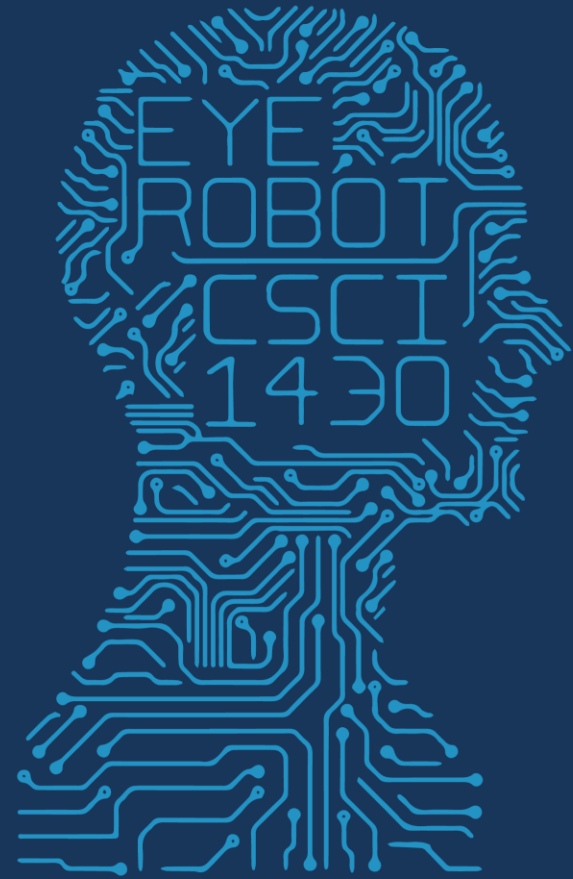




1950

FUTURE VISION



1 FEBRUARY 2019

COMPUTER VISION

WHAT IS AN IMAGE?

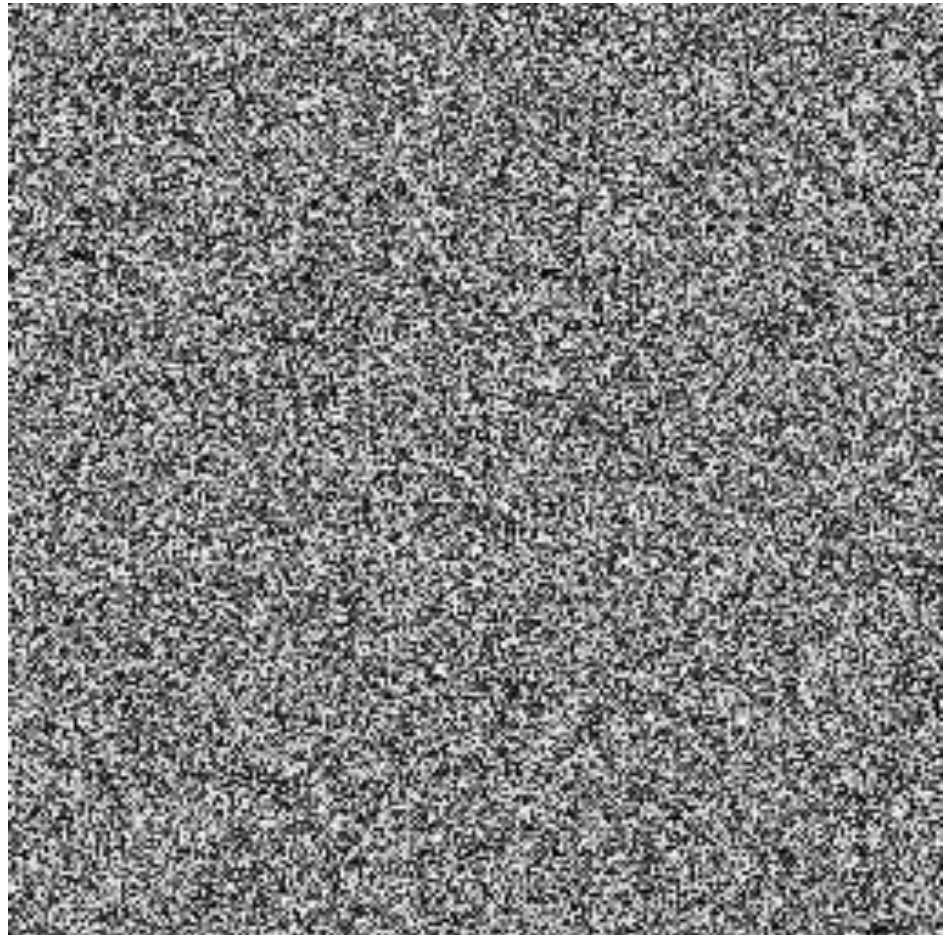
```
>>> from numpy import random as r
>>> I = r.rand(256,256);
```

Think-Pair-Share:

- What is this? What does it look like?
- Which values does it take?
- How many values can it take?

- Is it an image?

```
>>> from matplotlib import pyplot as p  
>>> l = r.rand(256,256);  
>>> p.imshow(l);  
>>> p.show();
```



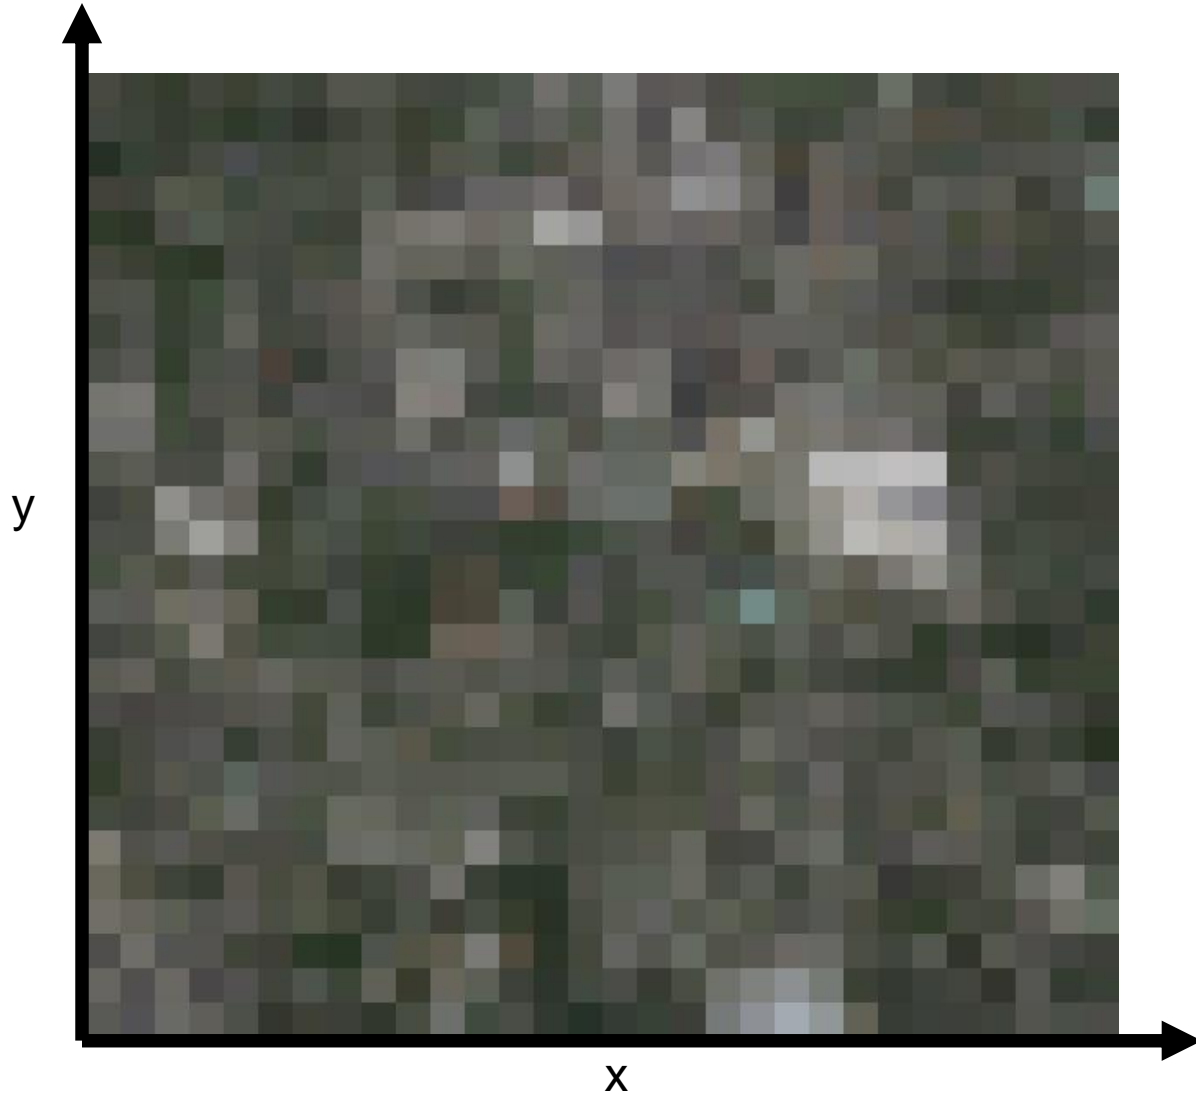
Dimensionality of an Image

- @ 8bit = 256 values $^$ 65,536
 - Computer says 'Inf' combinations.
- Some depiction of all possible scenes would fit into this memory.

Dimensionality of an Image

- @ 8bit = 256 values \wedge 65,536
 - Computer says ‘Inf’ combinations.
- Some depiction of all possible scenes would fit into this memory.
- Computer vision as making sense of an extremely high-dimensional space.
 - Subspace of ‘natural’ images.
 - Deriving low-dimensional, explainable models.

What is each part of an image?



What is each part of an image?

