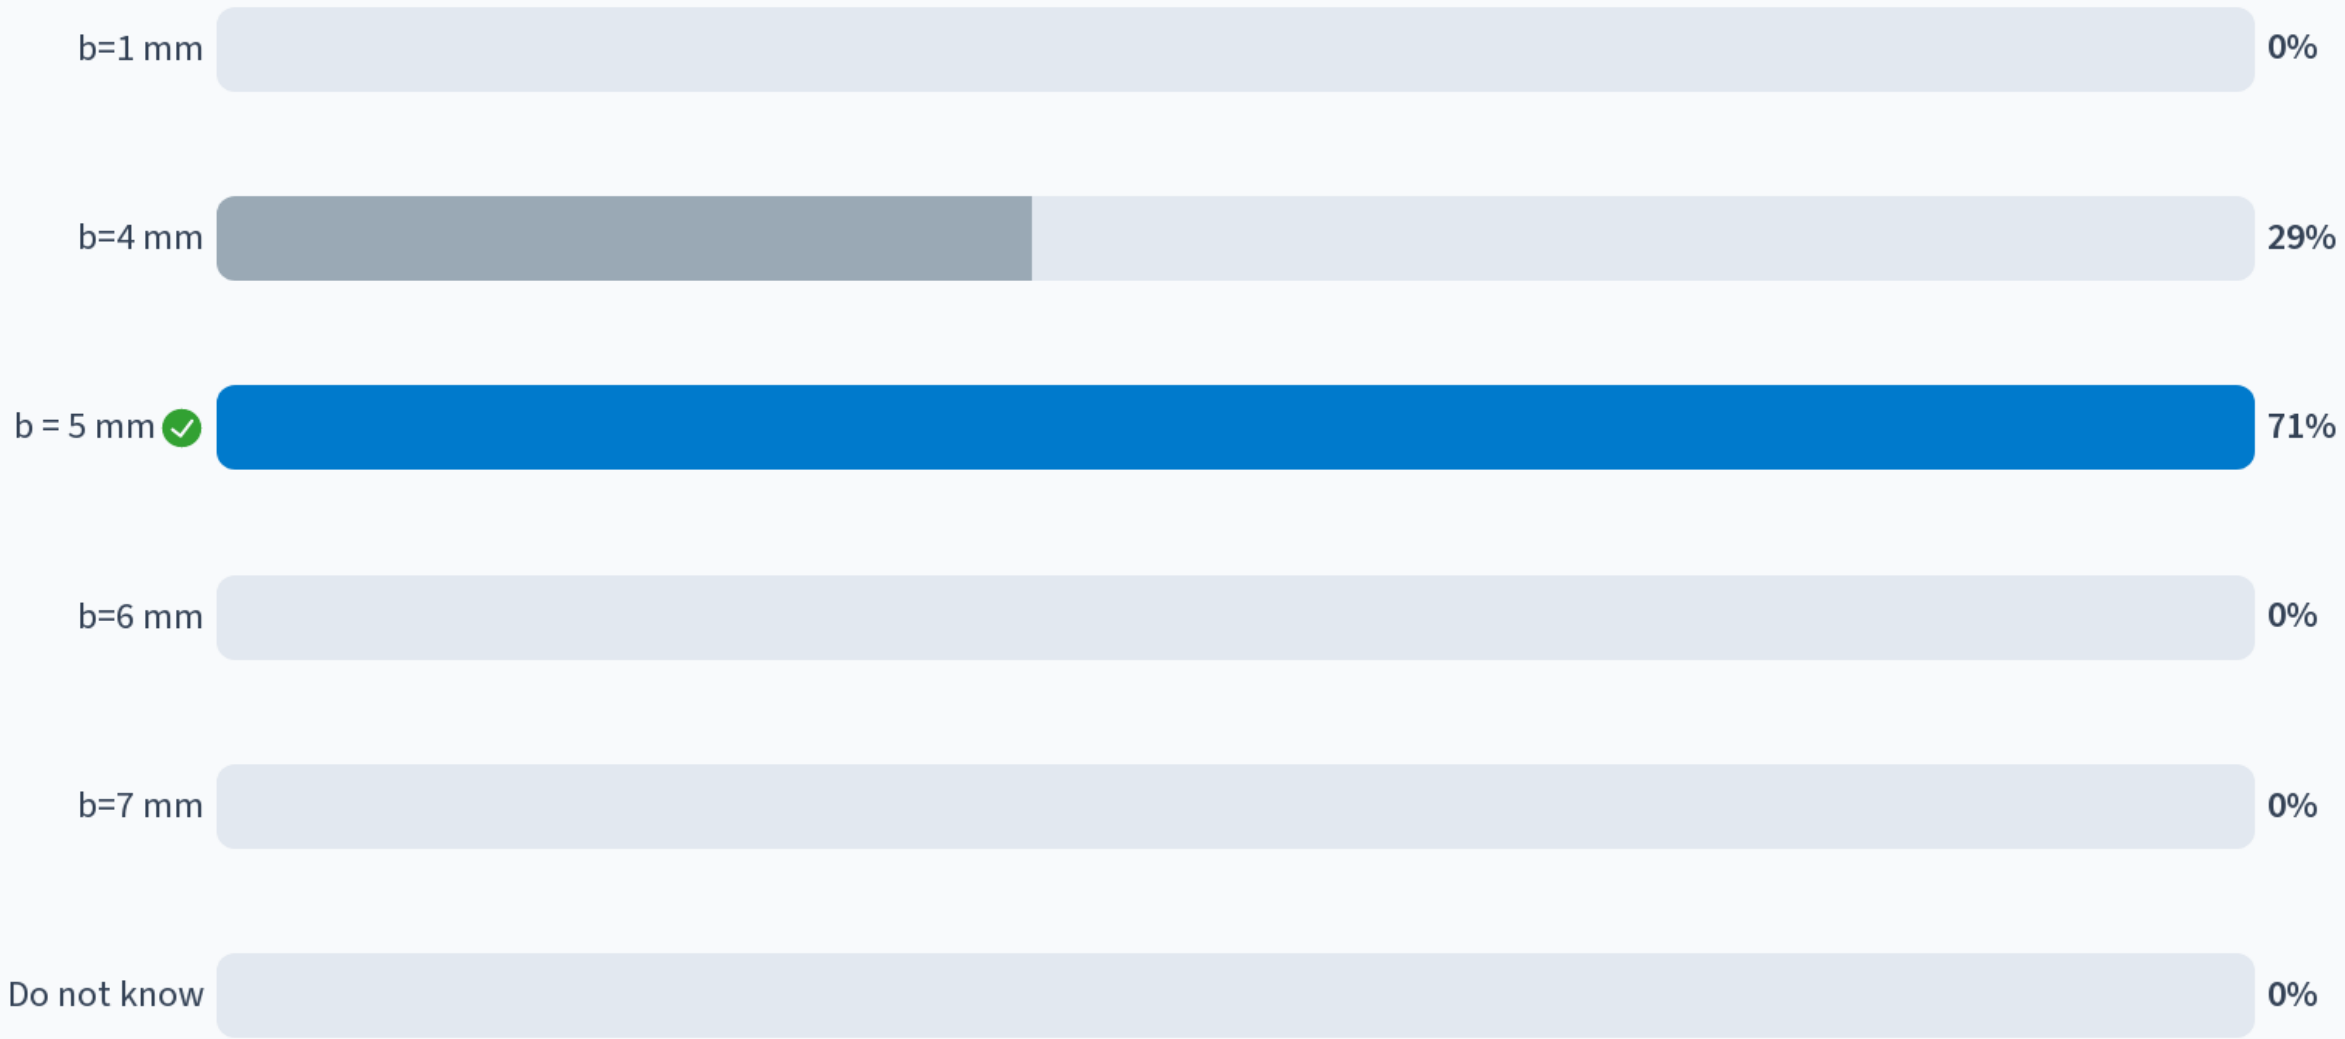

Start the presentation to see live content. For screen share software, share the entire screen. Get help at **pollev.com/app**

Focus or not to focus?

How do we make focused images? Placing the CCD right

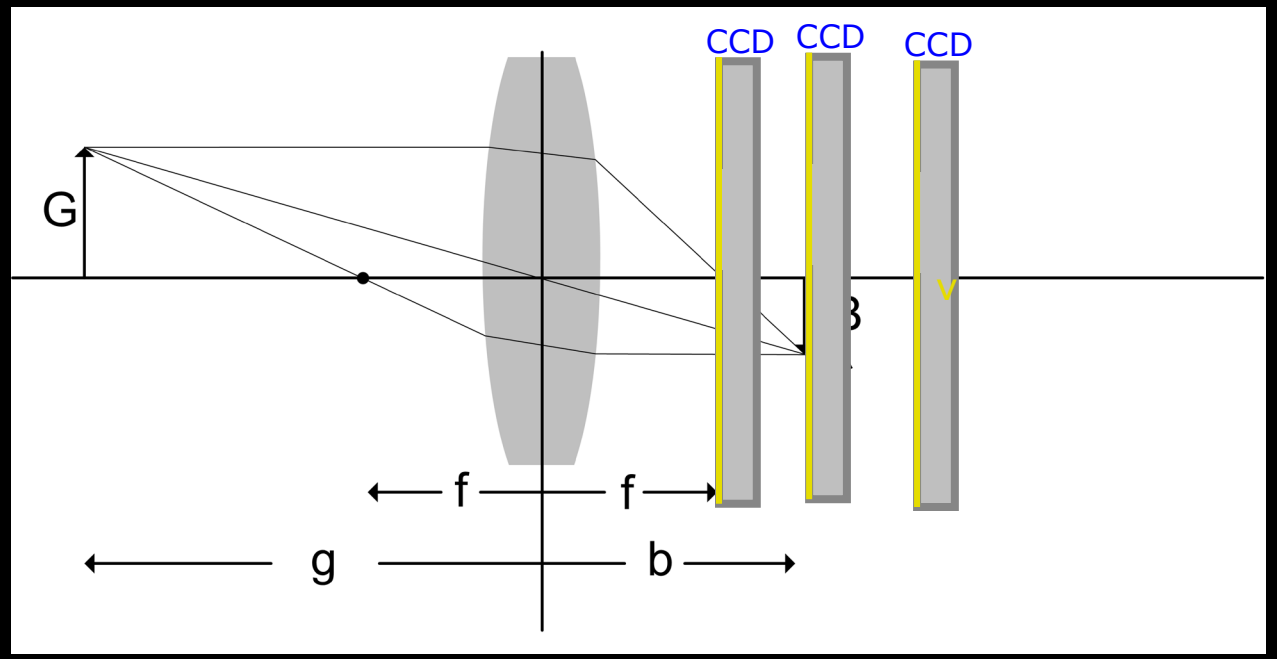

CCD should be placed at b!

World

Camera

g – distance to object

b – distance to intersection

Focusing

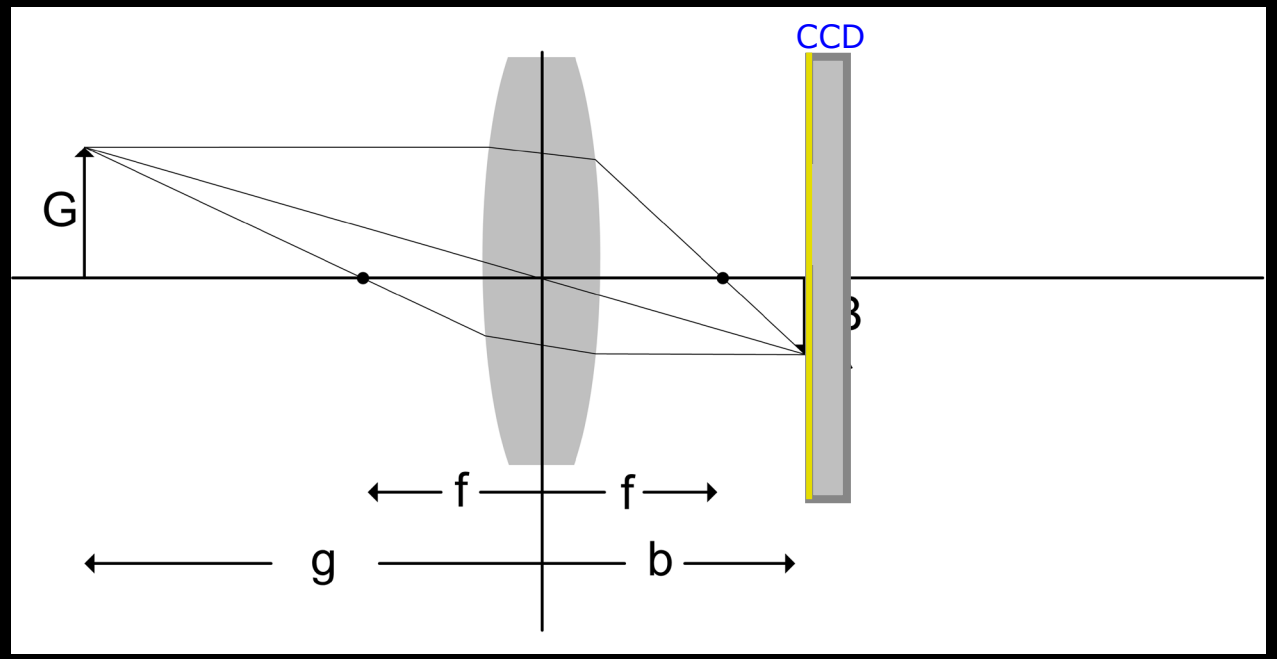

World

Camera

g – distance to object

b – distance to intersection

- We move the camera
- Distance to object (g) changes
- f is fixed
- b changes
- Move CCD to b
  - Focusing

