

Introduction to Computer Vision

Instructors: Jean Ponce and Matthew Trager
jean.ponce@inria.fr, matthew.trager@cims.nyu.edu


TAs: Jiachen (Jason) Zhu and Sahar Siddiqui
jiachen.zhu@nyu.edu, ss12414@nyu.edu

Outline

Image filtering

- Image filters in spatial domain
 - Filter is a mathematical operation on values of each patch
 - Smoothing, sharpening, measuring texture
- Image filters in the frequency domain
 - Filtering is a way to modify the frequencies of images
 - Denoising, sampling, image compression

What is a (digital) image?



An image is function $f : \Omega \rightarrow V$ defined on a rectangular array of pixels:


$$\Omega = \{(x, y) \mid 1 \leq x \leq N_{cols}, 1 \leq y \leq N_{rows}\} \subset \mathbb{Z}^2.$$

For scalar images, the range is usually a discrete set, $V = \{0, \dots, 2^a - 1\}$. Thus, f can also be viewed as a grid of integers.

The raster image (pixel matrix)



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



000 philg@mit.edu

Image filtering

For each pixel, compute function of a neighborhood and output a new value:

$$h[i, j] = g(f[i + k, j + l]_{k,l}) .$$

If g is linear, we talk about *linear filtering*:

$$h[i, j] = \sum_{k,l} g(k, l) \cdot f(i + k, j + l) .$$

- Same function applied at each position
- Output and input image are typically the same size

Applications

- Enhance images
 - Denoise, resize, increase contrast, etc.
- Extract information from images
 - Texture, edges, distinctive points, etc.
- Detect patterns
 - Template matching
- Convolutional Neural Networks

Example: box filter

- Replace each pixel with a weighted average of its neighborhood.
- The weights are called the filter kernel.
- What are the weights for the average of a 3x3 neighborhood?

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

1	1	1
1	1	1
1	1	1

“box filter”

Image filtering

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

1	1	1
1	1	1
1	1	1

 $f[\cdot, \cdot]$
 $h[\cdot, \cdot]$

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								

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

Image filtering

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

$f[\cdot, \cdot]$

$h[\cdot, \cdot]$

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							

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

Image filtering

$$g[\cdot, \cdot] \frac{1}{9} \begin{array}{|c|c|c|} \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline \end{array}$$

$f[\cdot, \cdot]$

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						

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

Image filtering

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

1	1	1
1	1	1
1	1	1

 $f[\cdot, \cdot]$
 $h[\cdot, \cdot]$

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	30					

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

Image filtering

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

1	1	1
1	1	1
1	1	1

 $f[\cdot, \cdot]$

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				

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

Image filtering

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

1	1	1
1	1	1
1	1	1

$$f[\cdot, \cdot]$$

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					

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

Image filtering

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

$$f[\cdot, \cdot]$$

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	

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


Box filter

What does it do?


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



Smoothing with box filter



Practice with linear filters




Original

0	0	0
0	1	0
0	0	0


?

Practice with linear filters




Original

0	0	0
0	1	0
0	0	0



Filtered
(no change)

Practice with linear filters




Original

0	0	0
0	0	1
0	0	0


?

Practice with linear filters




Original

0	0	0
0	0	1
0	0	0



Shifted left
By 1 pixel

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)

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

Sharpening




before




after

Other filters




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

Sobel




Vertical Edge
(absolute value)

Other filters



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

Sobel



Horizontal Edge
(absolute value)

Basic gradient filters

Horizontal Gradient

0	0	0
-1	0	1
0	0	0

or

-1	0	1
----	---	---

Vertical Gradient

0	-1	0
0	0	0
0	1	0

or

-1
0
1

Filter vs. convolution

Filtering (correlation):

$$h[i, j] = f \otimes g = \sum_{k, l} g[k, l] \cdot f[i + k, j + l].$$

Convolution:

$$h[i, j] = f \star g = \sum_{k, l} g[k, l] \cdot f[i - k, j - l].$$

Clearly equivalent if $g[i, j] = g[-i, -j]$, however in general there are differences.

Some properties

- Linearity:

$$g \star (f_1 + f_2) = g \star f_1 + g \star f_2, \quad g \otimes (f_1 + f_2) = g \otimes f_1 + g \otimes f_2$$

- Stationarity: if $T_{[u,v]}(f)[i,j] = f[i-u, j-v]$

$$g \star T_{[u,v]}(f) = T_{[u,v]}(g \star f), \quad g \otimes T_{[u,v]}(f) = T_{[u,v]}(g \otimes f)$$

Theorem: any linear shift-invariant operator can be represented as a convolution


- Associativity and commutativity:

$$g \star (h \star f) = (g \star h) \star f, \quad g \otimes (h \otimes f) \neq (g \otimes h) \otimes f$$

$$g \star f = f \star g, \quad g \otimes f \neq f \otimes g$$



Smoothing with box filter revisited

- What's wrong with this picture?
- What's the solution?



Gaussian Kernel

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




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, $\sigma = 1$


- Constant factor at front makes volume sum to 1 (can be ignored when computing the filter values, as we should renormalize weights to sum to 1 in any case)

Gaussian Kernel

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



$\sigma = 2$ with 30 x 30
kernel


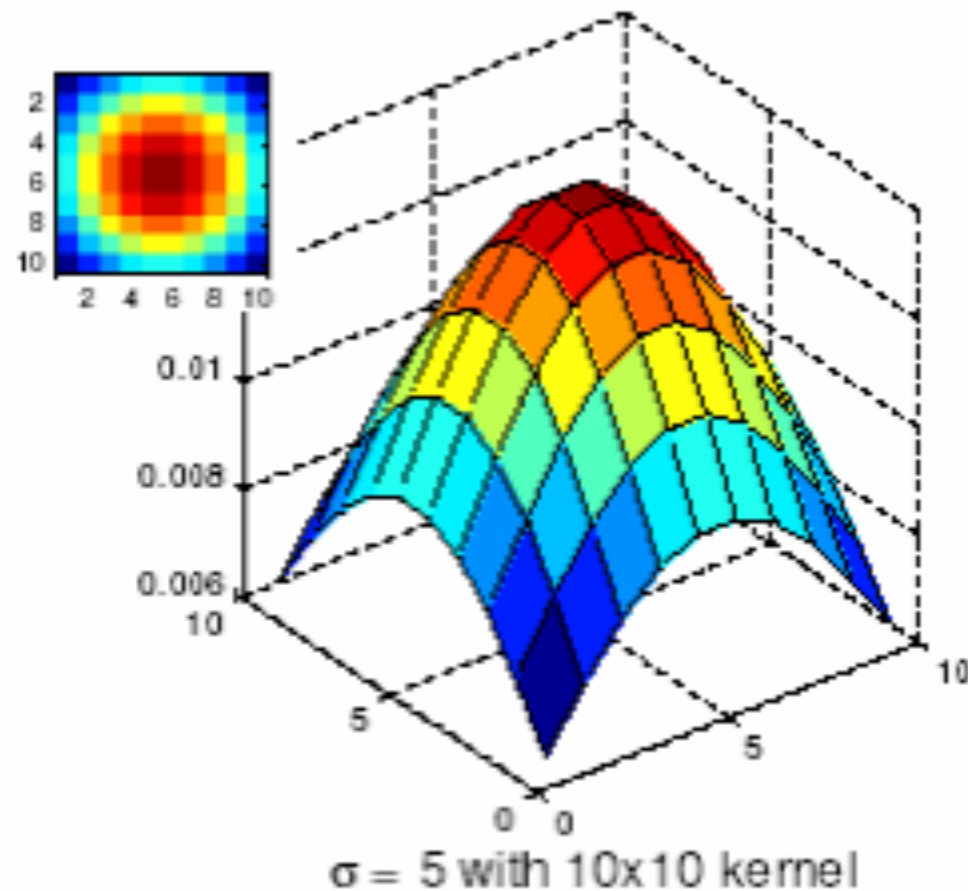


$\sigma = 5$ with 30 x 30
kernel

- Standard deviation σ : determines extent of smoothing


Choosing kernel width

- The Gaussian function has infinite support, but discrete filters use finite kernels





Choosing kernel width

Rule of thumb: set filter half-width to about 3σ



Gaussian vs. box filtering



Gaussian filters

- Convolution with self is another Gaussian
 - So can smooth with small- σ kernel, repeat, and get same result as larger- σ kernel would have
 - Convoluting two times with Gaussian kernel with std. dev. σ is same as convoluting once with kernel with std. dev. $\sigma\sqrt{2}$
- Separable kernel
 - Factors into product of two 1D Gaussians
 - Discrete example:

$$\begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix} = \begin{bmatrix} 1 \\ 2 \\ 1 \end{bmatrix} \begin{bmatrix} 1 & 2 & 1 \end{bmatrix}$$

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

Why is separability useful?

