

TSBB15 Computer Vision

Lecture 14 Image enhancement

Why image enhancement?


- Example of artifacts caused by image encoding



2016-04-19

2

Why image enhancement?




2016-04-19

3

Why image enhancement?

- Example of an image with sensor noise
 - ultrasound image of a beating heart



2016-04-19

4

Why image enhancement?

- IR-image
 - fixed pattern noise = spatial variations in gain and offset
 - Possibly even variations over time!
 - Hot/dead pixels
- A digital camera with short exposure time
 - Shot noise (photon noise)

2016-04-19

5

Methods for image enhancement

- Inverse filtering: the distortion process is modeled and estimated (e.g. motion blur) and the *inverse* process is applied to the image
- Image restoration: an *objective* quality (e.g. sharpness) is estimated in the image. The image is modified to increase the quality
- Image enhancement: modify the image to improve the visual quality, often with a subjective criterion

2016-04-19

6

Additive noise

- Some types of image distortion can be described as
 - Noise added on each pixel intensity
 - The noise has the identical distribution and is independent at each pixel (i.i.d.)
- Not all type of image distortion are of this type:
 - Multiplicative noise
 - Data dependent noise
 - Position dependent
- The methods discussed here assume additive i.i.d.-noise

What about pixel shot noise?

2016-04-19

7

Removing additive noise

- Image noise typically contains higher frequencies than images generally do
⇒ a low-pass filter can reduce the noise
- BUT: we also remove high-frequency signal components, e.g. at edges and lines
- HOWEVER: A low-pass filter works in regions without edges and lines

2016-04-19

8

Example: LP filter

Image with some noise

Filter, $\sigma = 1$

Filter, $\sigma = 2$

2016-04-19

9

Basic idea

The problem of low-pass filters is that we apply the same filter on the whole image

We need a filter that locally adapts to the image structures

A space variant filter

2016-04-19

10

Ordinary filtering / convolution

- Ordinary filtering can be described as a convolution of the signal f and the filter g :

$$h(\mathbf{x}) = (f * g)(\mathbf{x}) = \int f(\mathbf{x} - \mathbf{y}) g(\mathbf{y}) d\mathbf{y}$$

For each \mathbf{x} , we compute the integral between the filter g and a shifted signal f

2016-04-19

11

Adaptive filtering

- If we apply an adaptive (or position dependent, or space variant) filter $g_{\mathbf{x}}$, the operation cannot be expressed as a convolution, but instead as

$$h(\mathbf{x}) = \int f(\mathbf{x} - \mathbf{y}) g_{\mathbf{x}}(\mathbf{y}) d\mathbf{y}$$

For each \mathbf{x} , we compute the integral between a shifted signal f and the filter $g_{\mathbf{x}}$ where the filter depends on \mathbf{x}

2016-04-19

12

How to choose g_x ?

According to the previous discussion, we choose g_x such that:

- It contains a low-pass component that maintains the local image mean intensity Independent of x
- It contains a high-pass component that depends on the local signal structure Dependent of x
- Also: the resulting operation for computing h should be simple to implement Computational efficient

2016-04-19

13


High-frequency components in g_x

- If the signal is \approx 1D the filter can maintain the signal by reducing the frequency components orthogonal to the local structure
- The human visual system is less sensitive to noise along linear structures than to noise in the orthogonal direction
- Results in good subjective improvement of image quality

2016-04-19

14


Oriented noise



2016-04-19

15


Oriented noise



2016-04-19

16

Oriented noise



2016-04-19

17


Oriented noise



2016-04-19

18

Oriented noise



Edges and lines


- A. Without noise
- B. With oriented noise along
- C. With isotropic noise
- D. With oriented noise across

