Computer Vision

What is an image?

Slides Credit : James Tompkin

What is an image?

- Image
- Sampling
- Digital image representation
- Pixels and pixel manipulation
- Images in practice

What is an image?

```
>>> from numpy import random as r
>>> I = r.rand(256,256)
```

- What is this? What does it look like?
- Which values does it have?
- How many are there?

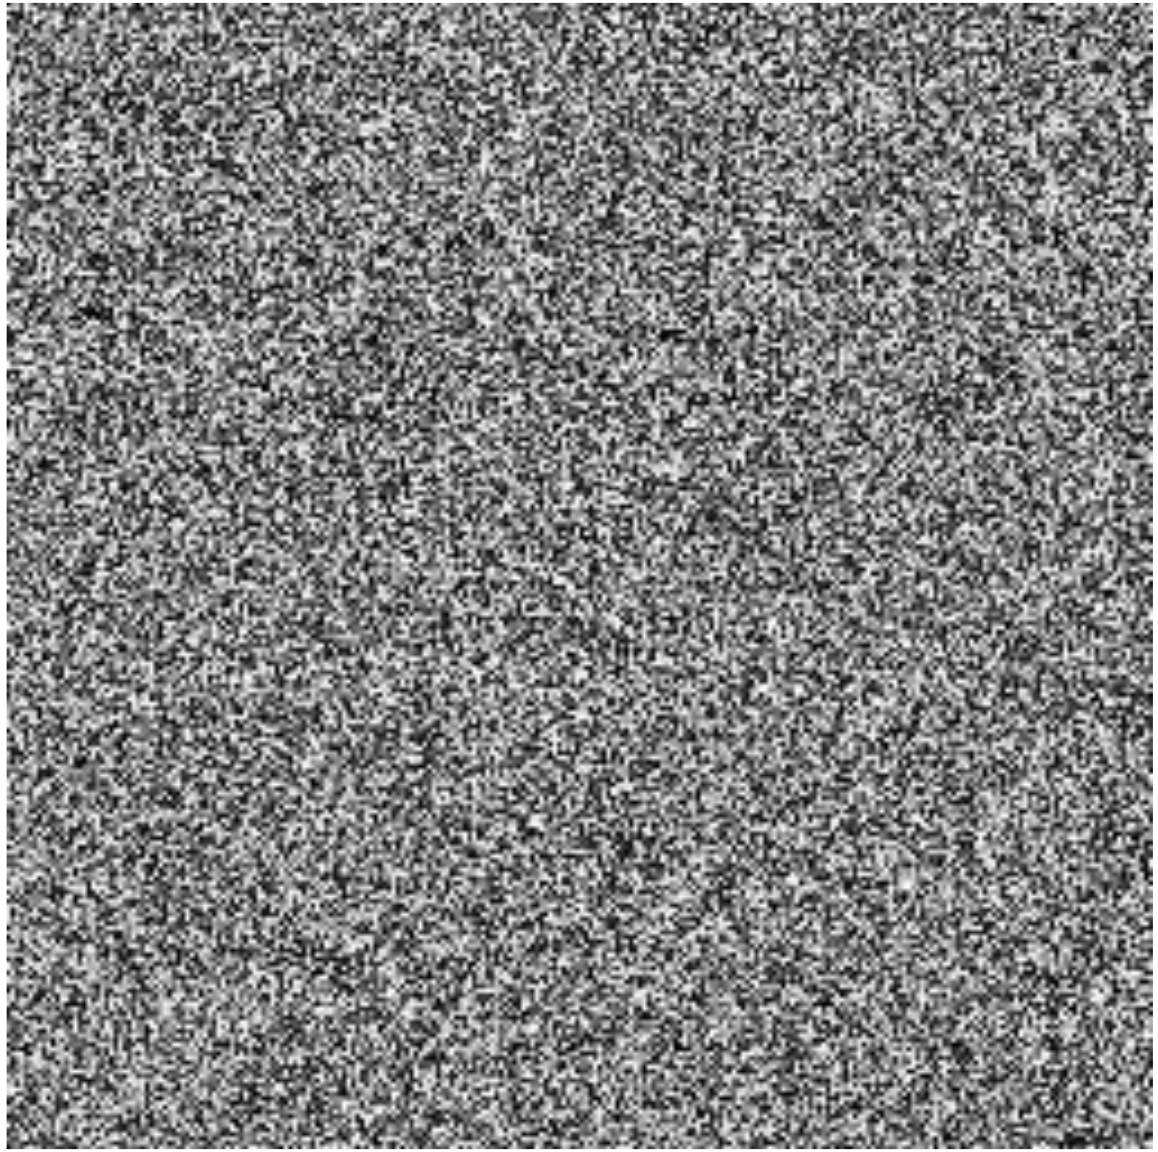

Dimensionality of an image

- 256x256 pixels @ 8bit = 256 value ^65536
  - There is absolutely no way to fit this into any memory

- Computer vision -> making sense of an extremely high-dimensional space.
  - Subspace of 'natural' images.
  - Deriving low-dimensional, explainable model.