$$\frac{1}{g} + \frac{1}{b} = \frac{1}{f}$$

Object size

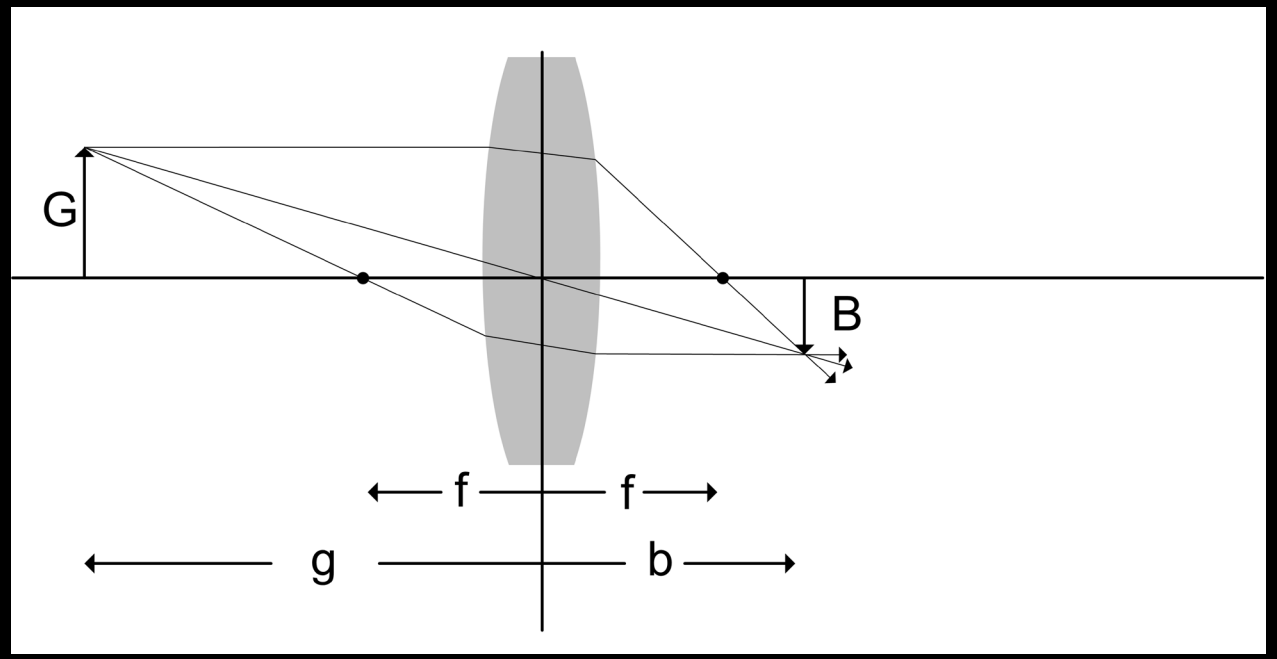

What is the size of an object on the CCD?

World

g – distance to object

G – Object height

Camera

b – distance to intersection

B – object height on CCD

An important relation!

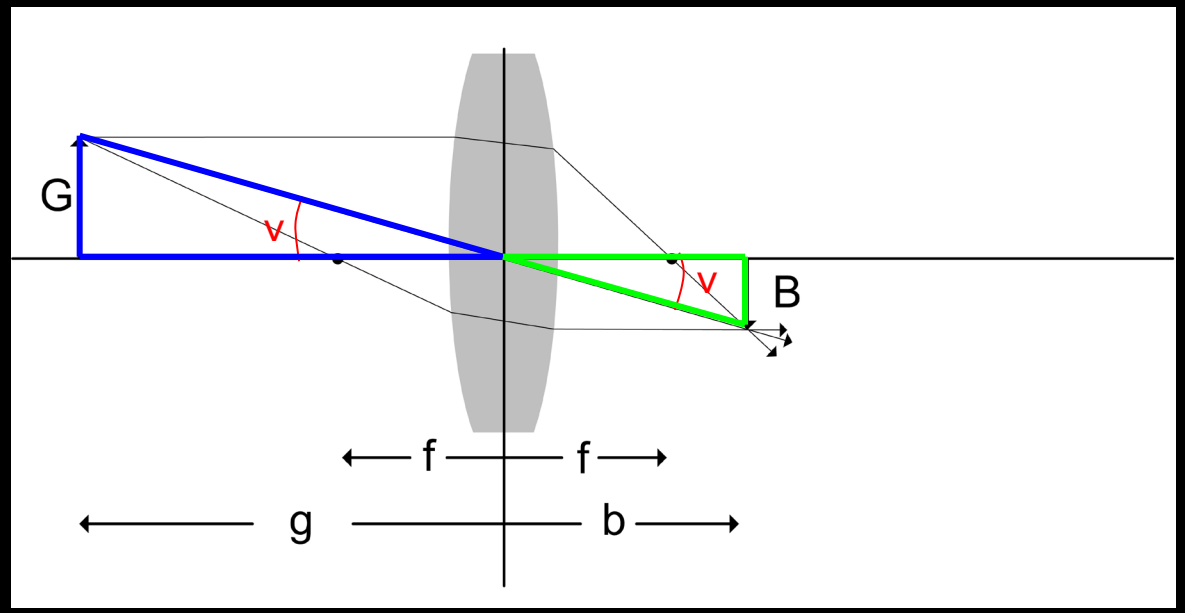

- Two triangles
- One with side length g and one with b
- B and G are related! – how?
- tangent

$$\frac{b}{B} = \frac{g}{G}$$

An important relation!

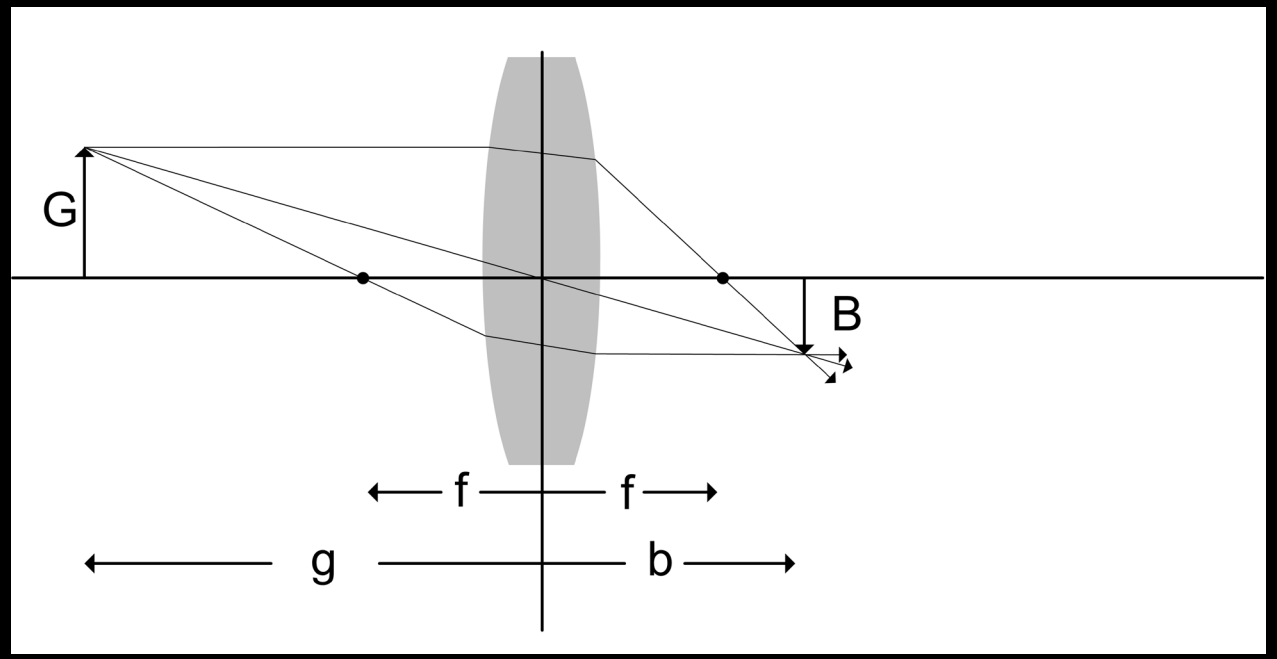

$$\frac{b}{B} = \frac{g}{G}$$

World

g – distance to object

G – Object height

Camera

b – distance to intersection

B – object height on CCD

How do we Zoom ?

We want to make B larger! How?

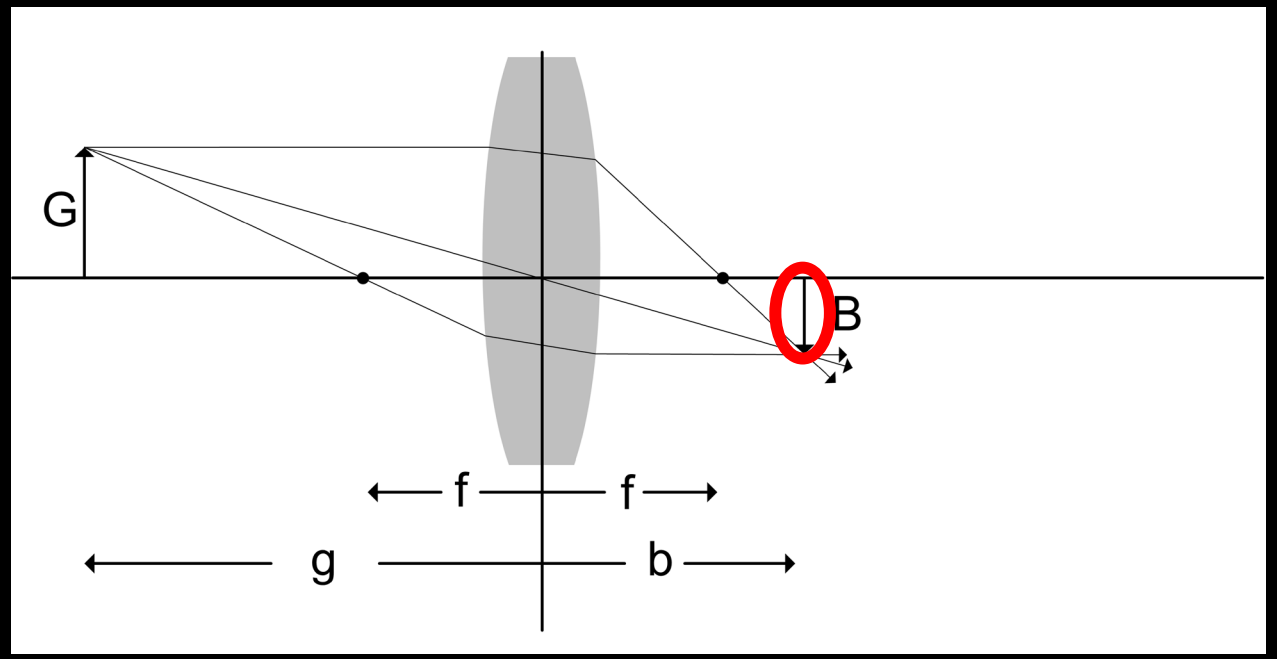

World

g – distance to object

G – Object height

Camera

b – distance to intersection

B – object height on CCD

Zoom

We want to make B larger! How?

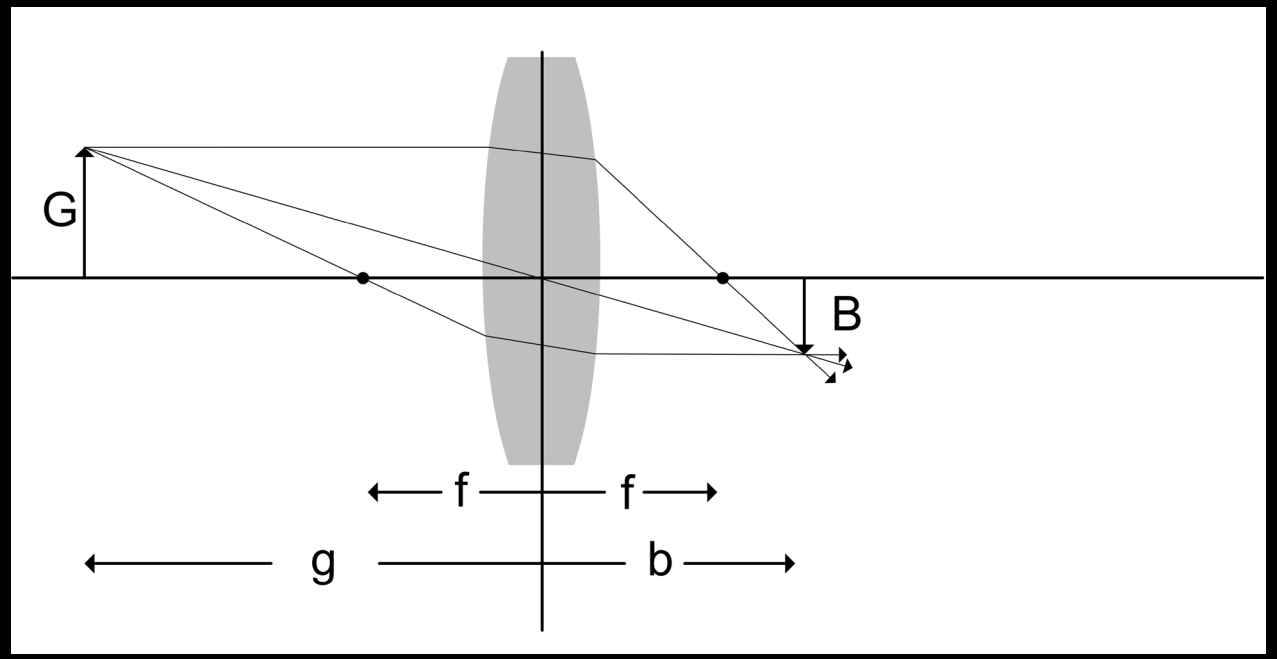

$$\frac{b}{B} = \frac{g}{G}$$

$$B = b\frac{G}{g}$$

Fixed

World

g – distance to object

G – Object height

Camera

b – distance to intersection

B – object height on CCD

Zoom

We want to make B larger – changing b!