19

Local structure information

- We compute the local orientation tensor $\mathbf{T}(\mathbf{x})$ at all points \mathbf{x} to control / steer $g_{\mathbf{x}}$
- At a point \mathbf{x} that lies in a locally 1D region, we obtain

$$\mathbf{T}(\mathbf{x}) = A \hat{\mathbf{e}} \hat{\mathbf{e}}^T$$



$\hat{\mathbf{e}}$ is normal to the linear structure

2016-04-19

20

Ansatz for g_x

We apply a filter that is given in the Fourier domain as

$$G_{HP}(\mathbf{u}) = G_\rho(u) (\hat{\mathbf{u}}^T \hat{\mathbf{e}})^2 \quad \mathbf{u} = u \hat{\mathbf{u}}$$

- G_{HP} is polar separable
- It attenuates frequency components that are \perp to $\hat{\mathbf{e}}$
- It maintains all frequency components that are \parallel to $\hat{\mathbf{e}}$

How to implement g_x ?

- We know that [EDUPACK – ORIENTATION]

$$(\hat{\mathbf{u}}^T \hat{\mathbf{e}})^2 = \langle \hat{\mathbf{u}} \hat{\mathbf{u}}^T | \hat{\mathbf{e}} \hat{\mathbf{e}}^T \rangle = \langle \hat{\mathbf{u}} \hat{\mathbf{u}}^T | \mathbf{T}(\mathbf{x}) \rangle$$

where $\mathbf{T}(\mathbf{x}) = \hat{\mathbf{e}} \hat{\mathbf{e}}^T$ (assume $A = 1!$)

- Using a N -D tensor basis $\hat{\mathbf{N}}_k = \hat{\mathbf{n}}_k \hat{\mathbf{n}}_k^T$ and its dual $\tilde{\mathbf{N}}_k$, we obtain:

$$\mathbf{T}(\mathbf{x}) = \sum_{k=1}^N \langle \mathbf{T}(\mathbf{x}) | \tilde{\mathbf{N}}_k \rangle \hat{\mathbf{N}}_k$$

How to implement g_x ?

$$\begin{aligned} (\hat{\mathbf{u}}^T \hat{\mathbf{e}})^2 &= \sum_{k=1}^N \langle \hat{\mathbf{u}} \hat{\mathbf{u}}^T | \hat{\mathbf{N}}_k \rangle \langle \mathbf{T}(\mathbf{x}) | \tilde{\mathbf{N}}_k \rangle \\ &= \sum_{k=1}^N \langle \hat{\mathbf{u}} \hat{\mathbf{u}}^T | \hat{\mathbf{n}}_k \hat{\mathbf{n}}_k^T \rangle \langle \mathbf{T}(\mathbf{x}) | \tilde{\mathbf{N}}_k \rangle \\ &= \sum_{k=1}^N (\hat{\mathbf{u}}^T \hat{\mathbf{n}}_k)^2 \langle \mathbf{T}(\mathbf{x}) | \tilde{\mathbf{N}}_k \rangle \end{aligned}$$

depends on \mathbf{u}
but not on $\hat{\mathbf{e}}$

depends on $\hat{\mathbf{e}}$
but not on \mathbf{u}

How to implement g_x ?

- Plug this into the expression for G_{HP} :

$$G_{HP}(\mathbf{u}) = \sum_{k=1}^N \langle \mathbf{T}(\mathbf{x}) | \tilde{\mathbf{N}}_k \rangle G_\rho(u) (\hat{\mathbf{u}}^T \hat{\mathbf{n}}_k)^2$$

depends on $\hat{\mathbf{e}}$
but not on \mathbf{u}

depends on \mathbf{u}
but not on $\hat{\mathbf{e}}$

How to implement g_x ?

Consequently, the filter G_{HP} is a **linear combination** of N filters, where each filter has a Fourier transform:

$$G_{HP,k}(\mathbf{u}) = G_\rho(u) (\hat{\mathbf{u}}^T \hat{\mathbf{n}}_k)^2$$

Independent of \mathbf{x}

and N scalars:

$$\langle \mathbf{T}(\mathbf{x}) | \tilde{\mathbf{N}}_k \rangle$$

Dependent of \mathbf{x}

2016-04-19

25

How to implement g_x ?