- Pixel -> picture element

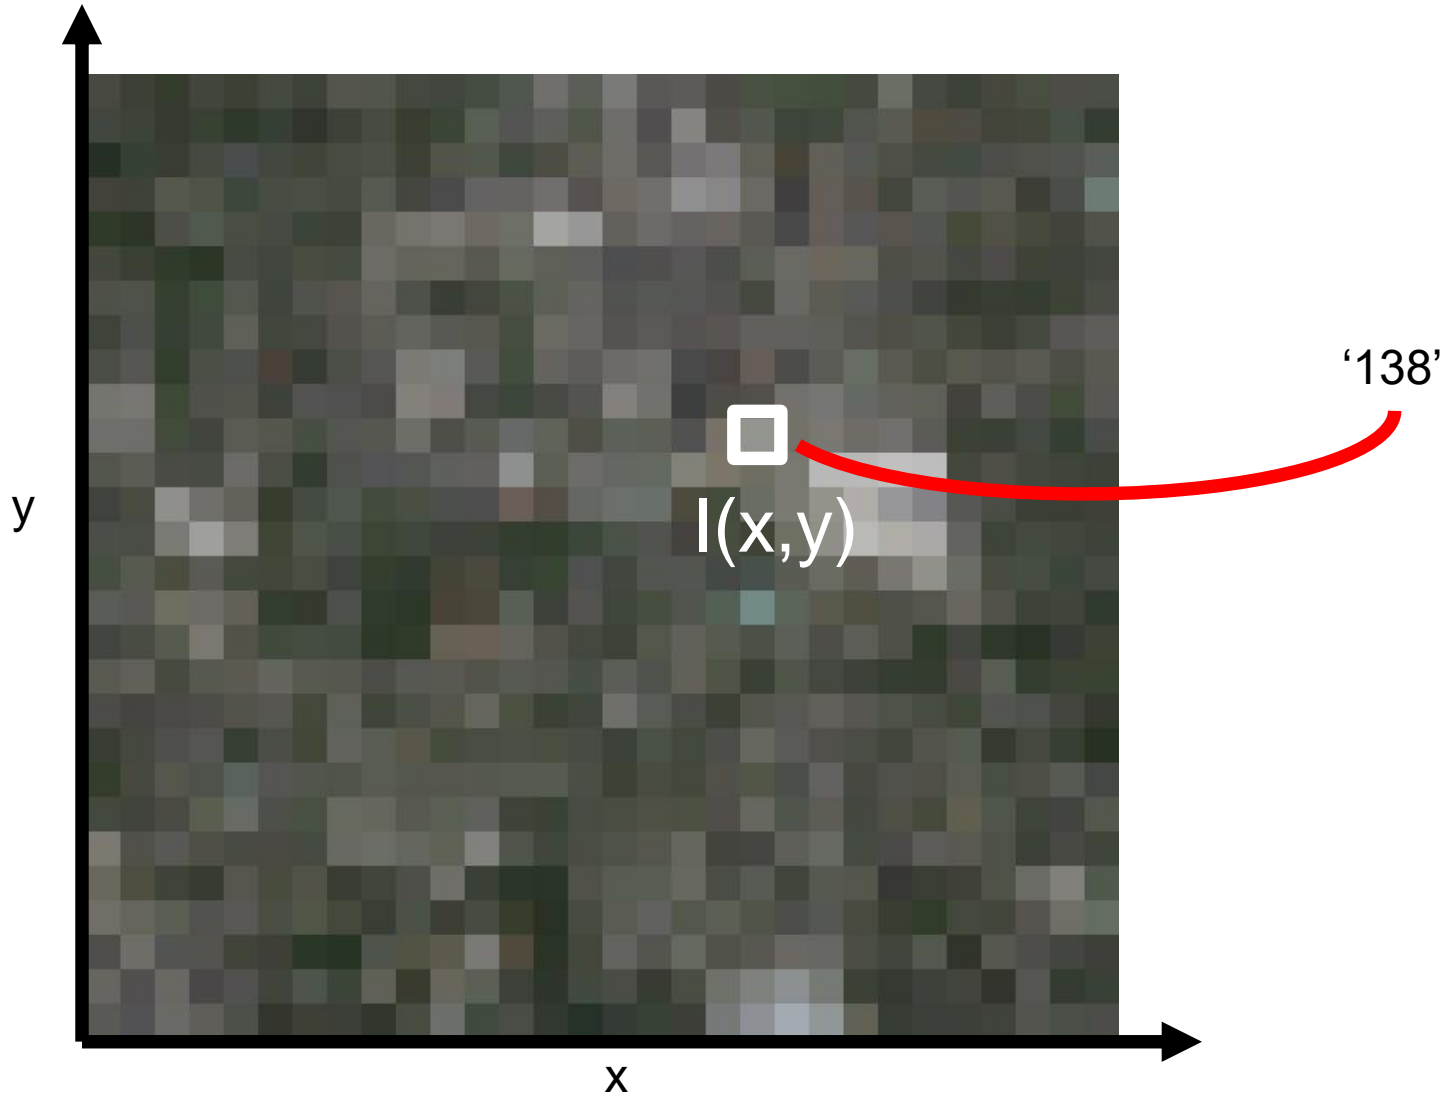
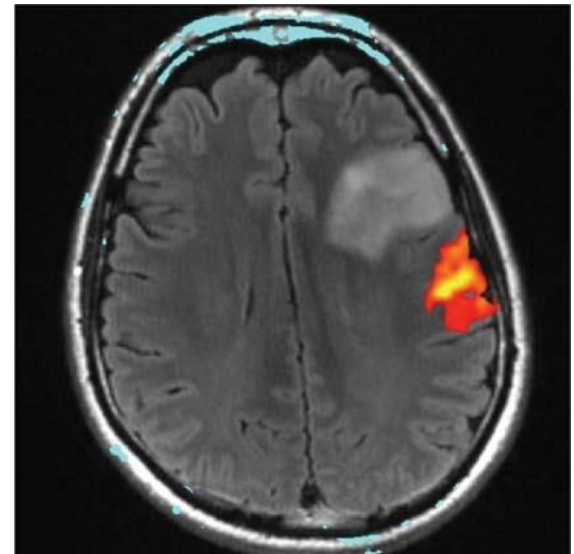
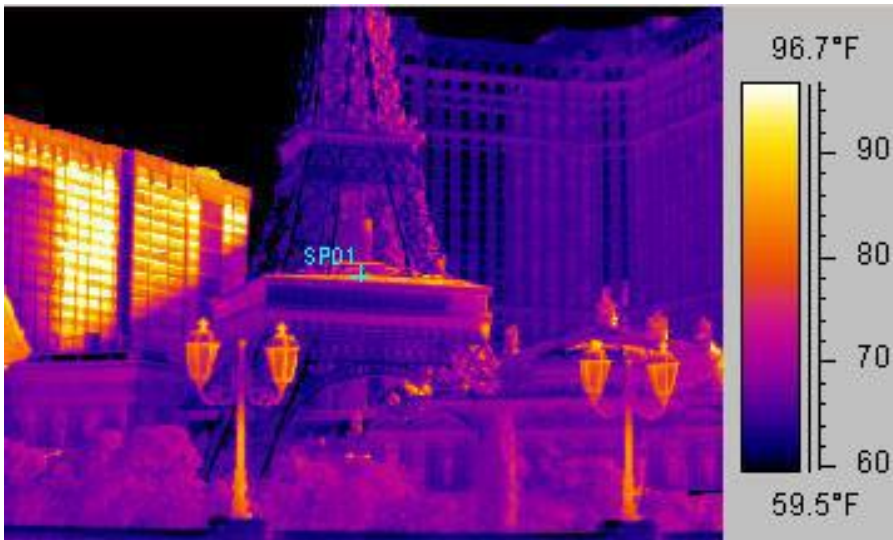


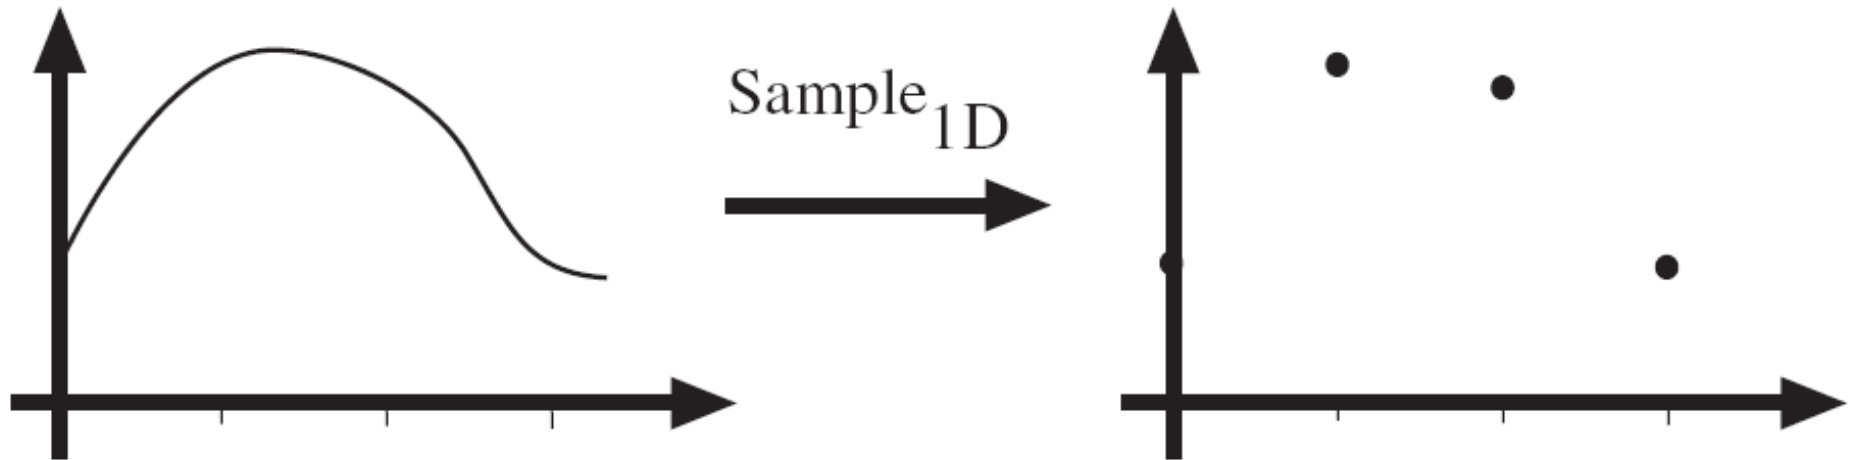
Image as a 2D sampling of signal

- Signal: function depending on some variable with physical meaning.
- Image: sampling of that function.
 - 2 variables: xy coordinates
 - 3 variables: xy + time (video)
 - ‘Brightness’ is the value of the function for visible light
- Can be other physical values too: temperature, pressure, depth ...

