## Filtering Images in the Spatial

 DomainChapter 3b G\&W
Ross Whitaker
(modified by Guido Gerig)
School of Computing
University of Utah

## Overview

- Correlation and convolution
- Linear filtering
- Smoothing, kernels, models
- Detection
- Derivatives
- Nonlinear filtering
- Median filtering
- Bilateral filtering
- Neighborhood statistics and nonlocal filtering


## Cross Correlation

- Operation on image neighborhood and small ...
- "mask", "filter", "stencil", "kernel"
- Linear operations within a moving window

| Input image |  |  |  |  | Filter |  |  |
| ---: | ---: | ---: | ---: | :--- | :--- | :--- | :--- | :--- |
| 100 | 130 | 104 | 99 | $\ldots$ |  |  |  |
| 87 | 95 | 103 | 150 | $\ldots$ | 0.0 | 0.1 | 0.0 |
| 50 | 36 | 150 | 104 | $\ldots$ | 0.1 | 0.6 | 0.1 |
| 20 | 47 | 205 | 77 | $\ldots$ | 0.0 | 0.1 | 0.0 |

## Cross Correlation

- Operation on image neighborhood and small ...
- "mask", "filter", "stencil", "kernel"
- Linear operations within a moving window
Input image

Filter

| 0.0 | 0.1 | 0.0 |
| :--- | :--- | :--- |
| 0.1 | 0.6 | 0.1 |
| 0.0 | 0.1 | 0.0 |

0.00 .10 .0

Output image

| 100 | 130 | 104 | 99 | $\ldots$ |
| ---: | ---: | ---: | ---: | ---: |
| 87 | 95 | 103 | 150 | $\ldots$ |
| 50 | 36 | 150 | 104 | $\ldots$ |
| 20 | 47 | 205 | 77 | $\ldots$ |

## Cross Correlation

- Operation on image neighborhood and small ...
- "mask", "filter", "stencil", "kernel"
- Linear operations within a moving window
Input image


## Filter

Output image

| 100 | 130 | 104 | 99 | $\ldots$ |
| ---: | ---: | ---: | ---: | ---: |
| 87 | 95 | 103 | 150 | $\ldots$ |
| 50 | 36 | 150 | 104 | $\ldots$ |
| 20 | 47 | 205 | 77 | $\ldots$ |


| 0.0 | 0.1 | 0.0 |
| :--- | :--- | :--- |
| 0.1 | 0.6 | 0.1 |
| 0.0 | 0.1 | 0.0 |

## Cross Correlation

- Operation on image neighborhood and small ...
- "mask", "filter", "stencil", "kernel"
- Linear operations within a moving window
Input image


## Filter

Output image

| 100 | 130 | 104 |
| ---: | ---: | ---: |
| 87 | 95 | 103 |
| 50 | 36 | 150 |
| 20 | 47 | 205 |


| 99 | $\ldots$ |  |  |  |
| ---: | :--- | :--- | :--- | :--- |
| 150 | $\ldots$ | 0.0 | 0.1 | 0.0 |
| 104 | $\ldots$ | 0.1 | 0.6 | 0.1 |
| 77 | $\ldots$ | 0.0 | 0.1 | 0.0 |

$$
\begin{aligned}
& 0.0 * 87+0.1 * 95+0.0 * 103 \\
+ & 0.1 * 50+0.6 * 36+0.1 * 150 \\
+ & 0.0 * 20+0.1 * 47+0.0 * 205=55.8
\end{aligned}
$$

## Cross Correlation

- Operation on image neighborhood and small ...
- "mask", "filter", "stencil", "kernel"
- Linear operations within a moving window
Input image


## Filter

## Output image

| 100 | 130 | 104 |
| ---: | ---: | ---: |
| 87 | 95 | 103 |
| 50 | 36 | 150 |
| 20 | 47 | 205 |

$99 \ldots$
$150 \ldots$
$104 \ldots$
$77 \ldots$

| 0.0 | 0.1 | 0.0 |
| :--- | :--- | :--- |
| 0.1 | 0.6 | 0.1 |
| 0.0 | 0.1 | 0.0 |



$$
\begin{aligned}
& 0.0 * 87+0.1 * 95+0.0 * 103 \\
+ & 0.1 * 50+0.6 * 36+0.1 * 150 \\
+ & 0.0 * 20+0.1 * 47+0.0 * 205=55.8
\end{aligned}
$$