Summarizing, the adaptive filter can be written as

$$g_x = g_{LP} + \sum_{k=1}^N \langle \mathbf{T}(\mathbf{x}) | \tilde{\mathbf{N}}_k \rangle g_{HP,k}$$

A fixed LP-filter

N fixed HP-filters

N position dependent scalars

2016-04-19

26

How to implement g_x ?

If the filter is applied to a signal, we obtain

$$\begin{aligned} h(\mathbf{x}) &= \int f(\mathbf{x} - \mathbf{y}) \left[g_{LP}(\mathbf{y}) + \sum_{k=1}^N \langle \mathbf{T}(\mathbf{x}) | \tilde{\mathbf{N}}_k \rangle g_{HP,k}(\mathbf{y}) \right] d\mathbf{y} \\ &= (f * g_{LP})(\mathbf{x}) + \sum_{k=1}^N \langle \mathbf{T}(\mathbf{x}) | \tilde{\mathbf{N}}_k \rangle (f * g_{HP,k})(\mathbf{x}) \end{aligned}$$

Standard convolution

Standard convolutions

Position dependent scalars

2016-04-19

27

Outline of method, version 1

1. Estimate the local orientation tensor $\mathbf{T}(\mathbf{x})$ at each image point \mathbf{x}
2. Apply a number of fixed filters to the image; one LP-filter g_{LP} and the N HP-filters $g_{HP,k}$
3. At each point \mathbf{x} :
 1. Compute the N scalars $\langle \mathbf{T}(\mathbf{x}) | \tilde{\mathbf{N}}_k \rangle$
 2. Form the linear combination of the N HP-filter responses and the N scalars and add the LP-filter response
4. At each point \mathbf{x} , the result is the filter response $h(\mathbf{x})$ of the locally adapted filter g_x

The filter g_x is also called a *steerable filter*

2016-04-19

28

Observation

- \mathbf{T} can be estimated for any image dimension
- The filters g_{LP} and $g_{HP,k}$ can be formulated for any image dimension

⇒ The method can be implemented for any dimension of the signal (2D, 3D, 4D, ...)

Remaining questions

1. What happens in regions that are not i1D, i.e., if \mathbf{T} has not rank 1?
2. What happens if $A \neq 1$?
3. How to choose the radial function G_ρ ?

Non i1D signals

- The tensor's eigenvectors with non-zero eigenvalues span the subspace of the Fourier domain that contains the signal's energy
- **Equivalent:** For a given local region with orientation tensor \mathbf{T} , let $\hat{\mathbf{u}}$ define an arbitrary orientation. The product $\hat{\mathbf{u}}^T \mathbf{T} \hat{\mathbf{u}}$ is a measure of how much of this orientation the region contains.

Non i1D signals

- But

$$\hat{\mathbf{u}}^T \mathbf{T} \hat{\mathbf{u}} = \langle \hat{\mathbf{u}} \hat{\mathbf{u}}^T | \mathbf{T} \rangle$$

which means that the adaptive filtering should work in general, even if the signal is non i1D

How about $A = 1$?

- Previously we assumed $A = 1$, but normally A depends on the local amplitude of the signal (depends on \mathbf{x})
- In order to achieve $A = 1$, \mathbf{T} must be pre-processed
- The resulting tensor is called the **control tensor \mathbf{C}**

2016-04-19

33

Pre-processing of \mathbf{T}

- The filter g_x is supposed to vary slowly with \mathbf{x} , but \mathbf{T} contains high-frequency noise that comes from the image noise
- This noise can be reduced by an initial LP-filtering of \mathbf{T} (i.e., of its elements)
- The result is denoted \mathbf{T}_{LP}

2016-04-19

34

Pre-processing of \mathbf{T}

\mathbf{T}_{LP} must then be *normalized*:

Eigen-decomposition of \mathbf{T}_{LP}

$$\mathbf{T}_{LP} = \sum_{k=1}^n \lambda_k \hat{\mathbf{e}}_k \hat{\mathbf{e}}_k^T \quad \lambda_k \geq \lambda_{k+1}$$

$$\mathbf{C} = \sum_{k=1}^n \gamma_k \hat{\mathbf{e}}_k \hat{\mathbf{e}}_k^T \quad \gamma_k \geq \gamma_{k+1}$$

$$\gamma_k = \gamma_k(\lambda_1, \dots, \lambda_n)$$


Same eigenvectors as \mathbf{T}_{LP} but different eigenvalues

2016-04-19

35

Modification of the eigenvalues

Examples of γ_1 as function of, e.g., $\|\mathbf{T}\| = \sqrt{\lambda_1^2 + \lambda_2^2 + \dots}$




2016-04-19

36


Modification of the eigenvalues

Examples of γ_{k+1}/γ_k as function of λ_{k+1}/λ_k




The radial function G_ρ

- Should “mainly” be equal to 1
- Should tend to 0 for $u = \pi$
- Together with the LP-filter g_{LP} : an all-pass filter



The adaptive filter in 2D




Examples of $G(\mathbf{u})$ for different $\mathbf{C}(\mathbf{x})$

Outline of method, version 2

1. Estimate the local tensor in each image point: $\mathbf{T}(\mathbf{x})$
2. LP-filter the tensor: $\mathbf{T}_{LP}(\mathbf{x})$
3. In each image point:
 1. Compute the eigenvalues and eigenvectors of $\mathbf{T}_{LP}(\mathbf{x})$.
 2. Map the eigenvalues λ_k to γ_k .
 3. Re-combine γ_k and the eigenvectors to form the control tensor \mathbf{C}
 4. Compute the scalars $\langle \mathbf{C} | \hat{\mathbf{N}}_k \rangle$ for all $k = 1, \dots, N$
4. Filter the image with g_{LP} and the N HP-filters $g_{HP,k}$
5. In each image point: form the linear combination of the filter responses and the scalars


Example

Original noisy image



2016-04-19

Image after enhancement



41

Example

QuickTime™ and a YUV420 codec decompressor are needed to see this picture.

2016-04-19

42

An iterative method

- Adaptive filtering can be iterated for reducing the noise
- If the filter size is reduced at the same time, a close-to continuous transition is achieved (**evolution**)
- This leads to another method for image enhancement: *anisotropic diffusion*

2016-04-19

43

Scale space recap (from lecture 2)

- The linear Gaussian scale space related to the image f is a family of images $L(x,y;s)$

$$L(x, y; s) = (g_s * f)(x, y)$$

Convolution over (x,y) only!

parameterized by the scale parameter s , where

$$g_s(x, y) = \frac{1}{2\pi s} e^{-\frac{1}{2s}(x^2+y^2)}$$

A Gaussian LP-filter with $\sigma^2 = s$

Note: $g_s(x,y) = \delta(x,y)$ for $s = 0$

2016-04-19

44

Scale space recap (from lecture 2)

- $L(x,y;s)$ can also be seen as the solution to the PDE

$$\frac{\partial}{\partial s} L = \frac{1}{2} \nabla^2 L$$

$$\frac{\partial}{\partial s} L = \frac{1}{2} \left(\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} \right) L$$

with boundary condition $L(x,y;0) = f(x,y)$

The diffusion equation

Example:

L = temperature

s = time

2016-04-19

45

Image enhancement based on linear diffusion

- This means that $L(x,y;s)$ is an LP-filtered version of $f(x,y)$ for $s > 0$.
- The larger s is, the more LP-filtered is f
 - High-frequency noise will be removed for larger s
- As before: also high-frequency image components (e.g. edges) will be removed
- We need to control the diffusion process such that edges remain
 - How?