Example 2D Images

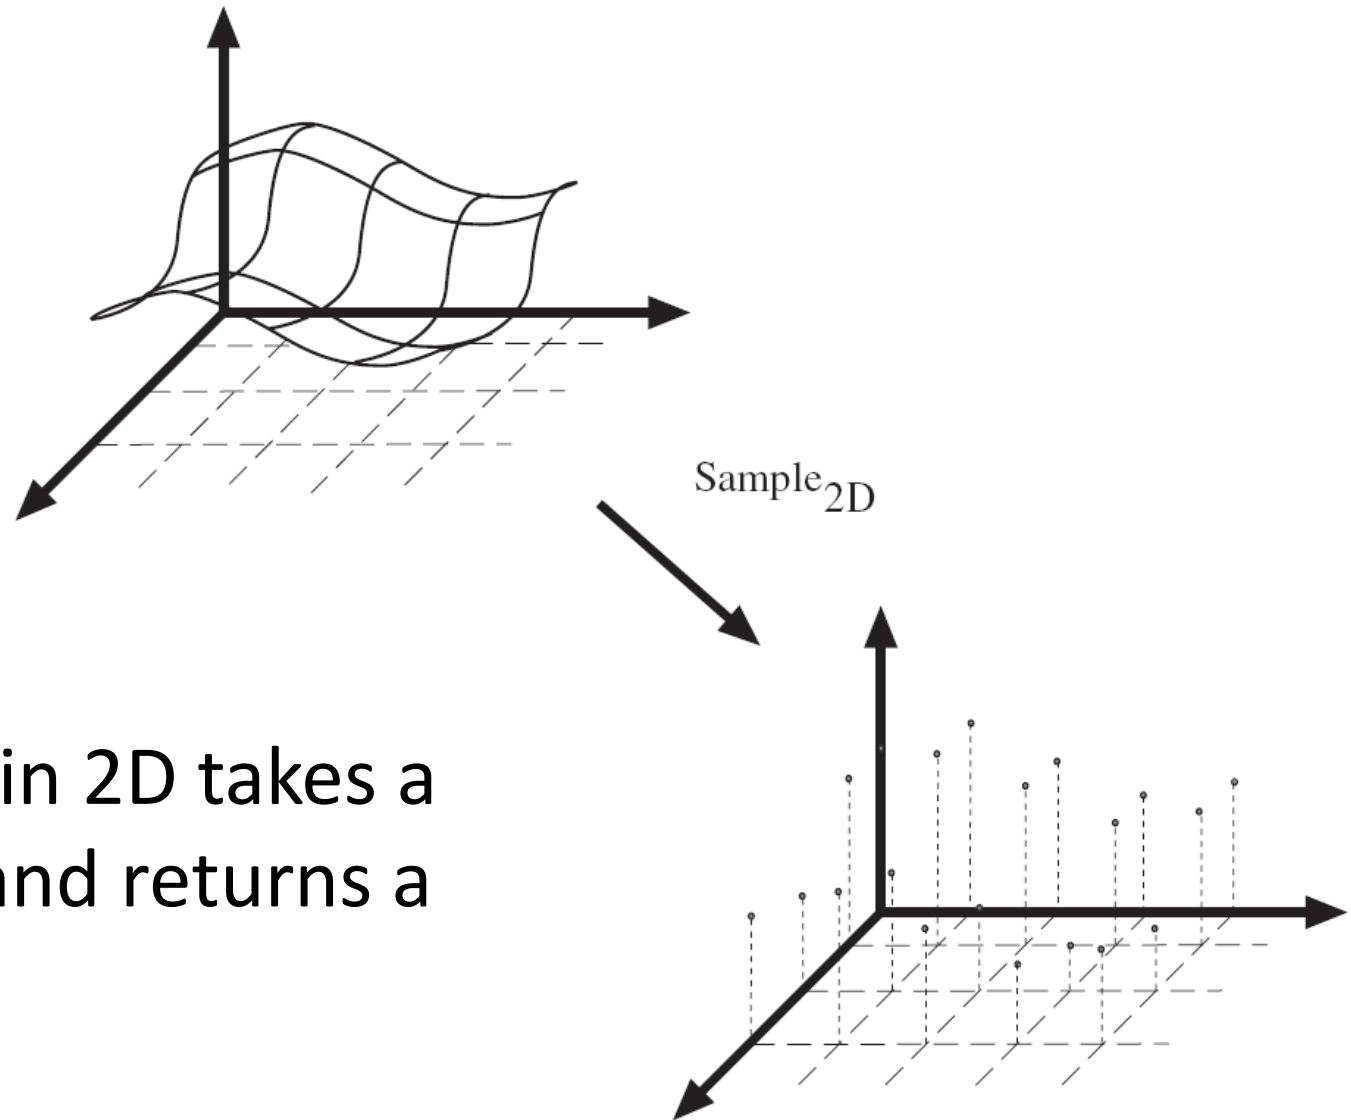


Sampling in 1D



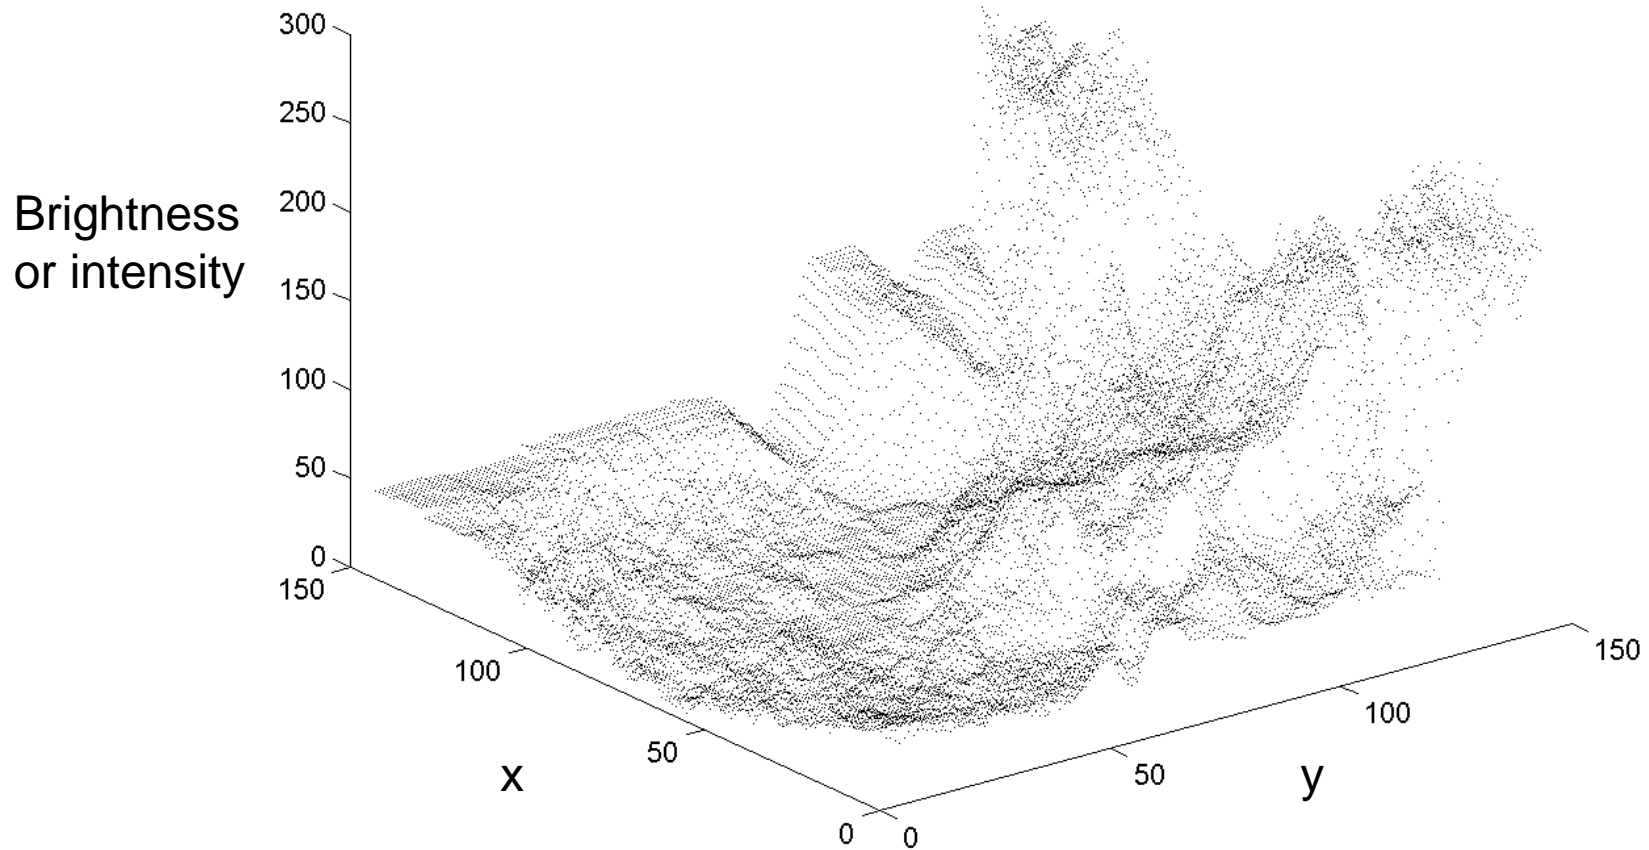
- Sampling in 1D takes a function, and returns a vector whose elements are values of that function at the sample points.

Sampling in 2D



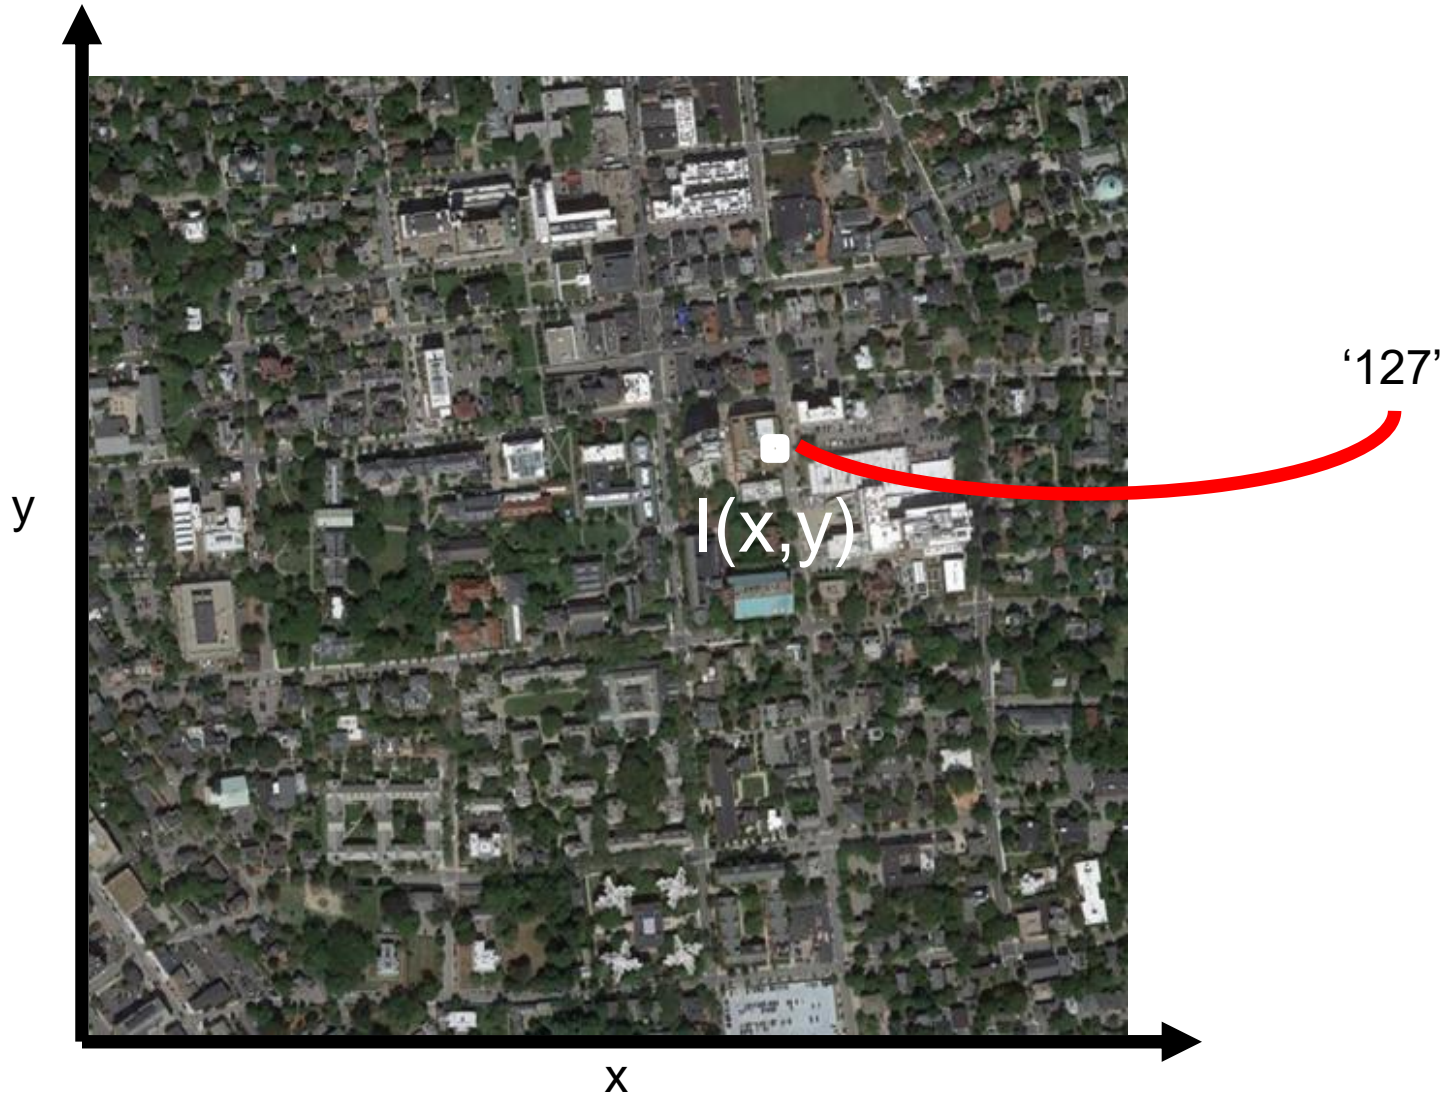
- Sampling in 2D takes a function and returns a matrix.

Grayscale Digital Image

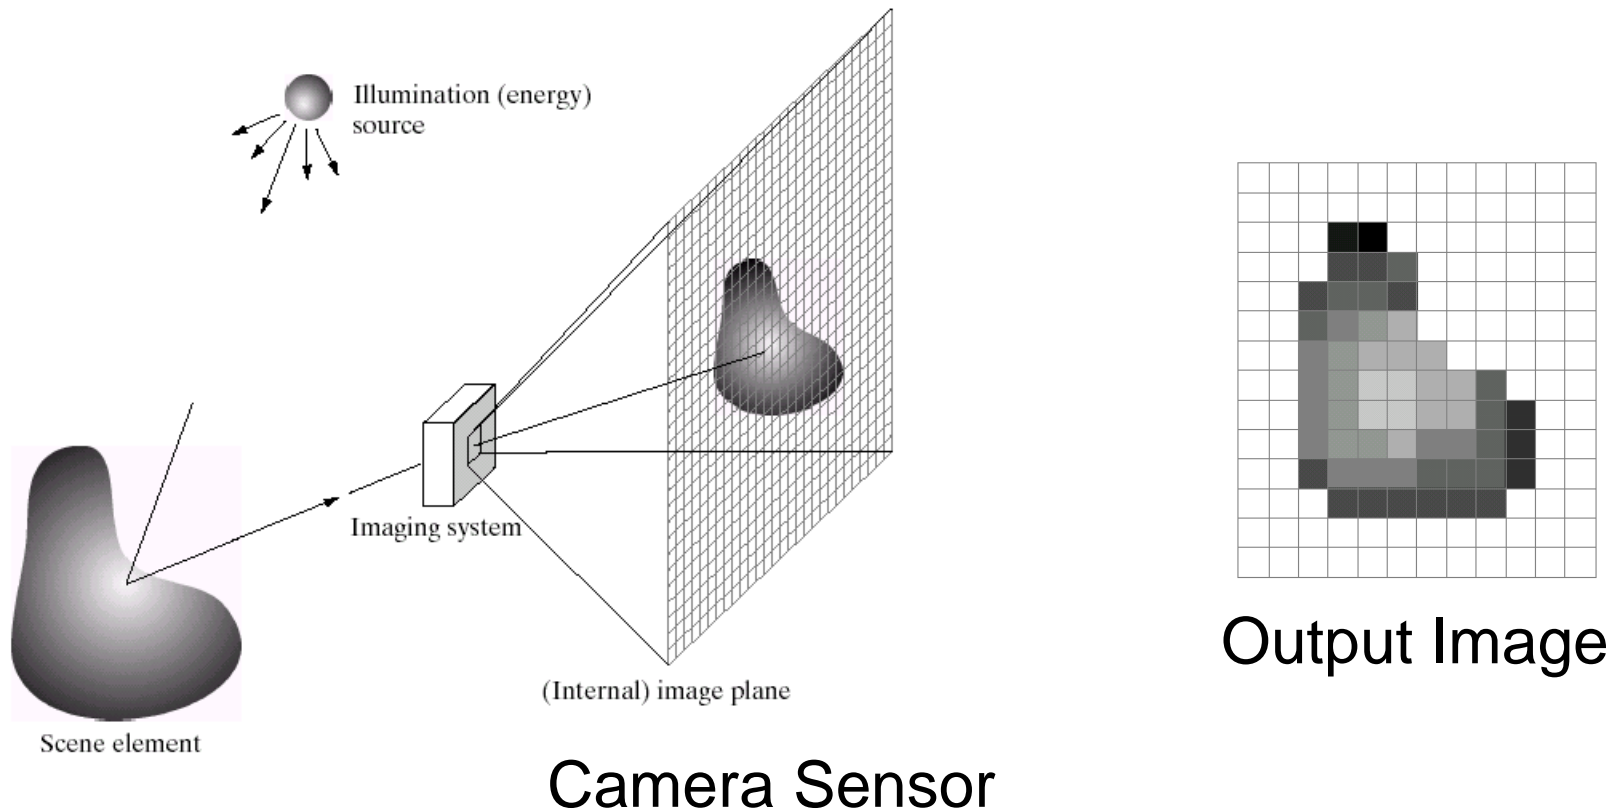


What is each part of a photograph?

