Topic 10:

Gaussian & Laplacian Pyramids

- The Gaussian pyramid
- Constructing the Gaussian pyramid
 - The REDUCE() function
- Constructing the Laplacian pyramid
 - The EXPAND() function
- Applications

The Gaussian Pyramid

Gaussian Pyramid Representation of the Photo



Original photo \mathbb{S}

The Gaussian Pyramid

Goal: Develop representation to decompose images into information at multiple scales, to extract feature or structure of interest, to attenuate noise.

Applications:

- Efficient image coding
- Progressive transmission
- Image blending
- Image enhancement
- Efficient Processing
- ...too many to list here!



Application: Pyramid Image Blending

Goal: Merge two images without visible seams







Horror Photo



© prof. dmartin

The Gaussian Pyramid

Goal: Develop representation to decompose images into information at multiple scales, to extract feature or structure of interest, to attenuate noise.

Input:

Image *I* of size $(2^N + 1) \times (2^N + 1)$

Output:

- *N* images $g_0, ..., g_{N-1}$
- g_l has size

$$(2^{N-l}+1) \times (2^{N-l}+1)$$



Why is it Called a Pyramid?

Idea: Representation can be pictured as a "pyramid" of $3x3, 5x5, 9x9, \dots, (2^{N}+1)x(2^{N}+1)$ images



The Gaussian Pyramid

The representation is based on 2 basic operations:

1. Smoothing

Smooth the image with a sequence of smoothing filters, each of which has twice the radius of the previous one.

2. Downsampling

Reduce image size by half after each smoothing.



Topic 10:

Gaussian & Laplacian Pyramids

- The Gaussian pyramid
- Constructing the Gaussian pyramid
 - The REDUCE() function
- Constructing the Laplacian pyramid
 - The EXPAND() function
- Applications

The Gaussian Pyramid

The representation is based on 2 basic operations:

1. Smoothing

Smooth the image with a sequence of smoothing filters, each of which has twice the radius of the previous one.

2. Downsampling

Reduce image size by half after each smoothing.



Original photo g_0

$$\hat{g}_1 = w^* g_0$$

$$\hat{g}_1 = x^* g_0$$

$$\hat{g}_1 = x^* g_0$$

$$\hat{g}_1 = x^* g_0$$

$$\hat{g}_1 = x^* g_0$$



$$\hat{g}_1(i,j) = \sum_{m=-2}^2 \sum_{n=-2}^2 w(m,n) \cdot g_0(i-m,j-n)$$

Operation #1: Smooth Image at N-1 Scales ĝ slide over $\hat{g}_2 = w * \hat{g}_1$ W image $= w * (w * g_0)$ $= (w * w) * g_0$ can be thought of

as a filter h = w * w whose radius is **twice** that of w.



$$\hat{g}_3 = w^* \hat{g}_2$$
$$= (w^* w^* w)^* g_0$$

ĝ₂ ↓ ↓ slide over image

radius is 4 times that of w



Operation #1: Smooth Image at N-1 Scales

$$\hat{g}_4 = w^* \hat{g}_3$$

 $= (w^* w^* w^* w)^* g_0$







Smoothing Filter in 1D: Derivation from 4 Criteria

- 1. \hat{w} always has 5 elements (aka "5-tap" filter)
- 2. \hat{w} symmetric about

$$\hat{w} = \begin{bmatrix} c \\ b \\ a \\ b \\ c \end{bmatrix} = \begin{bmatrix} 1/4 - a/2 \\ 1/4 \\ a \\ 1/4 \\ 1/4 - a/2 \end{bmatrix}$$

3. applying \hat{w} to a **constant** image does not change it $\sum_{m=-2} \hat{w}(m) = 1$ $\iff a + 2b + 2c = 1$



4. Equal contribution a + 2c = 2b = 1/2

To satisfy Criteria 1-4 we have a 2 equations & 3 unknowns \Rightarrow *a* remains a free parameter $\hat{w}(2) = \hat{w}(-2) = \frac{1}{4} - \frac{a}{2}, \hat{w}(-1) = \hat{w}(1) = \frac{1}{4}$ usually $a \in [0.3, 0.6]$

Defining the Smoothing Filter in 2D





Exploiting separability to compute \hat{g}_1 :

- 1. Convolve each row of g_0 with \hat{W} .
- 2. Convolve the columns fo the result with \hat{W} again.