## Cross Correlation

- Operation on image neighborhood and small ...
- "mask", "filter", "stencil", "kernel"
- Linear operations within a moving window



## Cross Correlation

- Operation on image neighborhood and small ...
- "mask", "filter", "stencil", "kernel"
- Linear operations within a moving window



## Cross Correlation

- Operation on image neighborhood and small ...
- "mask", "filter", "stencil", "kernel"
- Linear operations within a moving window

| Input image |  |  |  |  | Filter |  |  | Output image |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 100 | 130 | 104 | 99 | $\ldots$ |  |  | - | - - |
| 87 | 95 | 103 | 150 | ... | 0. | 0.10 .0 | - | - - |
| 50 | 36 | 150 | 104 | ... | 0.1 | 0.60 .1 | - | - - |
| 20 | 47 | 205 | 77 | ... | 0.0 | 0.10 .0 | - | - - |

## Cross Correlation

- Operation on image neighborhood and small ...
- "mask", "filter", "stencil", "kernel"
- Linear operations within a moving window

Input image

| 100 | 130 | 104 | 99 | $\ldots$ |
| ---: | ---: | ---: | ---: | ---: |
| 87 | 95 | 103 | 150 | $\ldots$ |
| 50 | 36 | 150 | 104 | $\ldots$ |
| 20 | 47 | 205 | 77 | $\ldots$ |

Output image

| 0.0 | 0.1 | 0.0 |
| :--- | :--- | :--- |
| 0.1 | 0.6 | 0.1 |
| 0.0 | 0.1 | 0.0 |

## Cross Correlation

- Operation on image neighborhood and small ...
- "mask", "filter", "stencil", "kernel"
- Linear operations within a moving window



## Cross Correlation

- Operation on image neighborhood and small ...
- "mask", "filter", "stencil", "kernel"
- Linear operations within a moving window



## Cross Correlation

-1D $g(x)=\sum_{s=-a}^{a} w(s) f(x+s)$
-2D $g(x, y)=\sum_{s=-a}^{a} \sum_{t=-b}^{b} w(s, t) f(x+s, y+t)$

$$
w(s, t)=\begin{array}{ccccc}
w(-a,-b) & \cdots & & \cdots & w(a,-b) \\
\vdots & & & & \vdots \\
\vdots & \cdots & w(0,0) & \cdots & \\
w(-a, b) & \cdots & & \vdots \\
& & \cdots & w(a, b) \\
\hline
\end{array}
$$

## Correlation: Technical Details

- How to filter boundary?



## Correlation: Technical Details

- Boundary conditions
- Boundary not filtered (keep it 0)
- Pad image with amount (a,b)
- Constant value or repeat edge values
- Cyclical boundary conditions
- Wrap or mirroring


| d | b | d |
| :--- | :--- | :--- |
| q | P | q |
| d | b | d |

## Correlation: Technical Details

- Boundaries
- Can also modify kernel - no longer correlation
- For analysis
- Image domains infinite
- Data compact (goes to zero far away from origin)

$$
g(x, y)=\sum_{s=-\infty}^{\infty} \sum_{t=-\infty}^{\infty} w(s, t) f(x+s, y+t)
$$

## Correlation: Properties

- Shift invariant

$$
\begin{gathered}
g=w \circ f \quad g(x, y)=w(x, y) \circ f(x, y) \\
w(x, y) \circ f\left(x-x_{0}, y-y_{0}\right)=\sum_{s=-\infty}^{\infty} \sum_{t=-\infty}^{\infty} w(s, t) f\left(x-x_{0}+s, y-y_{0}+t\right)=g\left(x-x_{0}, y-y_{0}\right)
\end{gathered}
$$

## Correlation: Properties

- Shift invariant

$$
\begin{gathered}
g=w \circ f \quad g(x, y)=w(x, y) \circ f(x, y) \\
w(x, y) \circ f\left(x-x_{0}, y-y_{0}\right)=\sum_{s=-\infty}^{\infty} \sum_{t=-\infty}^{\infty} w(s, t) f\left(x-x_{0}+s, y-y_{0}+t\right)=g\left(x-x_{0}, y-y_{0}\right)
\end{gathered}
$$