- Pixel -> picture element

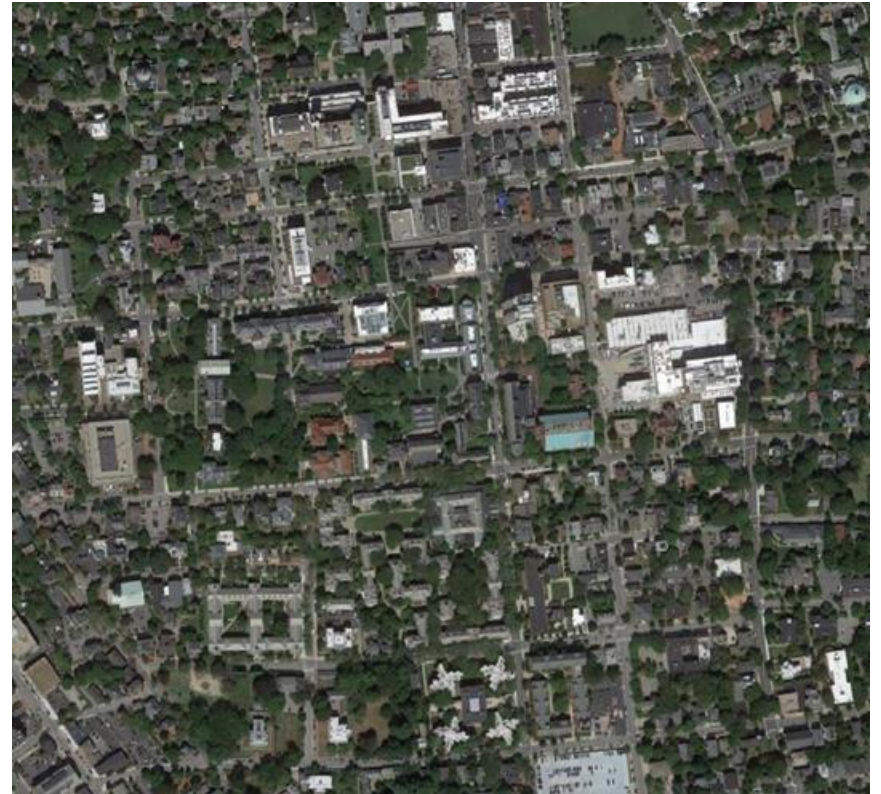


Integrating light over a range of angles.

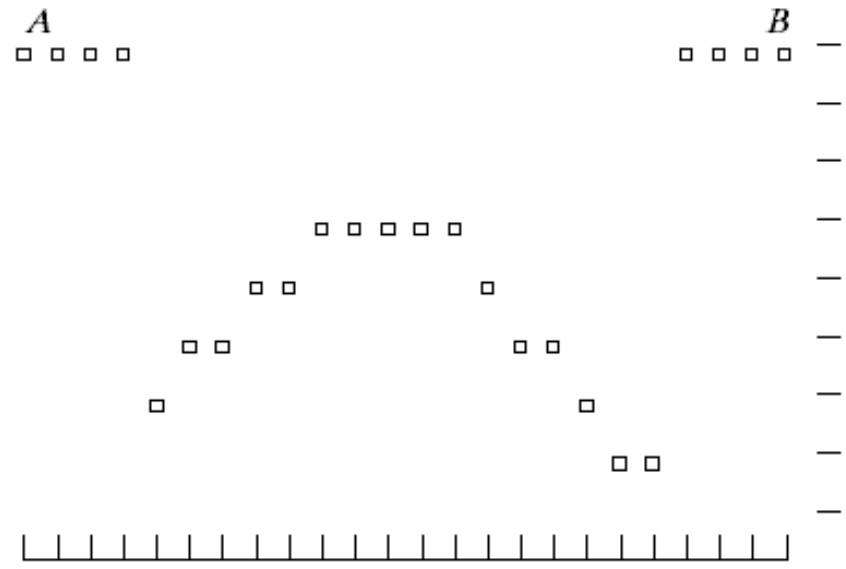
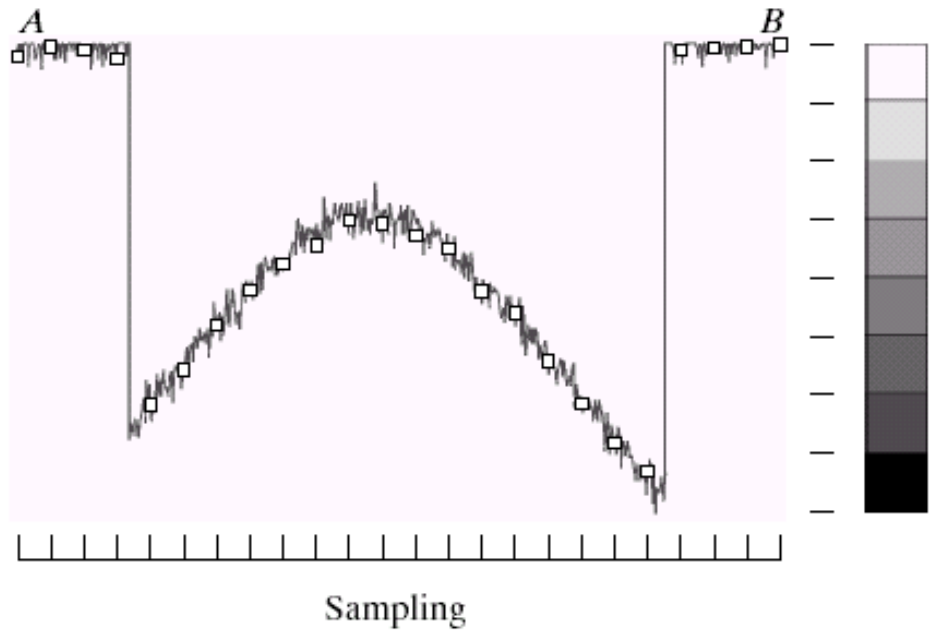
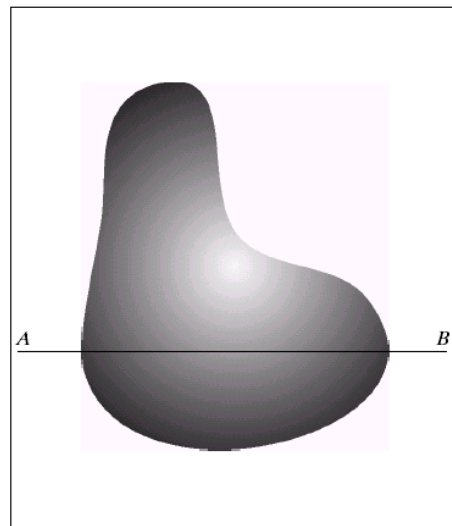


Resolution – geometric vs. spatial resolution

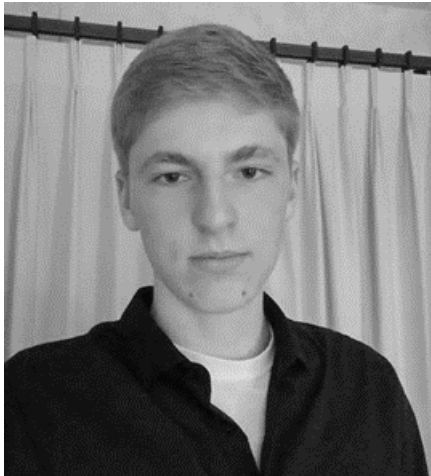
Both images are $\sim 500 \times 500$ pixels



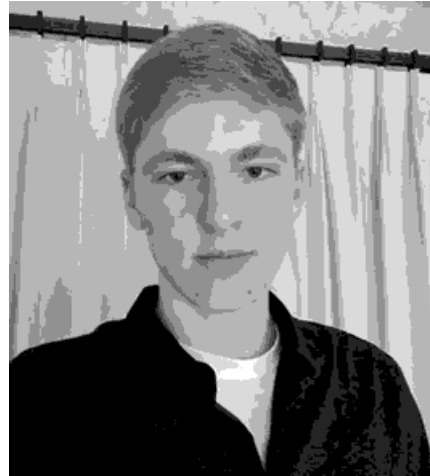
Quantization



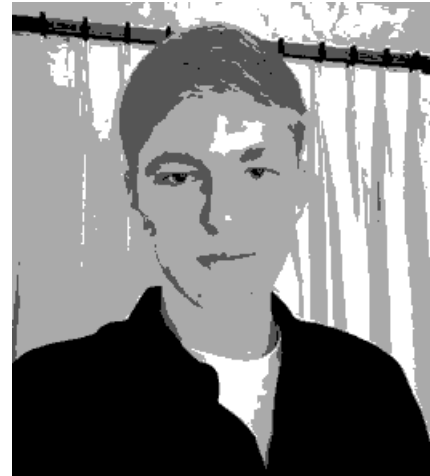
Quantization Effects – Radiometric Resolution



8 bit – 256 levels



4 bit – 16 levels



2 bit – 4 levels



1 bit – 2 levels

Images in Python Numpy

N x M grayscale image “im”

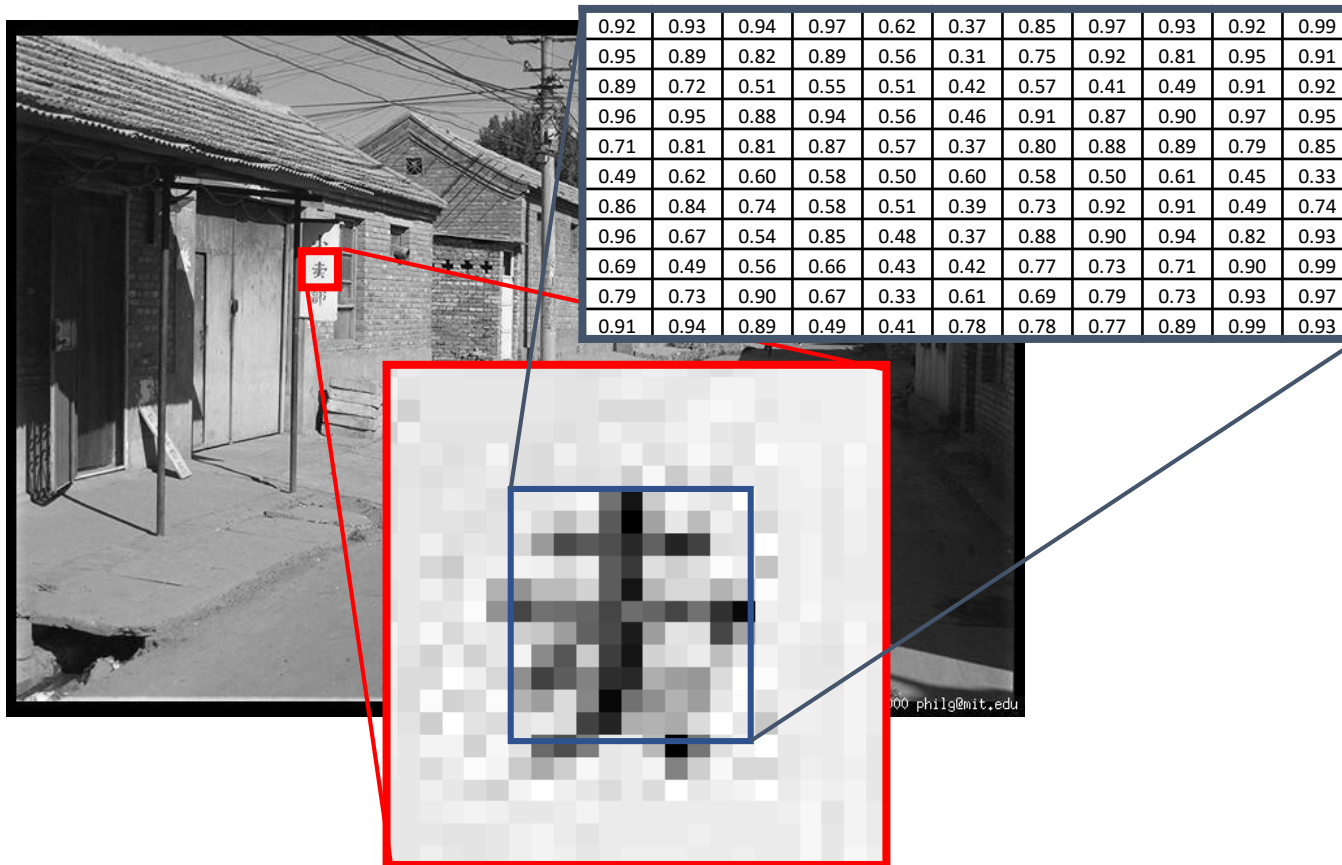
- `im[0,0]` = top-left pixel value
- `im[y, x]` = y pixels down, x pixels to right
- `im[N-1, M-1]` = bottom-right pixel

Row ↓

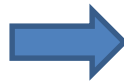
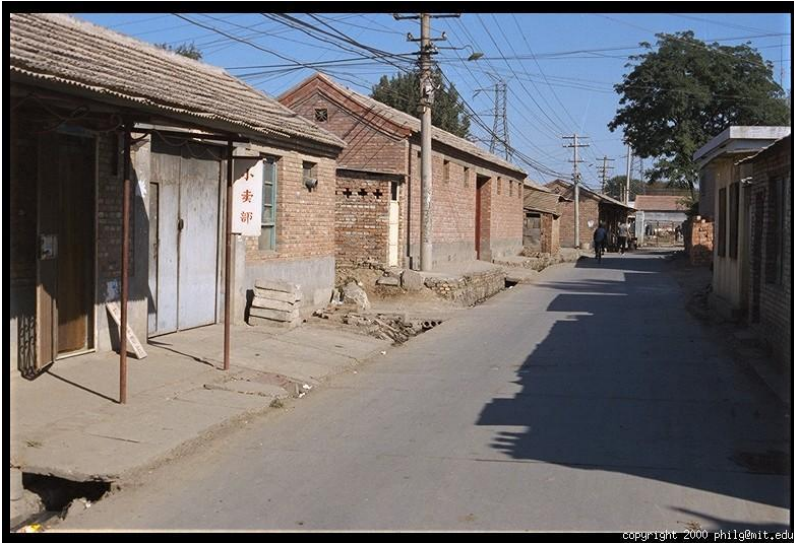
Column →