Operation #2: Downsample the Smoothed Image

Since \hat{g}_1 contains less image detail, we downsample it by 2 (ie. store every other pixel)

$$g_1(i,j) = \hat{g}_1(2i,2j)$$



 g_1

2^N + 1

 $2^{N-1} + 1$

Topic 10:

Gaussian & Laplacian Pyramids

- The Gaussian pyramid
- Constructing the Gaussian pyramid
 - The REDUCE() function
- Constructing the Laplacian pyramid
 - The EXPAND() function
- Applications

Operations #1 & #2: The REDUCE() Function

Smoothing & downsampling are combing into a single REDUCE function.

 $w * g_l$

 g_l



 $-2^{N-l}+1 \longrightarrow$

 $g_{l+1} = \text{REDUCE}(g_l)$

 $g_{l+1}(i,j) = \sum_{l=1}^{\infty} \sum_{l=1}^{\infty} w(m,n) \cdot g_l(2i-m,2j-n)$ m = -2 n = -2



| downsample $\times 2$



 g_{l+1}

 $\bullet 2^{N-l-1} + 1 \bullet$





REDUCE(g_0)











The 1D REDUCE() function in matrix notation

$$\begin{bmatrix} \vdots \\ g_1 \\ \vdots \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{bmatrix} \begin{bmatrix} c & b & a & b & c & 0 & \cdots & 0 \\ 0 & c & b & a & b & c & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & \cdots & c \end{bmatrix} \begin{bmatrix} \vdots \\ \vdots \\ g_0 \\ \vdots \\ \vdots \end{bmatrix}$$

downsampling matrix D₀

convolution matrix C₀

$$g_1 = D_0 \cdot C_0 \cdot g_0$$

General expression: $g_{l+1} = D_l \cdot C_l \cdot g_l$

The Gaussian Pyramid

Goal: Develop representation to decompose images into information at multiple scales, to extract feature or structure of interest, to attenuate noise.

Input:

Image *I* of size $(2^N + 1) \times (2^N + 1)$

Output:

- *N* images $g_0, ..., g_{N-1}$
- g_l has size

$$(2^{N-l}+1) \times (2^{N-l}+1)$$



Operations #1 & #2: The REDUCE() Function

Smoothing & downsampling are combing into a single REDUCE function.

 $w * g_l$

 g_l



 $-2^{N-l}+1 \longrightarrow$

 $g_{l+1} = \text{REDUCE}(g_l)$

 $g_{l+1}(i,j) = \sum_{l=1}^{\infty} \sum_{l=1}^{\infty} w(m,n) \cdot g_l(2i-m,2j-n)$ m = -2 n = -2



| downsample $\times 2$



 g_{l+1}

 $\bullet 2^{N-l-1} + 1 \bullet$

What Does Smoothing Take Away?

 $g_0 = I$ /
original
photo

What Does Smoothing Take Away?

smoothed photo

 $\hat{g}_1 = w * g_0$



What Does Smoothing Take Away?

 $g_0 - \hat{g}_1$



Details in g_0 that were **not** represented in \hat{g}_1 .

Topic 10:

Gaussian & Laplacian Pyramids

- The Gaussian pyramid
- Constructing the Gaussian pyramid
 - The REDUCE() function
- Constructing the Laplacian pyramid
 - The EXPAND() function
- Applications

The Laplacian Pyramid

Idea:

Rather than store the smoothed images, store only **the difference** between levels g_l and g_{l+1}


Idea:

Rather than store the smoothed images, store only **the difference** between levels g_l and g_{l+1}

most values are 0 (gray) ⇒ can be quantized and represented with few bits



Idea:

Rather than store the smoothed images, store only **the difference** between levels g_l and g_{l+1}