- Linear $w \circ(\alpha e+\beta f)=\alpha w \circ e+\beta w \circ f$

Compact notation

$$
C_{w f}=w \circ f
$$

## Filters: Considerations

- Normalize
- Sums to one
- Sums to zero (some cases, see later)
- Symmetry
- Left, right, up, down
- Rotational
- Special case: auto correlation

$$
C_{f f}=f \circ f
$$

## Examples 1


000
010
000

## Examples 1



## Examples 1



$$
\begin{array}{ccc}
1 / 9 * & 1 & 1 \\
1 & 1 & 1 \\
1 & 1 & 1
\end{array}
$$

## Examples 1



## Examples 1



Examples 2


## Smoothing and Noise

Noisy image

$5 \times 5$ box filter


## Noise Analysis

- Consider an a simple image I() with additive, uncorrelated, zero-mean noise of variance s
- What is the expected rms error of the corrupted image?
- If we process the image with a box filter of size $2 a+1$ what is the expected error of the filtered image?

$$
\operatorname{RMSE}=\left(\frac{1}{|\mathcal{D}|} \sum_{(\mathrm{x}, \mathrm{y}) \in \mathcal{D}}(\tilde{\mathrm{I}}(\mathrm{x}, \mathrm{y})-\mathrm{I}(\mathrm{x}, \mathrm{y}))^{2}\right)^{\frac{1}{2}}
$$

## Other Filters

- Disk
- Circularly symmetric, jagged in discrete case
- Gaussians
- Circularly symmetric, smooth for large enough stdev
- Must normalize in order to sum to one
- Derivatives - discrete/finite differences
- Operators


## Gaussian Kernel



Figure 2.1 The Gaussian kernel with unit standard deviation in 1D.

## Gaussian Kernel

$$
G_{1 D}(x ; \sigma)=\frac{1}{\sqrt{2 \pi} \sigma} e^{-\frac{x^{2}}{2 \sigma^{2}}}, G_{2 D}(x, y ; \sigma)=\frac{1}{2 \pi \sigma^{2}} e^{-\frac{x^{2}+y^{2}}{2 \sigma^{2}}}, G_{\mathrm{ND}}(\vec{x} ; \sigma)=\frac{1}{(\sqrt{2 \pi} \sigma)^{N}} e^{-\frac{|x|^{2}}{2 \sigma^{2}}}
$$

Normalization to 1.0




Figure 3.2 The Gaussian function at scales $\sigma=.3, \sigma=1$ and $\sigma=2$. The kernel is normalized, so the total area under the curve is always unity.

## Box versus Gaussian



## Convolution

Java demo: http://www.jhu.edu/signals/convolve/ http://www.jhu.edu/signals/discreteconv2/index.html

- Discrete

$$
g(x, y)=w(x, y) * f(x, y)=\sum_{s=-a}^{a} \sum_{t=-b}^{b} w(s, t) f(x-s, y-t)
$$

- Continuous

$$
g(x, y)=w(x, y) * f(x, y)=\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} w(s, t) f(x-s, y-t) d s d t
$$

- Same as cross correlation with kernel transposed around each axis
- The two operations (correlation and convolution) are the same if the kernel is symmetric about axes

$$
g=w \circ f=w^{*} * f \quad w^{*} \text { reflection of } w
$$

## Convolution: Properties

- Shift invariant, linear
- Commutative

$$
f * g=g * f
$$

- Associative

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

- Others (discussed later):
- Derivatives, convolution theorem, spectrum...