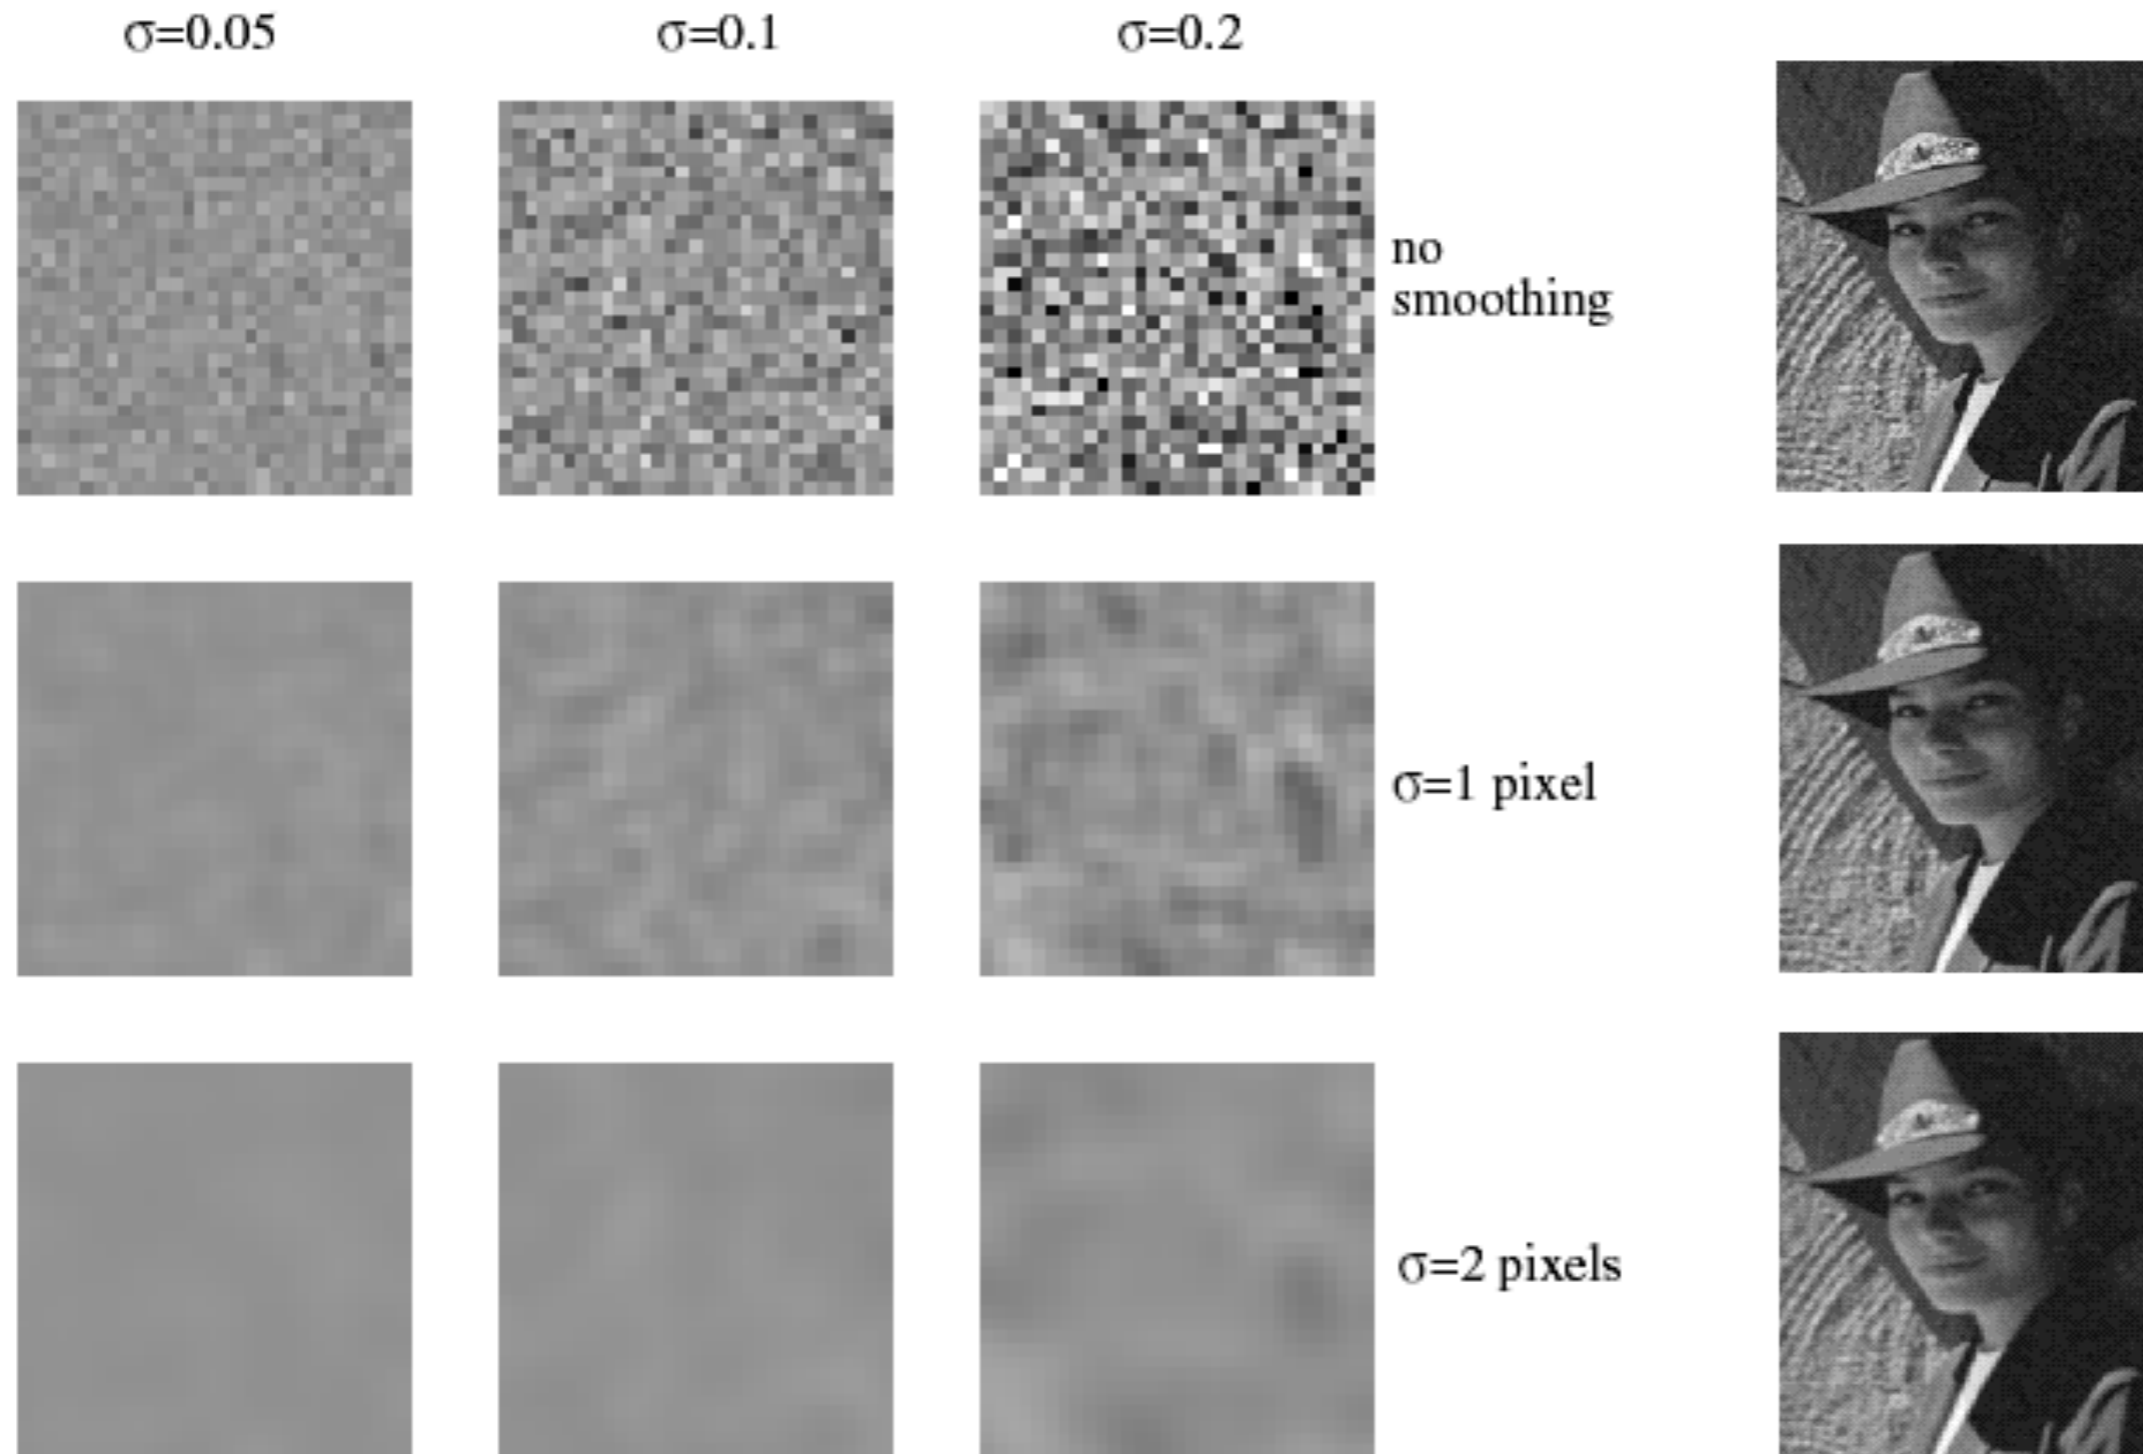
- Separability means that a 2D convolution can be reduced to two 1D convolutions (one along rows and one along columns)
- What is the complexity of filtering an $n \times n$ image with an $m \times m$ kernel?
 - $O(n^2 m^2)$
- What if the kernel is separable?
 - $O(n^2 m)$