2016-04-19

46

Step 1

- Modify the PDE by introducing a parameter μ :

$$\frac{\partial}{\partial s} L = \frac{\mu}{2} \nabla^2 L$$

- This PDE is solved by

$$L(x, y; s) = (g_s * f)(x, y)$$

Same as before

$$g_s(x, y) = \frac{1}{2\pi\mu s} e^{-\frac{1}{2\mu s}(x^2 + y^2)}$$

Slightly different

μ can be seen as a "diffusion speed":

Small μ : the diffusion process is slow when s increases

Large μ : the diffusion process is fast when s increases

2016-04-19

47

Step 2

- We want the image content to control μ
 - In flat regions: fast diffusion (large μ)
 - In non-flat region: slow diffusion (small μ)
- We need to do *space variant* diffusion
 - μ is a function of position (x,y)

Compare to the space variant filter g_x in adaptive filtering

2016-04-19

48

Inhomogeneous diffusion


- Perona & Malik suggested to use

$$\mu(x, y) = \frac{1}{1 + |\nabla f|^2 / \lambda^2}$$

where ∇f is the image gradient at (x, y) and λ is fixed a parameter

- Close to edges: $|\nabla f|$ is large $\Rightarrow \mu$ is small
- In flat regions: $|\nabla f|$ is small $\Rightarrow \mu$ is large

Inhomogeneous diffusion



Inhomogeneous diffusion

- Example



- Noise is effectively removed in flat region 😊
- Edges are preserved 😊
- Noise is preserved close to edges 😞

We want to be able to LP-filter along but not across edges, same as for adaptive filtering

Step 3

- The previous PDEs are all isotropic
⇒ The resulting filter g is isotropic
- The last PDE can be written:

$$\frac{\partial}{\partial s} L = \frac{\mu}{2} \nabla^2 L = \frac{1}{2} \operatorname{div}(\mu \operatorname{grad} L)$$

Divergence of (...)
maps 2D vector field to scalar field

Gradient of L ,
a 2D vector field

2016-04-19

53

Step 3

- Change μ from a scalar to a 2×2 symmetric matrix \mathbf{D}

$$\frac{\partial}{\partial s} L = \frac{1}{2} \operatorname{div}(\mathbf{D} \operatorname{grad} L)$$

- The solution is now given by

$$L(\mathbf{x}; s) = (g_s * f)(\mathbf{x})$$

← Same as before

$$g_s(\mathbf{x}) = \frac{1}{2\pi \det(\mathbf{D})^{1/2} s} e^{-\frac{1}{2s} \mathbf{x}^T \mathbf{D}^{-1} \mathbf{x}}$$

2016-04-19

54

Anisotropic diffusion

- The filter g is now anisotropic, i.e., not necessary circular symmetric
- The shape of g depends on eigensystem of \mathbf{D}
- \mathbf{D} is called a *diffusion tensor*
 - Can be given a physical interpretation, e.g. for anisotropic heat diffusion

2016-04-19

55

The diffusion tensor

- Since \mathbf{D} is symmetric 2×2 :

$$\mathbf{D} = \alpha_1 \mathbf{e}_1 \mathbf{e}_1^T + \alpha_2 \mathbf{e}_2 \mathbf{e}_2^T$$

where α_1, α_2 are the eigenvalues of \mathbf{D} , and \mathbf{e}_1 and \mathbf{e}_2 are corresponding eigenvectors

\mathbf{e}_1 and \mathbf{e}_2 form an ON-basis

2016-04-19


56

The filter g

- The corresponding shape of g is given by

The width of the filter in direction e_k is given by α_k

Iso-curves for $g \Rightarrow$



2016-04-19

57

Step 4

- We want g to be narrow across edges and wide along edges
- This means: \mathbf{D} should depend on (x,y)
 - A space variant anisotropic diffusion
- This is referred to as *anisotropic diffusion* in the literature
- Introduced by Weickert

2016-04-19

58

Anisotropic diffusion

- Information about edges and their orientation can be provided by an orientation tensor, e.g., the structure tensor \mathbf{T} in terms of its eigenvalues λ_1, λ_2
- However:
 - We want α_k to be close to 0 when λ_k is large
 - We want α_k to be close to 1 when λ_k is close to 0