## Computing Convolution

## Computing Convolution

- Compute time
- MxM mask
- NxN image


## Computing Convolution

- Compute time
- MxM mask
- NxN image $\} O\left(M^{2} N^{2}\right) \quad$ "tor" loops are nested 4 deep


## Computing Convolution

- Compute time
- MxM mask
$\mathrm{O}\left(\mathrm{M}^{2} \mathrm{~N}^{2}\right) \quad$ "for" loops are nested 4 deep
- NxN image $\int 0\left(\mathrm{~m}^{2} \mathrm{~N}^{2}\right)$
- Special case: separable

Two 1D kernels
$w=w_{x} * w_{y}$


## Separable Kernels

- Examples
- Box/rectangle
- Bilinear interpolation
- Combinations of partial derivatives
- d²f/dxdy
- Gaussian
- Only filter that is both circularly symmetric and separable
- Counter examples
- Disk
- Cone
- Pyramid


## Separability

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

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

```
DisplayTogetherArray[(Plot [gauss [x, |=1],{x, -3, 3)],
    Plot3D[gauss[x,\sigma=1] gauss[y,\sigma=1],{x,-3,3},{y,-3, 3}]},
    ImageSize-> 440];
```



Figure 3.7 A product of Gaussian functions gives a higher dimensional Gaussian function This is a consequence of the separability.

# Digital Images: Boundaries are "Lines" or "Discontinuities" 



## Example: Characterization of discontinuities?

Digital Images: Boundaries are "Lines" or "Discontinuities"


## Example: Characterization of discontinuities?

## Characterizing edges

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



## Differentiation and convolution

- Recall, for 2D function, $\mathrm{f}(\mathrm{x}, \mathrm{y})$ :
$\frac{\partial f}{\partial x}=\lim _{\varepsilon \rightarrow 0}\left(\frac{f(x+\varepsilon, y)}{\varepsilon}-\frac{f(x, y)}{\varepsilon}\right)$
- This is linear and shift invariant, so must be the result of a convolution.
- We could approximate this as
$\frac{\partial f}{\partial x} \approx \frac{f\left(x_{n+1}, y\right)-f\left(x_{n}, y\right)}{\Delta x}$
- (which is obviously a convolution)



## Derivatives: Finite Differences

$$
\frac{\partial f}{\partial x} \approx \frac{1}{2 h}(f(x+1, y)-f(x-1, y))
$$

$$
\begin{array}{ll}
\frac{\partial f}{\partial x} \approx w_{d x} \circ f & w_{d x}=\begin{array}{|c|c|c|}
-\frac{1}{2} & 0 & \frac{1}{2} \\
\frac{\partial f}{\partial y} \approx w_{d y} \circ f & w_{d y}=\begin{array}{|c|}
\hline-\frac{1}{2} \\
\hline
\end{array} \\
\hline \frac{1}{2} \\
\hline
\end{array}
\end{array}
$$

## Derivative Example



## Image gradient

- The gradient of an image: $\nabla f=\left[\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}\right]$
- The gradient points in the direction of most rapid change in intensity

$$
\left[\begin{array}{l}
\nabla f=\left[\frac{\partial f}{\partial x}, 0\right] \quad \\
\nabla f=\left[0, \frac{\partial f}{\partial_{j}}\right] \quad
\end{array}\right.
$$

- The gradient direction is given by:

$$
\theta=\tan ^{-1}\left(\frac{\partial f}{\partial y} / \frac{\partial f}{\partial x}\right)
$$

- how does this relate to the direction of the edge? perpendicular
- The edge strength is given by the gradient magnitude

$$
\|\nabla f\|=\sqrt{\left(\frac{\partial f}{\partial x}\right)^{2}+\left(\frac{\partial f}{\partial y}\right)^{2}}
$$

## Finite differences: example



- Which one is the gradient in the x-direction (resp. y-direction)?


## Finite difference filters

- Other approximations of derivative filters exist:

$$
\begin{aligned}
& \text { Prewitt: } M_{z}=\begin{array}{|r|r|r|}
\hline-1 & 0 & 1 \\
\hline-1 & 0 & 1 \\
\hline-1 & 0 & 1 \\
\hline
\end{array} ; \quad M_{y}=\begin{array}{|r|r|r|}
\hline 1 & 1 & 1 \\
\hline 0 & 0 & 0 \\
\hline-1 & -1 & -1 \\
\hline-1 & 0 & 1 \\
\hline-2 & 0 & 2 \\
\hline-1 & 0 & 1 \\
\text { Sobel: } \quad M_{x}=
\end{array} \quad M_{y}=\begin{array}{|r|r|r|}
\hline 1 & 2 & 1 \\
\hline 0 & 0 & 0 \\
\hline-1 & -2 & -1 \\
\hline
\end{array} \\
& \text { Roberts: } \quad M_{z}=\begin{array}{|r|l|}
\hline 0 & 1 \\
\hline-1 & 0 \\
\hline
\end{array} ; \quad M_{y}=\begin{array}{|r|r|}
\hline 1 & 0 \\
\hline 0 & -1 \\
\hline
\end{array}
\end{aligned}
$$