Noise



- **Salt and pepper noise:** contains random occurrences of black and white pixels
- **Impulse noise:** contains random occurrences of white pixels
- **Gaussian noise:** variations in intensity drawn from a Gaussian normal distribution

Reducing Gaussian noise



Smoothing with larger standard deviations suppresses noise, but also blurs the image

Reducing salt-and-pepper noise

What's wrong with the results?

3x3



5x5




7x7



Alternative idea: Median filtering

- A **median filter** operates over a window by selecting the median intensity in the window



Median filter


- Is median filtering linear?
- Let's try filtering

$$\begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 2 \\ 2 & 2 & 2 \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

Median filter

- What advantage does median filtering have over Gaussian filtering?
 - Robustness to outliers

filters have width 5 :



Gaussian vs. median filtering

3x3

5x5

7x7

Gaussian



Median




Other non-linear filters

- Weighted median (pixels further from center count less)
- Clipped mean (average, ignoring few brightest and darkest pixels)
- Max or min filter
- Bilateral filtering: to avoid blurring edges, only average with similar intensity values.

$$I_{\mathbf{p}}^b = \frac{1}{W_{\mathbf{p}}^b} \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_s}(\|\mathbf{p} - \mathbf{q}\|) G_{\sigma_r}(|I_{\mathbf{p}} - I_{\mathbf{q}}|) I_{\mathbf{q}}$$

$$\text{with } W_{\mathbf{p}}^b = \sum_{\mathbf{q} \in \mathcal{S}} G_{\sigma_s}(\|\mathbf{p} - \mathbf{q}\|) G_{\sigma_r}(|I_{\mathbf{p}} - I_{\mathbf{q}}|)$$

Bilateral filters




Border effects

- What about near the edge? need to extrapolate!
 - Methods:
 - clip filter (black)
 - wrap around
 - copy edge
 - reflect across edge



Images in frequency domain


- Wide range of applications: image analysis, image filtering, image reconstruction, and image compression.



A. Oliva, A. Torralba, P.G. Schyns,
“Hybrid Images,” SIGGRAPH 2006

Fourier analysis

- Joseph Fourier: “Any function is a weighted combination of sines and cosines.”



<https://youtu.be/-qgreAUpPwM?t=302>

The Fourier transform

- Continuous transform:

$$f: \mathbb{R} \rightarrow \mathbb{R}, \quad \hat{f}(\xi) = \int_{-\infty}^{+\infty} f(x) e^{-ix\xi} dx.$$

- Discrete transform:

$$f: [0, N-1] \rightarrow \mathbb{R}, \quad \hat{f}(k) = \frac{1}{N} \sum_{m=0}^{N-1} f(m) e^{-\frac{2\pi i}{N} km}.$$

- Intuition: \hat{f} collects coefficients in the representation of f in Fourier basis

$$f(x) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \hat{f}(\xi) e^{i\xi x} d\xi, \quad f(m) = \sum_{k=0}^{N-1} \hat{f}(k) e^{\frac{2\pi i}{N} km}.$$

Some useful properties

- f discrete $\Leftrightarrow \hat{f}$ periodic
- \hat{f} real $\Leftrightarrow f(x) = f(-x)$
- Convolution Theorem: $(f \star g) = \hat{f} \cdot \hat{g}$
- Differentiation: $(\hat{f}')(\xi) = i\xi \hat{f}(\xi)$
- Energy conservation: $\int |f(x)|^2 dx = \int |\hat{f}(\xi)|^2 d\xi$

Discrete 2D transform

- Transforms an $M \times N$ pixel grid $f(x, y)$ into an $M \times N$ grid of complex numbers $H(k_x, k_y)$:


$$H(k_x, k_y) = \frac{1}{MN} \sum_{x=0}^{M-1} \sum_{y=0}^{N-1} f(x, y) e^{-i2\pi \frac{2\pi}{MN} (k_x x + k_y y)}$$

- Usually, the log-magnitude at every pixel is usually plotted.
- Pixels in the center of the transform correspond to low frequencies/long wavelengths.