Elements of a Digital Image

Pixel: picture element

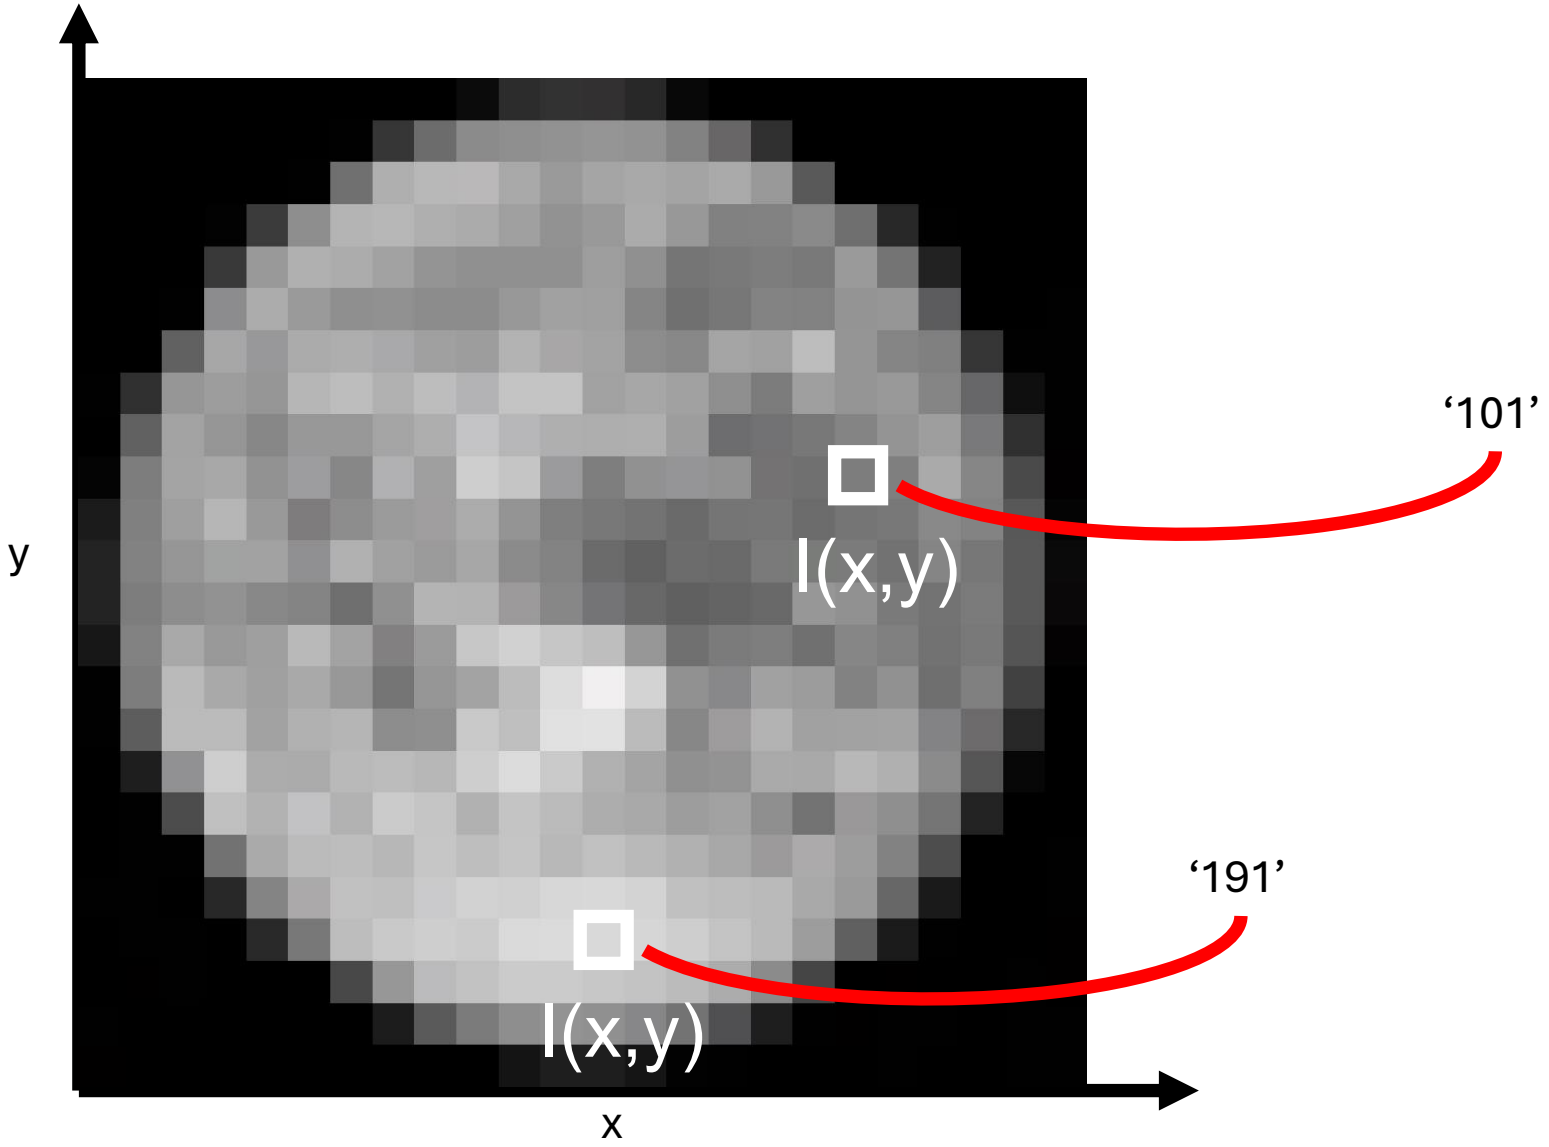

Image as a 2D sampling of signal

- Signal: function depending on some variable with physical meaning.

- Image: sampling of that function.
  - 2 variables: x, y coordinates.
  - 3 variables: x, y + time (video)
  - 'Brightness' is the value of the function for visible light.

- Other physical values too: temperature, pressure, depth,...

Examples of 2D images

(a) Natural Image

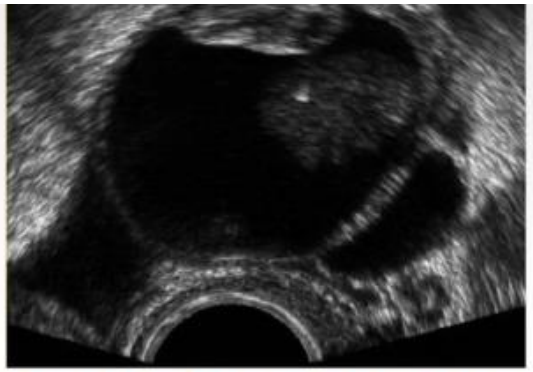

(c) Ultra Sound Image

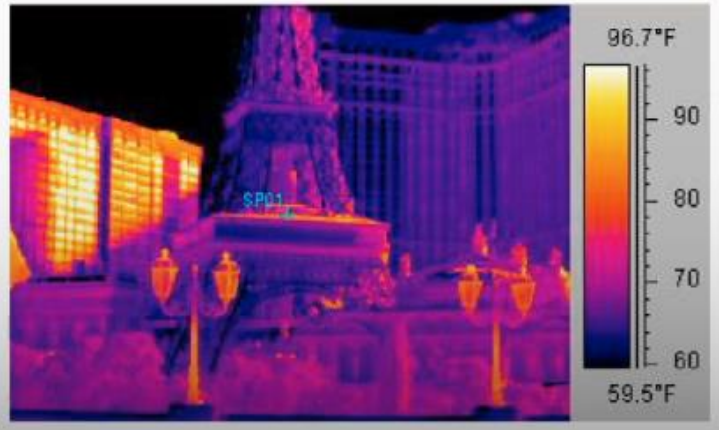

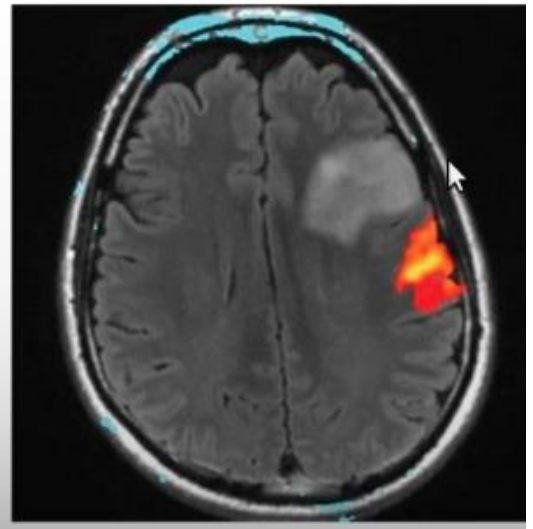

Sampling

Sampling in 1D takes a function and returns a vector whose elements are values of that function at the sample points.