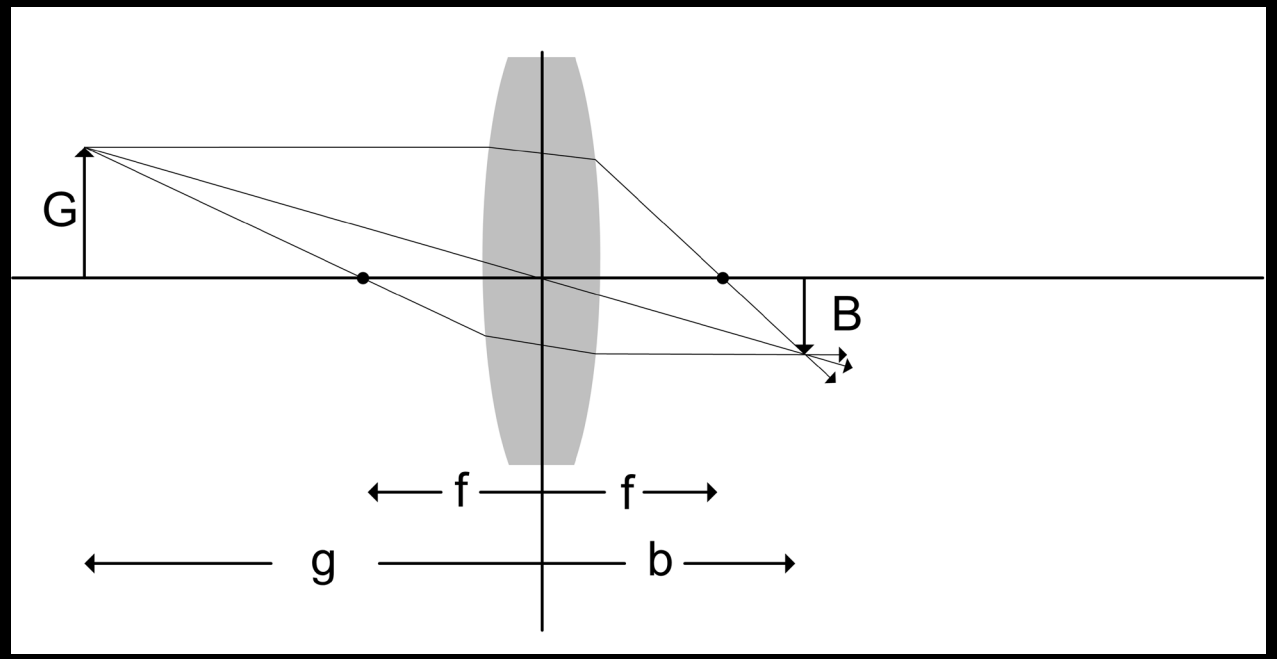

$$B = b\frac{G}{g}$$

$$\frac{1}{g} + \frac{1}{b} = \frac{1}{f}$$

constant

To change B we change the focal length!

World

g – distance to object

G – Object height

Camera

b – distance to intersection

B – object height on CCD

Changing the focal length?

- Not possible on a simple lens
- Need a "zoom lens"
- Several lenses together

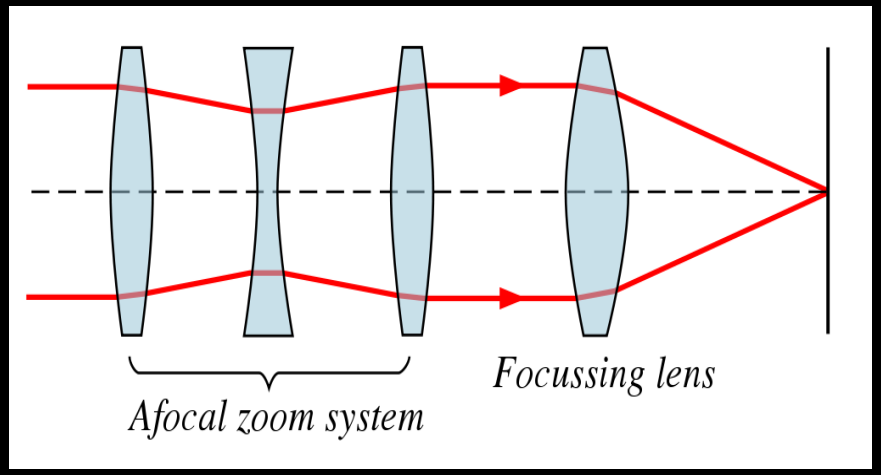

Afocal zoom principle

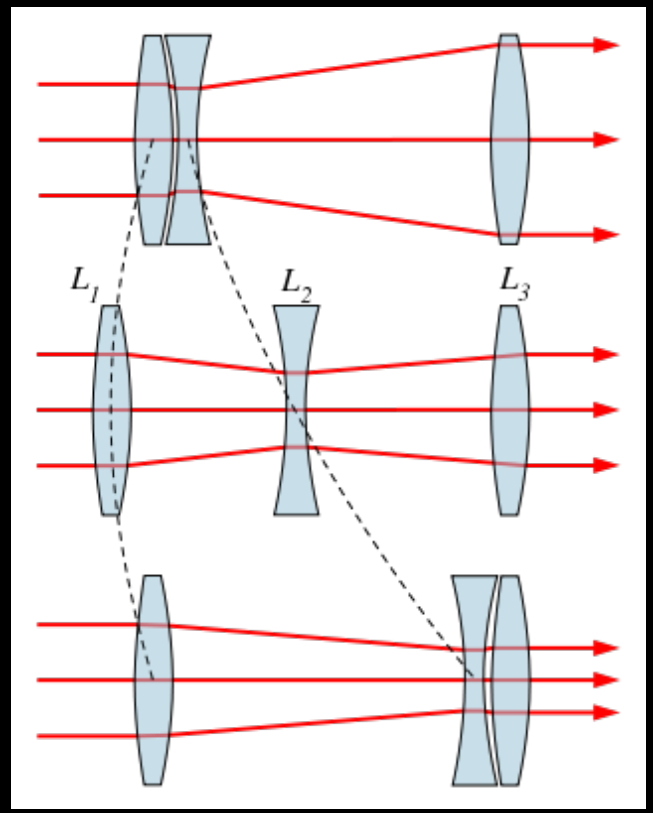

From WikiPedia: wikipedia.org/wiki/Zoom_lens

Field of view (FOV)

Two cameras with different focal length

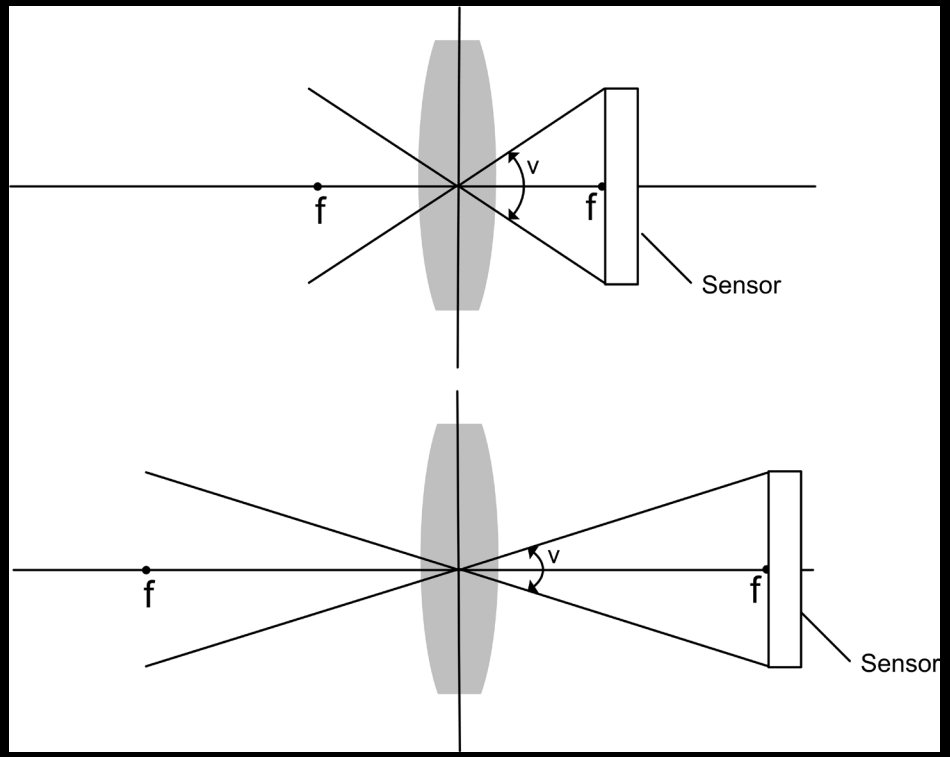

- Described by an angle
  - Large angle the larger FOV
- Depends on
  - CCD size
  - Focal length
- Fisheye lens
  - Small focal length
  - Large field of view
- CCD chip is a rectangle
  - Horizontal field of view
  - Vertical field of view
- Zoom changes field of view
  - Optical zoom
  - Digital zoom

Depth of field - dybdeskarphed

Depth of field

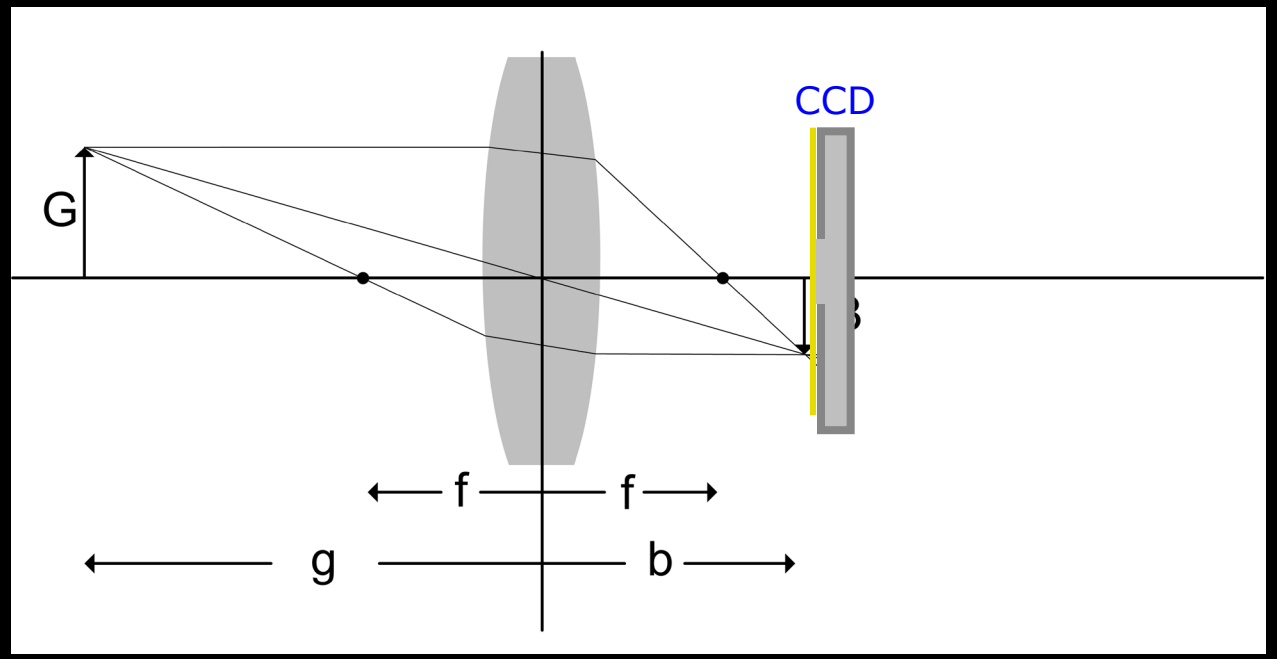

World

Camera

g – distance to object

b – distance to intersection

- CCD should be placed at b
- g is fixed – only focus at one distance!