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



Images in Python Numpy

Take care between 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"))`

Row ↓

Column →

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



IMAGE FILTERING

Image filtering

- Image filtering:
 - Compute function of local neighborhood at each position

$$h[m, n] = \sum_{k, l} f[k, l] I[m + k, n + l]$$

Image filtering

- Image filtering:
 - Compute function of local neighborhood at each position

`h=output`

`f=filter`

`I=image`

$$h[m,n] = \sum_{k,l} f[k,l] I[m+k,n+l]$$

`2d coords=k,l`

`2d coords=m,n`

[]

[]

[]

Example: box filter

$f[\cdot, \cdot]$

$\frac{1}{9}$	1	1	1
	1	1	1
	1	1	1

Image filtering

$$f[\cdot, \cdot] \frac{1}{9}$$

1	1	1
1	1	1
1	1	1

$I[\cdot, \cdot]$

$h[\cdot, \cdot]$

0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0	0
0	0	0	90	90	90	90	90	0	0	0
0	0	0	90	90	90	90	90	0	0	0
0	0	0	90	0	90	90	90	0	0	0
0	0	0	90	90	90	90	90	0	0	0
0	0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0

$$h[m, n] = \sum_{k, l} f[k, l] I[m + k, n + l]$$

$$m = 1, n = 1$$
$$k, l = [-1, 0, 1]$$

Image filtering

$$f[\cdot, \cdot] \frac{1}{9}$$

1	1	1
1	1	1
1	1	1

$I[\cdot, \cdot]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$h[\cdot, \cdot]$

	0	10							

$$h[m, n] = \sum_{k, l} f[k, l] I[m + k, n + l]$$

$$m = 2, n = 1$$
$$k, l = [-1, 0, 1]$$

Image filtering

$$f[\cdot, \cdot] \frac{1}{9}$$

1	1	1
1	1	1
1	1	1

$I[\cdot, \cdot]$

$h[\cdot, \cdot]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

	0	10	20						

$$h[m, n] = \sum_{k, l} f[k, l] I[m + k, n + l]$$

$$m = 3, n = 1$$

$$k, l = [-1, 0, 1]$$

Image filtering

$$f[\cdot, \cdot] \frac{1}{9}$$

1	1	1
1	1	1
1	1	1

$I[\cdot, \cdot]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$h[\cdot, \cdot]$

	0	10	20	30					

$$h[m, n] = \sum_{k, l} f[k, l] I[m + k, n + l]$$

$$m = 4, n = 1$$

$$k, l = [-1, 0, 1]$$

Image filtering

$$f[\cdot, \cdot] \frac{1}{9}$$

1	1	1
1	1	1
1	1	1

$I[\cdot, \cdot]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$h[\cdot, \cdot]$

	0	10	20	30	30				

$$h[m, n] = \sum_{k, l} f[k, l] I[m + k, n + l]$$

$$m = 5, n = 1$$
$$k, l = [-1, 0, 1]$$

Image filtering

$$f[\cdot, \cdot] \frac{1}{9}$$

1	1	1
1	1	1
1	1	1

$I[\cdot, \cdot]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$h[\cdot, \cdot]$

	0	10	20	30	30				
							?		
					50				

$$h[m, n] = \sum_{k, l} f[k, l] I[m + k, n + l]$$

$$m = 6, n = 4$$

$$k, l = [-1, 0, 1]$$

Image filtering

$$f[\cdot, \cdot] \frac{1}{9} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

$I[\cdot, \cdot]$

0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	90	0	90	90	90	0	0
0	0	0	90	90	90	90	90	0	0
0	0	0	0	0	0	0	0	0	0
0	0	90	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0

$h[\cdot, \cdot]$

	0	10	20	30	30	30	20	10	
	0	20	40	60	60	60	40	20	
	0	30	60	90	90	90	60	30	
	0	30	50	80	80	90	60	30	
	0	30	50	80	80	90	60	30	
	0	20	30	50	50	60	40	20	
	10	20	30	30	30	30	20	10	
	10	10	10	0	0	0	0	0	

$$h[m, n] = \sum_{k, l} f[k, l] I[m + k, n + l]$$

Box Filter

What does it do?

- Replaces each pixel with an average of its neighborhood
- Achieve smoothing effect (remove sharp features)

$$\frac{1}{9} f[\cdot, \cdot]$$

1	1	1
1	1	1
1	1	1

Box Filter

What does it do?

- Replaces each pixel with an average of its neighborhood
- Achieve smoothing effect (remove sharp features)
- Why does it sum to one?

$$\frac{1}{9} f[\cdot, \cdot]$$

1	1	1
1	1	1
1	1	1

Smoothing with box filter

$f[\cdot, \cdot]$

1 9	1	1	1
	1	1	1
	1	1	1

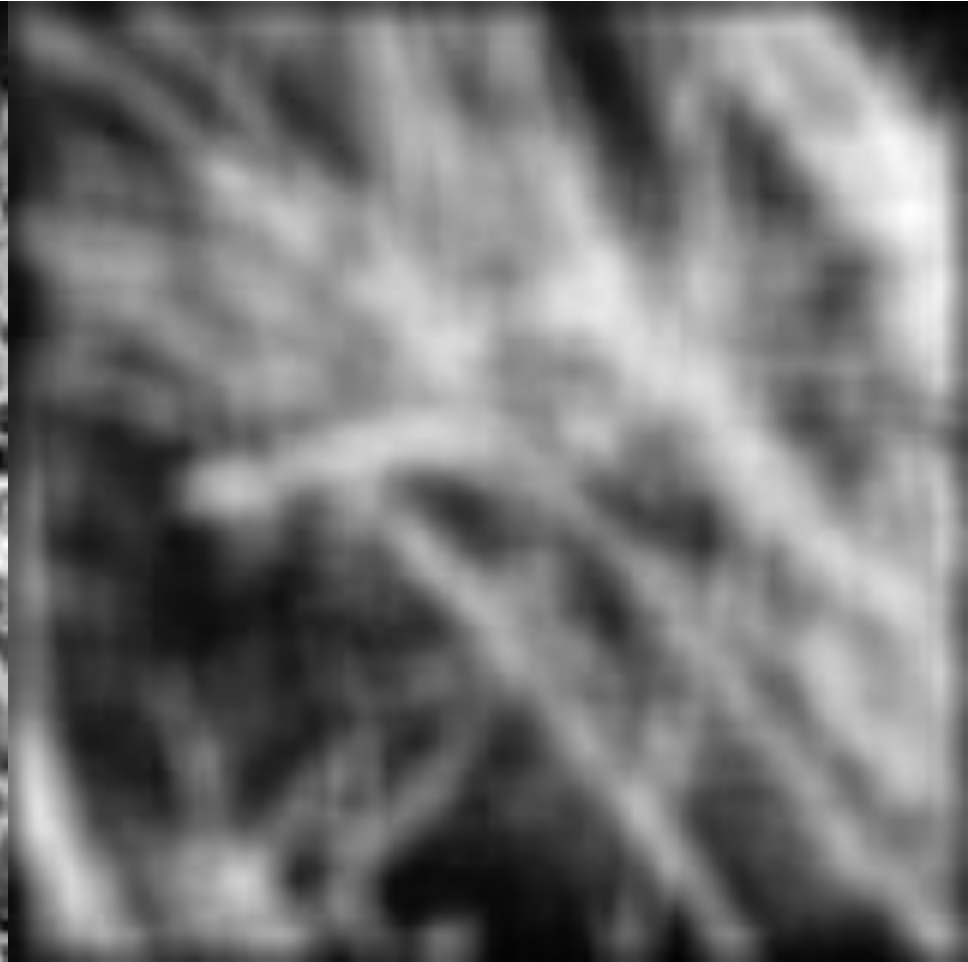


Image filtering

- Image filtering:
 - Compute function of local neighborhood at each position

$$h[m, n] = \sum_{k, l} f[k, l] I[m + k, n + l]$$

- Really important!
 - Enhance images
 - Denoise, resize, increase contrast, etc.
 - Extract information from images
 - Texture, edges, distinctive points, etc.
 - Detect patterns
 - Template matching

Think-Pair-Share time



1.

0	0	0
0	1	0
0	0	0

2.

0	0	0
0	0	1
0	0	0

3.

1	0	-1
2	0	-2
1	0	-1

4.

0	0	0
0	2	0
0	0	0

—

$\frac{1}{9}$

1	1	1
1	1	1
1	1	1

1. Practice with linear filters



Original

0	0	0
0	1	0
0	0	0

?

1. Practice with linear filters



Original

0	0	0
0	1	0
0	0	0



Filtered
(no change)

2. Practice with linear filters



Original

0	0	0
0	0	1
0	0	0

?

2. Practice with linear filters



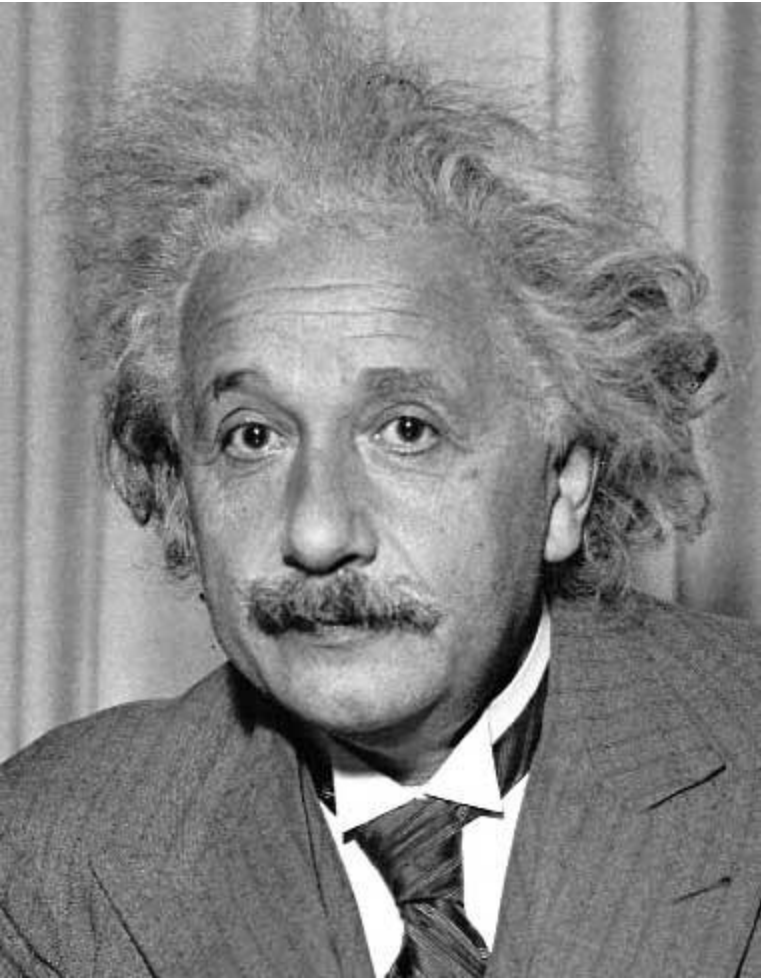
Original

0	0	0
0	0	1
0	0	0



Shifted left
By 1 pixel

3. Practice with linear filters



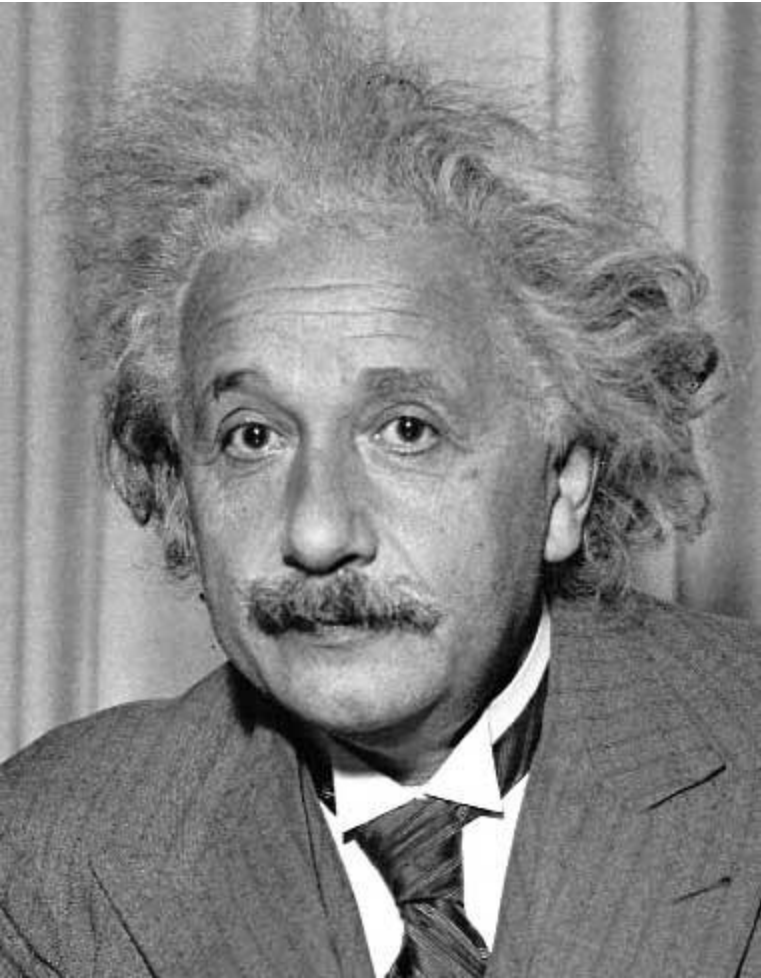
1	0	-1
2	0	-2
1	0	-1

Sobel



Vertical Edge
(absolute value)

3. Practice with linear filters



1	2	1
0	0	0
-1	-2	-1

Sobel



Horizontal Edge
(absolute value)

4. Practice with linear filters



Original

0	0	0
0	2	0
0	0	0

−

$\frac{1}{9}$

1	1	1
1	1	1
1	1	1

?