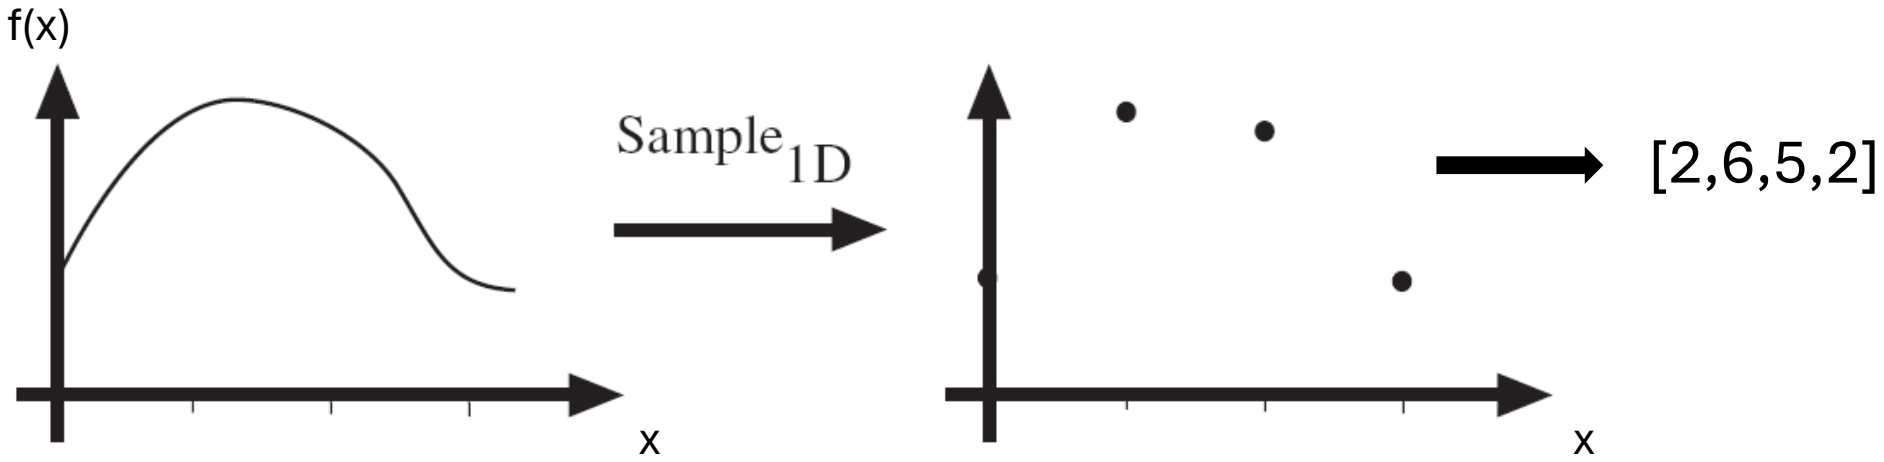

Sampling

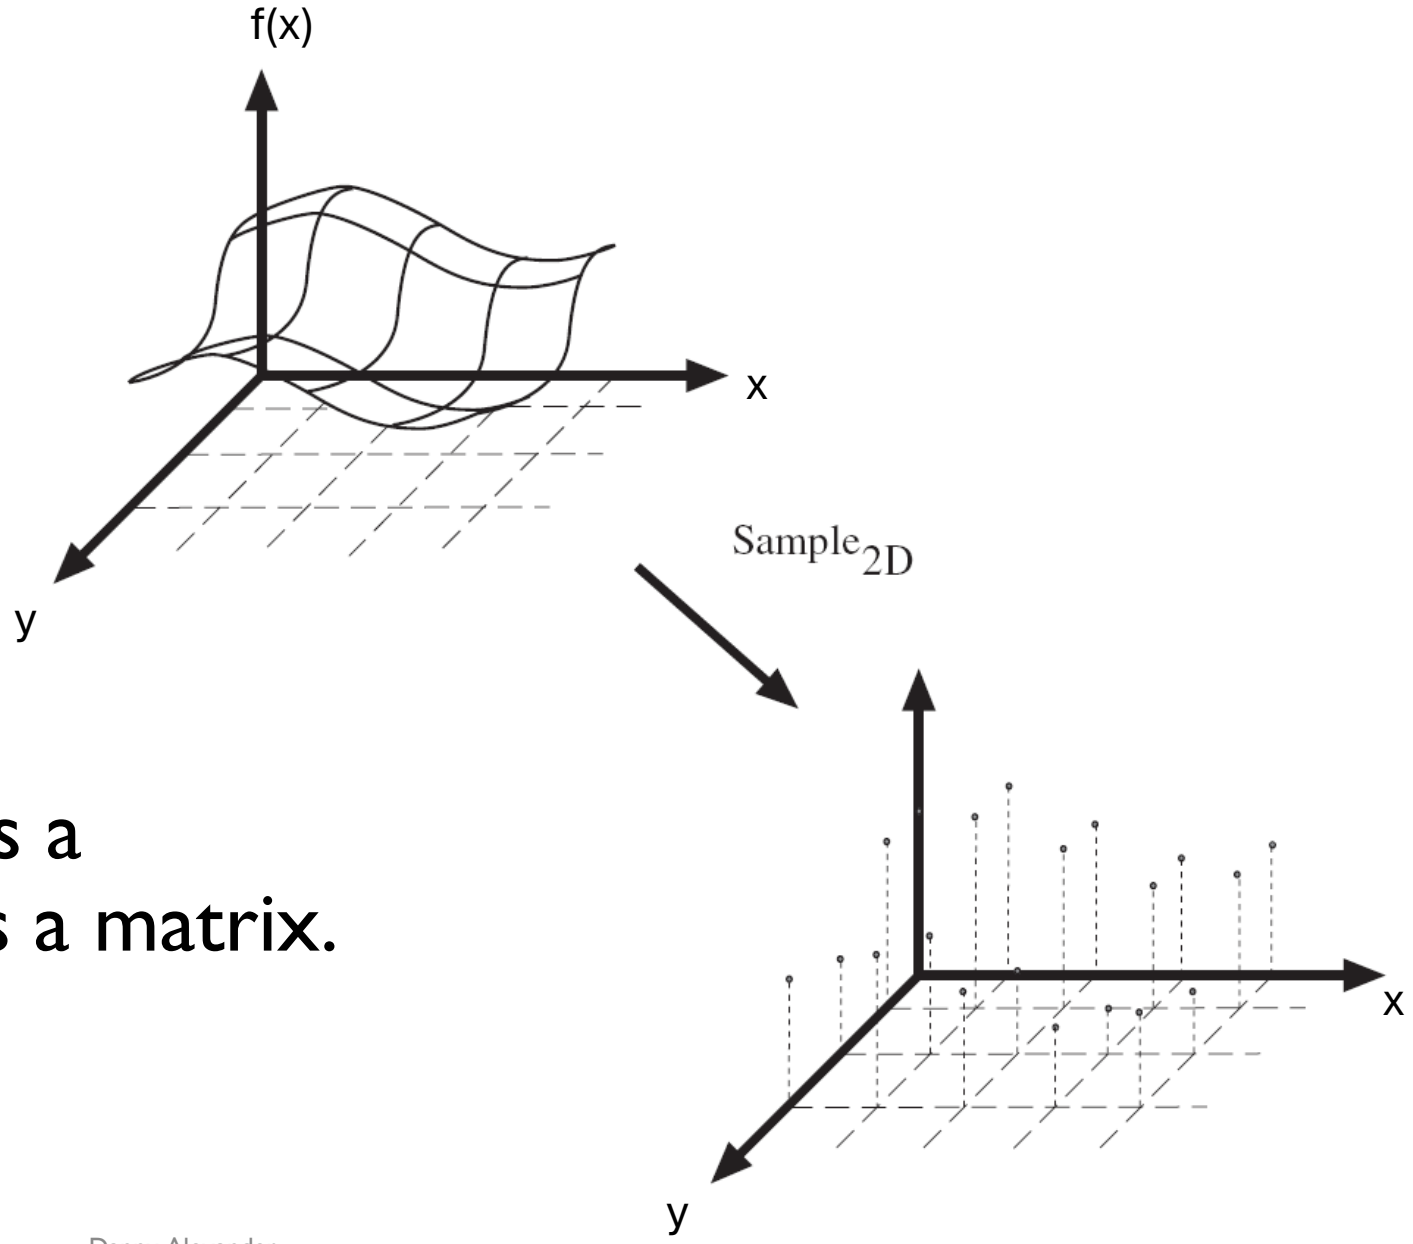

Sampling in 2D takes a function and returns a matrix.

Digital Image as ‘Height Map’

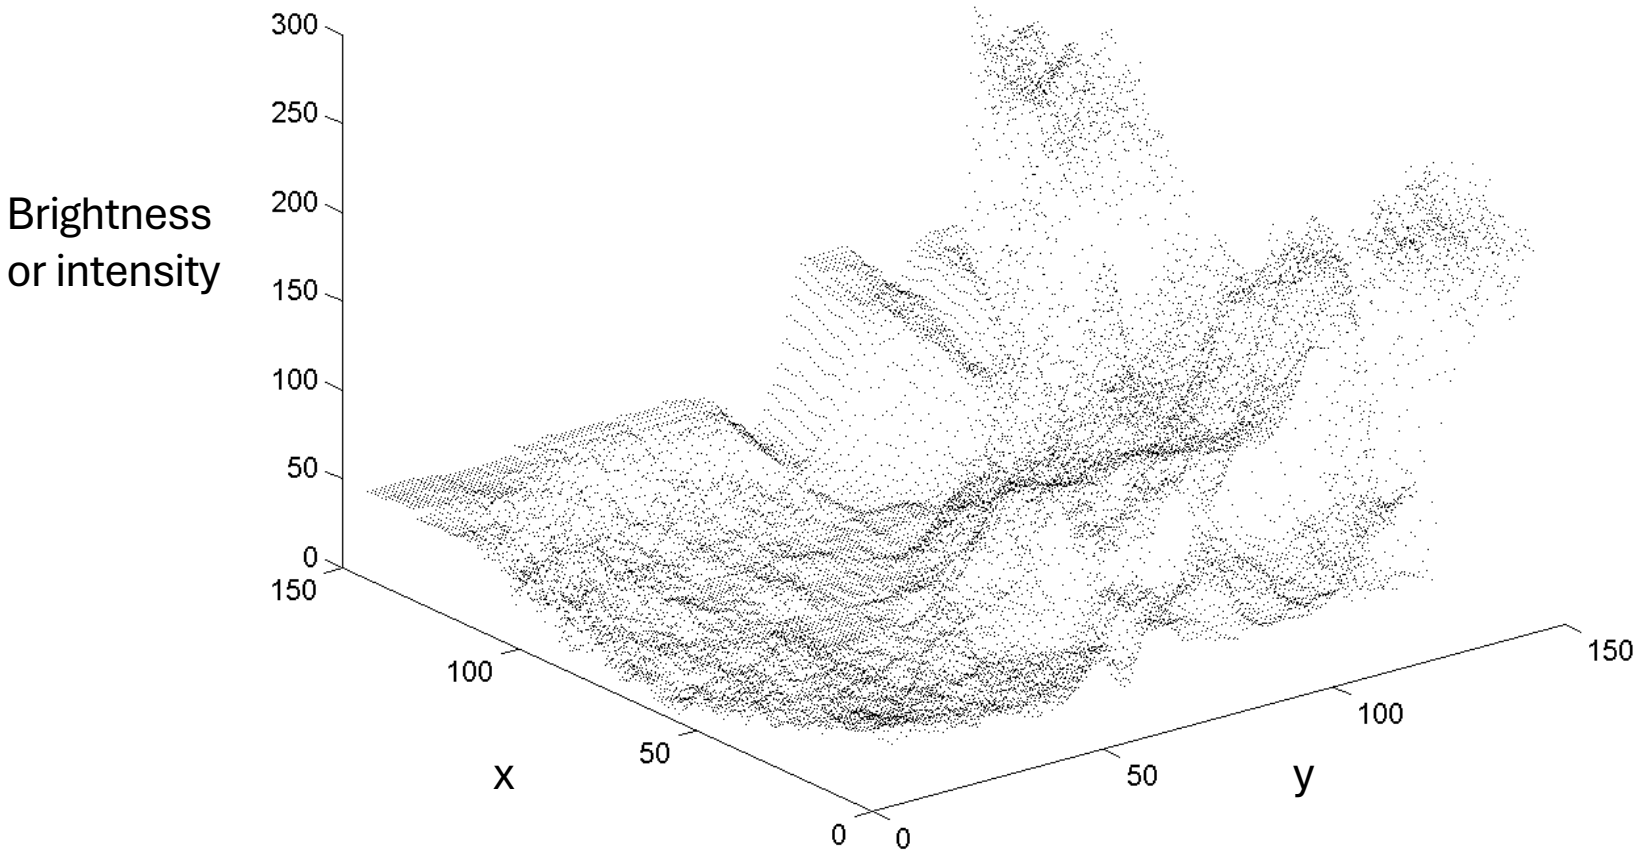

What is each part of a ptotograph?