Depth of field

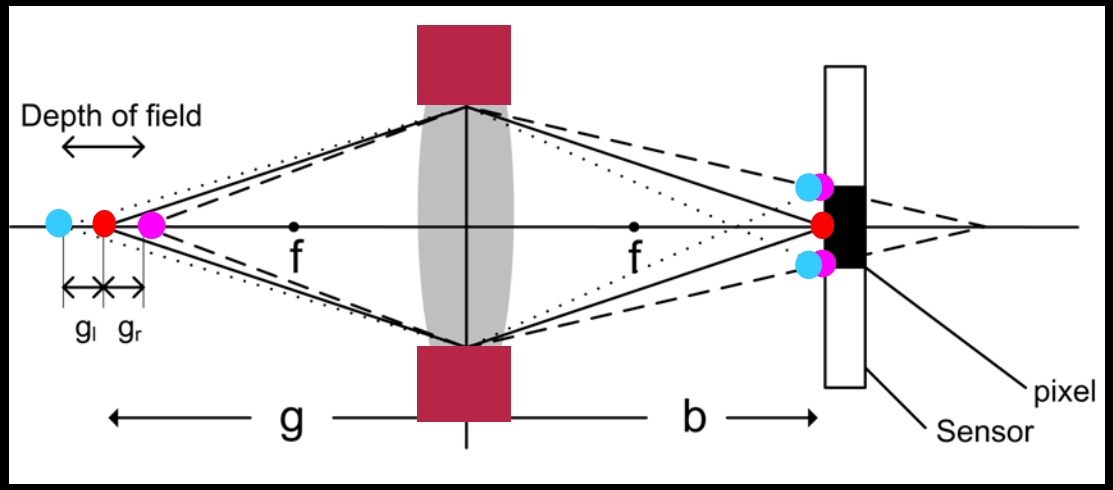

- Look at one pixel in the middle
- The object is placed at distance g
- How much can we move the object?
  - Light has to hit the same pixel
- Move it to the left (g_l)
- move it to the right (g_r) – still hit the same pixel (but twice)

Depth of field – Aperture (blænde)

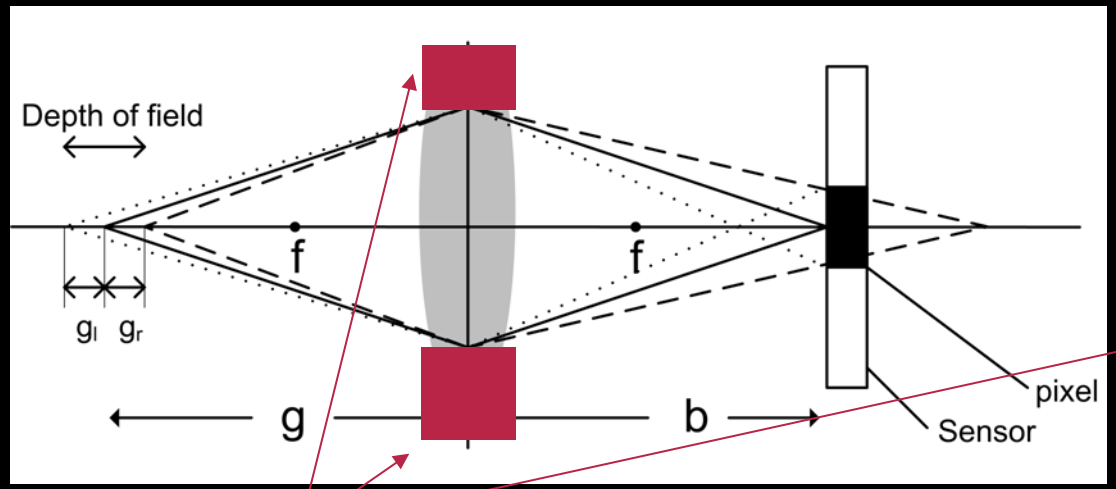

- The aperture controls the amount of light
- Small aperture
  - large depth of field
  - Less light -> longer exposure

How to acquire a good image?

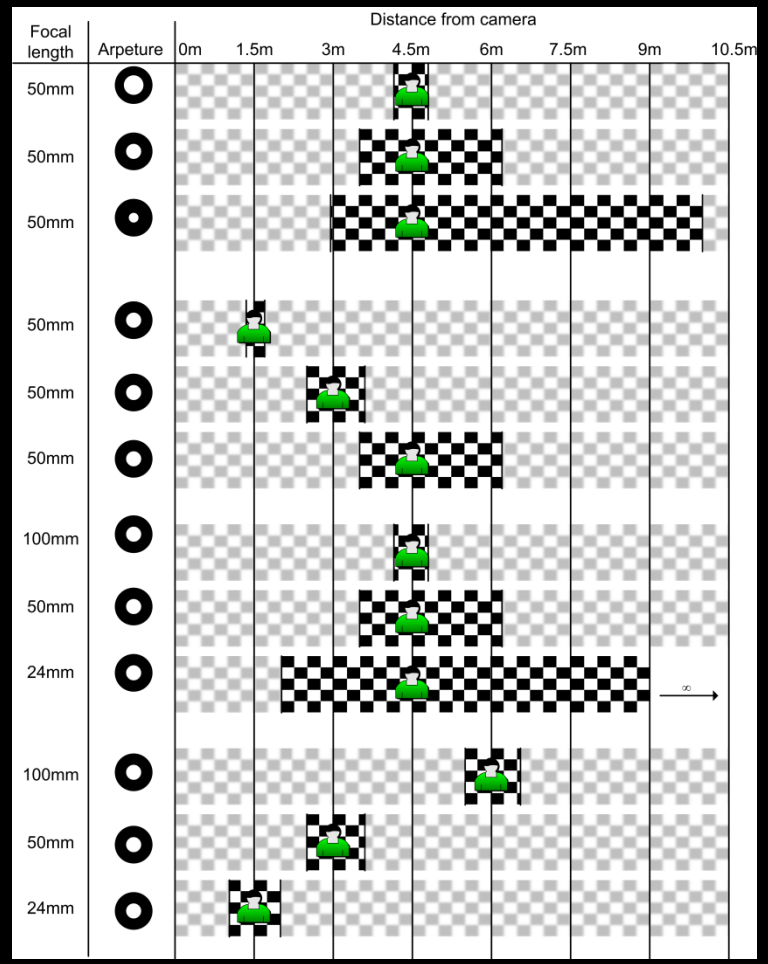

- Distance to object
- Motion of object
- Zoom
- Focus
- Depth-of-fields
- Focal length
- Shutter
- Field-of-view
- Aperture (DK: blænde)
- Sensor (size and type)

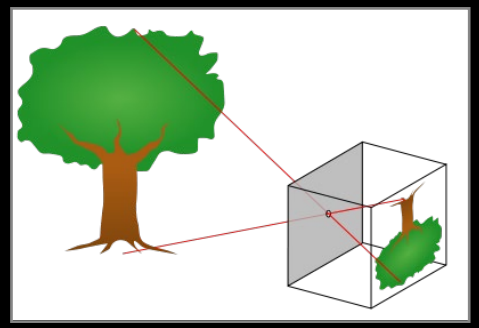

Image storage and compression

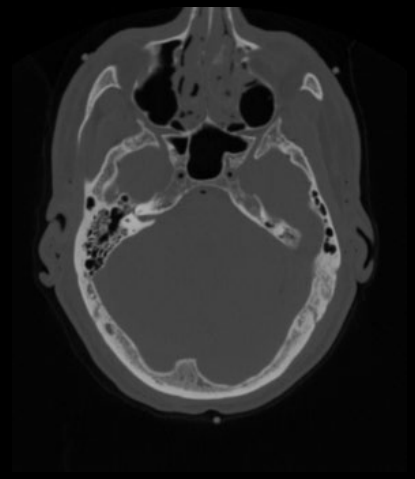

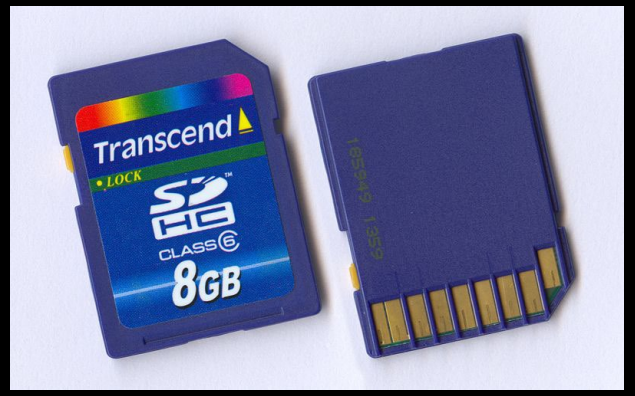

Learning objectives – image storage and compression

- Compute the run-length code of a grayscale image
- Compute the chain coding of a binary image
- Compute the compression ratio
- Describe the difference between a lossless and a lossy image format
- Decide if a given image should be stored using a lossless or a lossy image format

Hard disks, memory cards, CDs etc

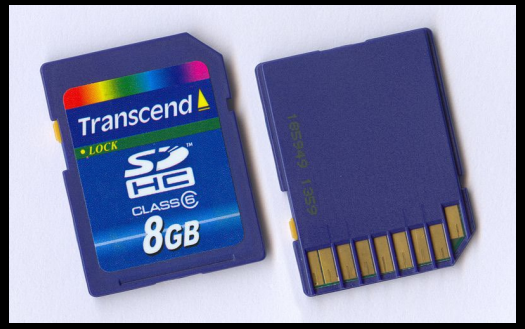

- Storage for bytes!
  - 500 GB?
  - 500 GigaBytes = 500.000.000.000 bytes!
- A hard disk does not know anything about images
- Stores data as lists of bytes
  - 17, 255, 1, 3, 87, 98, 11, …
- File on a hard disk
  - It has a length (in bytes, MB, GB)
  - Contains numbers! (Bytes)

We want to make an "image file"

Image as data

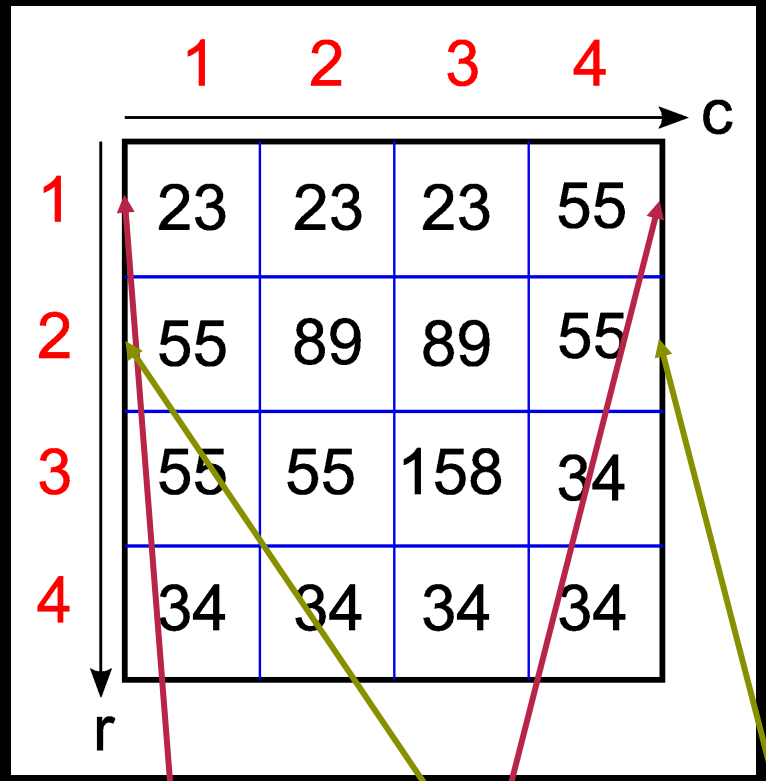

- How do we store this image as list of bytes?
- What do we need
  - Size of the image
    - Width as 2 bytes (0-65535)
    - Height as 2 bytes (0-65535)
  - The data

23,23,23,55,55,89,89,55,55,55,158,34,34,34,34,34

Simple image format

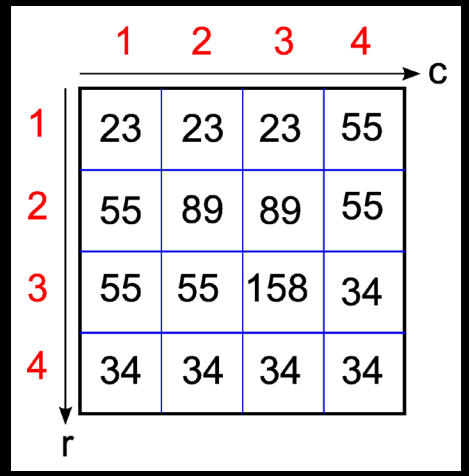

- Stores the image as
  - A header with information about size
  - Data with no compression
- Windows Bitmap Format (BMP)

Compression - make something smaller

- Is there a more "compact" way to represent the data below?
- Look for patterns
  - A series of numbers can be represented how?
    - The count and the value
- What is the "count and value" code?
  - Reduced from 16 to 12 values

Run length encoding

23,23,23,55,55,89,89,55,55,55,158,34,34,34,34,34

3,23, 2,55, 2,89, 3,55, 1,158, 5,34

Run length encoding

- Simple but useful data compression
- General – not only for images
- Is also used by the Windows Bitmap Format (BMP)

Run length coding of an image

1 1 3 5 2 3 1 2 2 3 2 201 3 19 2 147 4 130 3 147 2 88

1 1 3 5 2 3 1 3 2 201 3 19 2 147 4 130 3 147 2 88

1 3 5 2 2 3 3 4 1 2 2 3 2 201 3 19 2 147 4 130 3 147 2 88

1 1 3 5 2 3 1 2 2 4 4201 3 19 2 147 4 130 3 147 2 88

1 1 3 5 2 3 1 2 2 3 2 3 3 19 2 147 3 2 3 4 4 130 3 147 2 88

I do not know

Start the presentation to see live content. For screen share software, share the entire screen. Get help at **pollev.com/app**