## Pattern Matching



## Pattern Matching/Detection

- The optimal (highest) response from a filter is the autocorrelation evaluated at position zero

$$
\max _{\bar{x}} C_{f f}(\bar{x})=C_{f f}(0)=\int f(\bar{s}) f(\bar{s}) d \bar{s}
$$

- A filter responds best when it matches a pattern that looks itself
- Strategy
- Detect objects in images by correlation with "matched" filter


## Matched Filter Example

## w







盖

## Matched Filter Example:

## Correlation of template with image

## Matched Filter Example:

## Thresholding of correlation results

$$
1
$$

$$
1 \cdot 1
$$

$$
1 \cdot 1
$$

## Matched Filter Example: High correlation $\rightarrow$ template found



## Summary of 9/15

- Spatial Filtering
- Consider neighborhood information
- Special consideration at boundary




## Summary of $9 / 15$

- Common Filters



## Box versus Gaussian



## Gaussian Filtering



## Gaussian Filtering

$\sigma=2.72$

$\sigma=3.00$

$\sigma=16.44$

$\sigma=18.17$

$\sigma=20.09$


## Finite differences: example



- Which one is the gradient in the x-direction (resp. y-direction)?


## Cross-correlation and Convolution

- Cross-correlation

$$
g(x, y)=\sum_{s=-a}^{a} \sum_{t=-b}^{b} w(s, t) f(x+s, y+t)
$$

- Convolution

$$
g(x, y)=\sum_{s=-a}^{a} \sum_{t=-b}^{b} w(s, t) f(x-s, y-t)
$$

## Key Concepts

## - Separability



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

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



64

## Nonlinear Methods For Filtering

- Median filtering
- Bilateral filtering
- Neighborhood statistics and nonlocal filtering


## Median Filtering

- For each neighborhood in image
- Sliding window
- Usually odd size (symmetric) $5 \times 5,7 \times 7, \ldots$
- Sort the greyscale values
- Set the center pixel to the median
- Important: use "Jacobi" updates
- Separate input and output buffers
- All statistics on the original image old



## Median vs Gaussian



## Median Filter

- Issues
- Boundaries
- Compute on pixels that fall within window
- Computational efficiency
- What is the best algorithm?
- Properties
- Removes outliers (replacement noise - salt and pepper)
- Window size controls size of structures
- Preserves straight edges, but rounds corners and features


## Replacement Noise

- Also: "shot noise", "salt\&pepper"
- Replace certain \% of pixels with samples from pdf
- Best filtering strategy: filter to avoid outliers



## Smoothing of S\&P Noise

- It's not zero mean (locally)
- Averaging produces local biases



## Smoothing of S\&P Noise

- It's not zero mean (locally)
- Averaging produces local biases



## Median Filtering



Median $3 \times 3$

Median 5x5

## Median Filtering



Median $3 \times 3$


Median 5x5

## Median Filtering

- Iterate


Median $3 \times 3$


2x Median $3 \times 3$

## Median Filtering

- Image model: piecewise constant (flat) $\bigcirc \bigcirc \bigcirc \bigcirc \bigcirc \bigcirc \bigcirc \bigcirc \bigcirc \bigcirc$


## Median Filtering

- Image model: piecewise constant (flat)



## Median Filtering

- Image model: piecewise constant (flat) $\bigcirc \bigcirc \bigcirc \bigcirc \bigcirc \bigcirc \bigcirc \bigcirc \bigcirc \bigcirc$


## Median Filtering

- Image model: piecewise constant (flat)



## Order Statistics

- Median is special case of order-statistics filters
- Instead of weights based on neighborhoods, weights are based on ordering of data

$$
\begin{array}{cc}
\text { Neighborhood } & \text { Ordering } \\
X_{1}, X_{2}, \ldots, X_{N} & X_{(1)} \leq X_{(2)} \leq \ldots \leq X_{(N)}
\end{array}
$$

