


Cameras and Images

Pinhole Camera




[Source: A. Torralba]



- Make your own camera
- http://www.foundphotography.com/PhotoThoughts/archives/2005/04/pinhole_camera_2.html


Pinhole Camera




- How it works
- The pinhole camera only allows rays from one point in the scene to strike each point of the paper.

Pinhole Camera

. Torralba]



[Source: A. Torralba]



- Example

Pinhole Camera

[Source: A. Torralba]



- Example


[Source: A. Torralba]




- Example

Pinhole Camera


[Source: A. Torralba]



a) Input (occluder present)



b) Reference (occluder absent)



c) Difference image (b-a)



d) Crop upside down



e) True view

- Remember this example?
- In this case the window acts as a pinhole camera into the room

[Adopted from S. Seitz]



- A digital camera replaces film with a sensor array
- Each cell in the array is light-sensitive diode that converts photons to electrons
- <http://electronics.howstuffworks.com/cameras-photography/digital/digital-camera.htm>

Image Formation





Image formation process producing a particular image depends on:



- lighting conditions
- scene geometry
- surface properties
- camera optics

Digital Image

Continuous image projected to sensor array



Sampling and quantization



<http://pho.to/media/images/digital/digital-sensors.jpg>

- Sample the 2D space on a regular grid
- Quantize each sample (round to nearest integer)

Digital Image

- Image is a matrix with integer values
- We will typically denote it with I
- $I(i, j)$ is called **intensity**
- Matrix I can be $m \times n$ (grayscale)
- or $m \times n \times 3$ (color)