Run length coding of an image

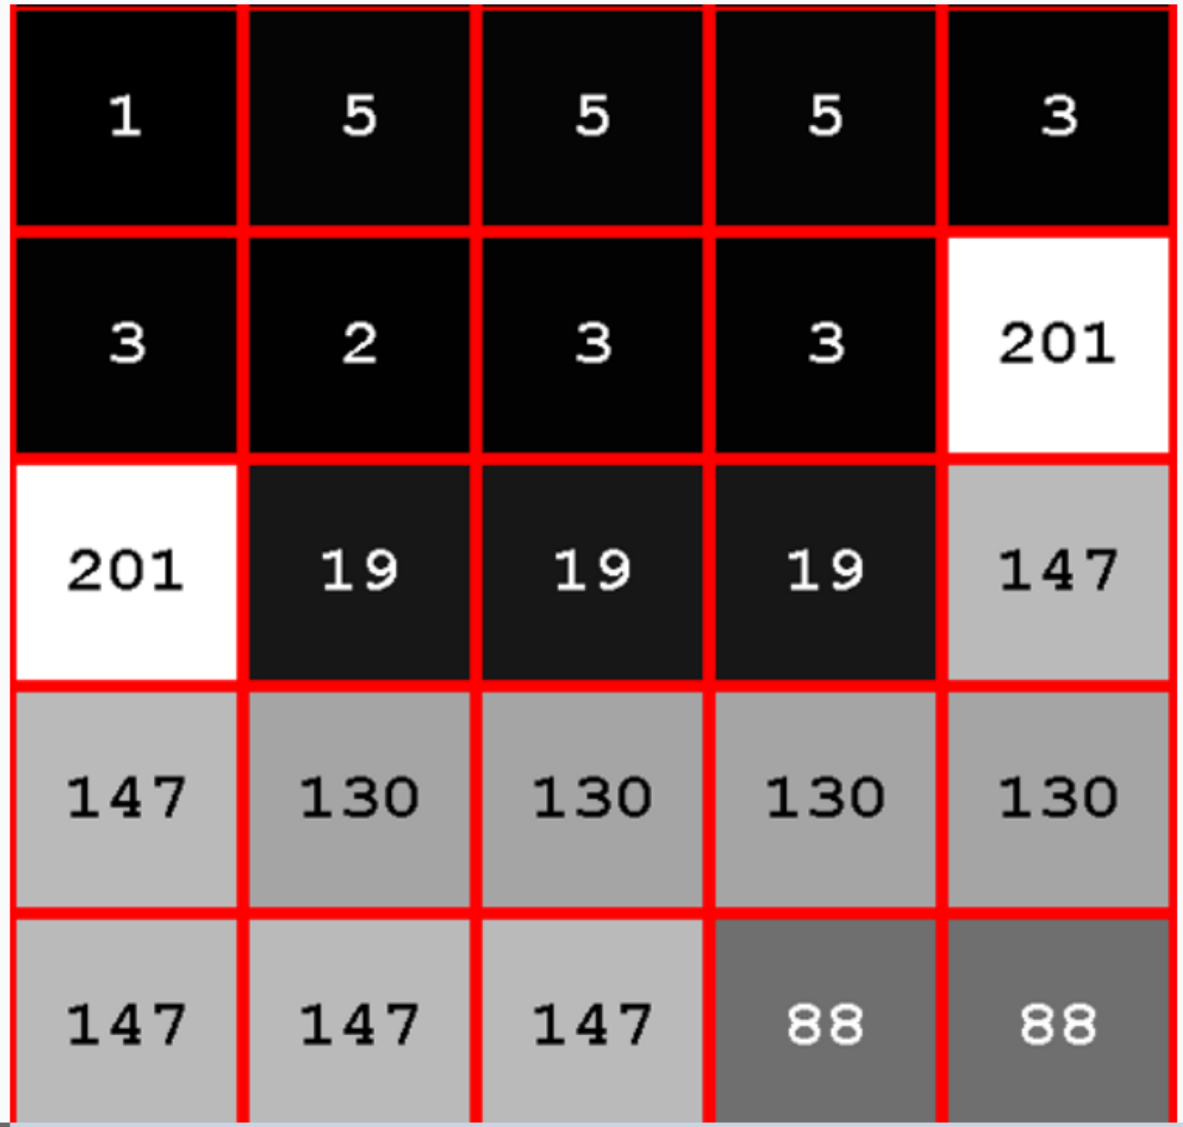

1 1 3 5 2 3 1 2 2 3 2 201 3 19 2 147 4 130 3 147 2 88 ✓	67%
1 1 3 5 2 3 1 3 2 201 3 19 2 147 4 130 3 147 2 88	0%
1 3 5 2 2 3 3 4 1 2 2 3 2 201 3 19 2 147 4 130 3 147 2 88	33%
1 1 3 5 2 3 1 2 2 4 4201 3 19 2 147 4 130 3 147 2 88	0%
1 1 3 5 2 3 1 2 2 3 2 3 3 19 2 147 3 2 3 4 4 130 3 147 2 88	0%
I do not know	0%

Start the presentation to see live content. For screen share software, share the entire screen. Get help at **pollev.com/app**

Run length coding of an image

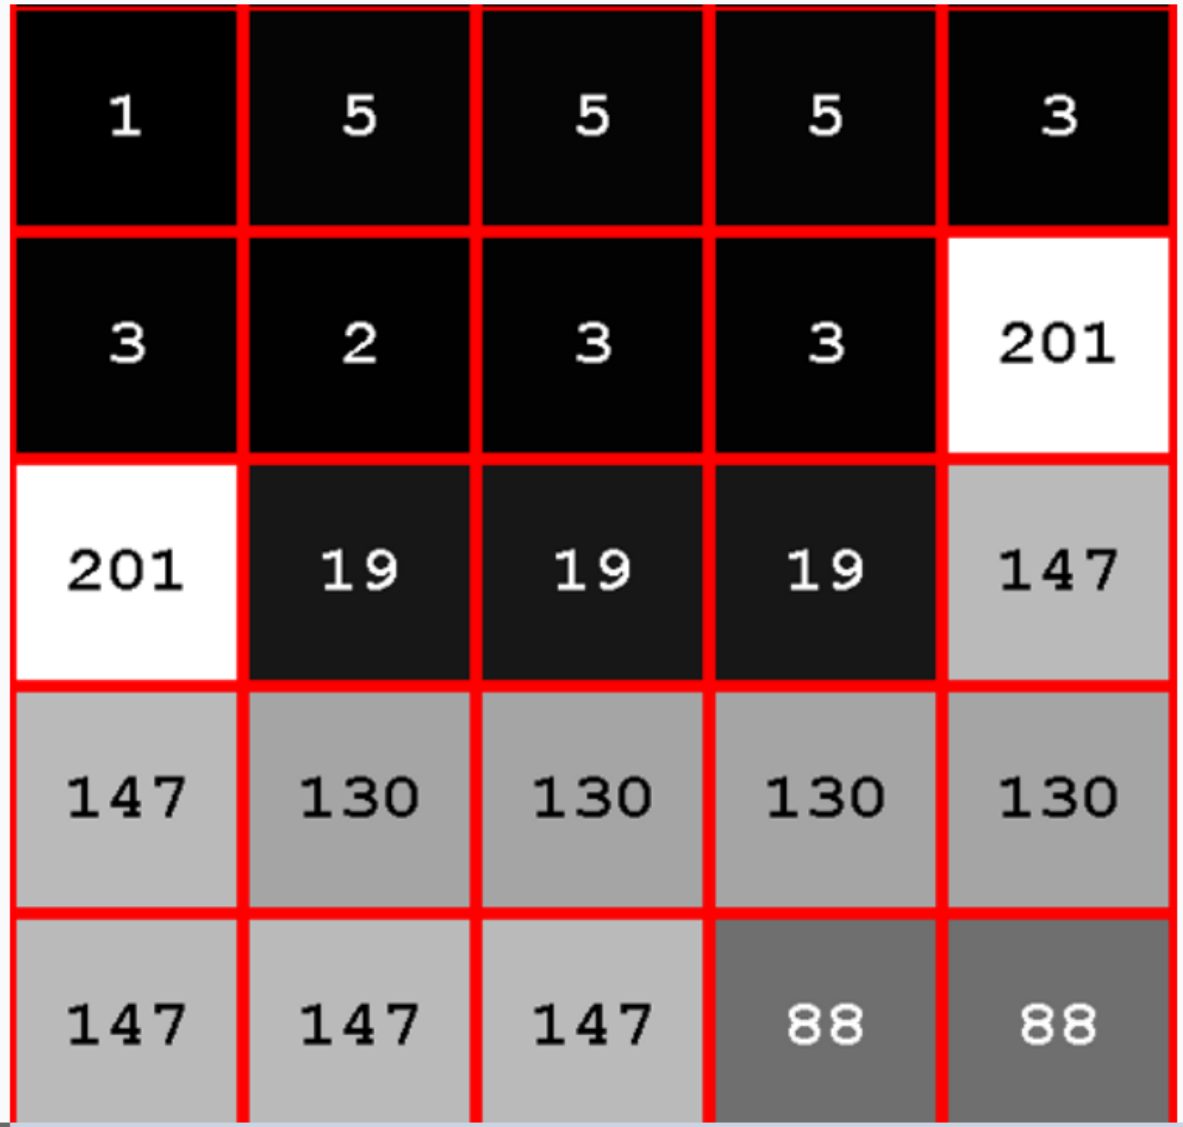

1 1 3 5 2 3 1 2 2 3 2 201 3 19 2 147 4 130 3 147 2 88 ✓	67%
1 1 3 5 2 3 1 3 2 201 3 19 2 147 4 130 3 147 2 88	0%
1 3 5 2 2 3 3 4 1 2 2 3 2 201 3 19 2 147 4 130 3 147 2 88	33%
1 1 3 5 2 3 1 2 2 4 4201 3 19 2 147 4 130 3 147 2 88	0%
1 1 3 5 2 3 1 2 2 3 2 3 3 19 2 147 3 2 3 4 4 130 3 147 2 88	0%
I do not know	0%

Start the presentation to see live content. For screen share software, share the entire screen. Get help at **pollev.com/app**

Compression ratio – how compressed?

- Gives a measure for how much data is compressed
- Our example
  - From 16 to 12
  - 16 : 12 = 4 : 3
  - Ratio 1.33

Compression ratio = uncompressed size / compressed size

Lossless image formats

- Do not throw away information
- Good for storing medical images
  - We do not want to destroy any information
- Not very effective for photos. Why?
  - To many changes in the image
- PNG (portable network graphics) is a good format

Noisy image

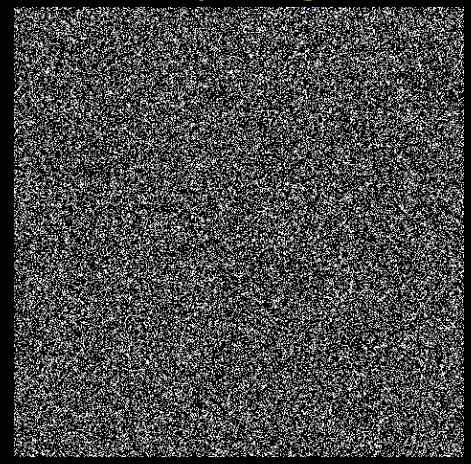

Lossy image formats

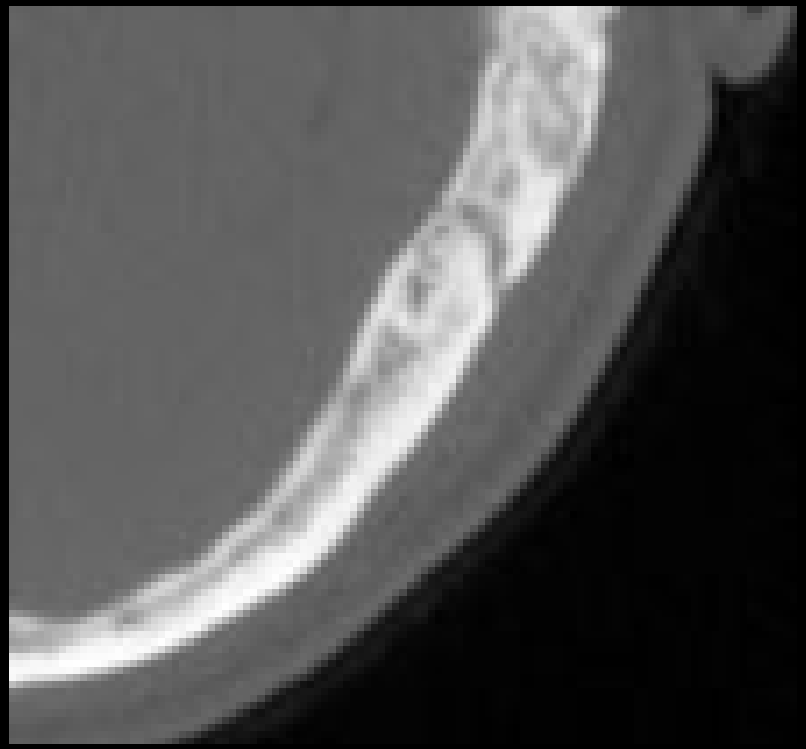

- Removes “unimportant” information
- JPEG is an example
- Removes the “high frequencies”
- Similar to the MP3 sound format

Compression artefacts

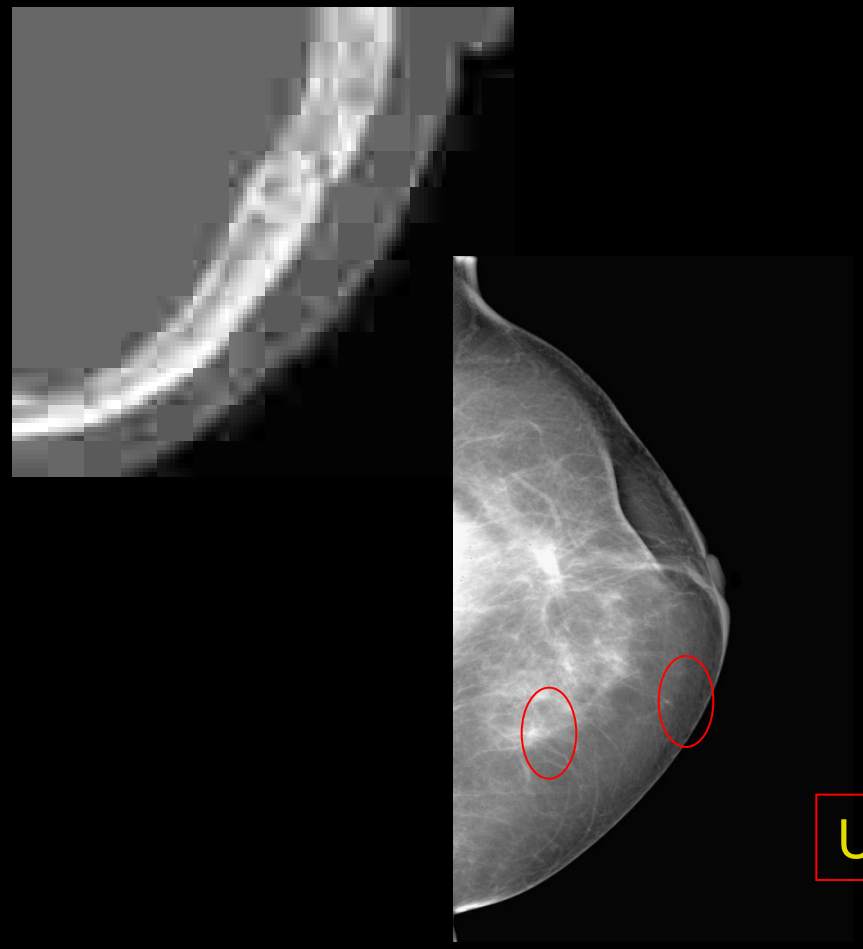

- Lossy compression changes the image
- Normally not a problem for photos
- BIG problem for medical images
- Mammogram
  - Looking for tiny bright spots
  - Would be changed by lossy compression