(Note that filter sums to 1)

4. Practice with linear filters



Original

0	0	0
0	2	0
0	0	0

−

$\frac{1}{9}$

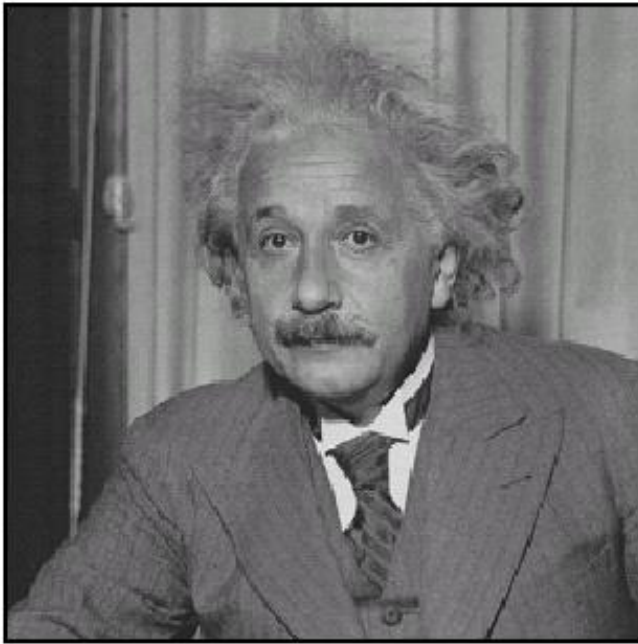
1	1	1
1	1	1
1	1	1



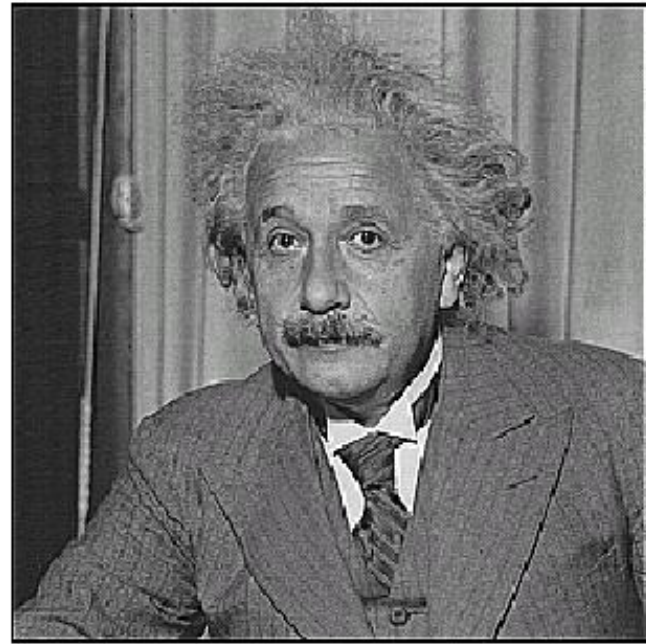
Sharpening filter

- Accentuates differences with local average

4. Practice with linear filters



before



after

Correlation and Convolution

- 2d correlation

$$h[m,n] = \sum_{k,l} f[k,l] I[m+k,n+l]$$

e.g., `h = scipy.signal.correlate2d(f, I)`

Correlation and Convolution

- 2d correlation

$$h[m,n] = \sum_{k,l} f[k,l] I[m+k,n+l]$$

e.g., `h = scipy.signal.correlate2d(f,I)`

- 2d convolution

$$h[m,n] = \sum_{k,l} f[k,l] I[m-k,n-l]$$

e.g., `h = scipy.signal.convolve2d(f,I)`

Convolution is the same as correlation with a 180° rotated filter kernel.
Correlation and convolution are identical when the filter kernel is symmetric.

Key properties of linear filters

Linearity:

$$\text{imfilter}(I, f_1 + f_2) = \text{imfilter}(I, f_1) + \text{imfilter}(I, f_2)$$

Shift invariance:

Same behavior given intensities regardless of pixel location m, n

$$\text{imfilter}(I, \text{shift}(f)) = \text{shift}(\text{imfilter}(I, f))$$

Any linear, shift-invariant operator can be represented as a convolution.

Convolution properties

Commutative: $a * b = b * a$

- Conceptually no difference between filter and signal
- But particular filtering implementations might break this equality, e.g., image edges

Associative: $a * (b * c) = (a * b) * c$

- Often apply several filters one after another: $((a * b_1) * b_2) * b_3$
- This is equivalent to applying one filter: $a * (b_1 * b_2 * b_3)$

Convolution properties

Commutative: $a * b = b * a$

- Conceptually no difference between filter and signal
- But particular filtering implementations might break this equality, e.g., image edges

Associative: $a * (b * c) = (a * b) * c$

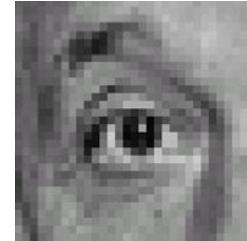
- Often apply several filters one after another: $((a * b_1) * b_2) * b_3$
- This is equivalent to applying one filter: $a * (b_1 * b_2 * b_3)$
- Correlation is not associative (rotation effect)
- Why important?

Recap of Monday

- Linear filtering (convolution)

$$h[m, n] = \sum_{k, l} f[k, l] I[m - k, n - l]$$

- Not a matrix multiplication
- Sum over Hadamard product
- Can smooth, sharpen, translate (among many other uses)



$\frac{1}{9}$

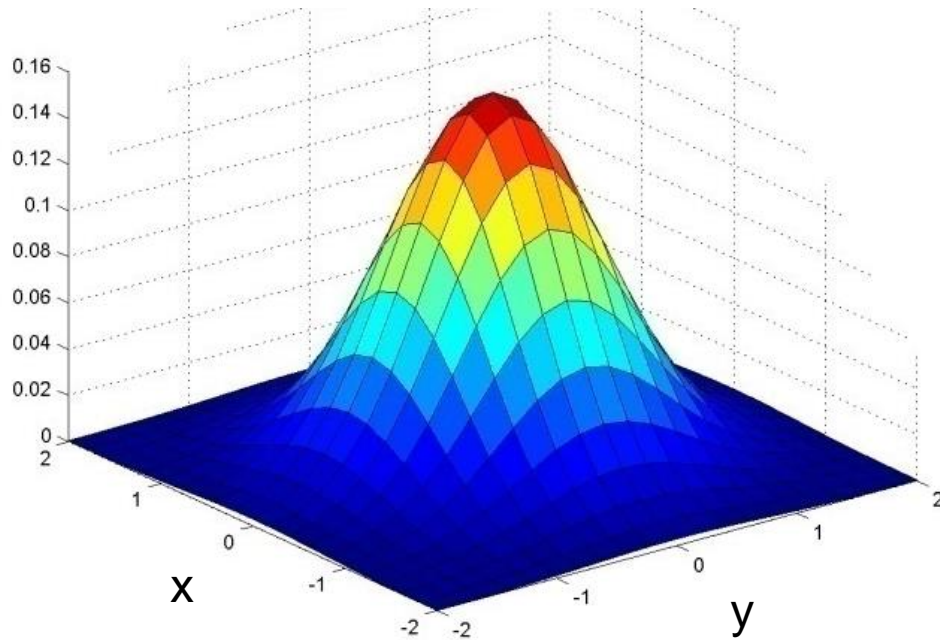
1	1	1
1	1	1
1	1	1

Convolution properties

- Commutative: $a * b = b * a$
 - Conceptually no difference between filter and signal
 - But particular filtering implementations might break this equality, e.g., image edges
- Associative: $a * (b * c) = (a * b) * c$
 - Often apply several filters one after another: $((a * b_1) * b_2) * b_3$
 - This is equivalent to applying one filter: $a * (b_1 * b_2 * b_3)$
 - Correlation is not associative (rotation effect)
 - Why important?
- Distributes over addition: $a * (b + c) = (a * b) + (a * c)$
- Scalars factor out: $ka * b = a * kb = k(a * b)$
- Identity: unit impulse $e = [0, 0, 1, 0, 0]$, $a * e = a$

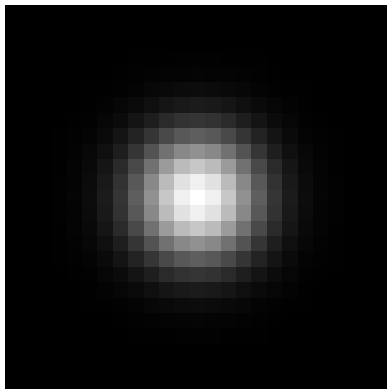
Important filter: Gaussian

Weight contributions of neighboring pixels by nearness



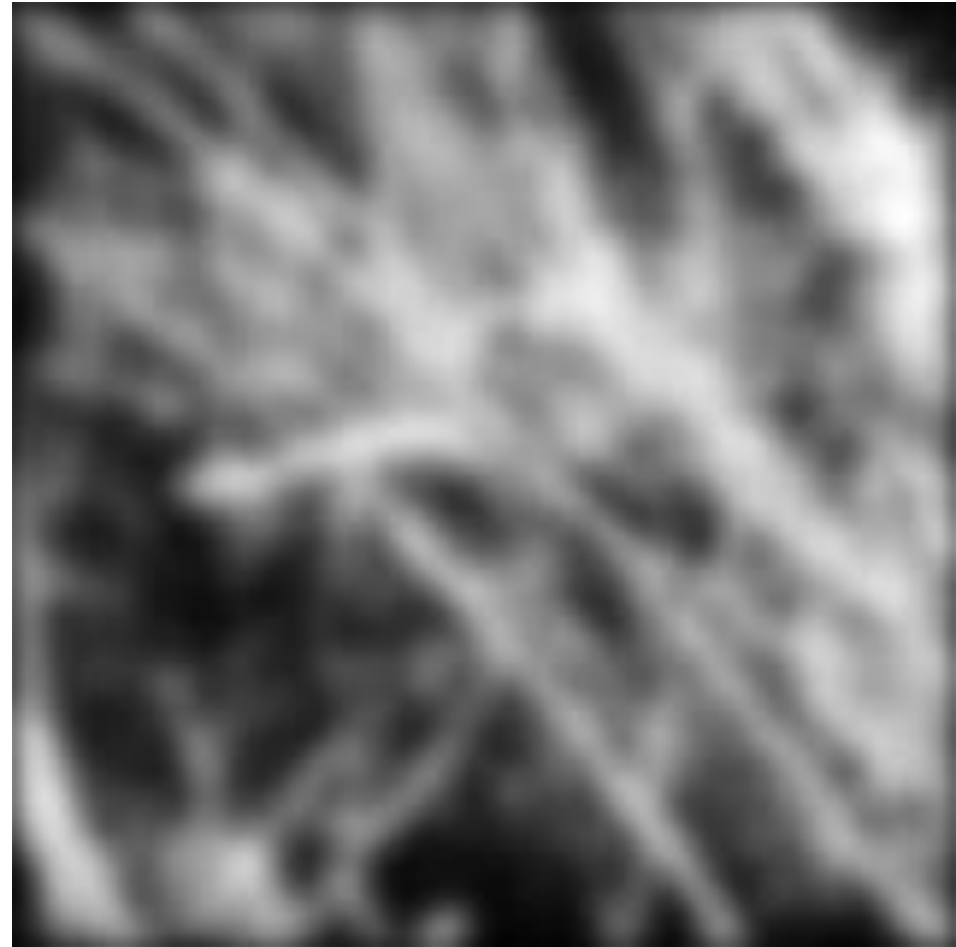
		x				
y	0.003	0.013	0.022	0.013	0.003	
	0.013	0.059	0.097	0.059	0.013	
	0.022	0.097	0.159	0.097	0.022	
	0.013	0.059	0.097	0.059	0.013	
	0.003	0.013	0.022	0.013	0.003	