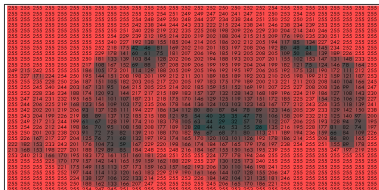
```

255 255 255 255 255 255 255 255 255 255 255 255 255 255 255 255 248 249 242 241 248 249 254 255 255 255 255 255 255 255
251 255 255 255 255 255 255 255 255 255 255 251 248 241 236 232 220 222 231 240 245 251 255 255 255 255 255 255
255 255 255 255 255 255 255 255 255 255 255 238 237 240 235 208 215 210 217 227 230 242 243 243 247 255 255 255 255
255 255 255 255 255 255 255 255 255 255 255 232 218 227 227 226 212 195 185 197 216 224 221 218 216 226 252 255 255 255 255
255 255 255 255 255 255 255 255 255 255 247 204 206 191 207 215 201 178 164 179 200 207 206 172 187 223 227 255 255 255 255
255 255 255 255 255 255 255 255 254 211 219 180 180 164 194 191 170 153 172 187 188 179 140 155 210 214 250 255 255 255
255 255 255 255 255 255 255 255 232 206 170 110 121 151 174 186 174 151 170 183 175 161 91 95 112 211 236 255 255 255
255 255 255 255 255 255 252 202 151 151 164 135 170 179 167 148 163 177 174 159 84 85 133 222 251 255 255 255
255 255 255 255 255 255 255 252 182 119 81 54 64 146 174 173 162 150 159 172 177 172 89 85 96 119 162 200 255 255
255 255 255 255 255 255 255 167 109 118 97 79 115 167 173 167 160 153 159 169 174 167 124 99 154 135 105 129 230 255
255 255 255 255 255 255 255 140 128 80 92 126 175 177 173 165 160 164 170 171 165 145 97 168 160 25 98 188 255
248 206 239 255 255 255 176 83 145 171 50 102 132 171 176 173 161 157 160 163 172 171 156 140 102 161 150 67 202 255
255 224 185 216 252 245 178 118 123 180 156 168 166 181 178 167 154 150 154 157 169 178 173 163 167 187 135 84 168 253
255 255 252 217 242 225 162 125 81 153 173 173 188 191 173 162 148 139 143 154 166 182 192 162 173 182 115 79 141 242
255 255 223 221 231 183 142 166 71 136 185 174 198 184 168 150 126 119 119 134 156 175 197 203 179 182 110 75 140 245
255 252 206 218 217 163 149 94 91 116 187 188 116 155 148 125 107 103 100 111 134 154 163 169 195 161 100 81 117 225
255 244 189 213 214 171 141 114 76 84 198 206 189 140 136 116 101 99 94 109 120 138 150 197 223 136 86 80 105 222
255 240 181 202 196 145 102 131 83 79 145 210 174 143 133 111 100 84 85 99 115 135 142 181 230 221 100 91 118 238
255 236 178 172 196 183 85 116 161 84 167 200 153 164 92 86 65 71 70 83 74 101 113 174 220 226 102 113 129 234
253 237 180 171 207 167 81 72 111 85 156 186 155 83 82 83 81 78 77 67 63 126 179 211 201 97 116 81 188
255 246 197 194 235 175 48 53 105 92 145 180 149 140 90 53 25 26 24 46 69 98 176 175 203 178 101 78 63 182
255 254 215 190 236 172 72 36 67 68 121 180 142 156 168 89 66 43 47 118 119 162 206 164 85 68 68 187
255 248 178 188 229 177 72 36 67 68 114 182 154 186 154 111 65 56 51 62 95 115 187 163 21 168 82 69 52 216
255 210 132 203 229 175 103 60 36 34 109 169 163 143 137 145 101 98 85 106 165 150 157 195 254 200 48 63 150 255
209 130 123 226 236 173 148 82 58 45 145 217 205 189 142 150 164 146 144 161 155 189 231 253 255 250 156 164 255
203 132 144 197 234 174 156 105 71 65 177 249 255 247 209 166 145 131 128 148 175 236 255 255 244 189 212 255
254 240 212 165 168 170 149 140 131 144 159 188 224 251 255 255 215 152 169 24 255 255 255 255 255 255 255 255
255 255 255 255 168 160 125 112 112 150 159 136 188 187 229 255 232 186 142 200 255 255 255 255 255 255 255 255
255 255 255 255 222 167 118 88 86 126 156 167 163 156 161 192 200 103 83 114 172 233 255 255 255 255 255 255 255
255 255 255 255 253 190 114 87 86 91 142 183 229 219 189 160 156 113 81 98 122 180 255 255 255 255 255 255 255
255 255 255 255 217 108 96 88 91 137 215 255 253 226 179 116 77 103 188 186 244 255 255 255 255 255 255 255 255
255 255 255 255 255 249 177 141 112 143 204 247 255 255 255 255 200 153 151 167 212 255 255 255 255 255 255 255
255 255 255 255 255 255 248 245 251 255 255 255 255 255 255 255 255 254 253 255 255 255 255 255 255 255 255



```

Digital Image

- Image is a matrix with integer values
- We will typically denote it with I
- $I(i, j)$ is called **intensity**
- Matrix I can be $m \times n$ (grayscale)
- or $m \times n \times 3$ (color)



Intensity



- We can think of a (grayscale) image as a function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ giving the intensity at position (i, j)
- Intensity 0 is black and 255 is white

Image Transformations

- As with any function, we can apply operators to an image, e.g.:



$I(i, j)$



$J(i, j) = I(i, j) + 50$



$J(i, j) = I(i, -j)$



$I(i, j) \cdot (I(i, j) < 250)$

- We'll talk about special kinds of operators, **correlation** and **convolution** (linear filtering)


[Adapted from: N. Snavely]

Linear Filters

Reading: Szeliski book, Chapter 3.2

Motivation: Finding Waldo

- How can we find Waldo?



[Source: R. Urtasun]

- Slide and compare!
- In formal language: **filtering**

Motivation: Noise reduction

- Given a camera and a still scene, how can you reduce noise?



[Source: S. Seitz]

Image Filtering

- Modify the pixels in an image based on some function of a local neighborhood of each pixel
- In other words... Filtering

10	5	3
4	5	1
1	1	7

Local image data

Some function



	7	

Modified image data


[Source: L. Zhang]

Applications of Filtering

- Detect patterns, e.g., **template matching**.
- Enhance an image, e.g., **denoise, resize**.
- Extract information, e.g., **texture, edges**.