Use JPEG (JPG) for photos only

Binary images

- Binary – means on or off
- Binary image – only two colors
- Background (0 = black)
- Foreground (1 = white)

Chain coding of binary images

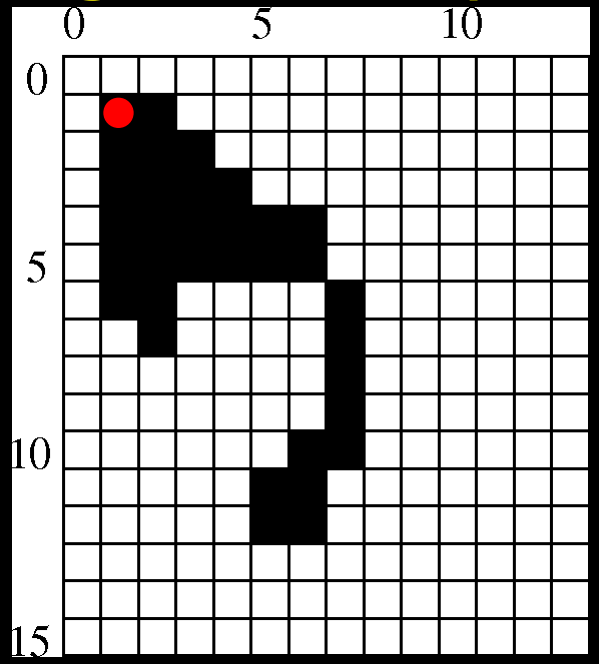

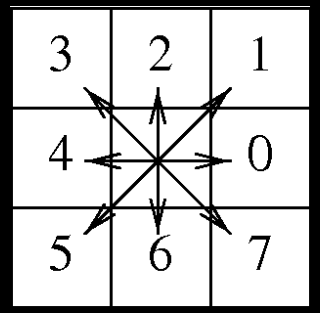

- Sufficient to describe the foreground
- Background given by the foreground
- The coordinates of the starting pixel is stored
- Secondly the sequence of step directions is stored

(1; 1)(07770676666564211222344456322222)

Chain code - what is in the image?

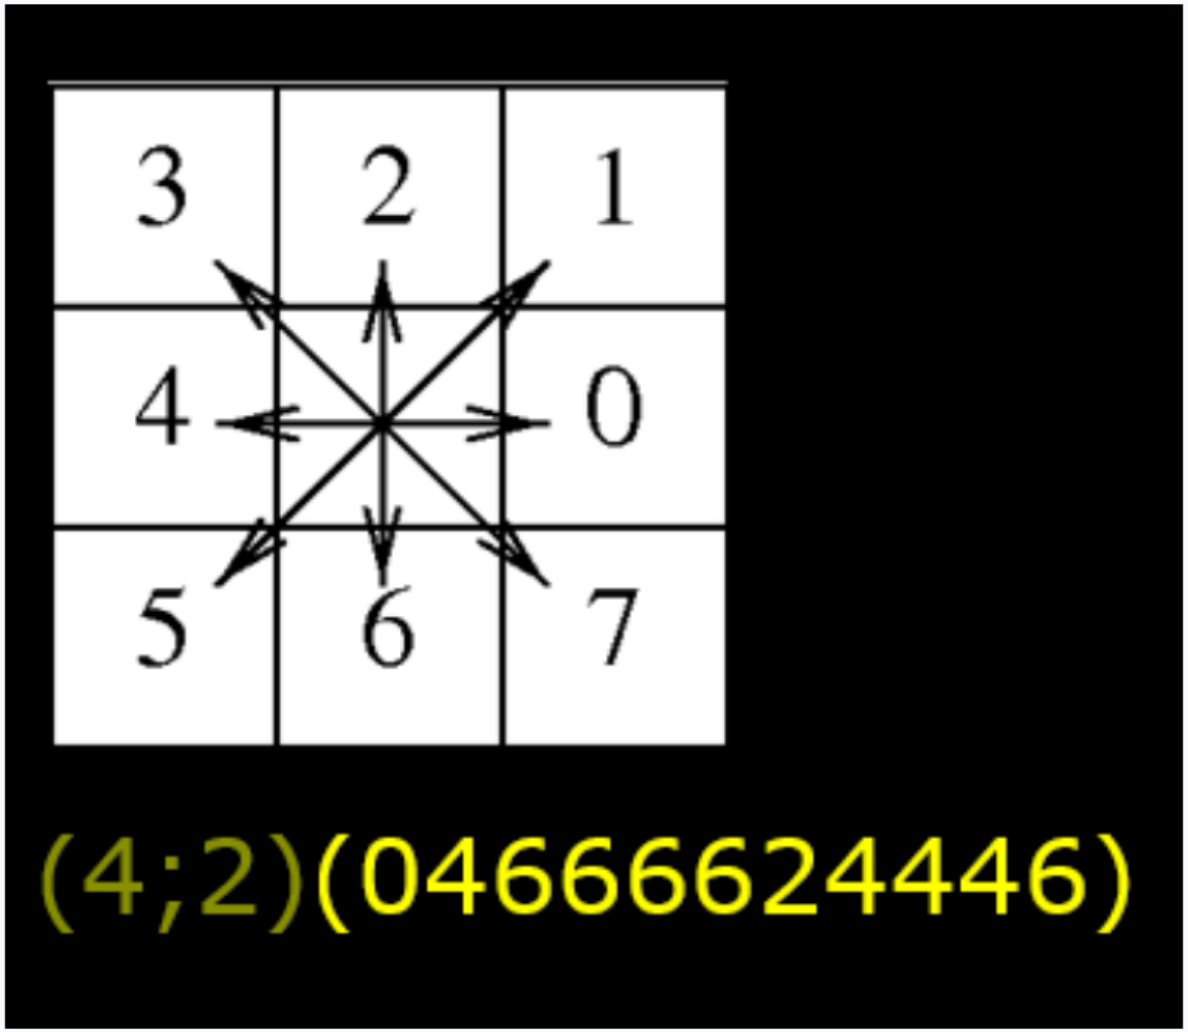

- House
- Flower
- Giraffe
- Dog
- Teapot
- Car
- Glass
- Bottle
- I do not know

Start the presentation to see live content. For screen share software, share the entire screen. Get help at **pollev.com/app**

Start the presentation to see live content. For screen share software, share the entire screen. Get help at **pollev.com/app**

Chain code - what is in the image?

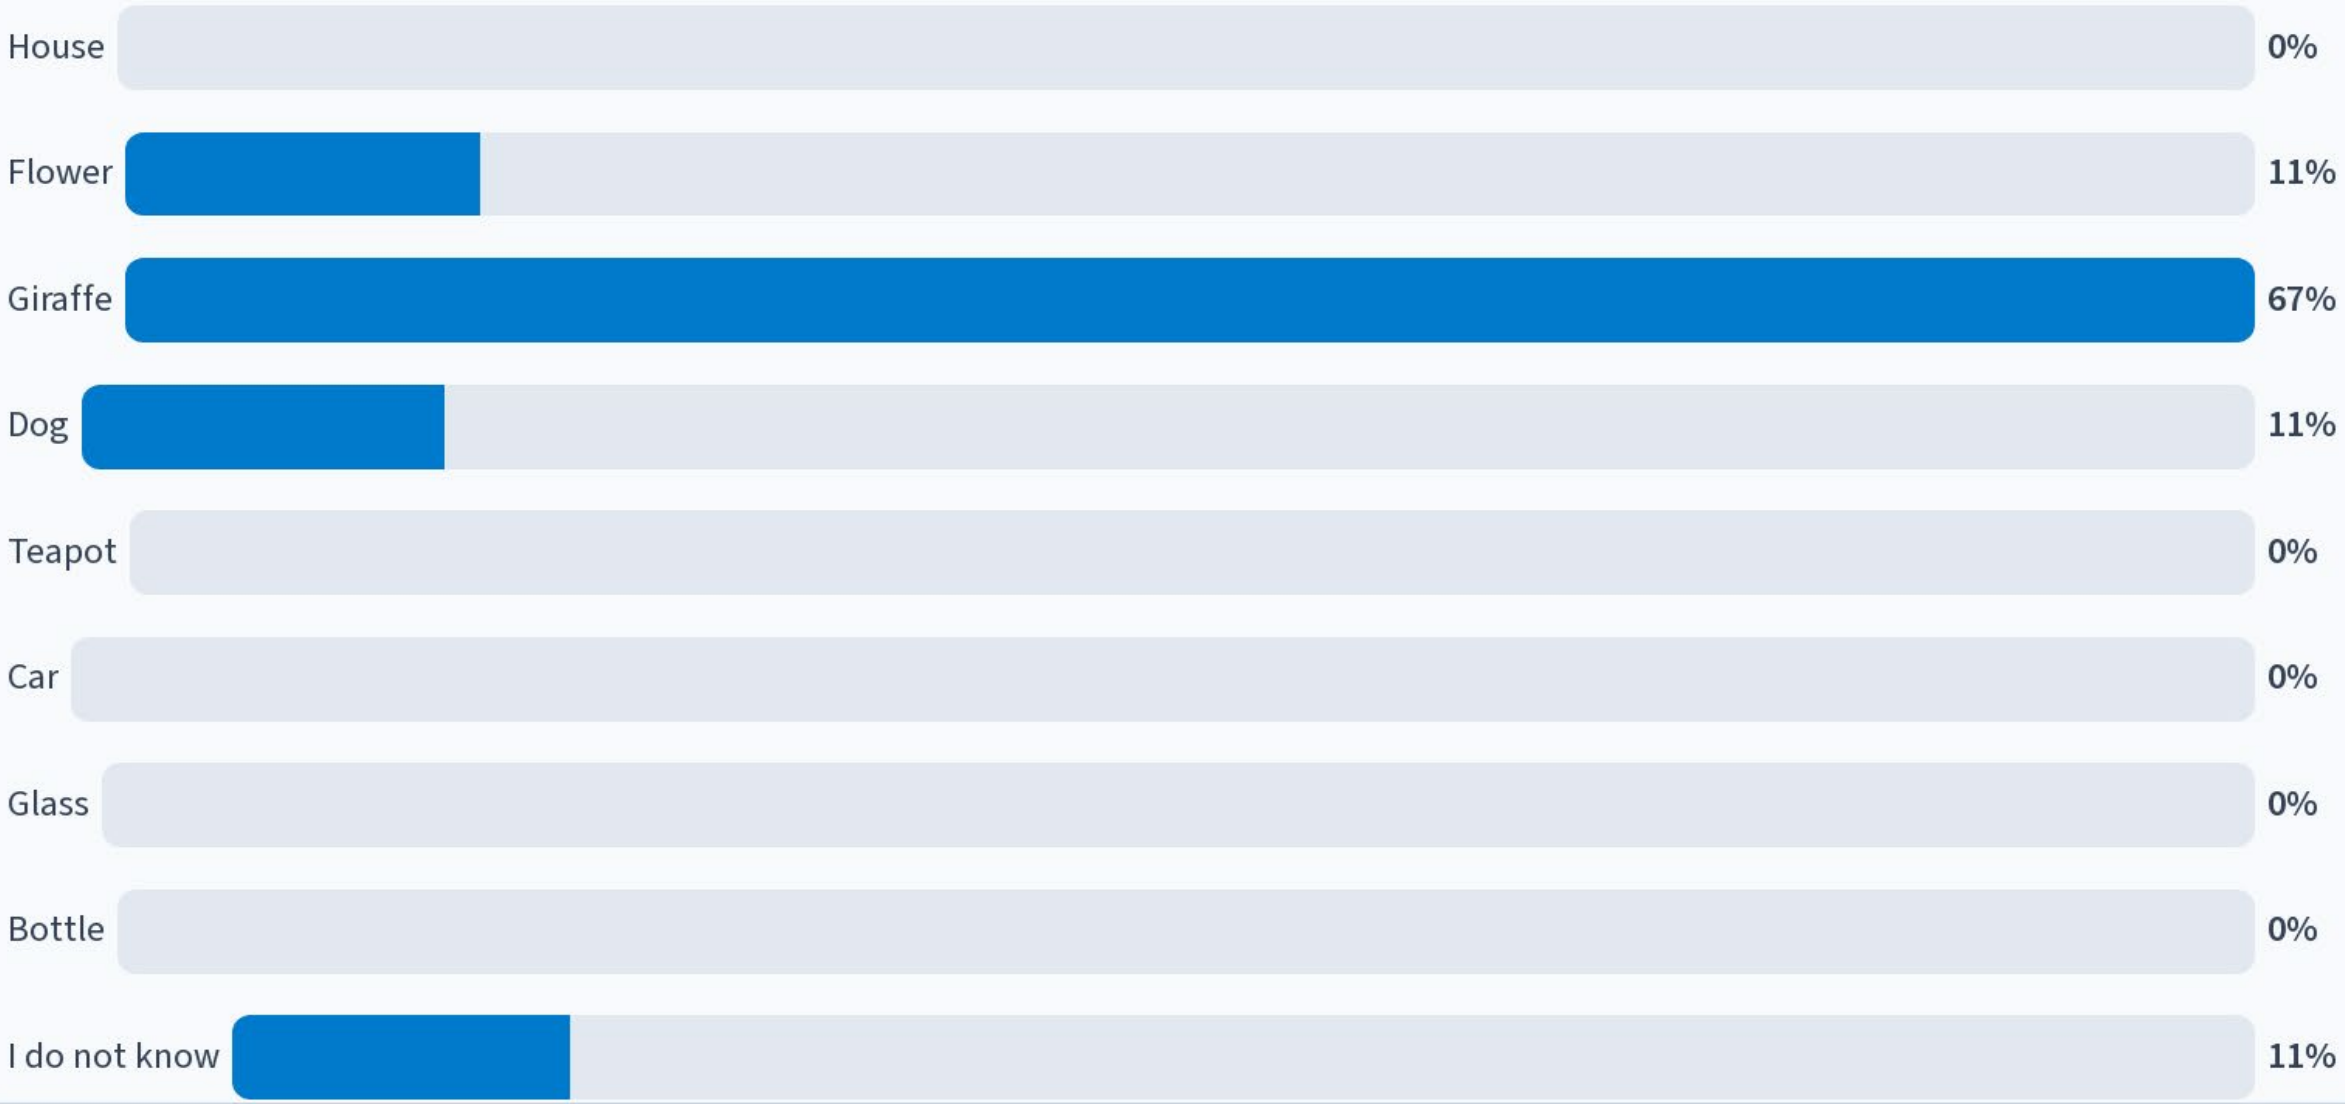

Start the presentation to see live content. For screen share software, share the entire screen. Get help at **pollev.com/app**