Filter $\quad F\left(X_{1}, X_{2}, \ldots, X_{N}\right)=\boldsymbol{\alpha}_{\mathbf{1}} X_{(\mathbf{1})}+\boldsymbol{\alpha}_{2} X_{(\mathbf{2})}+\ldots+\boldsymbol{\alpha}_{N} X_{(N)}$

Neighborhood average (box)
Median filter

$$
\alpha_{i}=1 / N
$$

$$
\alpha_{i}= \begin{cases}1 & i=(N+1) / 2 \\ 0 & \text { otherwise }\end{cases}
$$

Trimmed average (outlier removal)

$$
\alpha_{i}= \begin{cases}1 / M & (N-M+1) / 2 \leq i \leq(N+M+1) / 2 \\ 0 & \text { otherwise }\end{cases}
$$

## Median filter

- What advantage does median filtering have over Gaussian filtering?
- Robustness to outliers
filters have width 5 :



## Piecewise Flat Image Models

- Image piecewise flat -> average only within similar regions
- Problem: don't know region boundaries



## Piecewise-Flat Image Models

- Assign probabilities to other pixels in the image belonging to the same region
- Two considerations
- Distance: far away pixels are less likely to be same region
- Intensity: pixels with different intensities are less likely to be same region


## Gaussian: Blur Comes from Averaging across Edges



Same Gaussian kernel everywhere.
Source: http://people.csail.mit.edu/sparis/bf_course/slides/03_definition_bf.ppt

## Bilateral Filter

## No Averaging across Edges

[Aurich 95, Smith 97, Tomasi 98]


The kernel shape depends on the image content.

## Main Idea

## Distance (kernel/pdf)

$$
G\left(\mathbf{x}_{i}-\mathbf{x}_{j}\right)
$$




## Bilateral Filter Definition: an Additional Edge Term

Same idea: weighted average of pixels.


## Illustration a 1D Image

- 1 D image $=$ line of pixels

- Better visualized as a plot



## Gaussian Blur and Bilateral Filter

## Gaussian blur



$$
G B[I]_{\mathbf{p}}=\sum_{\mathbf{q} \in S} G_{\sigma}(\|\mathbf{p}-\mathbf{q}\|) I_{\mathbf{q}}
$$

## Gaussian Blur and Bilateral Filter

Gaussian blur


$$
G B[I]_{\mathbf{p}}=\sum_{\mathbf{q} \in S} G_{\sigma}(\|\mathbf{p}-\mathbf{q}\|) I_{\mathbf{q}}
$$

Bilateral filter
[Aurich 95, Smith 97, Tomasi 98]


$$
\begin{gathered}
B F[I]_{\mathbf{p}}=\frac{1}{W_{\mathbf{p}}} \sum_{\mathbf{q} \in S} G_{\sigma_{\mathrm{s}}}(\|\mathbf{p}-\mathbf{q}\|) G_{\sigma_{\mathrm{r}}}\left(\left|I_{\mathbf{p}}-I_{\mathbf{q}}\right|\right) I_{\mathbf{q}} \\
\text { normalization }
\end{gathered}
$$

## Bilateral Filter

- Neighborhood - sliding window
- Weight contribution of neighbors according to:

$$
\begin{aligned}
f_{i} \leftarrow k_{i}^{-1} \sum_{j \in N} f_{j} G\left(\mathbf{x}_{i}-\mathbf{x}_{j}\right) H\left(f_{i}-f_{j}\right) & \\
k_{i}=\sum_{j \in N} G\left(\mathbf{x}_{i}-\mathbf{x}_{j}\right) H\left(f_{i}-f_{j}\right) & \begin{array}{l}
\text { normalization: all } \\
\text { weights add up to } 1
\end{array}
\end{aligned}
$$

- G is a Gaussian (or lowpass), as is $\mathrm{H}, \mathrm{N}$ is neighborhood,
- Often use $G\left(r_{\mathrm{ij}}\right)$ where $\mathrm{r}_{\mathrm{ij}}$ is distance between pixels
- Update must be normalized for the samples used in this (particular) summation
- Spatial Gaussian with extra weighting for intensity
- Weighted average in neighborhood with downgrading of intensity outliers


## Bilateral Filter

Replaces the pixel value at $\mathbf{x}$ with an average of similar and nearby pixel values.