Noise reduction


- Simplest thing: replace each pixel by the average of its neighbors.
- This assumes that neighboring pixels are similar, and the noise to be independent from pixel to pixel.



[Source: S. Marschner]

Noise reduction


- Simplest thing: replace each pixel by the average of its neighbors.
- This assumes that neighboring pixels are similar, and the noise to be independent from pixel to pixel.



[Source: S. Marschner]

Noise reduction


- Simplest thing: replace each pixel by the average of its neighbors
- This assumes that neighboring pixels are similar, and the noise to be independent from pixel to pixel.
- **Moving average** in 1D: $[1, 1, 1, 1, 1]/5$



[Source: S. Marschner]

Noise reduction


- Simplest thing: replace each pixel by the average of its neighbors
- This assumes that neighboring pixels are similar, and the noise to be independent from pixel to pixel.
- Non-uniform weights $[1, 4, 6, 4, 1] / 16$




[Source: S. Marschner]

Moving Average in 2D

$I(i, j)$




$G(i, j)$




[Source: S. Seitz]

Moving Average in 2D

$$I(i, j)$$



$$G(i, j)$$



[Source: S. Seitz]

Moving Average in 2D

$$I(i, j)$$

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

$$G(i, j)$$

	0	10	20							

[Source: S. Seitz]

Moving Average in 2D

$I(i, j)$

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

$G(i, j)$

	0	10	20	30					

[Source: S. Seitz]

Moving Average in 2D

 $I(i, j)$

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

 $G(i, j)$

	0	10	20	30	30					

[Source: S. Seitz]

Moving Average in 2D

$I(i, j)$

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

$G(i, j)$

	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		

[Source: S. Seitz]

Linear Filtering: Correlation

- Involves weighted combinations of pixels in small neighborhoods:

$$G(i,j) = \frac{1}{(2k+1)^2} \sum_{u=-k}^k \sum_{v=-k}^k I(i+u, j+v)$$

- The output pixels value is determined as a weighted sum of input pixel values

$$G(i,j) = \sum_{u=-k}^k \sum_{v=-k}^k F(u,v) \cdot I(i+u, j+v)$$

Linear Filtering: Correlation

- Involves weighted combinations of pixels in small neighborhoods:

$$G(i, j) = \frac{1}{(2k + 1)^2} \sum_{u=-k}^k \sum_{v=-k}^k I(i + u, j + v)$$

- The output pixels value is determined as a weighted sum of input pixel values

$$G(i, j) = \sum_{u=-k}^k \sum_{v=-k}^k F(u, v) \cdot I(i + u, j + v)$$

- The entries of the weight **kernel** or **mask** $F(u, v)$ are often called the **filter coefficients**.

Linear Filtering: Correlation

- Involves weighted combinations of pixels in small neighborhoods:

$$G(i, j) = \frac{1}{(2k + 1)^2} \sum_{u=-k}^k \sum_{v=-k}^k I(i + u, j + v)$$

- The output pixels value is determined as a weighted sum of input pixel values

$$G(i, j) = \sum_{u=-k}^k \sum_{v=-k}^k F(u, v) \cdot I(i + u, j + v)$$

- The entries of the weight **kernel** or **mask** $F(u, v)$ are often called the **filter coefficients**.
- This operator is the **correlation** operator

$$G = F \otimes I$$

Linear Filtering: Correlation

- Involves weighted combinations of pixels in small neighborhoods:

$$G(i, j) = \frac{1}{(2k + 1)^2} \sum_{u=-k}^k \sum_{v=-k}^k I(i + u, j + v)$$

- The output pixels value is determined as a weighted sum of input pixel values



$$G(i, j) = \sum_{u=-k}^k \sum_{v=-k}^k F(u, v) \cdot I(i + u, j + v)$$

- The entries of the weight **kernel** or **mask** $F(u, v)$ are often called the **filter coefficients**.
- This operator is the **correlation** operator

$$G = F \otimes I$$

Linear Filtering: Correlation



- It's really easy!



filter F



Linear Filtering: Correlation

- It's really easy!





