

Computer Vision

ECE/CSE 576

Filters

Linda Shapiro

Professor of Computer Science & Engineering
Professor of Electrical & Computer Engineering

Let's do something interesting already!!

Want to make image smaller



448x448 -> 64x64



448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64




448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64




448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64




448x448 -> 64x64




448x448 -> 64x64





448x448 -> 64x64





448x448 -> 64x64




448x448 -> 64x64



448x448 -> 64x64



IS THIS ALL THERE IS??




THERE IS A BETTER WAY!



LOOK AT HOW MUCH BETTER




How do?




How do? Averaging!




How do? Averaging!



“interpolation




averaging



What is averaging?

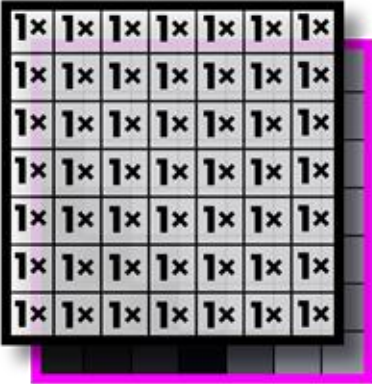


What is averaging? A weighted sum



What is averaging? A weighted sum

sum $\left[\frac{1}{49} \right]$




The diagram illustrates a uniform weight matrix for averaging. It consists of a 7x7 grid of 49 cells, each containing the text "1x". This grid is enclosed in large, stylized square brackets. To the left of the grid, the word "sum" is written in a bold, pink, sans-serif font. To the right of the grid, the fraction $\frac{1}{49}$ is written in a pink, sans-serif font, indicating that each element in the grid is multiplied by this weight to calculate the average.

Call this operation “*convolution*”

Filter or kernel

$\frac{1}{49}$

1×	1×	1×	1×	1×	1×	1×
1×	1×	1×	1×	1×	1×	1×
1×	1×	1×	1×	1×	1×	1×
1×	1×	1×	1×	1×	1×	1×
1×	1×	1×	1×	1×	1×	1×
1×	1×	1×	1×	1×	1×	1×
1×	1×	1×	1×	1×	1×	1×



Note: multiplying an image section by a filter is actually called “correlation” and convolution involves inverting the filter first, but since our filters are generally symmetric, we call everything convolution. This is what all computer vision people do.


Convolutions on larger images

$\frac{1}{49}$

1×	1×	1×	1×	1×	1×	1×
1×	1×	1×	1×	1×	1×	1×
1×	1×	1×	1×	1×	1×	1×
1×	1×	1×	1×	1×	1×	1×
1×	1×	1×	1×	1×	1×	1×
1×	1×	1×	1×	1×	1×	1×
1×	1×	1×	1×	1×	1×	1×



Kernel slides across image



Kernel slides across image


What happens at the edges of the images?