When the bilateral filter is centered, say, on a pixel on the bright side of the boundary, the similarity function $s$ assumes values close to one for pixels on the same side, and values close to zero for pixels on the dark side. The similarity function is shown in figure 1(b) for a $23 \times 23$ filter support centered two pixels to the right of the step in figure 1(a).

## Bilateral Filtering

Replaces the pixel value at $\mathbf{x}$ with an average of similar and nearby pixel values.


Gaussian Blurring
Bilateral

## Bilateral Filtering



Gaussian Blurring


Bilateral

## Nonlocal Averaging

- Recent algorithm
- NL-means, Baudes et al., 2005
- UINTA, Awate \& Whitaker, 2005
- Different model
- No need for piecewise-flat
- Images consist of some set of pixels with similar neighborhoods $\rightarrow$ average several of those
- Scattered around
- General area of a pixel
- All around
- Idea
- Average sets of pixels with similar neighborhoods


# UINTA: Unsupervised Information-Theoretic Adaptive Filtering : Excellent Introduction and Additional Readings (Suyash P. Awate) 


http://www.cs.utah.edu /~suyash/pubs/uinta/

Suyash P. Awate, Ross T. Whitaker Unsupervised, Information-Theoretic, Adaptive Image Filtering with
Applications to Image Restoration
IEEE Trans. Pattern Analysis \& Machine Intelligence (TPAMI) 2006, Vol. 28, Num. 3, pp. 364-376

## Nonlocal Averaging

- Strategy:
- Average pixels to alleviate noise
- Combine pixels with similar neighborhoods
- Formulation
$-n_{i, j}$ - vector of pixels values, indexed by $j$, from neighborhood around pixel i



## Nonlocal Averaging Formulation

- Distance between neighborhoods

$$
d_{i, k}=d\left(n_{i}, n_{k}\right)=\left\|n_{i}-n_{k}\right\|=\left(\sum_{j=1}^{N}\left(n_{i, j}-n_{k, j}\right)^{2}\right)^{\frac{1}{2}}
$$

- Kernel weights based on distances

$$
w_{i, k}=K\left(d_{i, k}\right)=e^{-\frac{d_{i, k}^{2}}{2 \sigma^{2}}}
$$

- Pixel values of $k$ neighborhoods: $f_{k}$


## Averaging Pixels Based on

 Weights- For each pixel, i, choose a set of pixel locations k :
$-k=1, \ldots, M$
- Average them together based on neighborhood weights (prop. to intensity pattern difference)

$$
g_{i} \leftarrow \frac{1}{\sum_{k=1}^{M} w_{i, k}} \sum_{k=1}^{M} w_{i, k} f_{k}
$$

## Nonlocal Averaging



## Some Details

- Window sizes: good range is $5 \times 5-$ $>11 \times 11$
- How to choose samples:
- Random samples from around the image
- UINTA, Awate\&Whitaker
- Block around pixel (bigger than window, e.g. 51x51)
- NL-means
- Iterate
- UNITA: smaller updates and iterate


## NL-Means Algorithm

- For each pixel, p
- Loop over set of pixels nearby
- Compare the neighorhoods of those pixels to the neighborhood of $p$ and construct a set of weights
- Replace the value of $p$ with a weighted combination of values of other pixels
- Repeat... but 1 iteration is pretty good


## Results



Noisy image (range 0.0-1.0)


Bilateral filter (3.0, 0.1)

## Results



Bilateral filter (3.0, 0.1)

## Results



Bilateral filter (3.0, 0.1)


NL means (7, 31, 1.0)

## Less Noisy Example



## Less Noisy Example



## Results



## Checkerboard With Noise



Original


Noisy


Filtered

## Quality of Denoising

- •, joint entropy, and RMS- error vs. number of iterations




## MRI Head



## MRI Head



## Fingerprint



## Fingerprint



## Results



Original


Noisy


Filtered

## Results



## Results



Original


Noisy


Filtered

## Fractal



## Piecewise Constant

- Several 10s of Iterations
- Tends to obliterate rare events (e.g. corners)



## Texture, Structure