Linear Filtering: Correlation

- It's really easy!



Linear Filtering: Correlation

- It's really easy!



$$G(i, j) = \sum_{u=-k}^k \sum_{v=-k}^k F(u, v) \cdot I(i + u, j + v)$$

$$G(i, j) = F(\text{cyan}) \cdot I(\text{cyan}) + F(\text{yellow}) \cdot I(\text{yellow}) + F(\text{orange}) \cdot I(\text{orange}) + \dots + F(\text{light blue}) \cdot I(\text{light blue})$$

Example

- What's the result?



Original

0	0	0
0	1	0
0	0	0

?

[Source: D. Lowe]

Example

- What's the result?



Original

0	0	0
0	1	0
0	0	0



**Filtered
(no change)**

[Source: D. Lowe]

Example

- What's the result?



Original

0	0	0
0	0	1
0	0	0

?

[Source: D. Lowe]

Example

- What's the result?



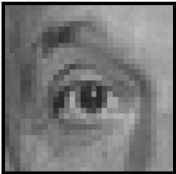
0	0	0
0	0	1
0	0	0



[Source: D. Lowe]

Example

- What's the result?




Original

$$* \left(\begin{array}{|c|c|c|} \hline 0 & 0 & 0 \\ \hline 0 & 2 & 0 \\ \hline 0 & 0 & 0 \\ \hline \end{array} - \frac{1}{9} \begin{array}{|c|c|c|} \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline \end{array} \right) =$$

[Source: D. Lowe]

Example

- What's the result?



Original


$$* \left(\begin{array}{|c|c|c|} \hline 0 & 0 & 0 \\ \hline 0 & 2 & 0 \\ \hline 0 & 0 & 0 \\ \hline \end{array} - \frac{1}{9} \begin{array}{|c|c|c|} \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline 1 & 1 & 1 \\ \hline \end{array} \right) =$$




Sharpening filter
(accentuates edges)

[Source: D. Lowe]

Sharpening



before



after

[Source: D. Lowe]

Example of Correlation

- What is the result of filtering the impulse signal (image) I with the arbitrary filter F ?

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	0	0	1	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	0	0	0
0	0	0	0	0	0	0	0

$I(i, j)$



a	b	c
d	e	f
g	h	i

$F(i, j)$

$G(i, j)$

[Source: K. Grauman]

Smoothing by averaging



depicts box filter:
white = high value, black = low value



original




filtered

- What if the filter size was 5×5 instead of 3×3 ?

[Source: K. Graumann]

Gaussian filter

- What if we want nearest neighboring pixels to have the most influence on the output?
- Removes high-frequency components from the image (low-pass filter).




$I(i, j)$

$$\frac{1}{16} \begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix}$$

$F(i, j)$

This kernel is an approximation of a 2d Gaussian function:

$$h(u, v) = \frac{1}{2\pi\sigma^2} e^{-\frac{u^2+v^2}{\sigma^2}}$$




[Source: S. Seitz]

Smoothing with a Gaussian




[Source: K. Grauman]

Mean vs Gaussian




Gaussian filter: Parameters

- **Size of filter or mask:** Gaussian function has infinite support, but discrete filters use finite kernels.



$\sigma = 5$ with
 10×10
kernel




$\sigma = 5$ with
 30×30
kernel


[Source: K. Grauman]

Gaussian filter: Parameters

- **Variance of the Gaussian:** determines extent of smoothing.




$\sigma = 2$ with
30 x 30
kernel



$\sigma = 5$ with
30 x 30
kernel

[Source: K. Grauman]

Gaussian filter: Parameters





```
for sigma=1:3:10
    h = fspecial('gaussian', fsize, sigma);
    out = imfilter(im, h);
    imshow(out);
    pause;
end
```

[Source: K. Grauman]

Is this the most general Gaussian?