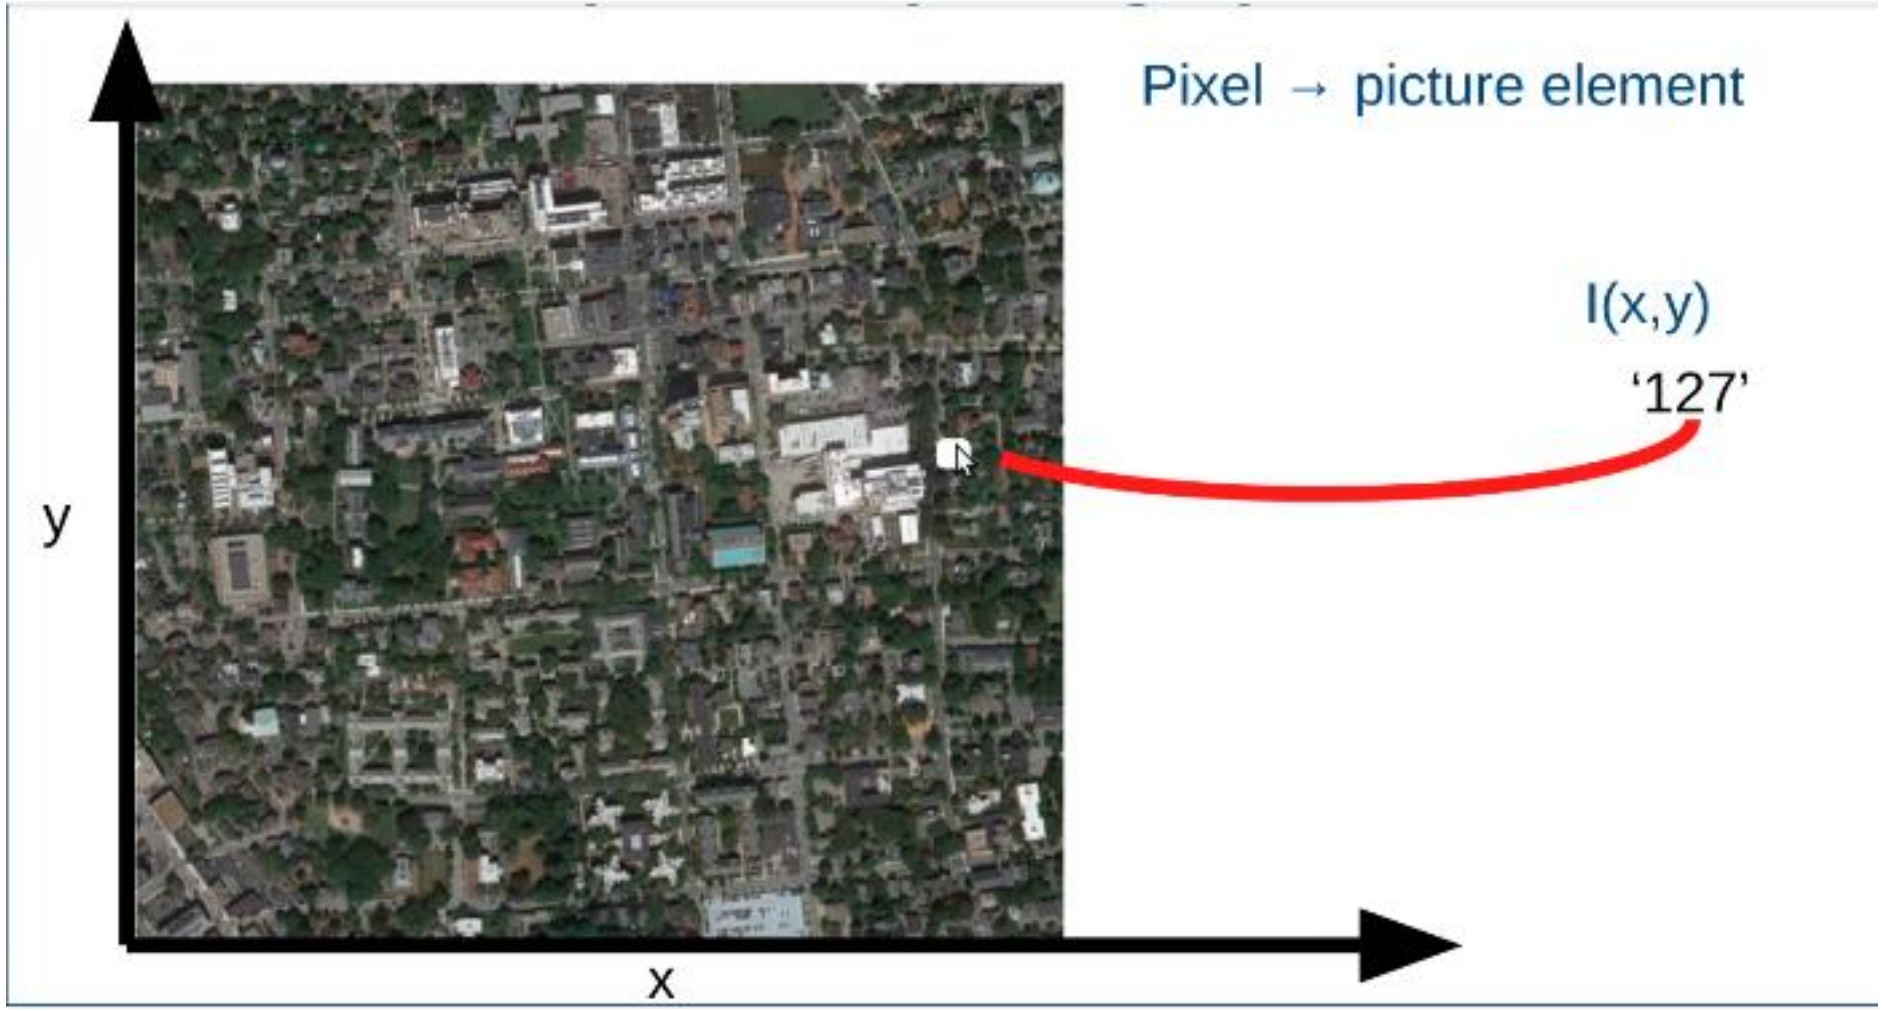

Light Integration Over the *Frustum*

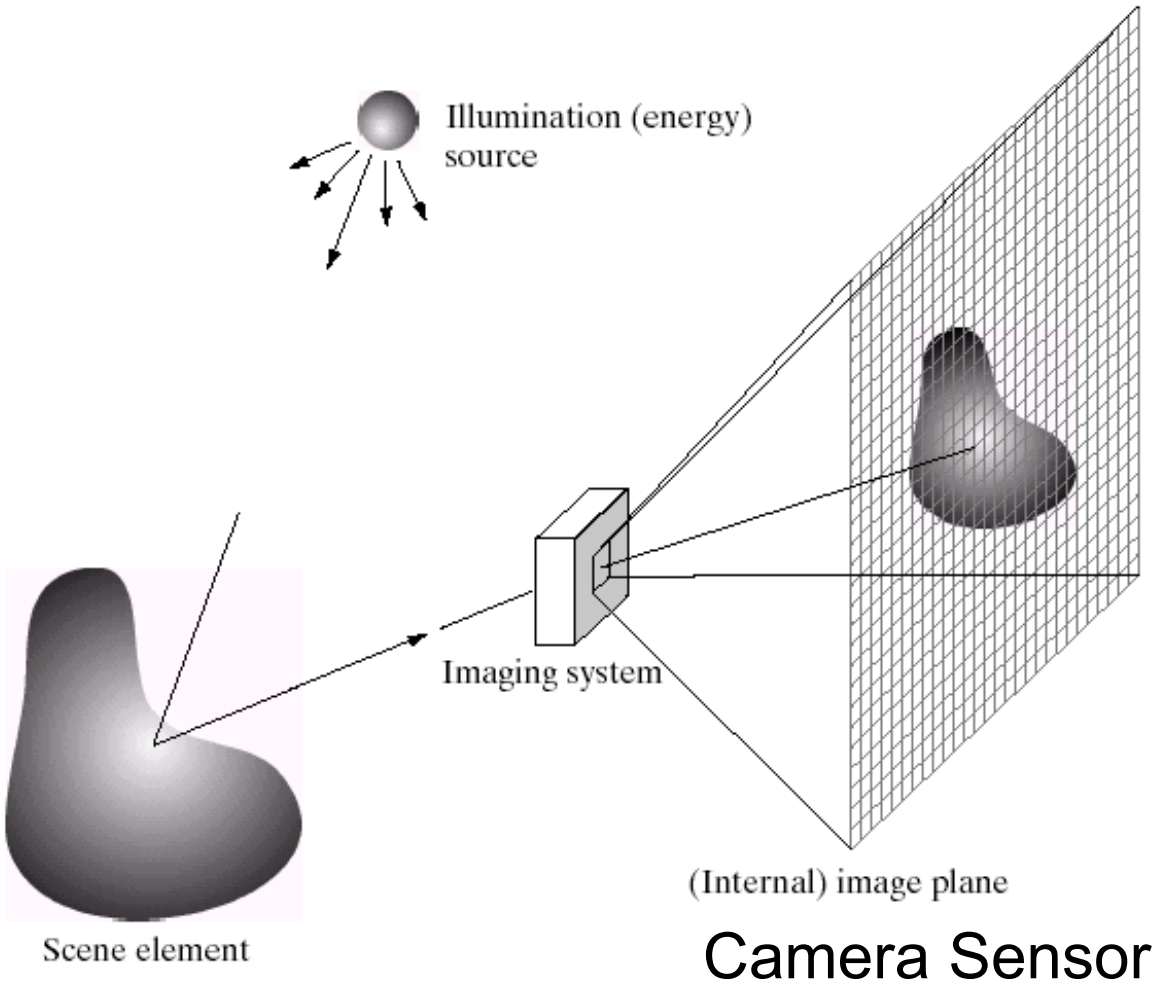

Camera Sensor

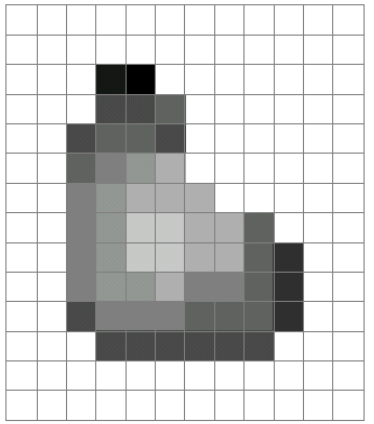

Output Image

Quantization

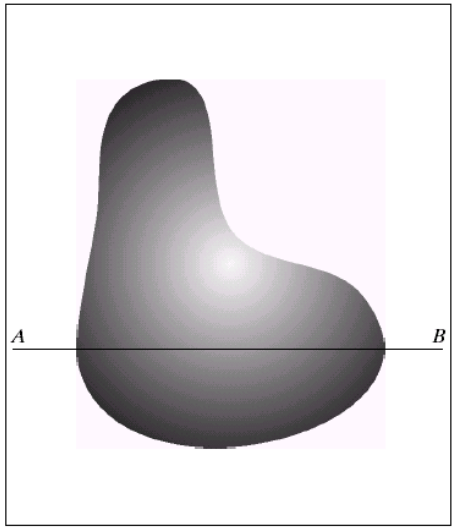

Underlying analog charge signal

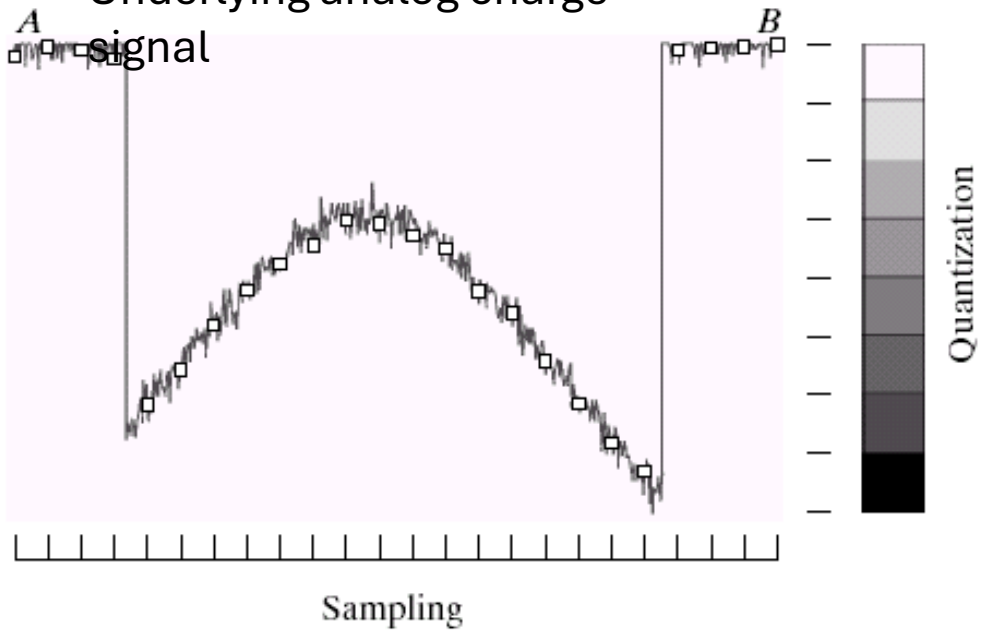

Quantized values