5 x 5, $\sigma = 1$

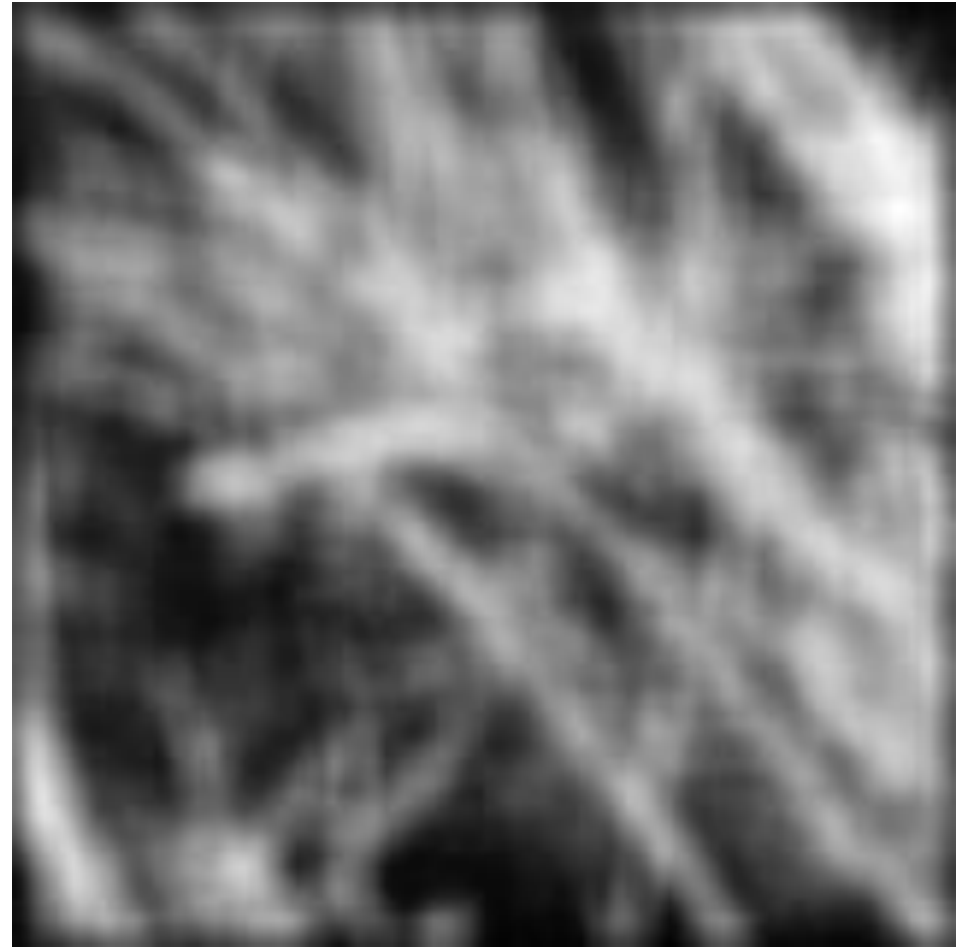


$$G_{\sigma} = \frac{1}{2\pi\sigma^2} e^{-\frac{(x^2+y^2)}{2\sigma^2}}$$

Smoothing with Gaussian filter



Smoothing with box filter



Gaussian filters

- Remove “high-frequency” components from the image (low-pass filter)
 - Images become more smooth
- Gaussian convolved with Gaussian...
 - ...is another Gaussian
 - So can smooth with small-width kernel, repeat, and get same result as larger-width kernel would have
 - Convolution twice with Gaussian kernel of width σ is same as convolution once with kernel of width $\sigma\sqrt{2}$
- *Separable* kernel
 - Factors into product of two 1D Gaussians

Separability of the Gaussian filter

$$\begin{aligned} G_{\sigma}(x, y) &= \frac{1}{2\pi\sigma^2} \exp\left(-\frac{x^2 + y^2}{2\sigma^2}\right) \\ &= \left(\frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{x^2}{2\sigma^2}\right) \right) \left(\frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{y^2}{2\sigma^2}\right) \right) \end{aligned}$$

The 2D Gaussian can be expressed as the product of two functions, one a function of x and the other a function of y

In this case, the two functions are the (identical) 1D Gaussian

Separability example

2D convolution
(center location only)

$$\begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline 2 & 4 & 2 \\ \hline 1 & 2 & 1 \\ \hline \end{array} * \begin{array}{|c|c|c|} \hline 2 & 3 & 3 \\ \hline 3 & 5 & 5 \\ \hline 4 & 4 & 6 \\ \hline \end{array}$$

The filter factors
into a product of 1D
filters:

$$\begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline 2 & 4 & 2 \\ \hline 1 & 2 & 1 \\ \hline \end{array} = \begin{array}{|c|} \hline 1 \\ \hline 2 \\ \hline 1 \\ \hline \end{array} \times \begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline \end{array}$$

Perform convolution
along rows:

$$\begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline \end{array} * \begin{array}{|c|c|c|} \hline 2 & 3 & 3 \\ \hline 3 & 5 & 5 \\ \hline 4 & 4 & 6 \\ \hline \end{array} = \begin{array}{|c|c|c|} \hline & 11 & \\ \hline & 18 & \\ \hline & 18 & \\ \hline \end{array}$$

Followed by convolution
along the remaining column:

Separability

Why is separability useful in practice?

MxN image, PxQ filter

- 2D convolution: $\sim MN PQ$ multiply-adds
- Separable 2D: $\sim MN(P+Q)$ multiply-adds

Speed up = $PQ/(P+Q)$

9x9 filter = $\sim 4.5x$ faster

Practical matters

How big should the filter be?

- Values at edges should be near zero
- Gaussians have infinite extent...
- Rule of thumb for Gaussian: set filter half-width to about 3σ

Practical matters

What about near the edge?

- The filter window falls off the edge of the image
- Need to extrapolate
- methods:
 - clip filter (black)
 - wrap around
 - copy edge
 - reflect across edge



Convolution in Convolutional Neural Networks

- Convolution is the basic operation in CNNs
- Learning convolution kernels allows us to learn which `features` provide useful information in images.

Sobel filter visualization

- What happens to negative numbers?
- For visualization:
 - Shift image + 0.5
 - If gradients are small, scale edge response

```
>> I = img_to_float32( io.imread( 'luke.jpg' ) );  
>> h = convolve2d( I, sobelKernel );
```

1	2	1
0	0	0
-1	-2	-1

Sobel

```
plt.imshow( h ); plt.imshow( h + 0.5 );
```



$$h(:, :, 1) < 0$$



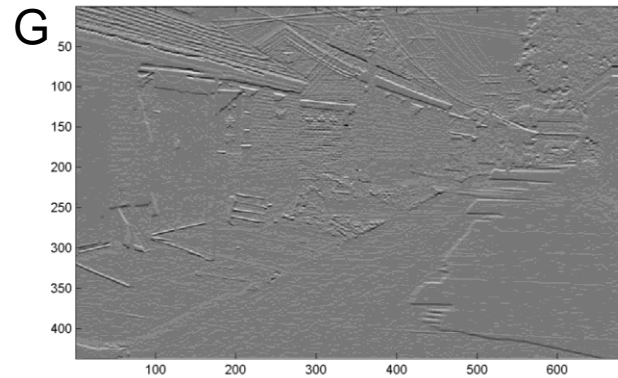
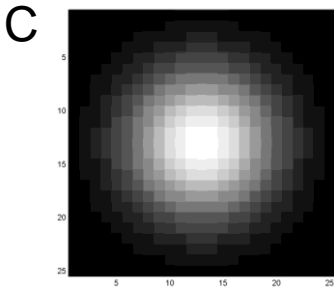
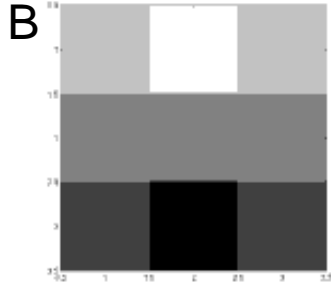
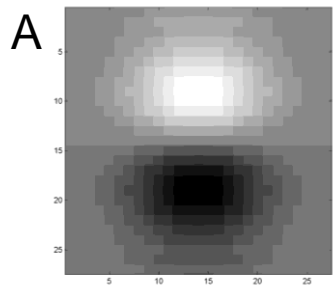
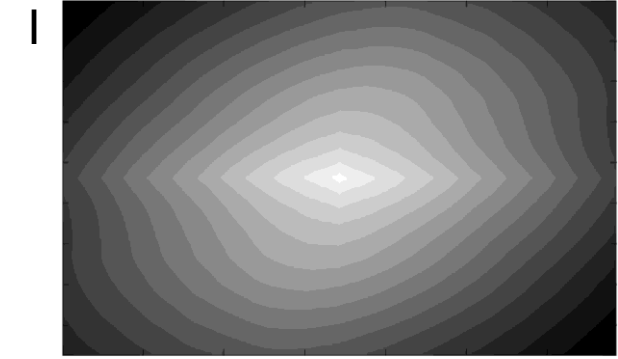
$$h(:, :, 1) > 0$$



Think-Pair-Share

* = Convolution operator

- a) $_ = D * B$
- b) $\bar{A} = _ * _$
- c) $F = D * _$
- d) $_ = D * D$

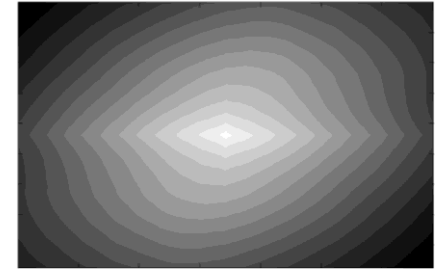


$$I = D * D$$

D (275 x 175 pixels)



I (from slide – 275 x 175)



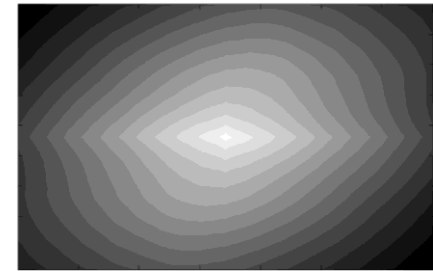
“...something to do with lack of content (black) at edges...”

$$I = D * D$$

D (275 x 175 pixels)



I (from slide – 275 x 175)



```
>> D = img_to_float32( io.imread( 'convexample.png' ) )
```

```
>> I = convolve2d( D, D )
```

```
>> np.max(I)
```

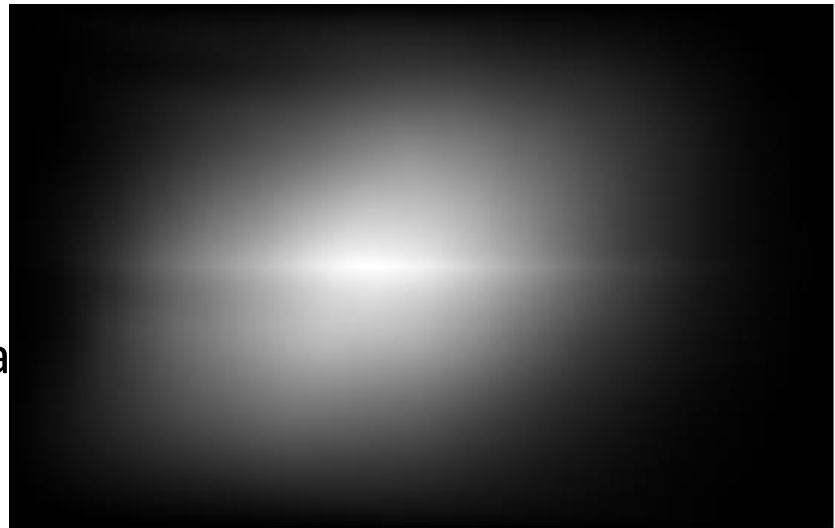
```
1.1021e+04
```

```
# Normalize for visualization
```

```
>> I_norm = (I - np.min(I)) / (np.ma
```

```
>> plt.imshow( I_norm )
```

I_norm

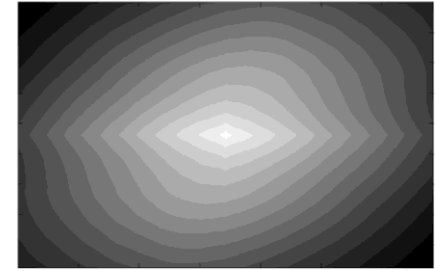


$$I = D * D$$

D (275 x 175 pixels)



I (from slide – 275 x 175)



$(275-1)/2$



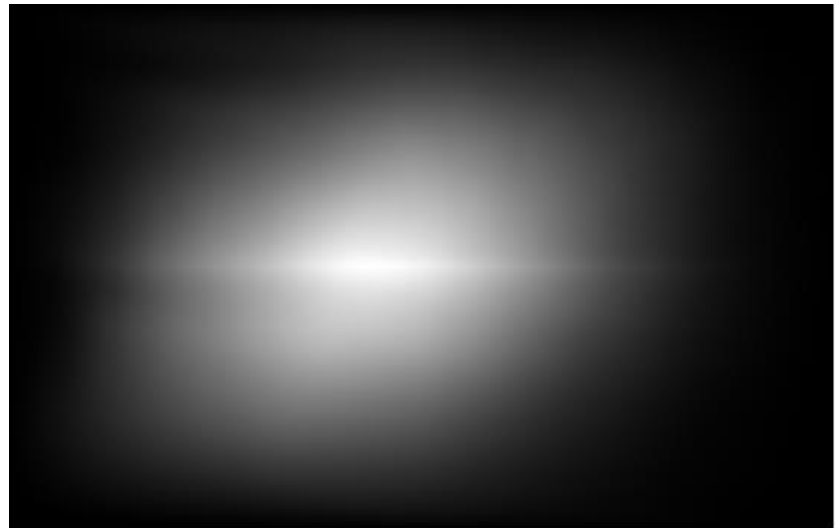
$(275-1)/2$



275

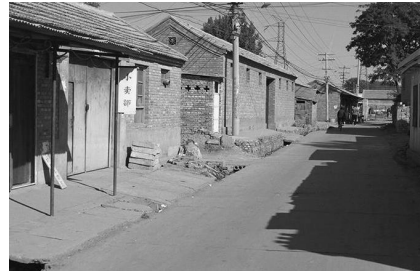
For x: $275 + (275-1)/2 + (275-1)/2$
 $= 549$

I_norm (549 x 349 pixels)

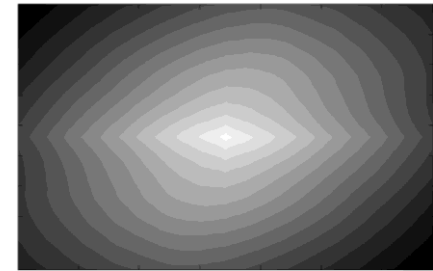


$$I = D * D$$

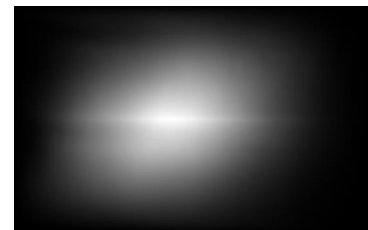
D (275 x 175 pixels)



I (from slide – 275 x 175)

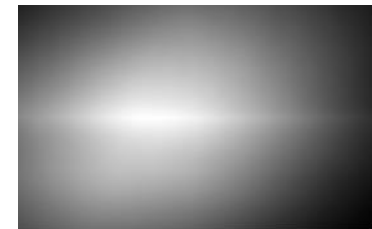


```
>> I = convolve2d( D, D, mode='full' )  
(Default; pad with zeros)
```



549 x 349

```
>> I = convolve2d( D, D, mode='same' )  
(Return same size as D)
```



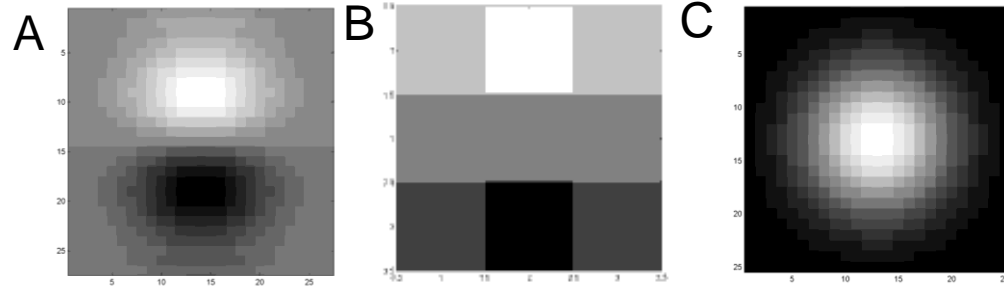
275 x 175

```
>> I = convolve2d( D, D, mode='valid' )  
(No padding)
```



Value = 10528.3
1x1

$A = B * C$ – “because it kind of looks like it.”