1. Zero padding: easiest but can give bad results
2. Duplicate the outermost row or column
3. Use wraparound

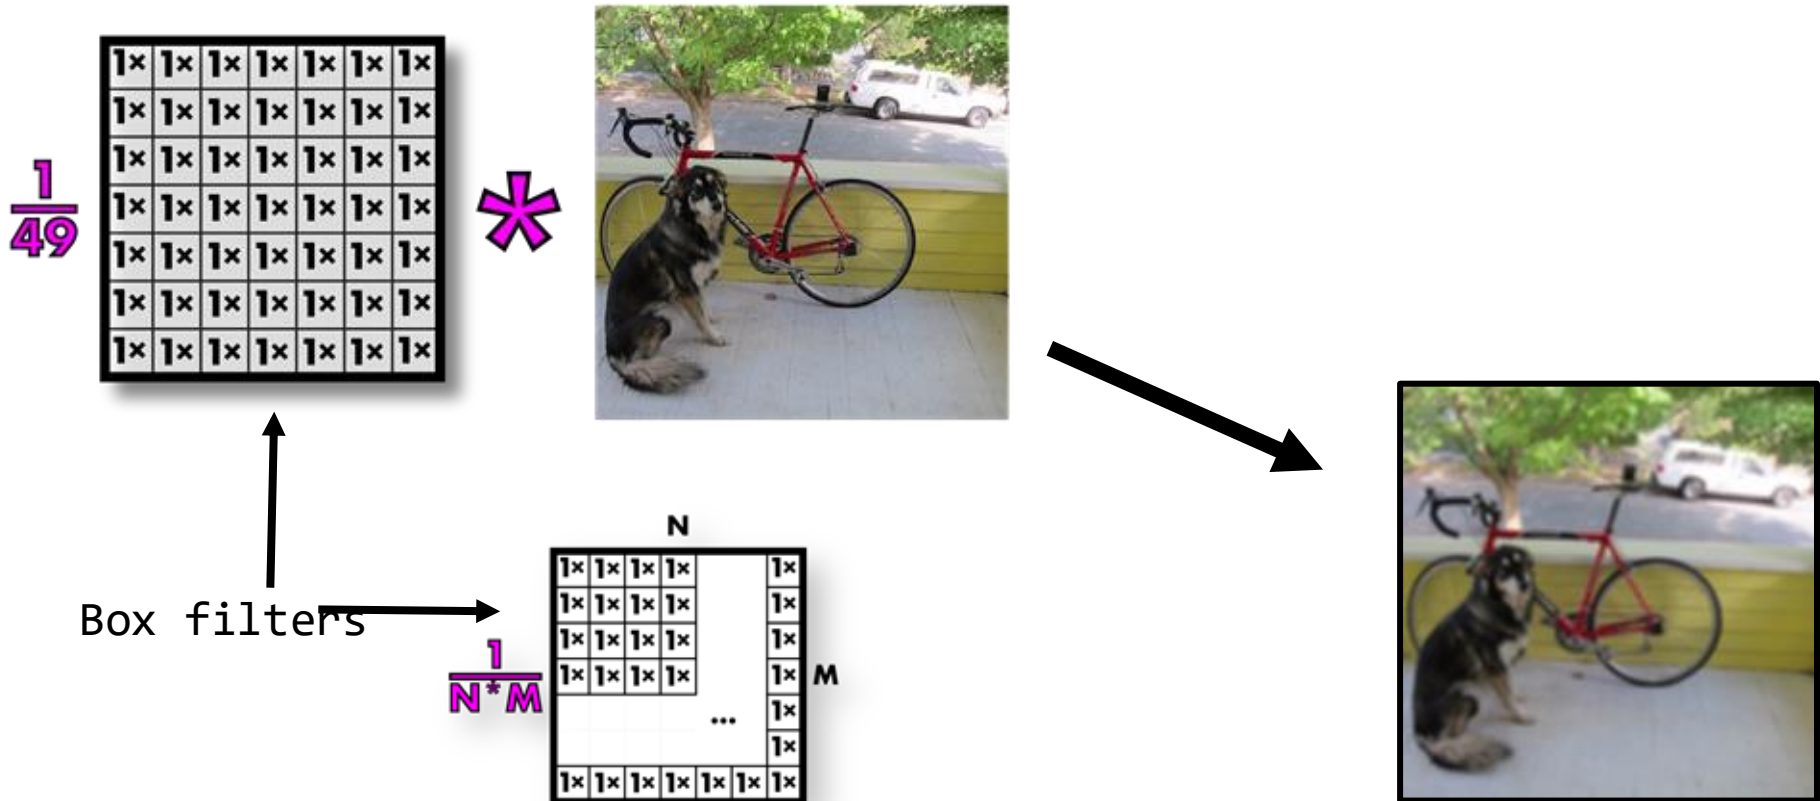
Convolutions on larger images

$\frac{1}{49}$