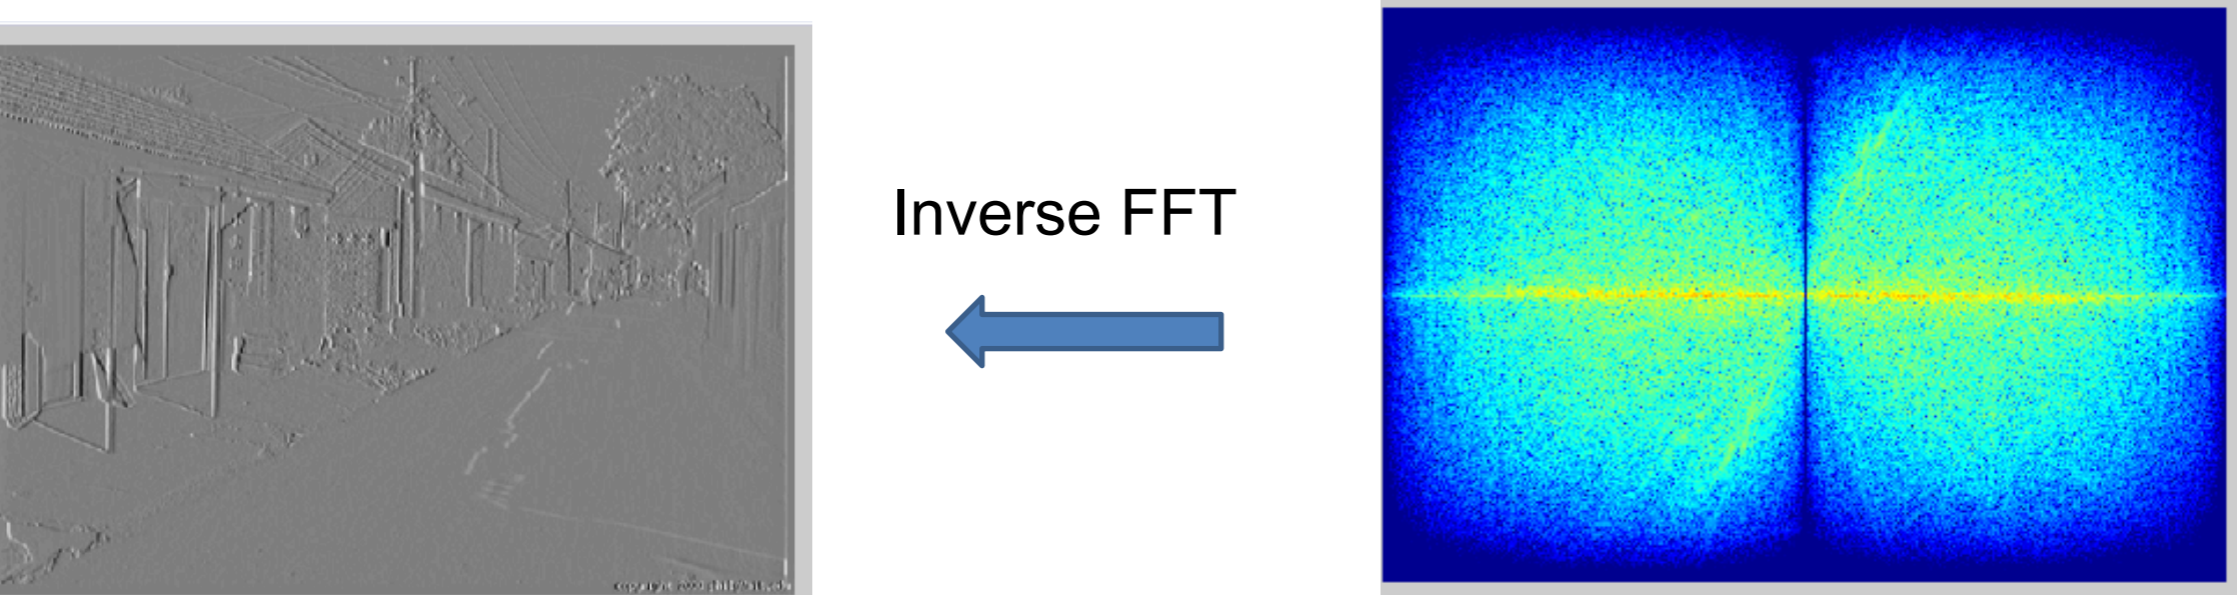
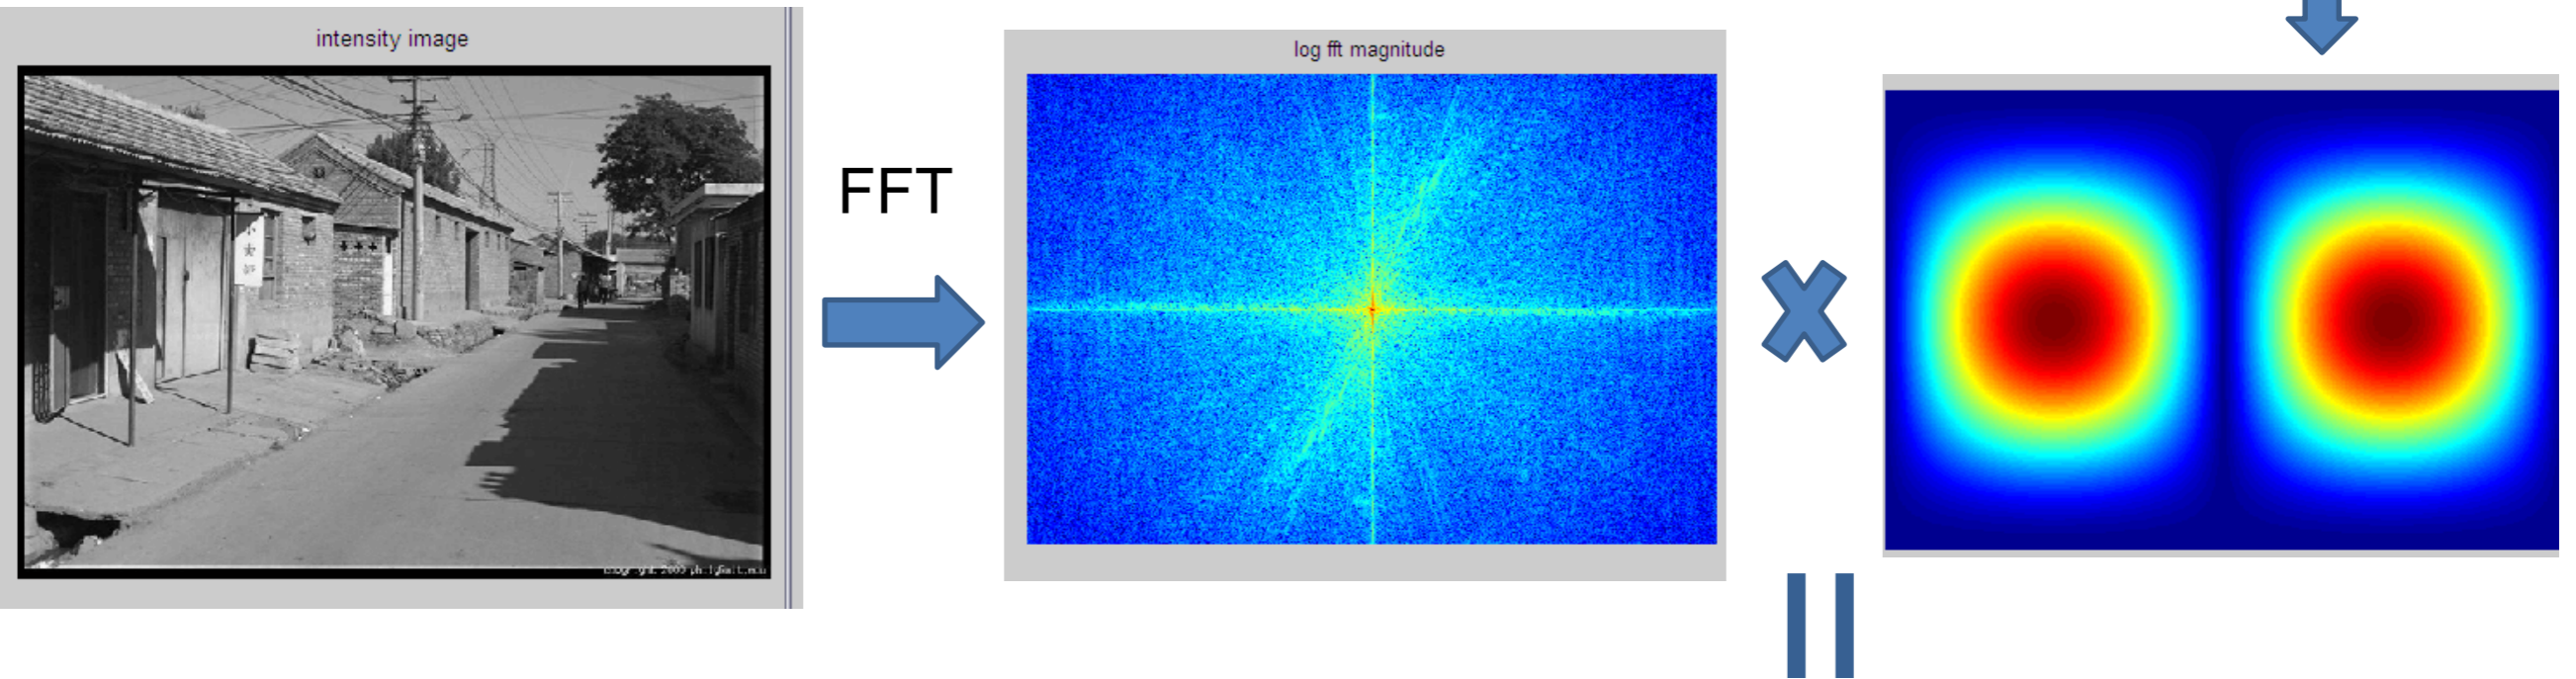
Filtering in spatial domain