- No, the most general form for $\mathbf{x} \in \mathbb{R}^d$

$$\mathcal{N}(\mathbf{x}; \mu, \Sigma) = \frac{1}{(2\pi)^{d/2} |\Sigma|^{1/2}} \exp\left(-\frac{1}{2}(\mathbf{x} - \mu)^T \Sigma^{-1}(\mathbf{x} - \mu)\right)$$



- But the simplified version is typically use for filtering.

Properties of the Smoothing

- All values are positive.
- They all sum to 1.

Properties of the Smoothing

- All values are positive.
- They all sum to 1.
- Amount of smoothing proportional to mask size.

Properties of the Smoothing

- All values are positive.
- They all sum to 1.
- Amount of smoothing proportional to mask size.
- Remove high-frequency components; low-pass filter.

Properties of the Smoothing

- All values are positive.
- They all sum to 1.
- Amount of smoothing proportional to mask size.
- Remove high-frequency components; low-pass filter.

Finding Waldo





image I



filter F


- Let's do normalized cross-correlation

Finding Waldo




- Result

Finding Waldo



- Result

Finding Waldo



Voila!

- **Homework:** Do it yourself! Code on class webpage. Don't cheat ;)

- **Convolution** operator


$$G(i, j) = \sum_{u=-k}^k \sum_{v=-k}^k F(u, v) \cdot I(i - u, j - v)$$

- **Equivalent** to flipping the filter in both dimensions (bottom to top, right to left) and apply correlation.

- **Convolution** operator

$$G(i, j) = \sum_{u=-k}^k \sum_{v=-k}^k F(u, v) \cdot I(i - u, j - v)$$


- **Equivalent** to flipping the filter in both dimensions (bottom to top, right to left) and apply correlation.
- If filter is **symmetric**, then correlation and convolution are **the same**





- **Convolution** operator

$$G(i, j) = \sum_{u=-k}^k \sum_{v=-k}^k F(u, v) \cdot I(i - u, j - v)$$

- **Equivalent** to flipping the filter in both dimensions (bottom to top, right to left) and apply correlation.
- If filter is **symmetric**, then correlation and convolution are **the same**



Correlation vs Convolution



- For a Gaussian or box filter, how will the outputs differ?
- If the input is an impulse signal, how will the outputs differ? $\delta * I?$, and $\delta \otimes I?$

"Optical" Convolution

- Camera Shake



Figure: Fergus, et al., SIGGRAPH 2006

- Blur in out-of-focus regions of an image.



Figure: Bokeh: <http://lullaby.homepage.dk/diy-camera/bokeh.html>
Click for more info

[Source: N. Snavely]

Properties of Convolution

Commutative : $f * g = g * f$

Associative : $f * (g * h) = (f * g) * h$

Distributive : $f * (g + h) = f * g + f * h$

Assoc. with scalar multiplier : $\lambda \cdot (f * g) = (\lambda \cdot f) * h$

- The Fourier transform of two convolved images is the product of their individual Fourier transforms:

$$\mathcal{F}(f * g) = \mathcal{F}(f) \cdot \mathcal{F}(g)$$

Properties of Convolution

Commutative : $f * g = g * f$

Associative : $f * (g * h) = (f * g) * h$

Distributive : $f * (g + h) = f * g + f * h$

Assoc. with scalar multiplier : $\lambda \cdot (f * g) = (\lambda \cdot f) * h$

- The Fourier transform of two convolved images is the product of their individual Fourier transforms:

$$\mathcal{F}(f * g) = \mathcal{F}(f) \cdot \mathcal{F}(g)$$

- **Homework:** Why is this good news?
- **Hint:** Think of complexity of convolution and Fourier Transform
- Both correlation and convolution are **linear shift-invariant (LSI) operators**: the effect of the operator is the same everywhere.