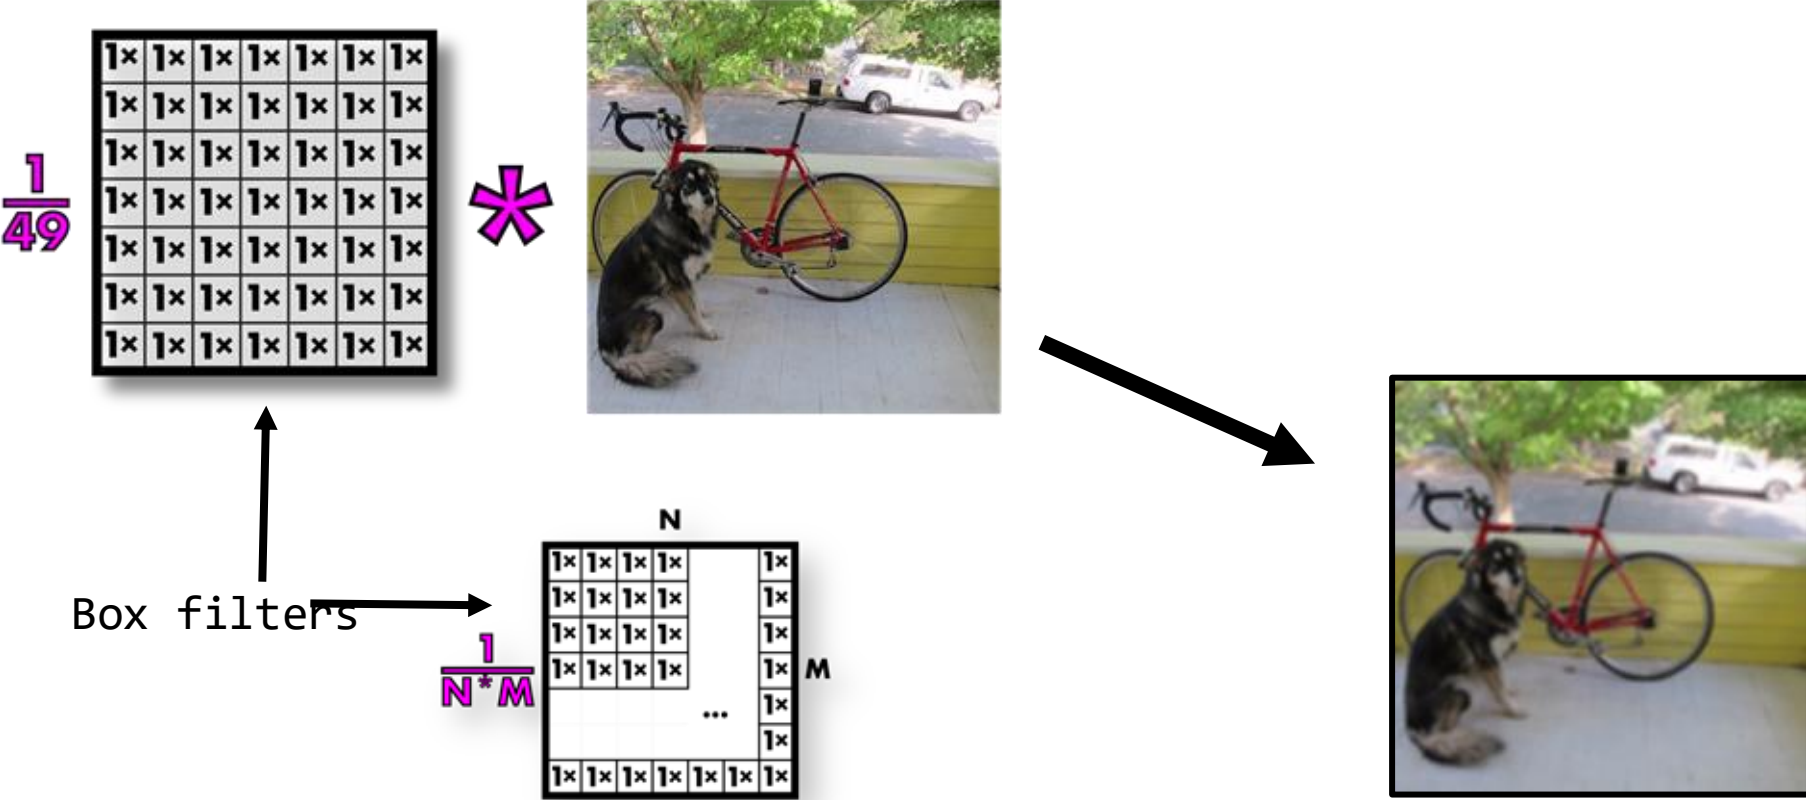
1x	1x	1x	1x	1x	1x	1x
1x	1x	1x	1x	1x	1x	1x
1x	1x	1x	1x	1x	1x	1x
1x	1x	1x	1x	1x	1x	1x
1x	1x	1x	1x	1x	1x	1x
1x	1x	1x	1x	1x	1x	1x
1x	1x	1x	1x	1x	1x	1x