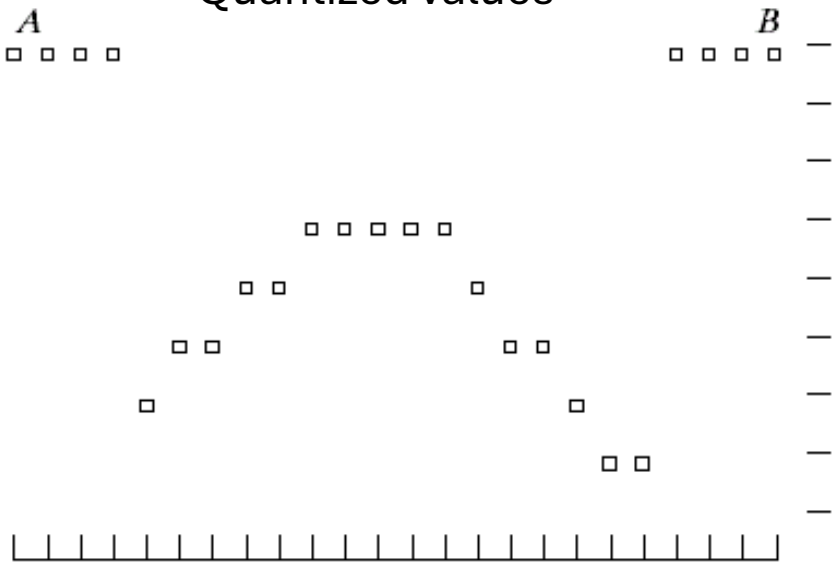

Quantization Effects – Radiometric Resolution

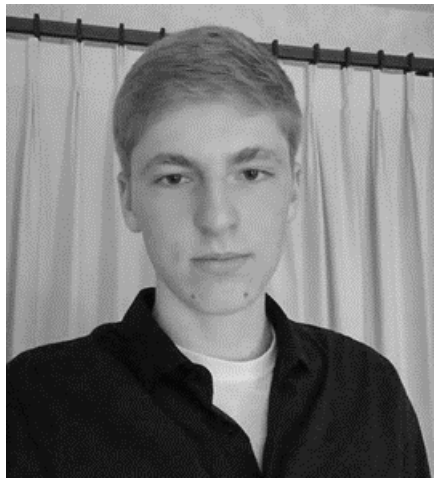

8 bit – 256 levels

4 bit – 16 levels

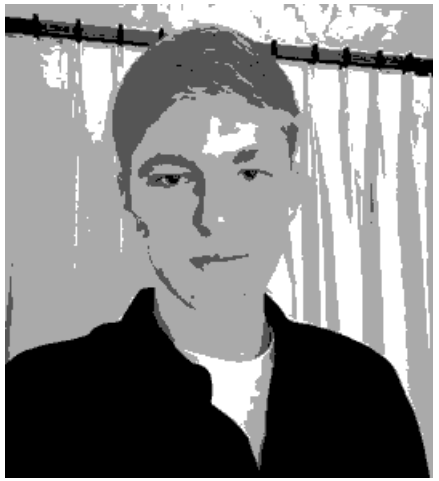

2 bit – 4 levels

1 bit – 2 levels

We often call this *bit depth*.
For photography, this is also related to *dynamic range*.