Properties of Convolution

Commutative : $f * g = g * f$

Associative : $f * (g * h) = (f * g) * h$

Distributive : $f * (g + h) = f * g + f * h$

Assoc. with scalar multiplier : $\lambda \cdot (f * g) = (\lambda \cdot f) * h$


- The Fourier transform of two convolved images is the product of their individual Fourier transforms:

$$\mathcal{F}(f * g) = \mathcal{F}(f) \cdot \mathcal{F}(g)$$

- **Homework:** Why is this good news?
- **Hint:** Think of complexity of convolution and Fourier Transform
- Both correlation and convolution are **linear shift-invariant (LSI) operators**: the effect of the operator is the same everywhere.

Gaussian Filter

- Convolution with itself is another Gaussian



- Convoluting twice with Gaussian kernel of width σ is the same as convoluting once with kernel of width $\sigma\sqrt{2}$
- We don't need to filter twice, just once with a bigger kernel

[Source: K. Grauman]


Boundary Effects

- What happens at the border of the image? What's the size of the output matrix?
- MATLAB: `FILTER2(G, F, SHAPE)`
- `shape = "full"` output size is sum of sizes of f and g
- `shape = "same"`: output size is same as f
- `shape = "valid"`: output size is difference of sizes of f and g

[Source: S. Lazebnik]

Boundary Effects

- What happens at the border of the image? What's the size of the output matrix?
- MATLAB: `FILTER2(G, F, SHAPE)`
- `shape = "full"`: output size is sum of sizes of f and g
- `shape = "same"`: output size is same as f
- `shape = "valid"`: output size is difference of sizes of f and g



[Source: S. Lazebnik]

- Remember correlation:

$$G(i, j) = \sum_{u=-k}^k \sum_{v=-k}^k F(u, v) \cdot I(i + u, j + v)$$

- Can we write that in a more compact form (with vectors)?

- Remember correlation:

$$G(i, j) = \sum_{u=-k}^k \sum_{v=-k}^k F(u, v) \cdot I(i + u, j + v)$$

- Can we write that in a more compact form (with vectors)?
- Define $\mathbf{f} = F(:)$, $T_{ij} = I(i - k : i + k, j - k : j + k)$, and $\mathbf{t}_{ij} = T_{ij}(:)$

$$G(i, j) = \mathbf{f}^T \cdot \mathbf{t}_{ij}$$

where \cdot is a dot product

- Remember correlation:

$$G(i, j) = \sum_{u=-k}^k \sum_{v=-k}^k F(u, v) \cdot I(i + u, j + v)$$

- Can we write that in a more compact form (with vectors)?
- Define $\mathbf{f} = F(:)$, $T_{ij} = I(i - k : i + k, j - k : j + k)$, and $\mathbf{t}_{ij} = T_{ij}(:)$

$$G(i, j) = \mathbf{f}^T \cdot \mathbf{t}_{ij}$$

where \cdot is a dot product

- Homework:** Can we write full correlation $G = F \otimes I$ in matrix form?

- Remember correlation:

$$G(i, j) = \sum_{u=-k}^k \sum_{v=-k}^k F(u, v) \cdot I(i + u, j + v)$$

- Can we write that in a more compact form (with vectors)?
- Define $\mathbf{f} = F(:)$, $T_{ij} = I(i - k : i + k, j - k : j + k)$, and $\mathbf{t}_{ij} = T_{ij}(:)$

$$G(i, j) = \mathbf{f}^T \cdot \mathbf{t}_{ij}$$

where \cdot is a dot product

- Homework:** Can we write full correlation $G = F \otimes I$ in matrix form?

Separable Filters: Speed-up Trick!

- The process of performing a convolution requires K^2 operations per pixel, where K is the size (width or height) of the convolution filter.
- Can we do faster?

Separable Filters: Speed-up Trick!

- The process of performing a convolution requires K^2 operations per pixel, where K is the size (width or height) of the convolution filter.
- Can we do faster?
- In many cases (**not all!**), this operation can be speed up by first performing a 1D horizontal convolution followed by a 1D vertical convolution, **requiring only $2K$ operations.**

Separable Filters: Speed-up Trick!

