# Images and Filters 

ECE P 596<br>Linda Shapiro

## What is an image?




$$
\begin{aligned}
& P=f(x, y) \\
& f: R^{2} \quad R
\end{aligned}
$$

1. We sample the image to get a discrete set of pixels with quantized values.
2. For a gray tone image there is one band $F(r, c)$, with values usually between 0 and 255.
3. For a color image there are 3 bands $R(r, c), G(r, c), B(r, c)$

## Image Operations

(functions of functions)


## Image Operations

(functions of functions)


## Image Operations

(functions of functions)


## Local image functions



Image filtering

$$
\mathrm{g}[\cdot, \cdot] \frac{1}{\frac{1}{9}} \begin{aligned}
& \frac{1}{2} \\
& \frac{1}{2} \\
& \hline
\end{aligned}
$$

## $f[. .$,

## $h[.,$.

| 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 |


|  | $A$ |  |  |  |  |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
|  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |

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

Image filtering

## $f[. .$,

## $h[.,$.

| 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 |


|  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | ${ }^{\circ}$ |  |  |  |  |  |
|  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |

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

Image filtering

$$
\bar{\sigma}\left[\bullet, \quad \cdot \frac{1}{9} \begin{array}{|l|l|l|}
\hline 1 & 1 & 1 \\
\hline 1 & 1 & 1 \\
\hline 1 & 1 & 1 \\
\hline
\end{array}\right.
$$

## $f[.,$.

## $h[.,$.

| 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 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |


|  |  | - |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | 20 |  |  |  |  |  |
|  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |

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

Image filtering

$$
\bar{\sigma}\left[\bullet, \quad \cdot \frac{1}{9} \begin{array}{|l|l|l|}
\hline 1 & 1 & 1 \\
\hline 1 & 1 & 1 \\
\hline 1 & 1 & 1 \\
\hline
\end{array}\right.
$$

$f[. .$,
$h[.,$.

| 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 |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |

$$
M\left[m, m=\sum_{k, l} O[k, l] f[m+m+m]\right.
$$

Image filtering

## $f[. .$,

## $h[.,$.

| 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 | 30 |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |

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

Image filtering

$$
\mathrm{g}[\cdot, \cdot] \frac{1}{\frac{1}{9}} \begin{aligned}
& \frac{1}{2} \\
& \frac{1}{2} \\
& \hline
\end{aligned}
$$

## $f[.,$.

$h[.,$.

| 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 |


|  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| - | ${ }^{10}$ | ${ }^{20}$ | ${ }^{30}$ | ${ }^{8}$ |  |  |  |
|  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |
|  |  |  | ? |  |  |  |  |
|  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |

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

Image filtering

## $f[.,$.

$h[.$, .]

| 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 | 30 |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  | ? |  |  |  |
|  |  |  |  | 50 |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |
|  |  |  |  |  |  |  |  |  |  |

$$
M\left[m, m=\sum_{k, l} O[k, l] f[m+m+m]\right.
$$

Image filtering

$$
f[. .,]
$$

$$
h[., .]
$$

| 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 | 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 | 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)

$$
\mathrm{g}[\cdot, \cdot]
$$



## Smoothing with box filter



## Practice with linear filters



| 0 | 0 | 0 |
| :--- | :--- | :--- |
| 0 | 1 | 0 |
| 0 | 0 | 0 |

?

Original

## Practice with linear filters



Original

| 0 | 0 | 0 |
| :--- | :--- | :--- |
| 0 | 1 | 0 |
| 0 | 0 | 0 |

Filtered
(no change)

## Practice with linear filters



| 0 | 0 | 0 |
| :--- | :--- | :--- |
| 0 | 0 | 1 |
| 0 | 0 | 0 |

?

Original

## Practice with linear filters



Original

| 0 | 0 | 0 |
| :--- | :--- | :--- |
| 0 | 0 | 1 |
| 0 | 0 | 0 |

Shifted left By 1 pixel

## Practice with linear filters



| 0 | 0 | 0 |
| :--- | :--- | :--- |
| 0 | 2 | 0 |
| 0 | 0 | 0 |$-\frac{1}{9}$| 1 | 1 | 1 |
| :--- | :--- | :--- |
| 1 | 1 | 1 |
| 1 | 1 | 1 |

Original

## Practice with linear filters



| 0 | 0 | 0 |
| :--- | :--- | :--- |
| 0 | 2 | 0 |
| 0 | 0 | 0 |

Original


Sharpening filter

- Accentuates differences with local average


## Sharpening


before

after

## Other filters



Vertical Edge
(absolute value)

## Other filters

| 1 | 2 | 1 |
| :---: | :---: | :---: |
| 0 | 0 | 0 |
| -1 | -2 | -1 |



Horizontal Edge (absolute value)

## Basic gradient filters

Horizontal Gradient

| 0 | 0 | 0 |
| :---: | :---: | :---: |
| -1 | 0 | 1 |
| 0 | 0 | 0 |

or

| -1 | 0 | 1 |
| :--- | :--- | :--- |

## Gaussian filter

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



Input image $f$


Filter $h$


Output image $g$

## Gaussian vs. mean filters



What does real blur look like?

## Important filter: Gaussian

- Spatially-weighted average



## Smoothing with Gaussian filter



## Smoothing with box filter



## Gaussian filters

- What parameters matter here?
- Variance of Gaussian: determines extent of smoothing




## Smoothing with a Gaussian

Parameter $\sigma$ is the "scale" / "width" / "spread" of the Gaussian kernel, and controls the amount of smoothing.