Resolution: geometric vs. spatial

Both images are 1000x1000 pixels

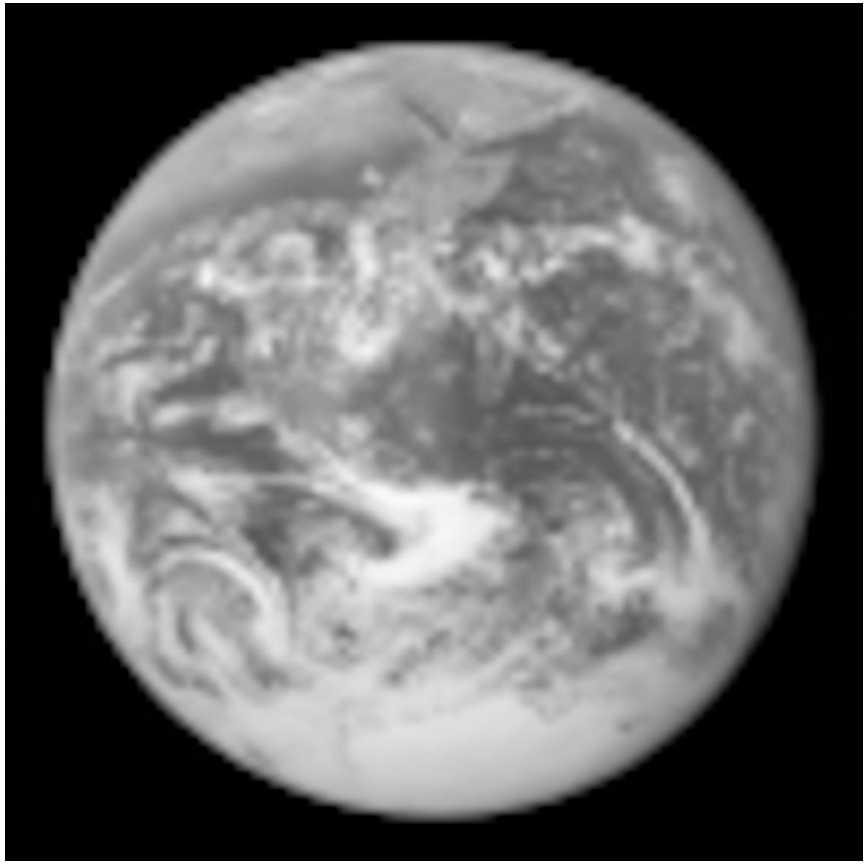

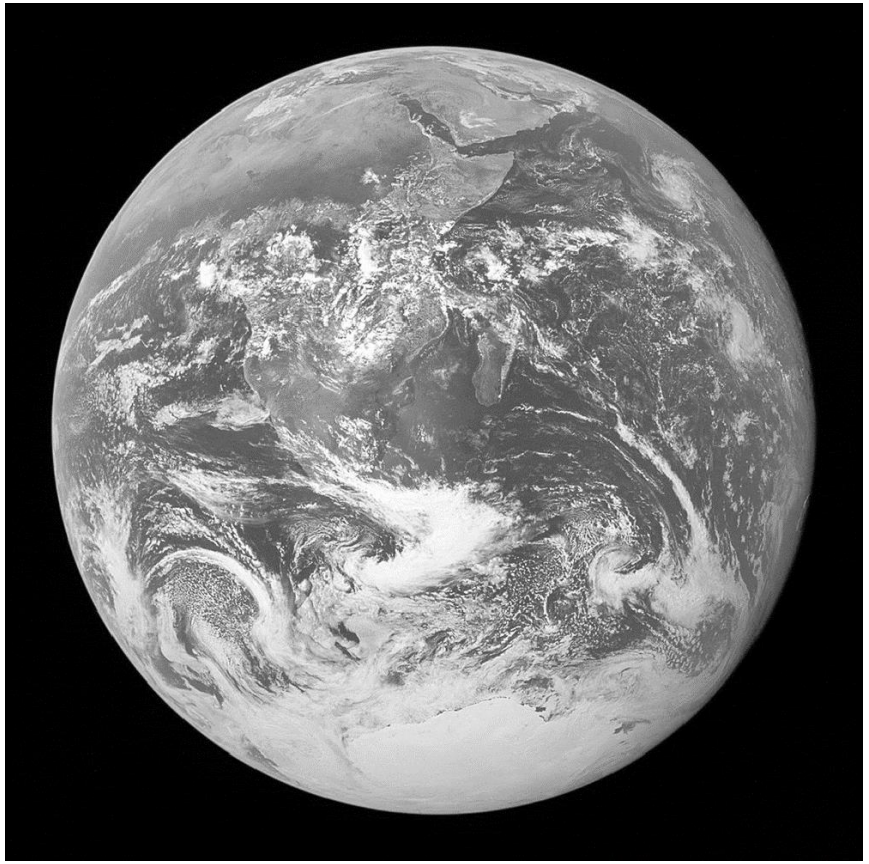

Images in Python (import numpy)

`N x M` grayscale image "`im`"

- `im[0,0]` = top-left pixel value
- `im[y,x]` = y pixels down, x pixels to right
- `im[N-1,M-1]` = bottom-right pixel

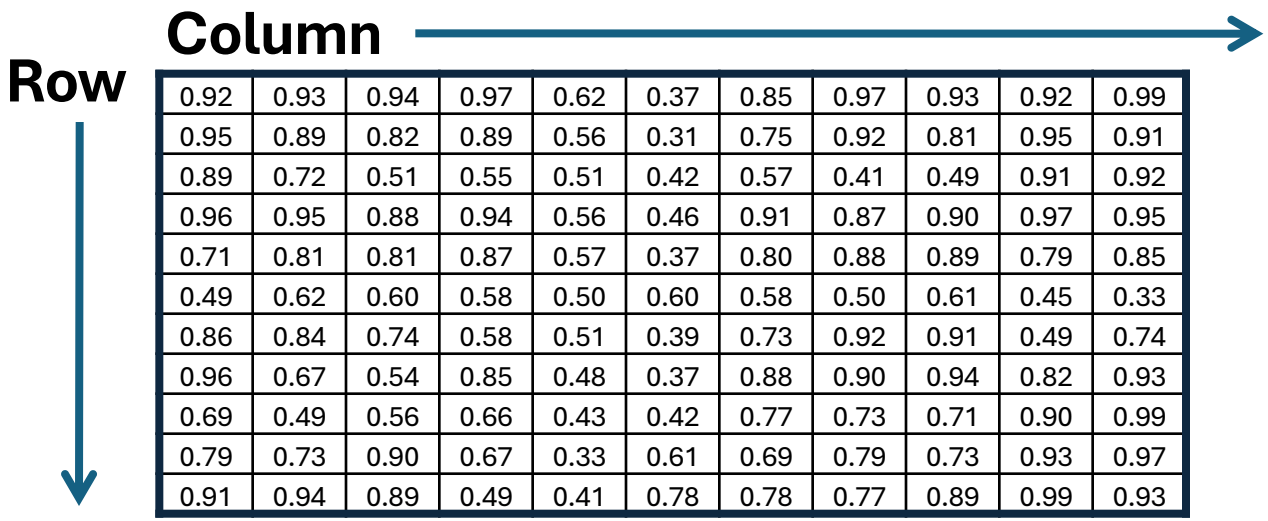

0.92	0.93	0.94	0.97	0.62	0.37	0.85	0.97	0.93	0.92	0.99
0.95	0.89	0.82	0.89	0.56	0.31	0.75	0.92	0.81	0.95	0.91
0.89	0.72	0.51	0.55	0.51	0.42	0.57	0.41	0.49	0.91	0.92
0.96	0.95	0.88	0.94	0.56	0.46	0.91	0.87	0.90	0.97	0.95
0.71	0.81	0.81	0.87	0.57	0.37	0.80	0.88	0.89	0.79	0.85
0.49	0.62	0.60	0.58	0.50	0.60	0.58	0.50	0.61	0.45	0.33
0.86	0.84	0.74	0.58	0.51	0.39	0.73	0.92	0.91	0.49	0.74
0.96	0.67	0.54	0.85	0.48	0.37	0.88	0.90	0.94	0.82	0.93
0.69	0.49	0.56	0.66	0.43	0.42	0.77	0.73	0.71	0.90	0.99
0.79	0.73	0.90	0.67	0.33	0.61	0.69	0.79	0.73	0.93	0.97
0.91	0.94	0.89	0.49	0.41	0.78	0.78	0.77	0.89	0.99	0.93