- The process of performing a convolution requires K^2 operations per pixel, where K is the size (width or height) of the convolution filter.
- Can we do faster?
- In many cases (**not all!**), this operation can be speed up by first performing a 1D horizontal convolution followed by a 1D vertical convolution, **requiring only $2K$ operations.**
- If this is possible, then the convolution filter is called **separable**.

Separable Filters: Speed-up Trick!

- The process of performing a convolution requires K^2 operations per pixel, where K is the size (width or height) of the convolution filter.
- Can we do faster?
- In many cases (**not all!**), this operation can be speed up by first performing a 1D horizontal convolution followed by a 1D vertical convolution, **requiring only $2K$ operations**.
- If this is possible, then the convolution filter is called **separable**.
- And it is the outer product of two filters:

$$\mathbf{F} = \mathbf{v} \mathbf{h}^T$$

- **Homework:** Think **why** in the case of separable filters 2D convolution is the same as two 1D convolutions

[Source: R. Urtasun]

Separable Filters: Speed-up Trick!


- The process of performing a convolution requires K^2 operations per pixel, where K is the size (width or height) of the convolution filter.
- Can we do faster?
- In many cases (**not all!**), this operation can be speed up by first performing a 1D horizontal convolution followed by a 1D vertical convolution, **requiring only $2K$ operations**.
- If this is possible, then the convolution filter is called **separable**.
- And it is the outer product of two filters:

$$\mathbf{F} = \mathbf{v} \mathbf{h}^T$$

- **Homework:** Think **why** in the case of separable filters 2D convolution is the same as two 1D convolutions

[Source: R. Urtasun]

How it Works



filter






image I


How it Works



How it Works




How it Works



Separable Filters: Gaussian filters

- One famous separable filter we already know:

$$\begin{aligned}\text{Gaussian} : f(x, y) &= \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{\sigma^2}} \\ &= \left(\frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{x^2}{\sigma^2}} \right) \cdot \left(\frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{y^2}{\sigma^2}} \right)\end{aligned}$$



Let's play a game...

Is this separable? If yes, what's the separable version?

$$\frac{1}{K^2} \begin{bmatrix} 1 & 1 & \dots & 1 \\ 1 & 1 & \dots & 1 \\ \vdots & \vdots & 1 & \vdots \\ 1 & 1 & \dots & 1 \end{bmatrix}$$

[Source: R. Urtasun]

Let's play a game...

Is this separable? If yes, what's the separable version?

$$\frac{1}{K^2} \begin{array}{|c|c|c|c|} \hline 1 & 1 & \cdots & 1 \\ \hline 1 & 1 & \cdots & 1 \\ \hline \vdots & \vdots & 1 & \vdots \\ \hline 1 & 1 & \cdots & 1 \\ \hline \end{array}$$

$$\frac{1}{K} \begin{array}{|c|c|c|c|} \hline 1 & 1 & \cdots & 1 \\ \hline \end{array}$$

What does this filter do?

[Source: R. Urtasun]

Let's play a game...

Is this separable? If yes, what's the separable version?

$$\frac{1}{16} \begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline 2 & 4 & 2 \\ \hline 1 & 2 & 1 \\ \hline \end{array}$$

[Source: R. Urtasun]

Let's play a game...

Is this separable? If yes, what's the separable version?

$$\frac{1}{16} \begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline 2 & 4 & 2 \\ \hline 1 & 2 & 1 \\ \hline \end{array}$$
$$\frac{1}{4} \begin{array}{|c|c|c|} \hline 1 & 2 & 1 \\ \hline \end{array}$$

What does this filter do?

[Source: R. Urtasun]

Let's play a game...

Is this separable? If yes, what's the separable version?

$$\frac{1}{8} \begin{array}{|c|c|c|} \hline -1 & 0 & 1 \\ \hline -2 & 0 & 2 \\ \hline -1 & 0 & 1 \\ \hline \end{array}$$

[Source: R. Urtasun]

Let's play a game...