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


intensity image



Filtering in frequency domain



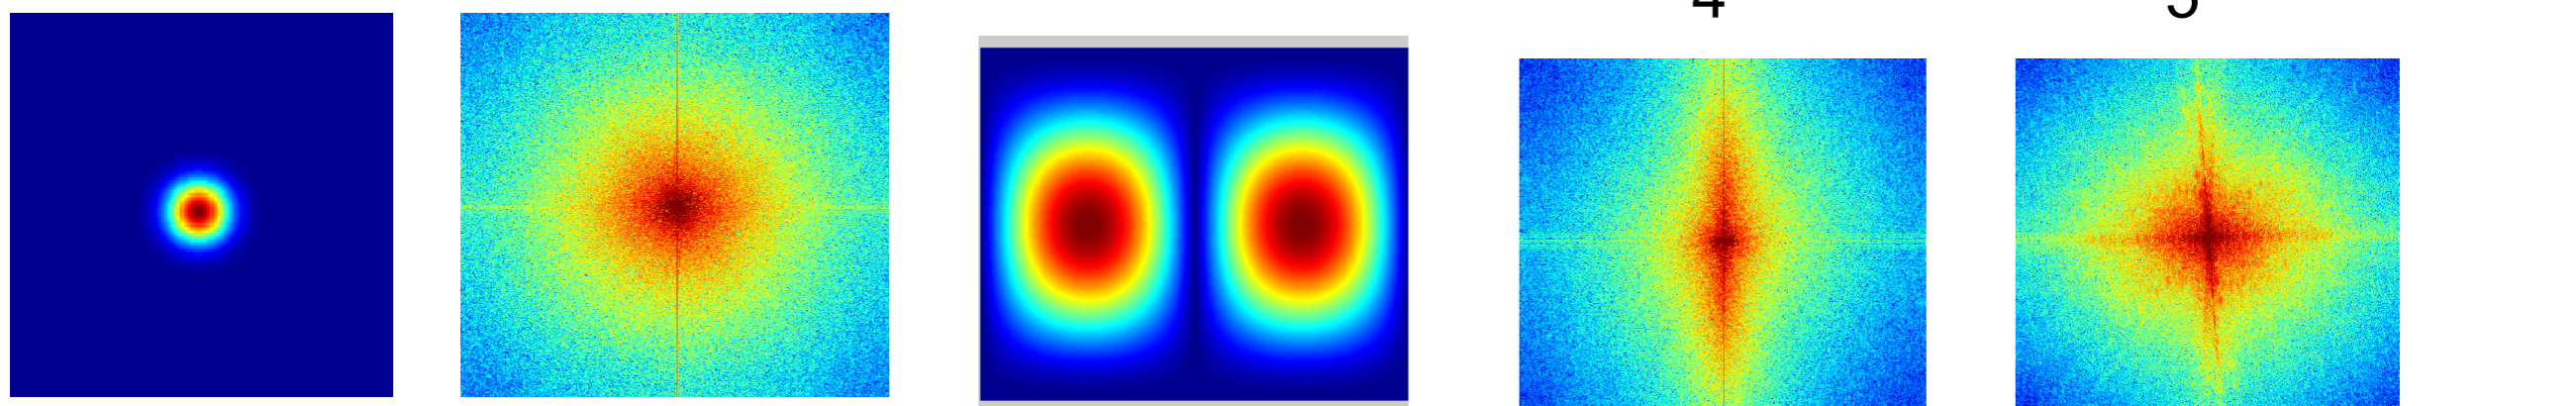
Fourier transform of a scene



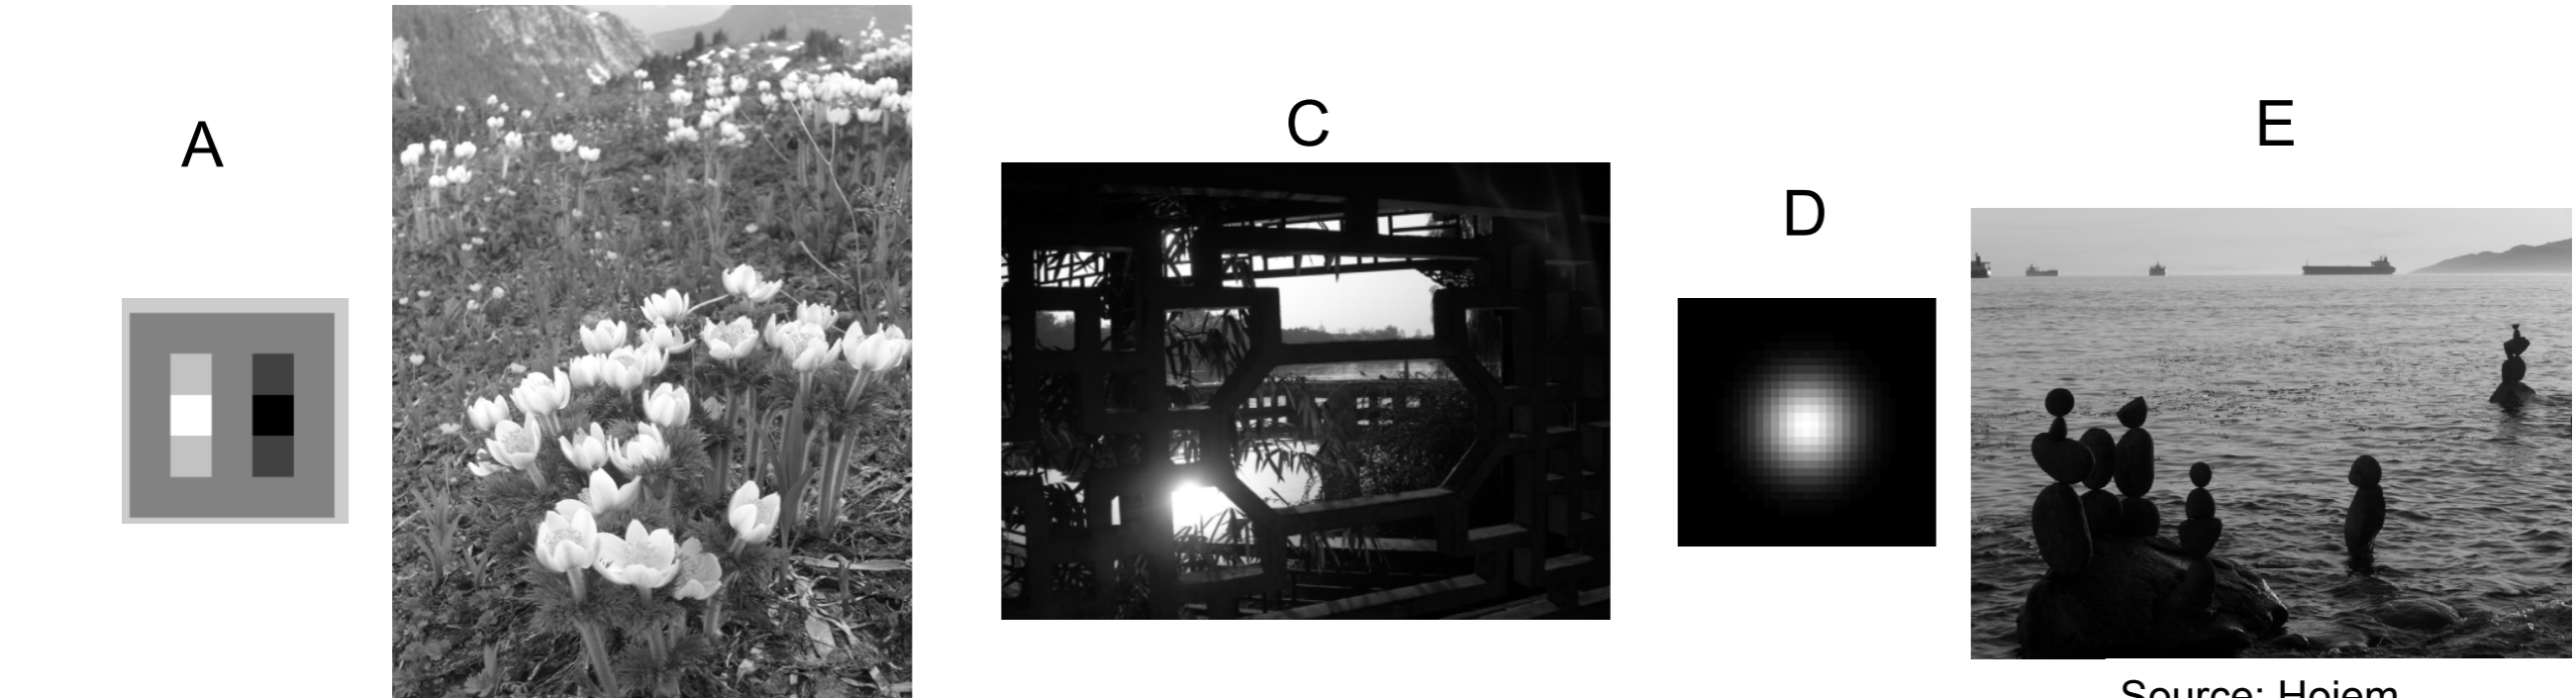
Question

1. Match the spatial domain image to the Fourier magnitude image



1 2 3 4 5



A B C D E




Low and high pass filtering



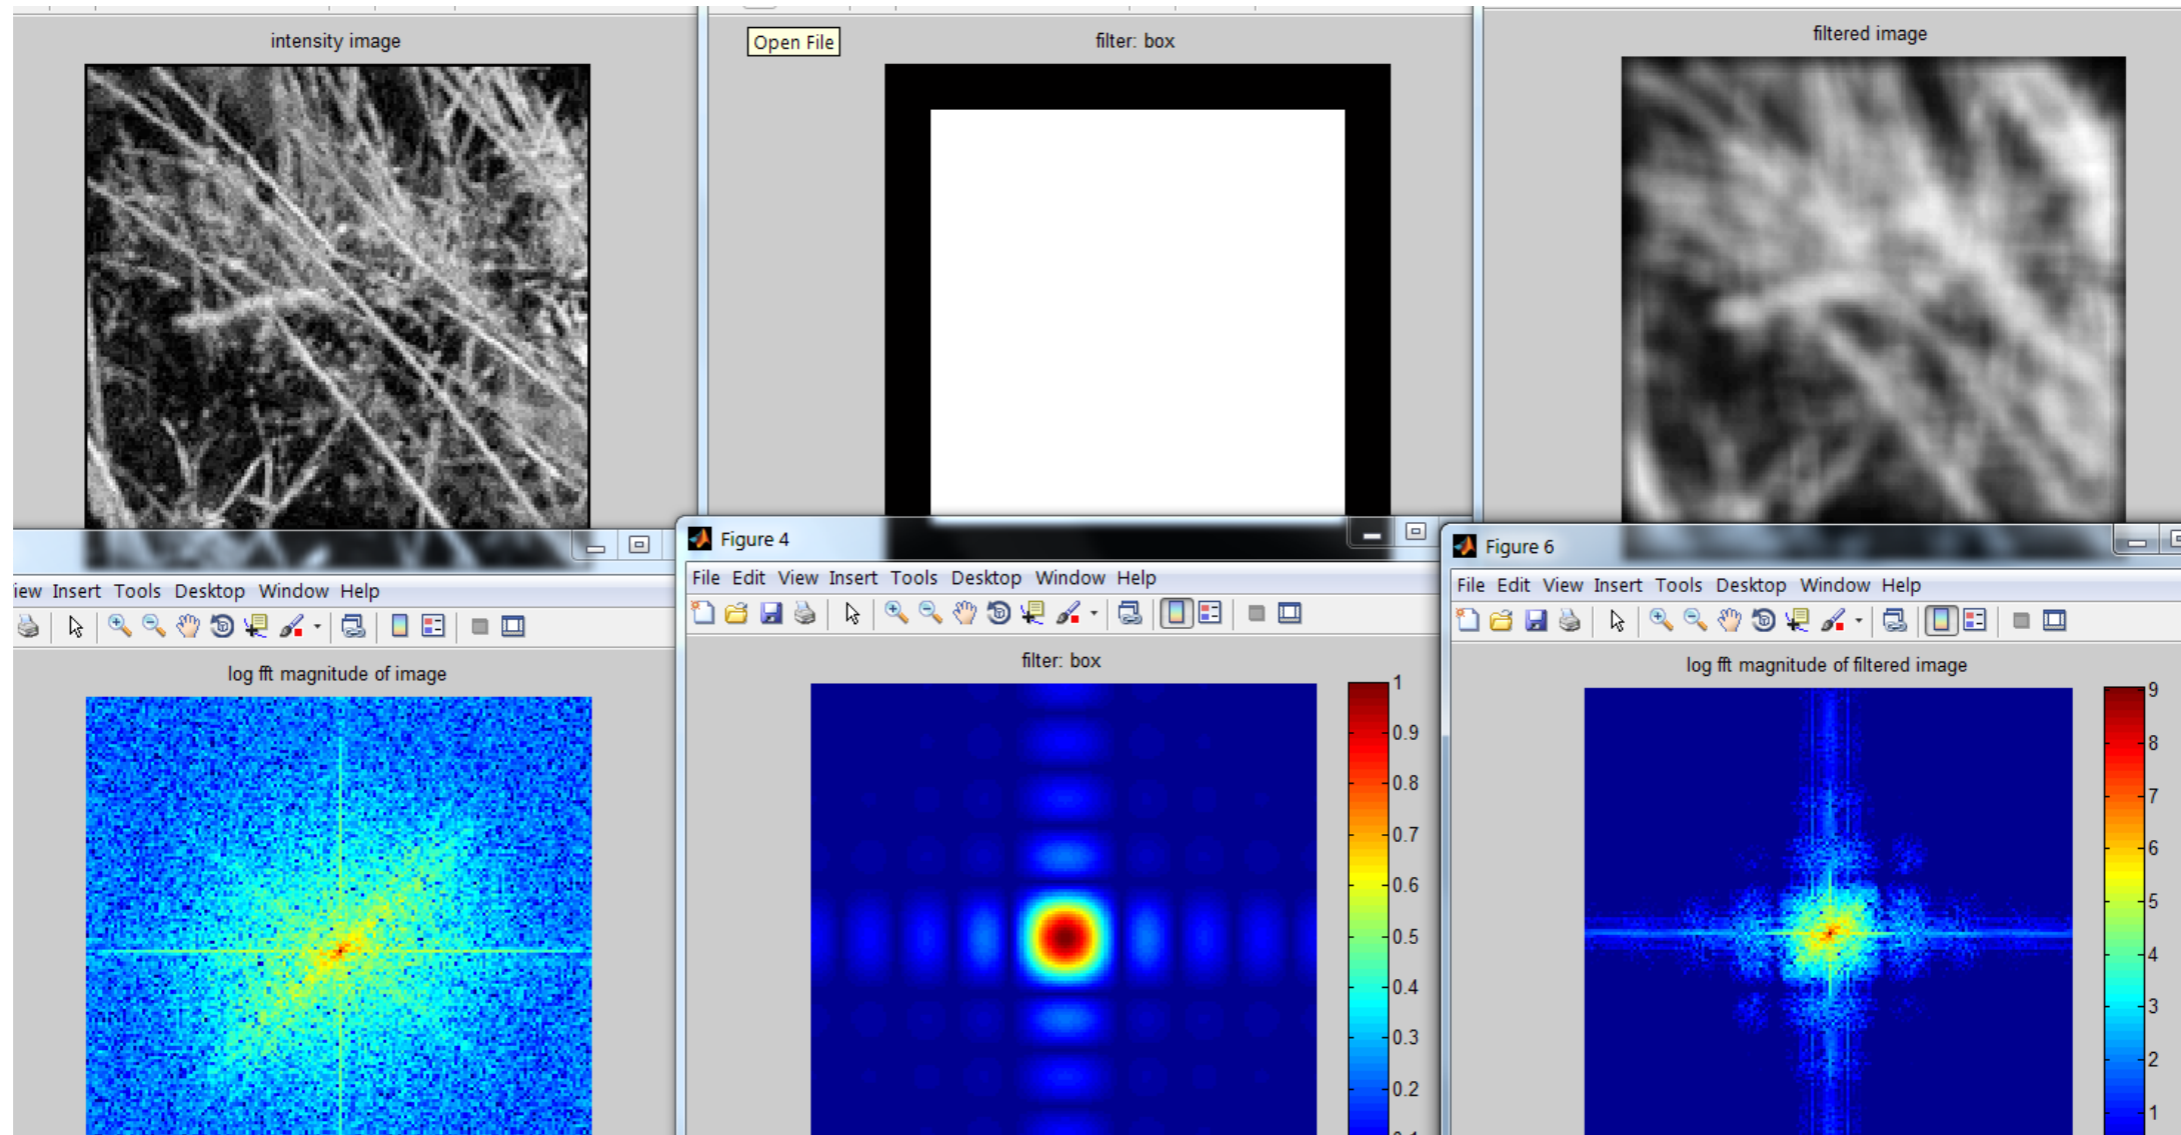
Filters in frequency domain

Gaussian




Filters in frequency domain

Box




Subsampling by a factor of 2



Throw away every other row and column to create a 1/2 size image

Aliasing




- Sub-sampling may be dangerous!
- Characteristic errors may appear:
 - Wagon wheels rolling the wrong way in movies
 - Checkerboards disintegrate in ray tracing
 - Striped shirts look funny on color television

Aliasing in video

Imagine a spoked wheel moving to the right (rotating clockwise).
Mark wheel with dot so we can see what's happening.


If camera shutter is only open for a fraction of a frame time (frame time = 1/30 sec. for video, 1/24 sec. for film):



Without dot, wheel appears to be rotating slowly backwards!
(counterclockwise)

How to sample

- When sampling a signal at discrete intervals, the sampling frequency must be $\geq 2 \times f_{\max}$
- f_{\max} = max frequency of the input signal
- This will allow to reconstruct the original perfectly from the sampled version



Anti-aliasing

Solutions:

- Sample more often
- Get rid of all frequencies that are greater than half the new sampling frequency
 - Will lose information
 - But it's better than aliasing
 - Apply a smoothing filter


Summary

- Filters are useful tools for manipulating images (denoising, sharpening, etc.)
- Spacial domain: linear filters (box, Gaussian), median filter, bilateral filter.
- Frequency domain: high and low pass filtering, aliasing.