## 2D edge detection filters



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

Laplacian of Gaussian or LoG filter

$$
\frac{\partial}{\partial x} h_{\sigma}(u, v)
$$



$$
\nabla^{2} h_{\sigma}(u, v)
$$

$\nabla^{2}$ is the Laplacian operator (sum of $2^{\text {nd }}$ derivatives):

$$
\nabla^{2} f=\frac{\partial^{2} f}{\partial x^{2}}+\frac{\partial^{2} f}{\partial y^{2}}
$$

Often approximated by

$$
\begin{array}{ccc}
0 & 1 & 0 \\
1 & -4 & 1 \\
0 & 1 & 0
\end{array}
$$

## First and second derivatives



Original


What are these good for?


First Derivative x



Second Derivative $\mathrm{x}, \mathrm{y}$


## Subtracting filters

$$
\operatorname{Sharpen}(x, y)=f(x, y)-\alpha\left(f * \nabla^{2} \mathcal{G}_{\sigma}(x, y)\right)
$$



Original


Second Derivative


Sharpened

## Combining filters

$f * g * g^{\prime}=f * h$ for some $h$


| 0 | 0 | 0 | 0 | 0 |
| :---: | :---: | :---: | :---: | :---: |
| 0 | 0 | -1 | 0 | 0 |
| 0 | -1 | 4 | -1 | 0 |
| 0 | 0 | -1 | 0 | 0 |
| 0 | 0 | 0 | 0 | 0 |



It's also true: $f *(g * h)=(f * g) * h$

$$
f * g=g * f
$$

## Combining Gaussian filters

$$
\begin{gathered}
* *=? \\
f * \mathcal{G}_{\sigma} * \mathcal{G}_{\sigma^{\prime}}=f * \mathcal{G}_{\sigma^{\prime \prime}} \\
\sigma^{\prime \prime}=\sqrt{\sigma^{2}+\sigma^{\prime 2}}
\end{gathered}
$$

More blur than either individually (but less than $\sigma^{\prime \prime}=\sigma+\sigma^{\prime}$ )

## Separable filters

$$
\begin{array}{r}
\mathcal{G}_{\sigma}=\mathcal{G}_{\sigma}^{x} * \mathcal{G}_{\sigma}^{y} \\
\mathcal{G}_{\sigma}^{x}(x, y)=\frac{1}{Z} e^{\frac{-\left(x^{2}\right)}{2 \sigma^{2}}} \\
\mathcal{G}_{\sigma}^{y}(x, y)=\frac{1}{Z} e^{\frac{-\left(y^{2}\right)}{2 \sigma^{2}}}
\end{array}
$$

Compute Gaussian in horizontal direction, followed by the vertical direction. Much faster!


Not all filters are separable.
Freeman and Adelson, 1991

## Linear vs. Non-Linear Fitters


(a)

(e)

(b)

(f)

(c)

(g)

(d)

(h)
a. original image with Gaussian noise, b. Gaussian filtered, c. median filtered, d. bilateral filtered e. original image with shot noise, f. Gaussian filtered, g. median filtered, h. bilateral filtered

## Spatially varying filters

- Some filters vary spatially.
- The bilateral filter is the product of a domain kernel (Gaussian) and a data dependent range kernel.
- $\mathrm{d}(\mathrm{i}, \mathrm{j}, \mathrm{k}, \mathrm{I})=\exp \left[\left(-(\mathrm{i}-\mathrm{k})^{2}+(\mathrm{j}-\mathrm{I})^{2}\right) / 2 \sigma_{\mathrm{d}}{ }^{2}\right]$ is the domain kernel
- $r(i, j, k, I)=\exp \left[-||f(i, j)-f(k, l)||^{2} / 2 \sigma_{r}^{2}\right]$ is the range kernel
- $w(i, j, k, I)=d(i, j, k, I) * r(i, j, k, I)$ is their product
- $g(i, j)=\Sigma_{k, I} f(k, l) w(i, j, k, I) / \Sigma_{k, l} w(i, j, k, I)$ is the bilateral filter


## Constant blur: same kernel everywhere



Same Gaussian kernel everywhere.
Slides courtesy of Sylvian Paris

## Bilateral filter: kernel depends on intensity

## Maintains edges when blurring!



The kernel shape depends on the image content.
Slides courtesy of Sylvian Paris

## Borders

## What to do about image borders:


black

fixed

periodic

reflected

## Image Sampling



F(

) =


## Image Scaling

This image is too big to fit on the screen. How can we reduce it?

How to generate a halfsized version?


## Image sub-sampling



1/8

1/4

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

- called image sub-sampling


## Image sub-sampling



Why does this look so bad?

## Down-sampling



- Aliasing can arise when you sample a continuous signal or image
- occurs when your sampling rate is not high enough to capture the amount of detail in your image
- Can give you the wrong signal/image-an alias
- formally, the image contains structure at different scales
- called "frequencies" in the Fourier domain
- the sampling rate must be high enough to capture the highest frequency in the image


## Subsampling with Gaussian pre-filtering



G $1 / 8$
G 1/4

Gaussian 1/2
Solution: filter the image, then subsample

- Filter size should double for each $1 / 2$ size reduction.


## Finale

- Filtering is just applying a mask to an image.
- Computer vision people call the linear form of these operations "convolutions". They are actually "correlations," since the true convolution inverts the mask.
- There are many nonlinear filters, too, such as median filters and morphological filters.
- Filtering is the lowest level of image analysis and is taught heavily in image processing courses.