Is this separable? If yes, what's the separable version?

$$\frac{1}{8} \begin{array}{|c|c|c|} \hline -1 & 0 & 1 \\ \hline -2 & 0 & 2 \\ \hline -1 & 0 & 1 \\ \hline \end{array}$$
$$\frac{1}{2} \begin{array}{|c|c|c|} \hline -1 & 0 & 1 \\ \hline \end{array}$$

What does this filter do?

[Source: R. Urtasun]

How can we tell if a given filter F is indeed separable?

- Inspection... this is what we were doing.
- Looking at the analytic form of it.

How can we tell if a given filter F is indeed separable?

- Inspection... this is what we were doing.
- Looking at the analytic form of it.
- Look at the **singular value decomposition (SVD)**, and if only one singular value is non-zero, then it is separable

$$F = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T = \sum_i \sigma_i u_i v_i^T$$

with $\mathbf{\Sigma} = \text{diag}(\sigma_i)$.

How can we tell if a given filter F is indeed separable?

- Inspection... this is what we were doing.
- Looking at the analytic form of it.
- Look at the **singular value decomposition (SVD)**, and if only one singular value is non-zero, then it is separable

$$F = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T = \sum_i \sigma_i u_i v_i^T$$

with $\mathbf{\Sigma} = \text{diag}(\sigma_i)$.

- Matlab: `[U,S,V] = svd(F);`

How can we tell if a given filter F is indeed separable?

- Inspection... this is what we were doing.
- Looking at the analytic form of it.
- Look at the **singular value decomposition (SVD)**, and if only one singular value is non-zero, then it is separable

$$F = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T = \sum_i \sigma_i \mathbf{u}_i \mathbf{v}_i^T$$

with $\mathbf{\Sigma} = \text{diag}(\sigma_i)$.

- Matlab: $[\mathbf{U}, \mathbf{S}, \mathbf{V}] = \text{SVD}(\mathbf{F})$;
- $\sqrt{\sigma_1} \mathbf{u}_1$ and $\sqrt{\sigma_1} \mathbf{v}_1^T$ are the vertical and horizontal filter.

[Source: R. Urtasun]

How can we tell if a given filter F is indeed separable?

- Inspection... this is what we were doing.
- Looking at the analytic form of it.
- Look at the **singular value decomposition (SVD)**, and if only one singular value is non-zero, then it is separable

$$F = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T = \sum_i \sigma_i \mathbf{u}_i \mathbf{v}_i^T$$


with $\mathbf{\Sigma} = \text{diag}(\sigma_i)$.

- Matlab: $[\mathbf{U}, \mathbf{S}, \mathbf{V}] = \text{SVD}(F)$;
- $\sqrt{\sigma_1} \mathbf{u}_1$ and $\sqrt{\sigma_1} \mathbf{v}_1^T$ are the vertical and horizontal filter.


[Source: R. Urtasun]

Sharpening revisited

- What does blurring take away?



- Let's add it back



[Source: S. Lazebnik]

Sharpening



[Source: N. Snavely]

Summary – Stuff You Should Know

- **Correlation:** Slide a filter across image and compare (via dot product)
- **Convolution:** Flip the filter to the right and down and do correlation
- **Smooth** image with a Gaussian kernel: bigger σ means more blurring
- **Some** filters (like Gaussian) are **separable**: you can filter faster. First apply 1D convolution to each row, followed by another 1D conv. to each column
- Applying first a Gaussian filter with σ_1 and then another Gaussian with σ_2 is the same as applying one Gaussian filter with $\sigma = \sqrt{\sigma_1^2 + \sigma_2^2}$

Matlab functions:

- `IMFILTER`: can do both correlation and convolution
- `CORR2`, `FILTER2`: correlation, `NORMXCORR2` normalized correlation
- `CONV2`: does correlation
- `FSPECIAL`: creates special filters including a Gaussian

Next time:

Edge Detection