2016-04-19

59


From \mathbf{T} to \mathbf{D}

- The diffusion tensor \mathbf{D} is obtained from the orientation tensor \mathbf{T} by modifying the eigenvalues and keeping the eigenvectors
 - similar to the control tensor \mathbf{C} , e.g.

$$\alpha_k = \exp(-\lambda_k/k)$$

For example

k is a control parameter



2016-04-19

60

Anisotropic diffusion: putting things together

1. At all points:
 1. compute a local orientation tensor $\mathbf{T}(\mathbf{x})$
 2. compute $\mathbf{D}(\mathbf{x})$ from $\mathbf{T}(\mathbf{x})$
2. Apply anisotropic diffusion onto the image by locally iterating

Left hand side:
the change in L
at (x,y) between
 s and $s+\partial s$

$$\frac{\partial}{\partial s} L = \frac{1}{2} \operatorname{div}(\mathbf{D} \nabla L)$$

Right hand side:
can be computed
locally at each
point (x,y)

This defines how scale space level
 $L(x,y;s+\partial s)$ is generated from $L(x,y;s)$

61

Implementation aspects

- The anisotropic diffusion iterations can be done with a constant diffusion tensor field $\mathbf{D}(\mathbf{x})$, computed once from the original image (faster)
- Alternatively: re-compute $\mathbf{D}(\mathbf{x})$ between every iteration (slower)

2016-04-19

62

Regularization

- We assume \mathbf{D} to have a slow variation with respect to \mathbf{x} (cf. adaptive filtering)
- This means

$$\frac{\partial}{\partial s} L \approx \frac{1}{2} \operatorname{tr} [\mathbf{D} (\operatorname{div} \operatorname{grad} L)] = \frac{1}{2} \operatorname{tr} [\mathbf{D} (\mathbf{H} L)]$$

The Hessian of L = second order derivatives of L

$$\mathbf{H} L = \begin{pmatrix} \frac{\partial^2}{\partial x^2} L & \frac{\partial^2}{\partial x \partial y} L \\ \frac{\partial^2}{\partial x \partial y} L & \frac{\partial^2}{\partial y^2} L \end{pmatrix}$$

2016-04-19

63

Numerical implementation

- Several numerical schemes for implementing anisotropic diffusion exist
- Simplest one:
 - Replace all partial differentials with finite differences

$$L(x, y; s + \Delta s) = L(x, y; s) + \Delta s \operatorname{tr} [\mathbf{D}(\mathbf{H}L)]$$

The Hessian of L
can be approximated by
convolving L with:

$$H_{11} : \begin{pmatrix} 0 & 0 & 0 \\ 1 & -2 & 1 \\ 0 & 0 & 0 \end{pmatrix} \quad H_{12} : \begin{pmatrix} \frac{1}{4} & 0 & -\frac{1}{4} \\ 0 & 0 & 0 \\ -\frac{1}{4} & 0 & \frac{1}{4} \end{pmatrix} \quad H_{22} : \begin{pmatrix} 0 & 1 & 0 \\ 0 & -2 & 0 \\ 0 & 1 & 0 \end{pmatrix}$$

2016-04-19

64

Algorithm Outline

1. Set parameters
e.g.: k , Δs , number of iterations, ...
2. Iterate
 1. Compute orientation tensor \mathbf{T}
 2. Modify eigenvalues $\Rightarrow \mathbf{D}$
 3. Computer Hessian $\mathbf{H} L$
 4. Update L according to:

$$L(x, y; s + \Delta s) = L(x, y; s) + \Delta s \operatorname{tr} [\mathbf{D}(\mathbf{H}L)]$$

2016-04-19

65

A note

- The image f is never convolved by the space variant anisotropic filter g
- Instead, the effect of g is generated incrementally based on the diffusion eq.
- In adaptive filtering: we never convolve f with g_x , instead several fixed filters are applied onto f and their results are combined in a non-linear fashion

2016-04-19

66

Comparison

Inhomogenous diffusion



Anisotropic diffusion