This is called box filter




Box filters smooth image



Box filters smooth image




Now we resize our smoothed image



So much better!




Box filters have artifacts




Box filters have artifacts



We want a smoothly weighted kernel





Gaussians



2d Gaussian


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




Better smoothing with Gaussians



*



=



Better smoothing with Gaussians



Better smoothing with Gaussians




Box Filtered




Gaussian Filtered

Wow, so what was that convolution thing??





$$q = a \times r + b \times s + c \times t + d \times u + e \times v + f \times w + g \times x + h \times y + i \times z$$

Wow, so what was that convolution thing??



$$q = a \times r + b \times s + c \times t + d \times u + e \times v + f \times w + g \times x + h \times y + i \times z$$

Calculate it, go!




Calculate it, go!

-1	-1	-1					
-1	8	-1	...				
-1	-1	-1	16	105	153		
...			15	104	113	...	
			18	111	136		
			...				




			...				
...			q			...	
			...				

Guess that kernel!




Highpass Kernel: finds edges


(applied to the graytone image!)




Guess that kernel!




Identity Kernel: Does nothing!



Guess that kernel!




Sharpen Kernel: sharpens! (applied to all three bands)




Note: sharpen = highpass + identity!



Guess that kernel!







Emboss Kernel: stylin' (applied to all three bands)






Guess those kernels!



Sobel Kernels: edges (applied to a graytone image and thresholded)



Sobel Kernels: edges and gradient!



Sobel Kernels: edges and gradient!



This visualization is showing the magnitude and direction of the gradient. We will talk further about this when we discuss edges.

And so much more!!

Assignment 2

Image resizing and filtering

First things first!

- Remember that you might need some of your code from the previous hw (e.g. `set_pixel`) for this hw as well. Have your code from hw1 in your src folder.

Assignment 2

Nearest Neighbor Interpolation and Resizing

- `def nn_interpolate(image im, float x, float y, int c) -> float`
- `def nn_resize(image im, int w, int h) -> image`

Bilinear Interpolation and Resizing

- `def bilinear_interpolate(image im, float x, float y, int c) -> float`
- `def bilinear_resize(image im, int w, int h) -> image`

Assignment 2

Box Filter

- `def l1_normalize(image im) -> None`
- `def make_box_filter(int w) -> image`

Convolution

- `def convolve_image(image im, image filter, int preserve) -> image`
- `def make_highpass_filter() -> image`
- `def make_sharpen_filter() -> image`
- `def make_emboss_filter() -> image`

Assignment 2

Gaussian

- `def make_gaussian_filter(float sigma) -> image`

Hybrid Images

- `def add_image(image a, image b) -> image`
- `def sub_image(image a, image b) -> image`

Assignment 2

Sobel Operator (next lecture)

- `def make_gx_filter() -> image`
- `def make_gy_filter() -> image`
- `def feature_normalize(image im) -> None`
- `def *sobel_image(image im) -> image`
- `def colorize_sobel(image im) -> image`