Video Analysis

- Video – images coming in a stream from a camera
- Automated video analysis applications
  - Industrial / agricultural sorting machines
  - Activity alerts for surveillance cameras
  - Sports tracking
  - Self-driving vehicles
  - Driver awareness tracking / alerts
  - Space-ship navigation
  - Tracking of surgical instruments
  - And many more..

Change detection in videos

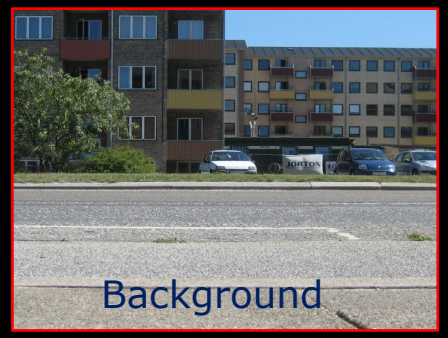

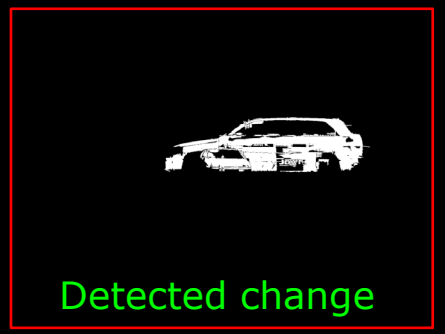

- Automatically detects changes in video stream
- The basis for many processing steps
  - Human pose tracking
  - Vehicle tracking
  - Alert systems
  - Cell tracking
  - …

Learning objectives – Video change detection

- Describe the concept of change detection
- Describe the camera, the processing and the total system frame rate
- Compute the maximum frame rate based on an algorithm processing time
- Compute a background/reference image when the scene is static or slowly changing
- Use pre-processing steps like color conversion and resizing to make images from a video stream ready to be analyzed
- Use image differencing to compute changes in a video stream
- Use background subtraction to compute changes in a video stream
- Use a threshold to create a binary image from a difference image
- Describe alternative approaches for background/reference image estimation
- Describe different scenarios where an action can be taken based on detected changes in a video stream

Exercise 2b

- The goal of exercise 2b is to implement and test a small change detection system

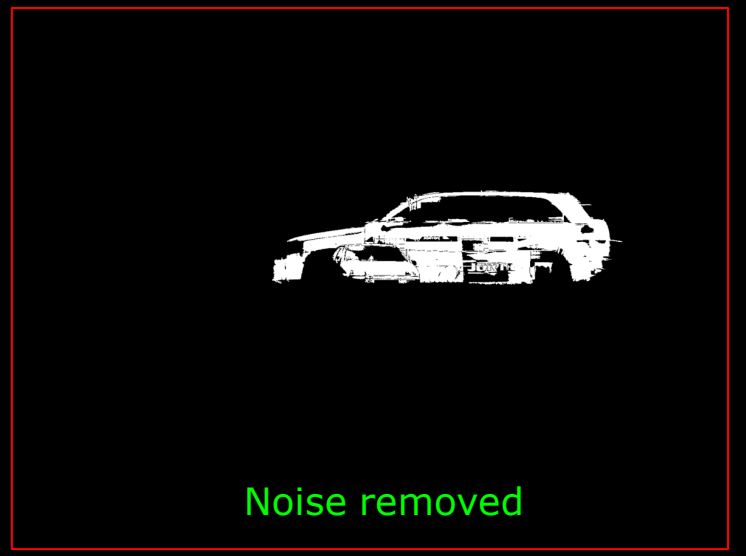

Cameras and videos – frame rate

- A camera delivers video in the forms of a stream of images (also called frames)
- The frame rate is the amount of frames per second. For example 20 frames/s (measured in Hz)
- For video processing we have two frame rates
  - How many frames can the camera deliver per second
  - How many frames can we analyze per second
- The *system frame* rate is the minimum of the camera and processing frame rate

Camera frame rate

Your camera is attached to your computer using a USB-2 connection. On your system, the maximum transfer speed is 30 megabytes per second (MB/s). The images are RGB images (3 bytes per pixel) and their size is 640 x 480. What is the maximum camera frame rate on your system?

- 16
- 25
- 32
- 61
- 103
- I do not know

Start the presentation to see live content. For screen share software, share the entire screen. Get help at **pollev.com/app**

Camera frame rate

Your camera is attached to your computer using a USB-2 connection. On your system, the maximum transfer speed is 30 megabytes per second (MB/s). The images are RGB images (3 bytes per pixel) and their size is 640 x 480. What is the maximum camera frame rate on your system?

Start the presentation to see live content. For screen share software, share the entire screen. Get help at **pollev.com/app**

Camera frame rate

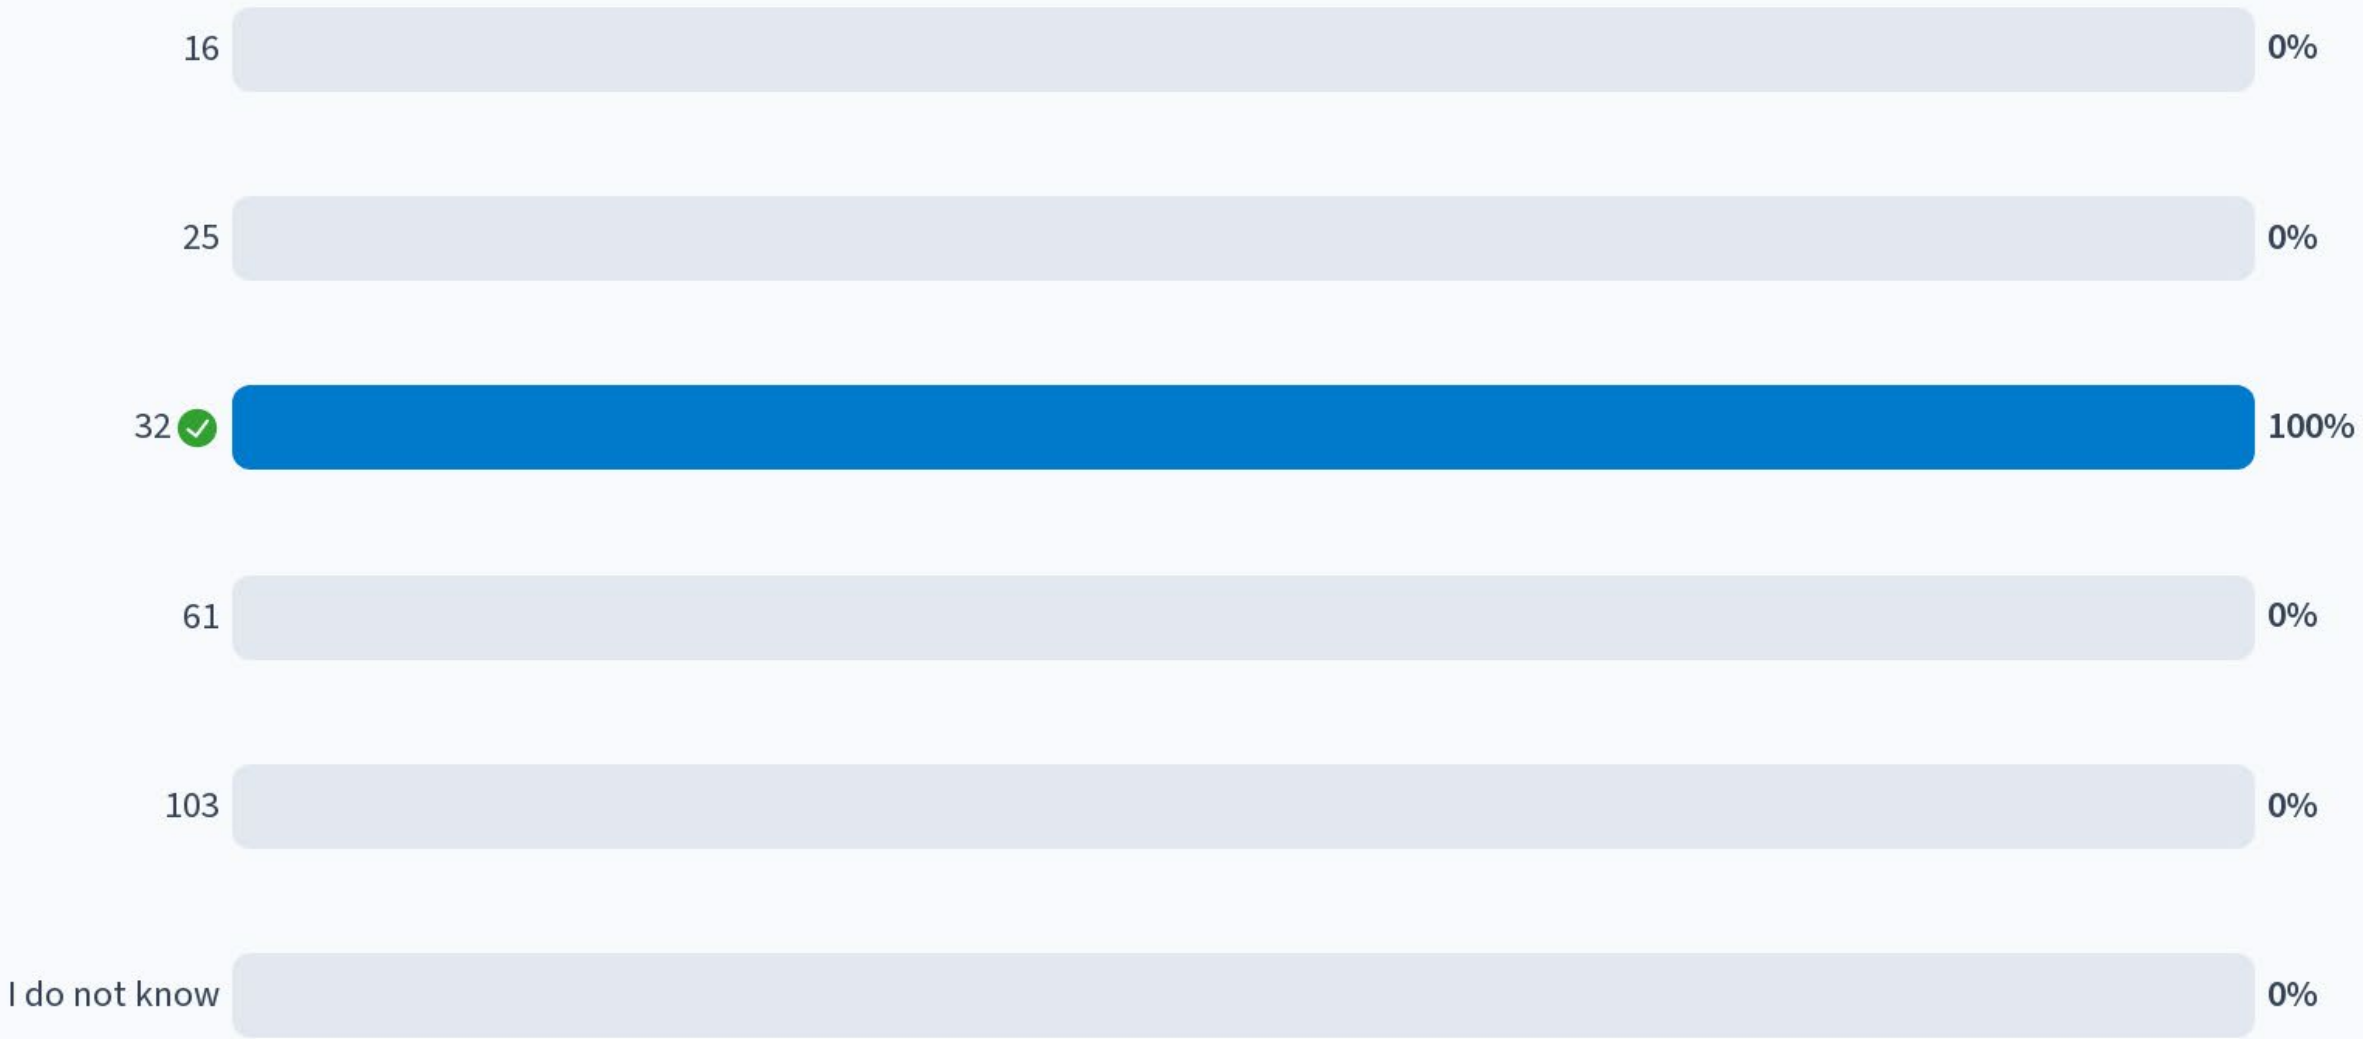

Start the presentation to see live content. For screen share software, share the entire screen. Get help at **pollev.com/app**

Your fancy image analysis algorithm uses 24 milliseconds to analyze one image. The camera delivers 60 frames per second. What is the maximum frame rate of your system?