Grayscale Intensity

Color

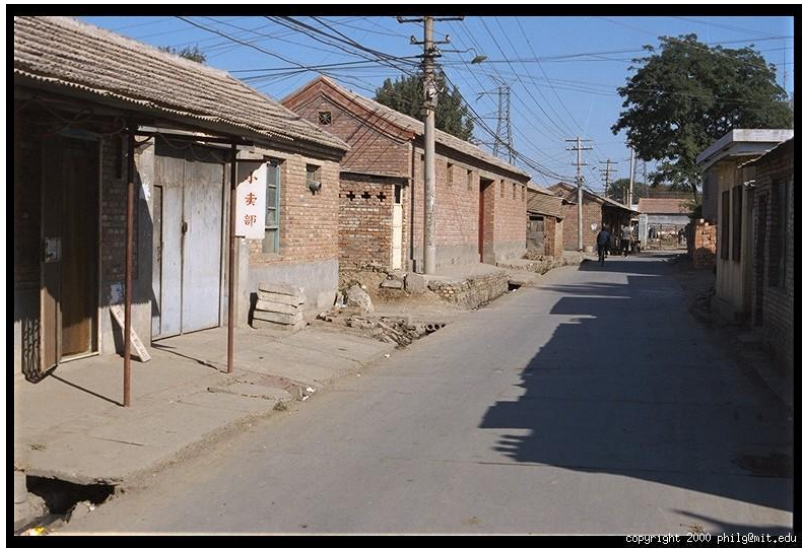

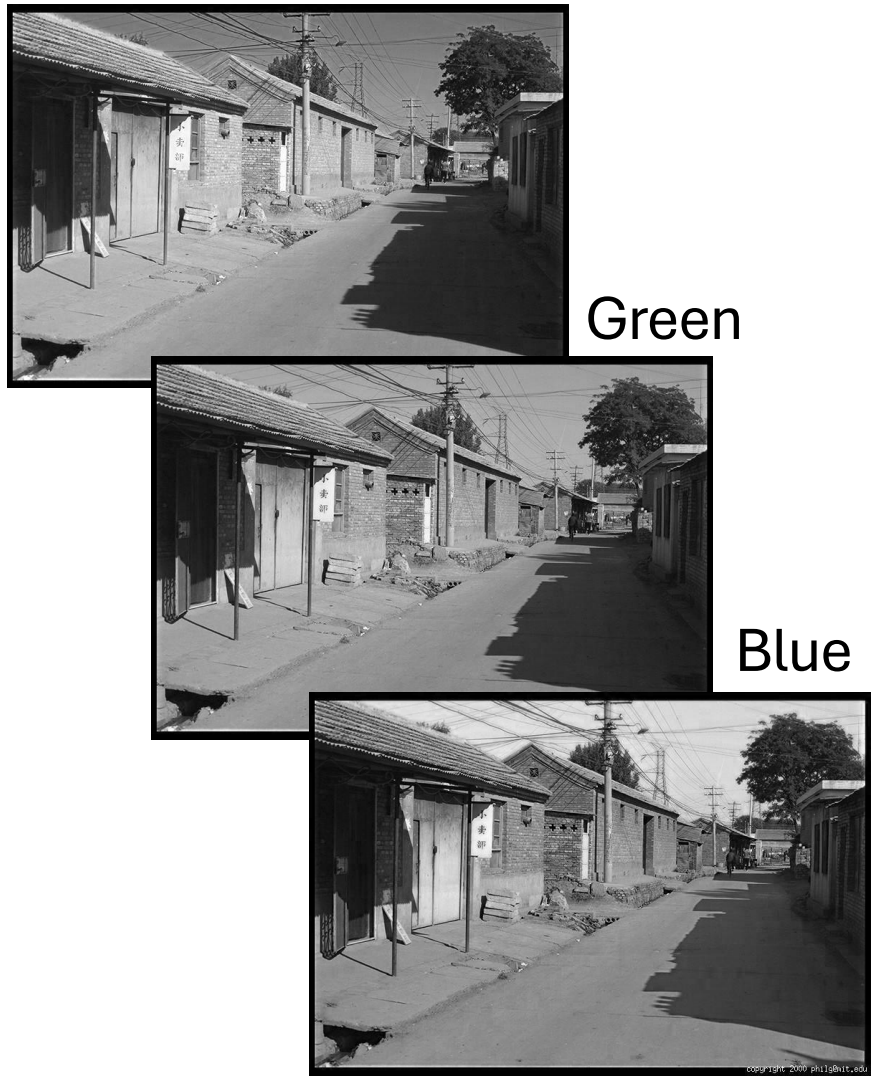

Color

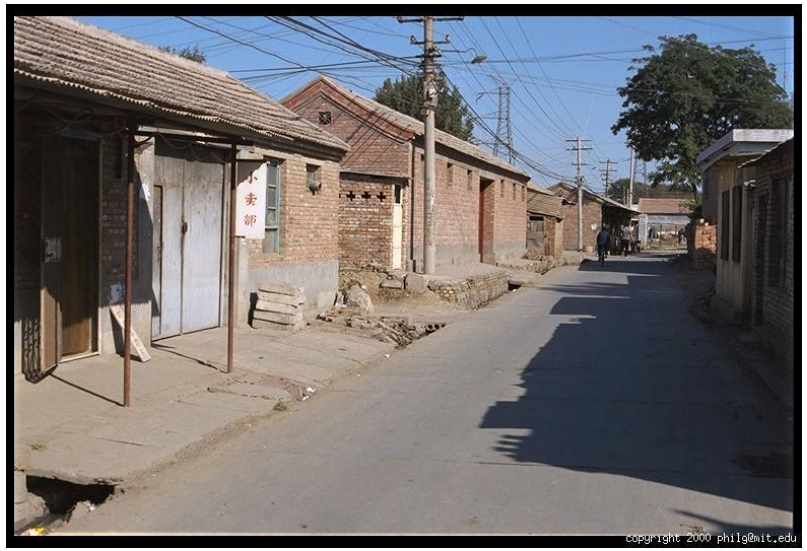

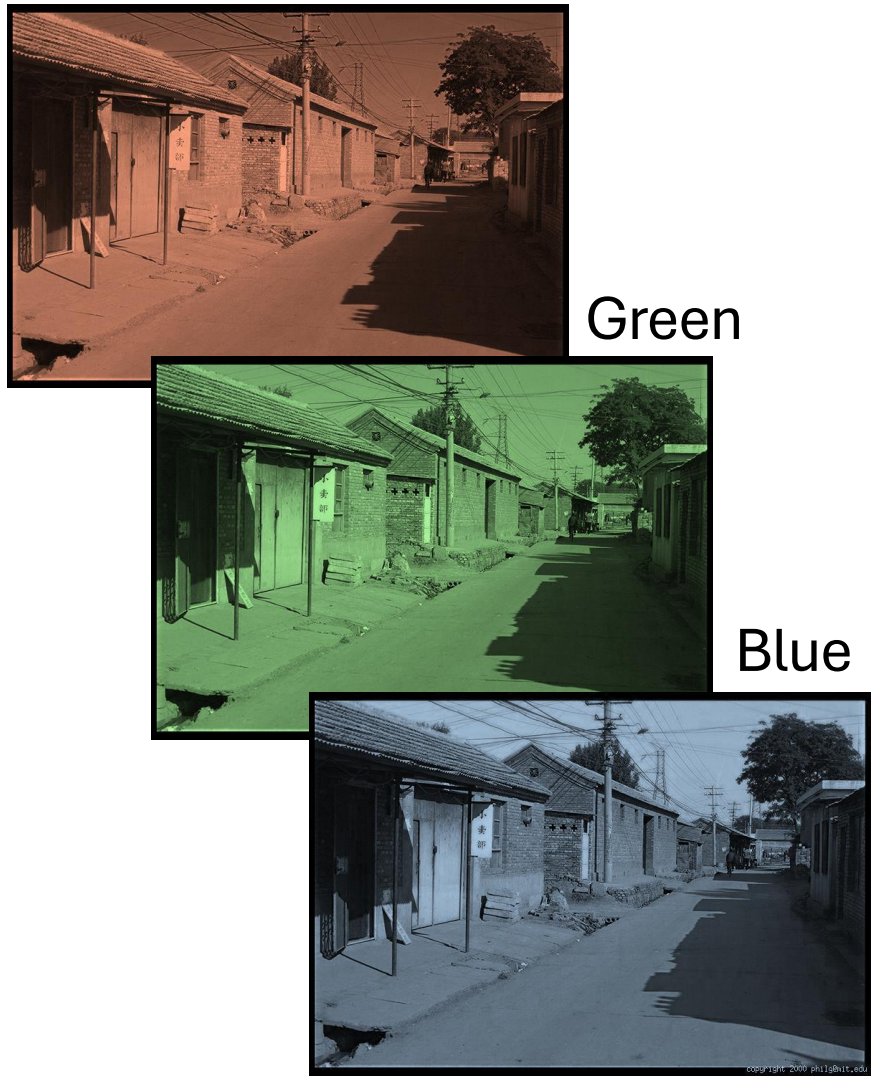

Images in Python (import numpy)