Edge detection

- **Goal:** Identify sudden changes (discontinuities) in an image
- Intuitively, edges carry most of the semantic and shape information from the image

Edges are caused by a variety of factors



Edge detection

- **Ideal:** artist's line drawing




- **Reality:**



Closeup of edges



Closeup of edges



Closeup of edges




Closeup of edges




Edge detection

- An edge is a place of rapid change in the image intensity function


image



intensity function
(along horizontal scanline)





first derivative



edges correspond to
extrema of derivative

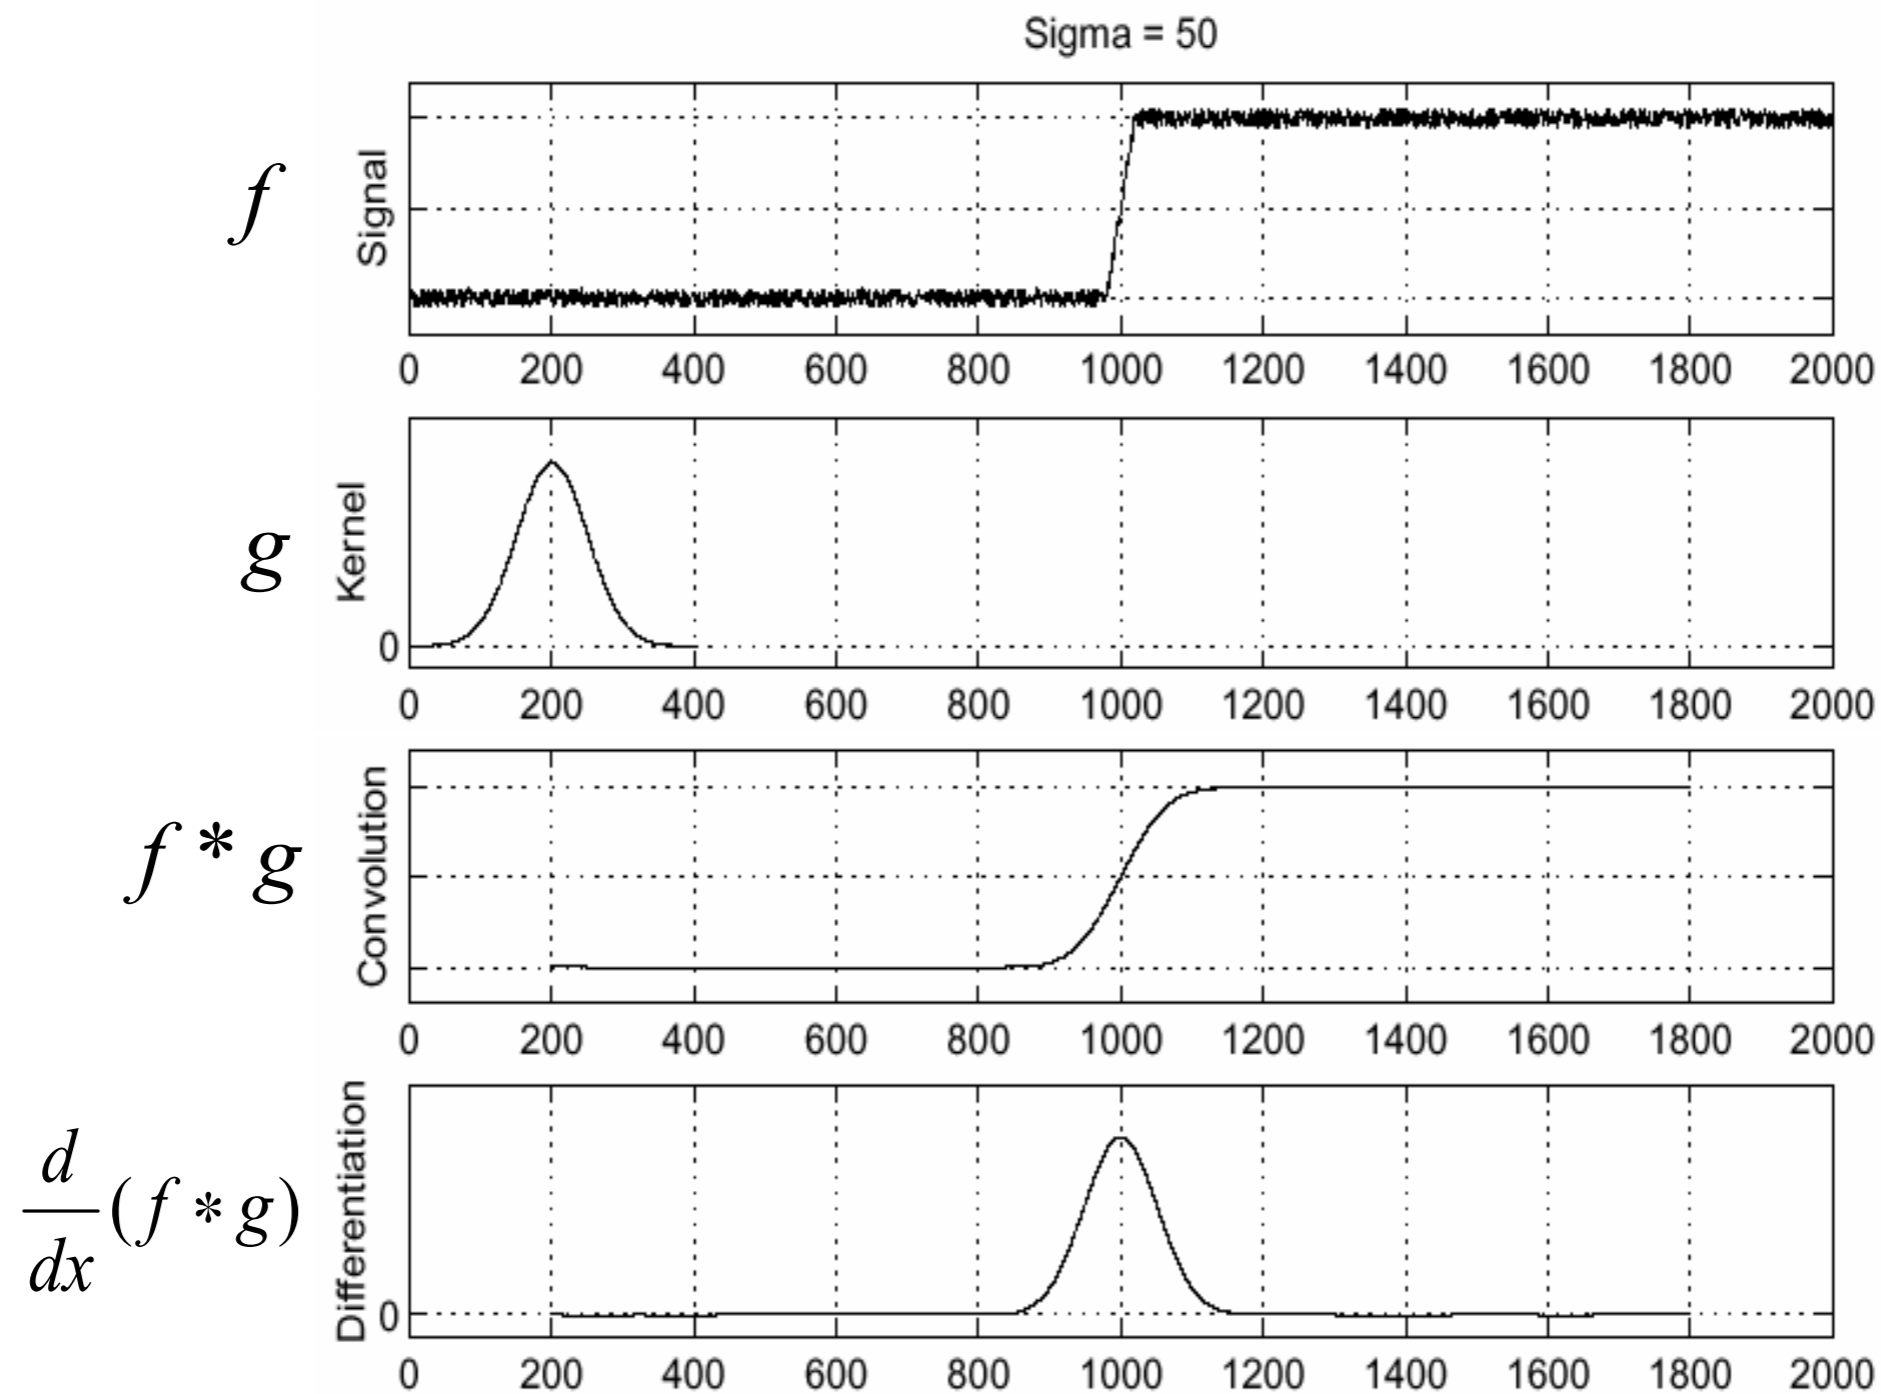
Effects of noise

Consider a single row or column of the image



Where is the edge?