- 12 Hz
- 27 Hz
- 38 Hz
- 41 Hz
- 67 Hz
- I do not know

Start the presentation to see live content. For screen share software, share the entire screen. Get help at **pollev.com/app**

Your fancy image analysis algorithm uses 24 milliseconds to analyze one image. The camera delivers 60 frames per second. What is the maximum frame rate of your system?

Start the presentation to see live content. For screen share software, share the entire screen. Get help at **pollev.com/app**

Your fancy image analysis algorithm uses 24 milliseconds to analyze one image. The camera delivers 60 frames per second. What is the maximum frame rate of your system?

Start the presentation to see live content. For screen share software, share the entire screen. Get help at **pollev.com/app**

How do we detect changes in a video stream?

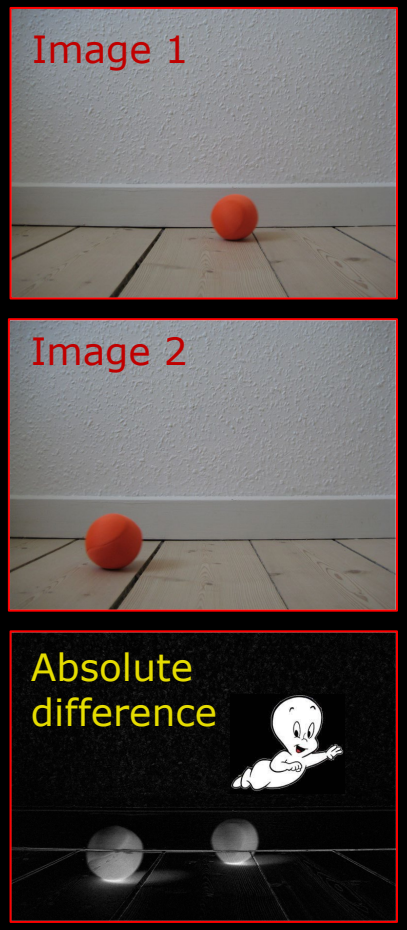

- One solution (out of many):
  - Subtract the previous image from the current image and take the absolute value in each pixel
  - *Image differencing*
- Several drawbacks:
  - Looses track of for example cars stopped for red light
  - *Ghost differences*

A change detection program

- Estimate and save a background/reference image
- Stop = False
- Do
  - Capture and pre-process one image
  - Compare with reference image (perhaps just subtraction)
  - Threshold difference image
  - Filter noise
  - Decide if something should be done
  - If ‘q’ key pressed:
    - Stop = True
- While not Stop

A change detection program

- Estimate and save a background/reference image
- Stop = False
- Do
  - Capture and pre-process one image
  - Compare with reference image (perhaps just subtraction)
  - Threshold difference image
  - Filter noise
  - Decide if something should be done
  - If ‘q’ key pressed:
    - Stop = True
- While not Stop

Background estimation

- Estimate a robust background image that can be used to detect significant changes

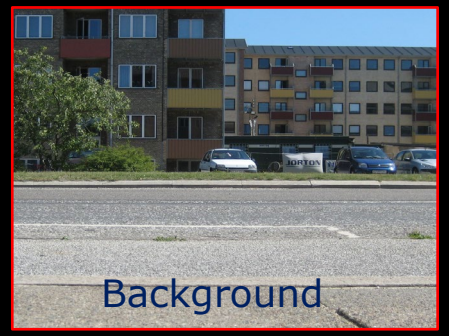

- The scene can be more or less complicated
  - A static scene – very controlled light and no moving objects
  - Slowly changing scene – light changes due to the movement of the sun
  - Rapidly changing scene – Fast movement of leaves due to wind

Naturally occurring changes

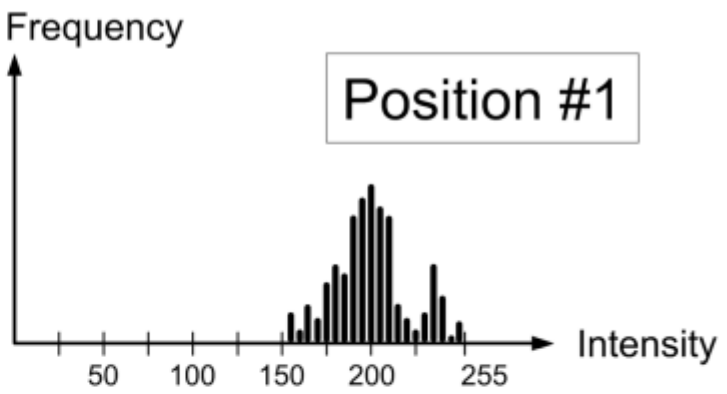

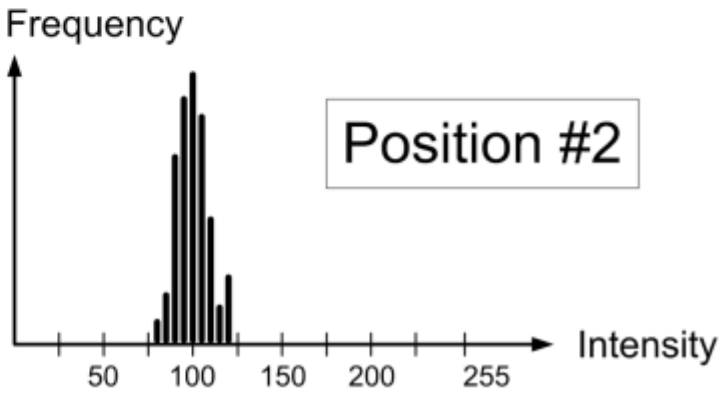

A camera is mounted on the parking space at building 101 The goal is to count the cars coming in and out. What could cause smaller rapid changes that should not be analyzed?

Start the presentation to see live content. For screen share software, share the entire screen. Get help at **pollev.com/app**

Background estimation – slowly changing scene

- Estimating a slowly changing background/reference image

```
ref_image = get_image_from_camera()
stop = False
alpha = 0.95
do
        new_image = get_image_from_camera()
        old_ref = ref_image
        ref_image = alpha * old_ref + (1 – alpha) * new_image
        …(do something more)
while not stop
```

A change detection program

- Estimate and save a background/reference image
- Stop = False
- Do
  - Capture and pre-process one image
  - Compare with reference image (perhaps just subtraction)
  - Threshold difference image
  - Filter noise
  - Decide if something should be done
  - If ‘q’ key pressed:
    - Stop = True
- While not Stop

Get new image and make it ready for processing

4032 x 3024

RGB (3 bytes per pixel)

4032 x 3024

Gray (1 byte per pixel)

640 x 480

Gray (1 byte per pixel)

A change detection program

- Estimate and save a background/reference image
- Stop = False
- Do
  - Capture and pre-process one image
  - Compare with reference image (perhaps just subtraction)
  - Threshold difference image
  - Filter noise
  - Decide if something should be done
  - If ‘q’ key pressed:
    - Stop = True
- While not Stop

Compare with reference image

- Simplest approach:
  - Absolute difference between background and new image
- More advanced approaches based on pixel-wise statistics exists

A change detection program

- Estimate and save a background/reference image
- Stop = False
- Do
  - Capture and pre-process one image
  - Compare with reference image (perhaps just subtraction)
  - Threshold difference image
  - Filter noise
  - Decide if something should be done
  - If ‘q’ key pressed:
    - Stop = True
- While not Stop

Threshold difference image

- Identify the pixel that have significantly changed
  - Set a threshold, T, in the difference image
  - Pixels with a value higher than the threshold is set to 1 the rest to 0

- Complicated to choose the correct threshold

A change detection program

- Estimate and save a background/reference image
- Stop = False
- Do
  - Capture and pre-process one image
  - Compare with reference image (perhaps just subtraction)
  - Threshold difference image
  - Filter noise
  - Decide if something should be done
  - If ‘q’ key pressed:
    - Stop = True
- While not Stop

Remove noise from binary image

- Remove pixels that can be considered noise
  - Isolated pixels
  - Pixels in small groups

- Filtering, morphological operations, BLOB analysis – more about this later in the course

Absolute difference

Thresholded image (Binary)

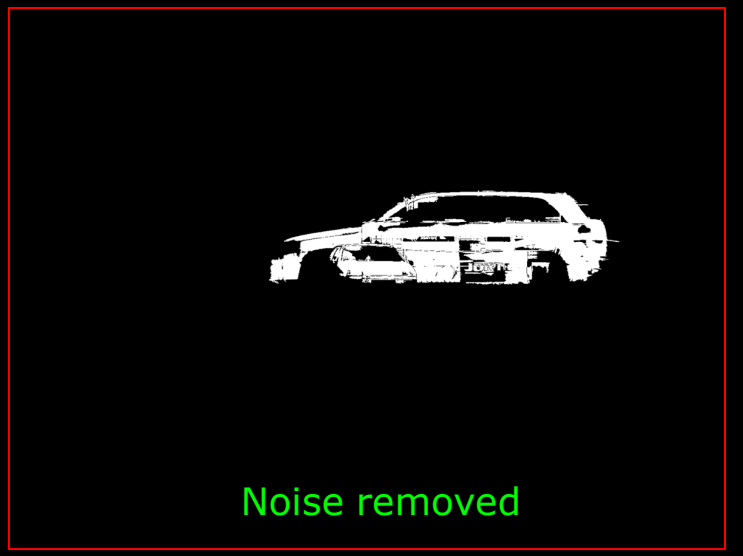
Noise removed

A change detection program

- Estimate and save a background/reference image
- Stop = False
- Do
  - Capture and pre-process one image
  - Compare with reference image (perhaps just subtraction)
  - Threshold difference image
  - Filter noise
  - Decide if something should be done
  - If ‘q’ key pressed:
    - Stop = True
- While not Stop

Decide if something should be done?

- Depends on the application and the scene
- Certain percentage of the total amount of pixel have changed
  - Sound an alarm?

- The changed pixels has the same size and shape as a car
  - Tell that a car is here or start analyzing the car

- The changed pixels look like a face or a person
  - Recognize the face
  - Track the human

- …

Noise removed

Advanced change detection techniques

- Active research and development for 30+ years
- Advanced reference image estimation
  - Pixel wise multi-class estimation
  - Statistical testing per-pixel to detect changes
  - Other color spaces
  - …

Next week
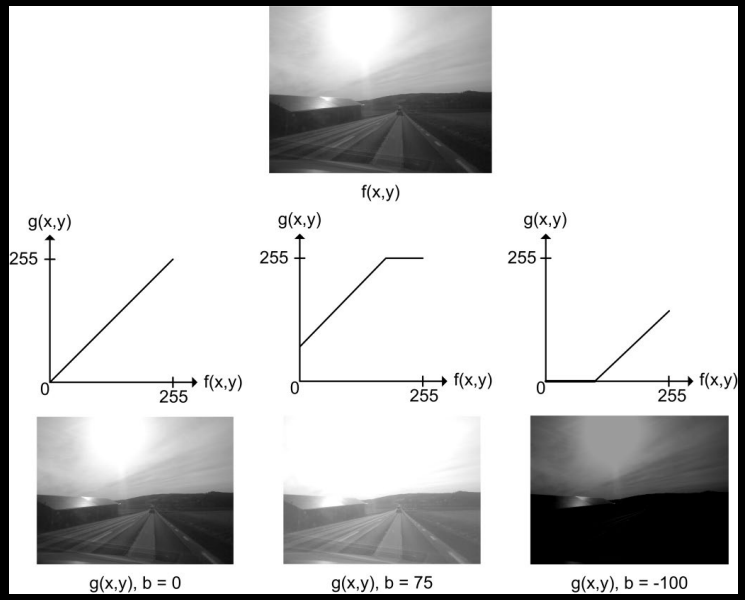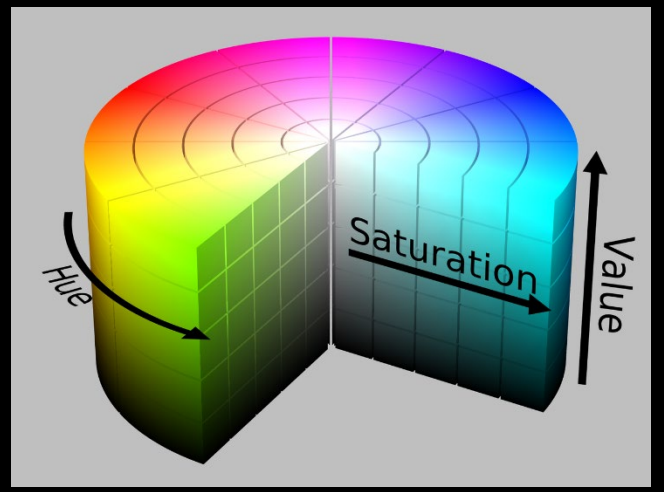
- Pixel wise operations
- Colour images
- PCA on images