`N x M` grayscale image "`im`"

- `im[0,0,0]` = top-left pixel value, red channel
- `im[y,x,1]` = y pixels down, x pixels to right, green channel
- `im[N-1,M-1,2]` = bottom-right pixel, blue channel

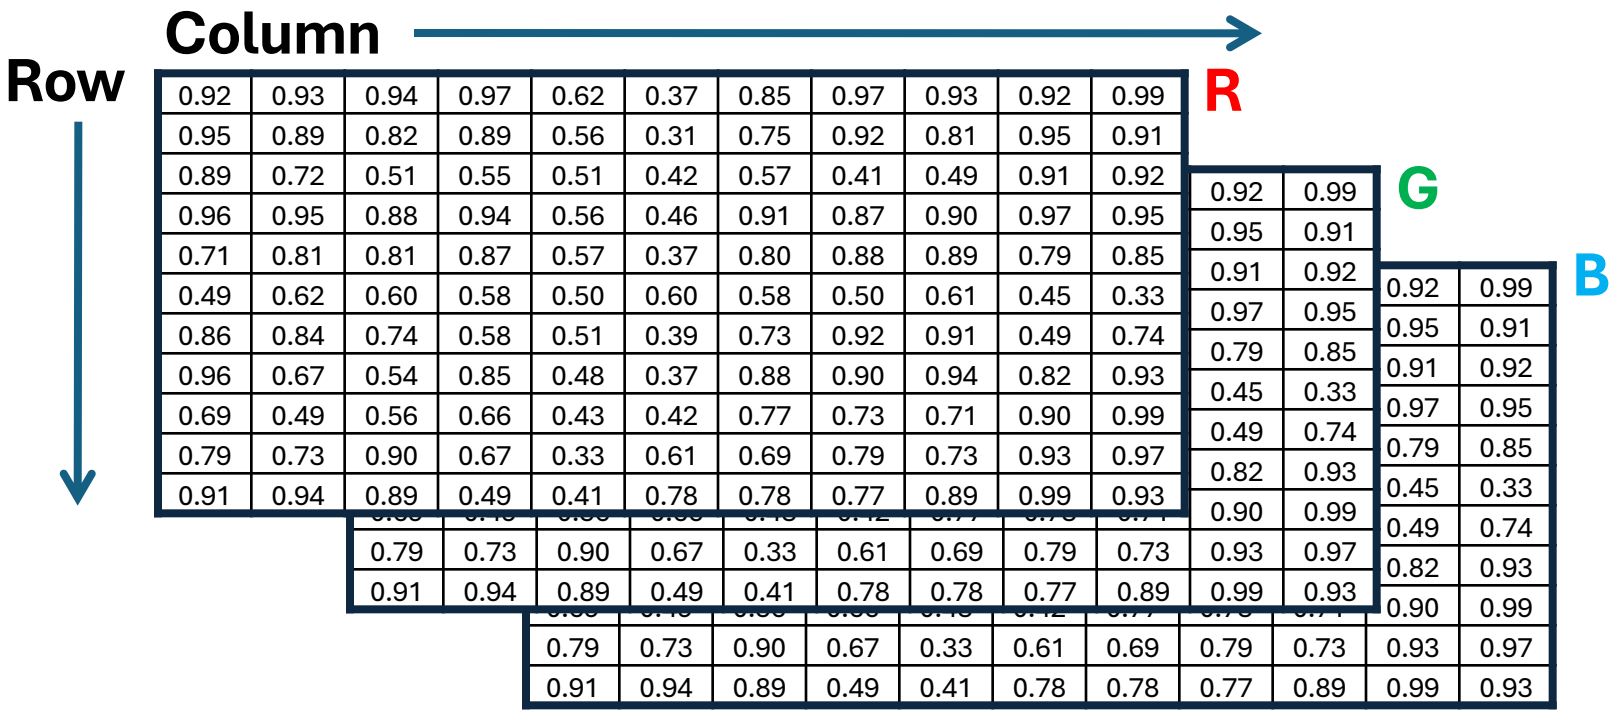

Images in Python (import `numpy`, `scikit`)

Take care of types!

- `uint8` (values 0 to 255) – `io.imread("file.jpg")`
- `float32` (values 0 to 255) – `io.imread("file.jpg").astype(np.float32)`
- `float32` (values 0 to 1) – `img_as_float32(io.imread("file.jpg"))`

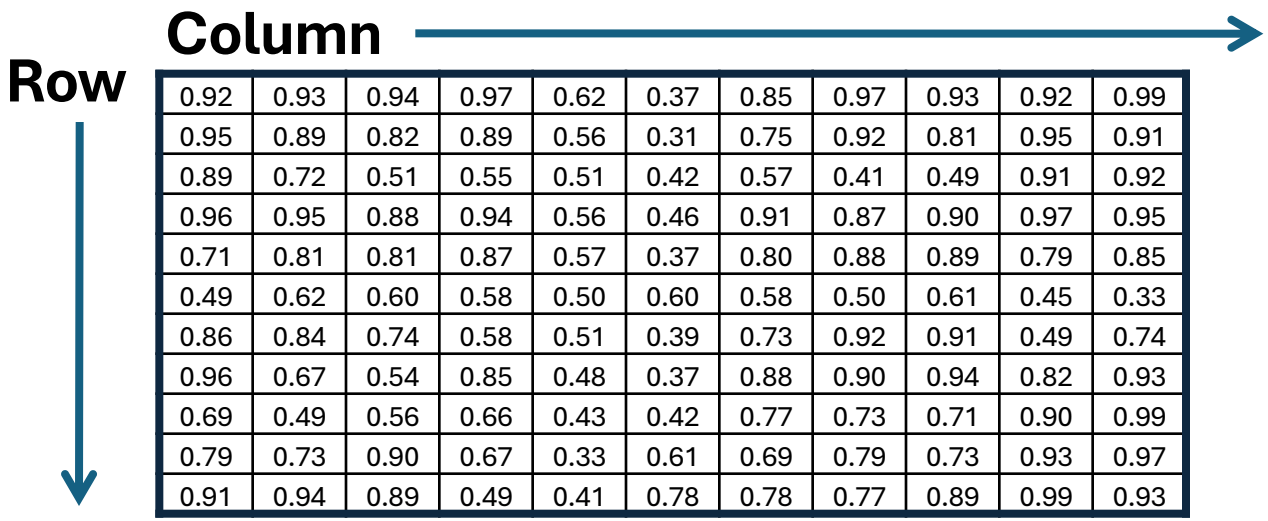

0.92	0.93	0.94	0.97	0.62	0.37	0.85	0.97	0.93	0.92	0.99
0.95	0.89	0.82	0.89	0.56	0.31	0.75	0.92	0.81	0.95	0.91
0.89	0.72	0.51	0.55	0.51	0.42	0.57	0.41	0.49	0.91	0.92
0.96	0.95	0.88	0.94	0.56	0.46	0.91	0.87	0.90	0.97	0.95
0.71	0.81	0.81	0.87	0.57	0.37	0.80	0.88	0.89	0.79	0.85
0.49	0.62	0.60	0.58	0.50	0.60	0.58	0.50	0.61	0.45	0.33
0.86	0.84	0.74	0.58	0.51	0.39	0.73	0.92	0.91	0.49	0.74
0.96	0.67	0.54	0.85	0.48	0.37	0.88	0.90	0.94	0.82	0.93
0.69	0.49	0.56	0.66	0.43	0.42	0.77	0.73	0.71	0.90	0.99
0.79	0.73	0.90	0.67	0.33	0.61	0.69	0.79	0.73	0.93	0.97
0.91	0.94	0.89	0.49	0.41	0.78	0.78	0.77	0.89	0.99	0.93

Thanks

- Next lecture: image Filtering