Solution: smooth first




- To find edges, look for peaks in $\frac{d}{dx}(f * g)$

Derivative theorem of convolution


- Differentiation is convolution, and convolution is associative:
$$\frac{d}{dx}(f * g) = f * \frac{d}{dx}g$$
- This saves us one operation:



Derivative of Gaussian filter




$$* [1 \ 0 \ -1] =$$



- Is this filter separable?

Tradeoff between smoothing and localization



1 pixel

3 pixels

7 pixels

- Smoothed derivative removes noise, but blurs edge. Also finds edges at different “scales”.

Designing an edge detector


- Criteria for a good edge detector:
 - **Good detection:** find all real edges, ignoring noise or other artifacts
 - **Good localization**
 - detect edges as close as possible to the true edges
 - return one point only for each true edge point
- Cues of edge detection
 - Differences in color, intensity, or texture across the boundary
 - Continuity and closure
 - High-level knowledge

Canny edge detector



- This is probably the most widely used edge detector in computer vision
- Theoretical model: step-edges corrupted by additive Gaussian noise
- Canny has shown that the first derivative of the Gaussian closely approximates the operator that optimizes the product of *signal-to-noise ratio* and localization

J. Canny, [**A Computational Approach To Edge Detection**](#), IEEE Trans. Pattern Analysis and Machine Intelligence, 8:679-714, 1986.


Derivative of Gaussian filters



x-direction



y-direction




Compute Gradients (DoG)



X-Derivative of Gaussian



Y-Derivative of Gaussian



Gradient Magnitude

Building an edge detector





How to turn these thick regions of the gradient into curves?

Thresholded norm of the gradient

Non-maximum suppression

- For each location q above threshold, check that the gradient magnitude is higher than at neighbors p and r along the direction of the gradient
- May need to interpolate to get the magnitudes at p and r



Before Non-max Suppression




After non-max suppression

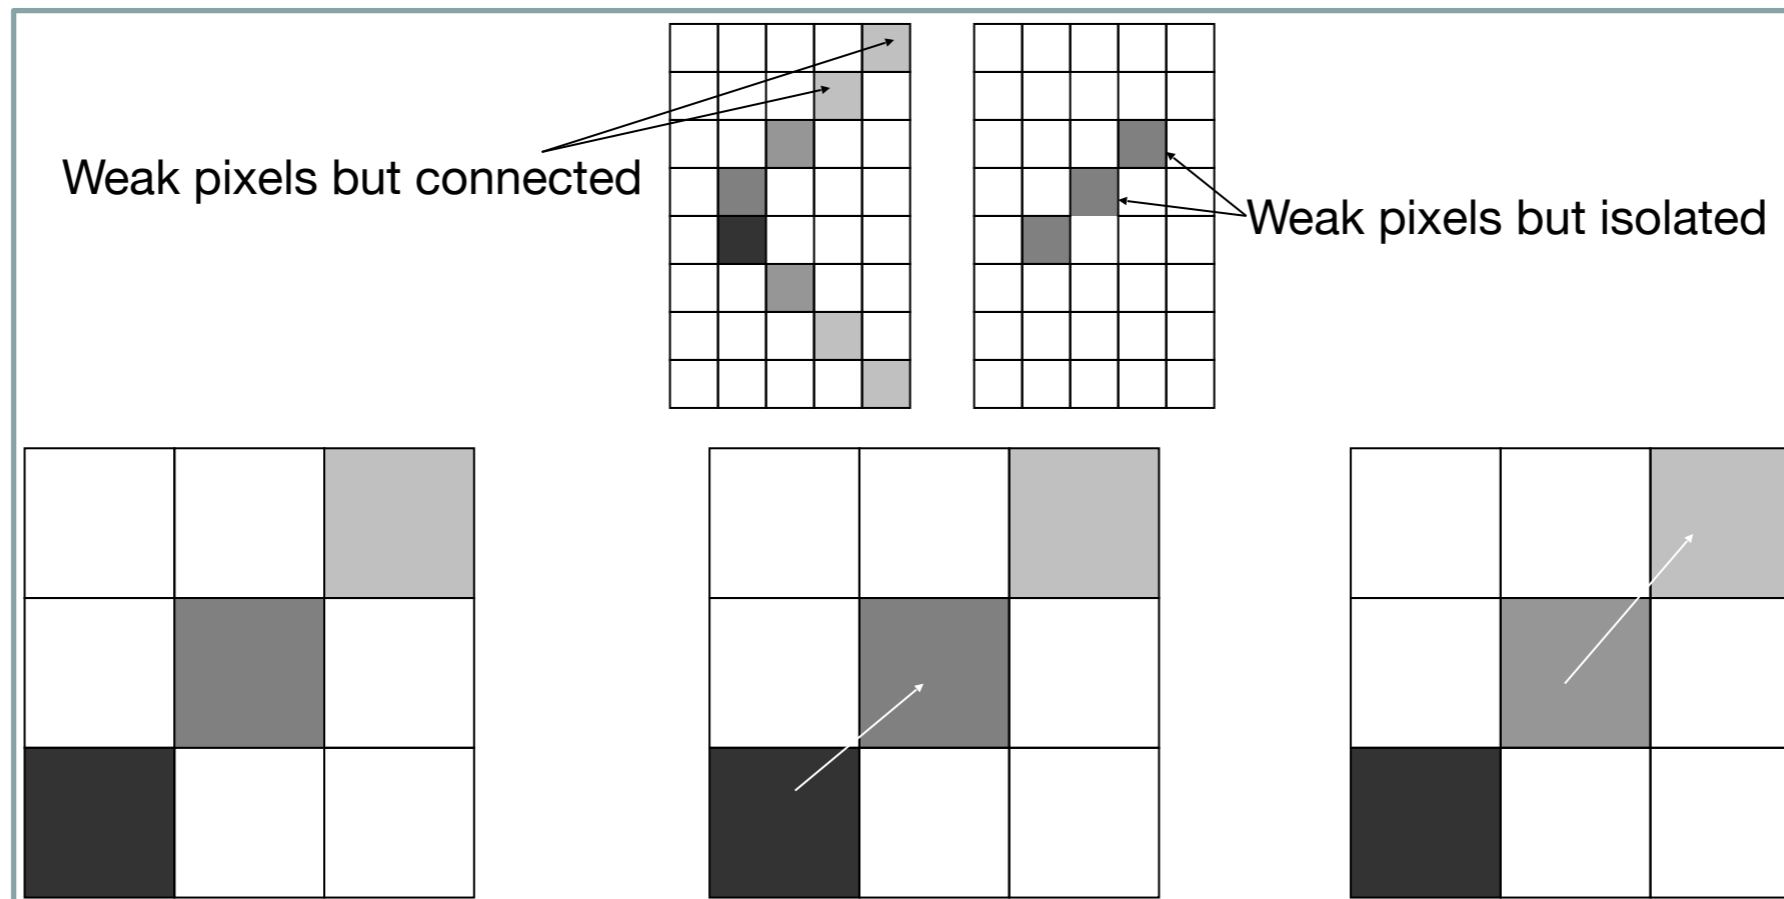


Hysteresis thresholding

- Check that maximum value of gradient value is sufficiently large
 - drop-outs? use **hysteresis**
 - use a high threshold to start edge curves and a low threshold to continue them.



Canny edge detector



Very strong edge response.
Let's start here

Weaker response but it is
connected to a confirmed
edge point. Let's keep it.

Continue...

Final Canny Edges



Effect of σ (Gaussian kernel spread/size)



original

Canny with $\sigma = 1$

Canny with $\sigma = 2$

The choice of σ depends on desired behavior

- large σ detects large scale edges
- small σ detects fine features

Canny edge detector

1. Filter image with x , y derivatives of Gaussian
2. Find magnitude and orientation of gradient
3. Non-maximum suppression:
 - Thin multi-pixel wide “ridges” down to single pixel width
4. Thresholding and linking (hysteresis):
 - Define two thresholds: low and high
 - Use the high threshold to start edge curves and the low threshold to continue them