C is a Gaussian filter
(or something close to it it),
and we know that it ‘blurs’.

When the filter ‘looks like’ the image = ‘template matching’

Filtering viewed as comparing an image of
what you want to find against all image regions.

For symmetric filters: use either convolution or correlation.

For nonsymmetric filters: correlation is template matching.

Filtering: Correlation and Convolution

- 2d correlation

$$h[m,n] = \sum_{k,l} f[k,l] I[m+k,n+l]$$

e.g., `h = scipy.signal.correlate2d(f,I)`

- 2d convolution

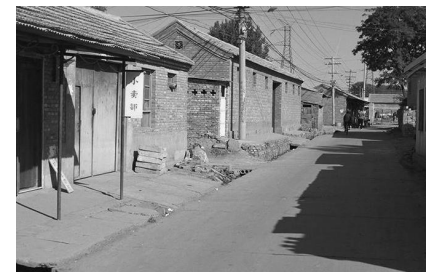
$$h[m,n] = \sum_{k,l} f[k,l] I[m-k,n-l]$$

e.g., `h = scipy.signal.convolve2d(f,I)`

Convolution is the same as correlation with a 180° rotated filter kernel.
Correlation and convolution are identical when the filter kernel is symmetric.

OK, so let's test this idea. Let's see if we can use correlation to 'find' the parts of the image that look like the filter.

D (275 x 175 pixels)



```
>> f = D[ 57:117, 107:167 ]
```

Expect response 'peak' in middle of I

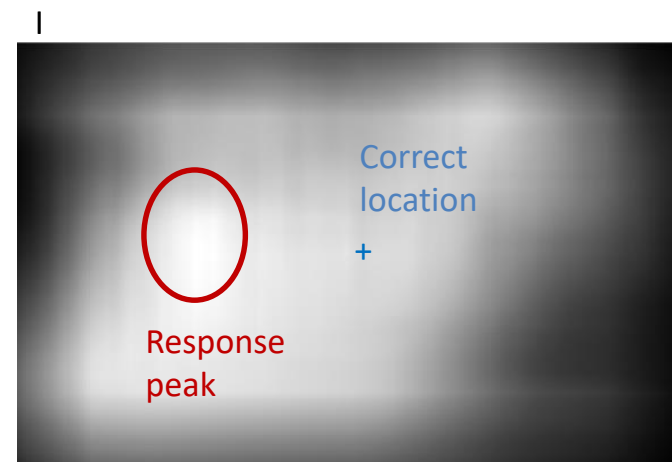


f
61 x 61

```
>> I = correlate2d( D, f, 'same' )
```

Hmm...

That didn't work – why not?



Correlation

$$h[m,n] = \sum_{k,l} f[k,l] I[m+k,n+l]$$

e.g., `h = scipy.signal.correlate2d(f, I)`

As brightness in I increases, the response in h will increase, as long as f is positive.

Overall brighter regions will give higher correlation response -> not useful!

OK, so let's subtract the mean

```
>> f = D[ 57:117, 107:167 ]  
>> f2 = f - np.mean(f)  
>> D2 = D - np.mean(D)
```

*Now zero centered.
Score is higher only when dark parts
match and when light parts match.*

```
>> I2 = correlate2d( D2, f2, 'same' )
```

D2 (275 x 175 pixels)



f2
61 x 61

I2



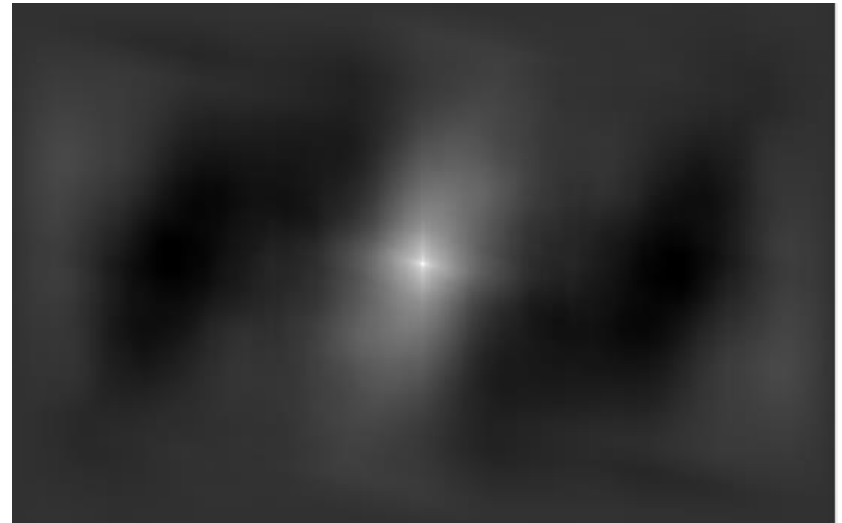
Or even

```
>> I3 = correlate2d( D2, D2, 'full' )
```

D2 (275 x 175 pixels)



I3



What happens with convolution?

```
>> f = D[ 57:117, 107:167 ]
```

```
>> f2 = f - np.mean(f)
```

```
>> D2 = D - np.mean(D)
```

```
>> I2 = convolve2d( D2, f2, 'same' )
```

D2 (275 x 175 pixels)



f2
61 x 61

I2



NON-LINEAR FILTERS

Median filters

- Operates over a window by selecting the median intensity in the window.
- ‘Rank’ filter as based on ordering of gray levels
 - E.G., min, max, range filters

Median filters

- Operates over a window by selecting the median intensity in the window.
- What advantage does a median filter have over a mean filter?

Noise – Salt and Pepper Jack



Mean Jack – 3 x 3 filter



Very Mean Jack – 11 x 11 filter



Noisy Jack – Salt and Pepper



Median Jack – 3 x 3



Very Median Jack – 11 x 11



Median filters

- Operates over a window by selecting the median intensity in the window.
- What advantage does a median filter have over a mean filter?
- Is a median filter a kind of convolution?

Median filters

- Operates over a window by selecting the median intensity in the window.
- What advantage does a median filter have over a mean filter?
- Is a median filter a kind of convolution?

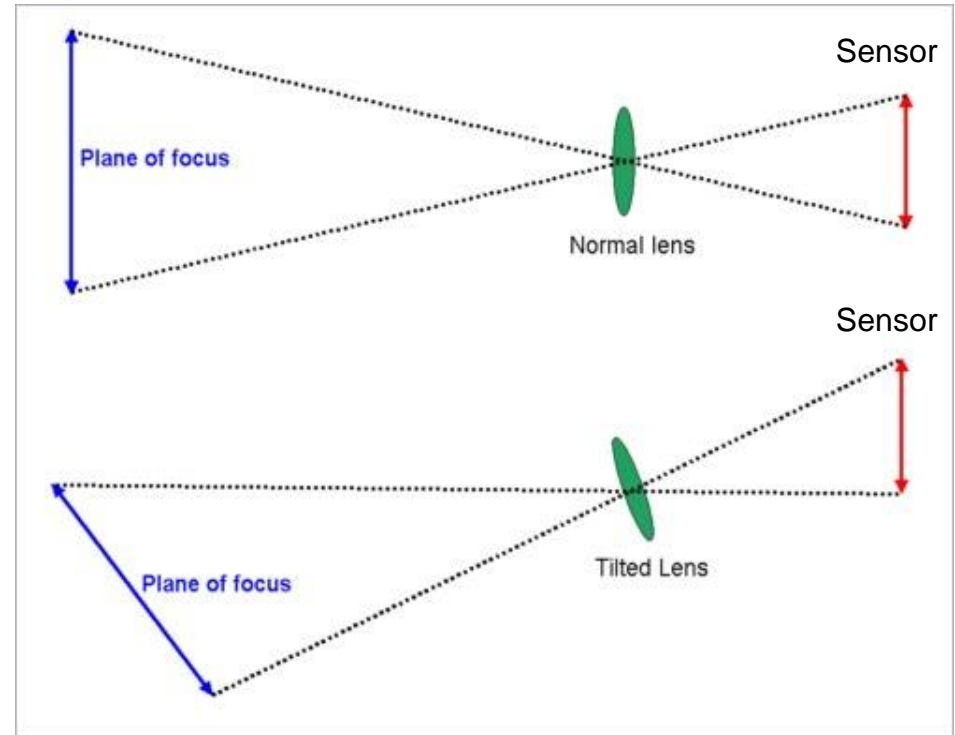
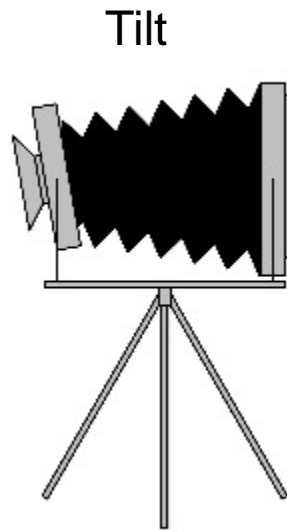
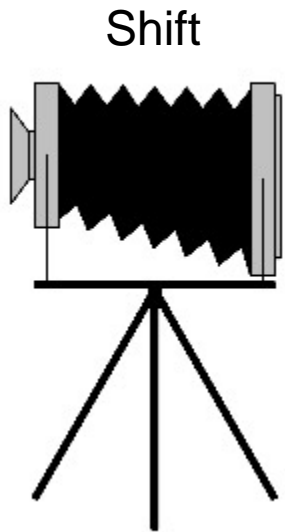
Interpretation: Median filtering is sorting.



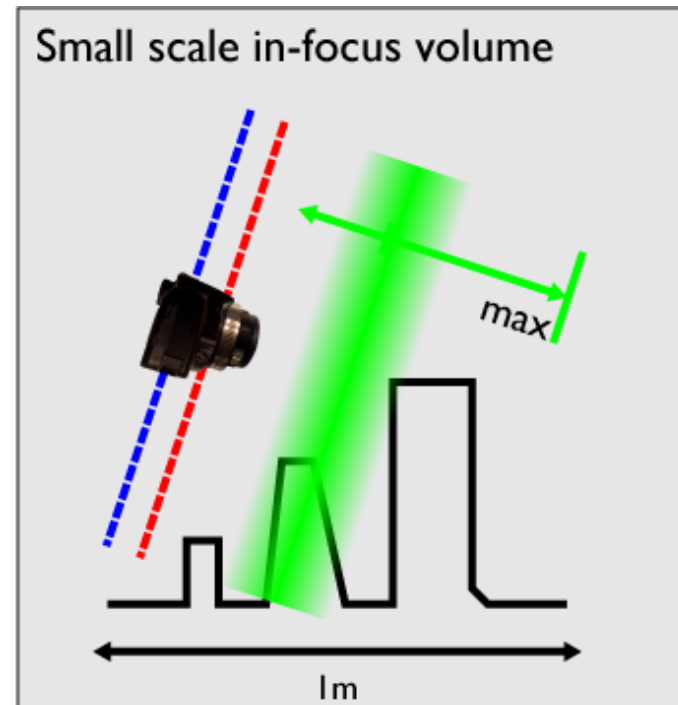
Tilt-shift photography



Tilt shift camera



Macro photography



Can we fake tilt shift?

- We need to blur the image
 - OK, now we know how to do that.

Can we fake tilt shift?

- We need to blur the image
 - OK, now we know how to do that.
- We need to blur progressively more away from our 'fake' focal point



But can I make it look more like a toy?

- Boost saturation – toys are very colorful
- We'll learn how to do this when we discuss color
- For now: transform to Hue, Saturation, Value instead of RGB



Next class: Thinking in Frequency