- To do this, we must compare adjacent levels g_l and g_{l+1}
- But g_l, g_{l+1} are not the same size!
 - \Rightarrow **Expand** g_{I+1} to

make it equal size to g_l



Operation #3: The EXPAND() Function



 $\text{EXPAND}(g_l)$

EXPAND
$$(g_l) = 4 \sum_{m=-2}^{2} \sum_{n=-2}^{2} w(m,n) \cdot g_l\left(\frac{i-m}{2}, \frac{j-n}{2}\right)$$

The EXPAND() function

The function upsamples level g_l by doubling its size from $(2^{N-1}+1) \times (2^{N-1}+1)$ to $(2^{N-1}+1) \times (2^{N-1}+1)$



Generalizing expression 1D

The EXPAND() function

The function upsamples level g_l by doubling its size from $(2^{N-l}+1) \times (2^{N-l}+1)$ to $(2^{N-l+1}+1) \times (2^{N-l+1}+1)$





$L_1 = g_1 - \text{EXPAND}(g_2)$



$L_2 = g_2 - \text{EXPAND}(g_3)$



Often we use a truncated Laplacian pyramid by storing images L_0 , L_1 , ..., L_k , g_{k+1} for k + 1 < N

Base case: store g_N



The Laplacian Image Pyramid

Idea: Represent g_{I} image by a g_{I+1} image and a detail image (called the Laplacian L₁ image) whose size is equal to the g_1 image

g



 $EXPAND(g_{I+1})$

The Laplacian Image Pyramid

Idea: This decomposition can be repeated several times!!

g





Laplacian L_I



+

The Laplacian Image Pyramid

Idea: This decomposition can be repeated several times!!



Given an input image of size $(2^{N}+1) \times (2^{N}+1)$, the Laplacian pyramid representation consists of • L₀, ..., L_{k-1}

• g_k for some $k \le N$



Transmission using EXPAND



5. Transmit L_{k-4}

4. Transmit L_{k-3}

3. Transmit L_{k-2}

2. Transmit L_{k-1}

Topic 10:

Gaussian & Laplacian Pyramids

- The Gaussian pyramid
- Constructing the Gaussian pyramid
 - The REDUCE() function
- Constructing the Laplacian pyramid
 - The EXPAND() function
- Applications
 - Multi-resolution image blending
 - Multi-resolution image editing

Application #2: Multi-Resolution Image Editing

How can we achieve the painting operations below?

Original photo



Adding makeup

Adding a glint

Final photo

(Berman, Bartell, Salesin, SIGGRAPH'94)

Image Editing

Approach #1

 Modify the image pixels directly, using a brush tool



Image Editing

Approach #1

 Modify the image pixels directly, using a brush tool



Difficulty: Manipulating pixel colors directly will distort/flatten fine scale detail

Image Editing

Approach #1

 Modify the image pixels directly, using a brush tool



Difficulty: Manipulating pixel colors directly will distort/flatten fine scale detail Approach #2

- Edit a lowerresolution version of the photo (e.g. level g_k of the Gauss pyramid for k > 0)
- How do we create an edited full-res photo?

Pyramid-Based Image Editing

1. Edit g_1 to obtain \tilde{g}_1 .

2. Add the details using $L_{
m 0}$

Definition of L_0 :

$$L_0 = g_0 - \text{EXPAND}(g_1)$$



Pyramid-Based Image Editing





Multi-Resolution Image Editing



Adding makeup

The procedure allows us to apply "paint" to an image while preserving image detail.

Original photo

Multi-Resolution Image Editing

(from Berman et al)



PLATE 4 Adding smog to an image with a range in scale of 100,000 to 1: (a) a sunny day in the park; (b) zooming out by a factor of 100,000; (c) adding smog over Chicago; (d) smog affects the park view. (See Chapter 4: Multiresolution Image Editing.)

Note: Authors used a wavelet pyramid rather than a Gaussian pyramid, but the principle is the same.

Topic 10:

Hierarchical image representations

- The Gaussian pyramid
- Constructing the Gaussian pyramid
 - The REDUCE() function
- Constructing the Laplacian pyramid
 - The EXPAND() function
- Applications
 - Multi-resolution image blending
 - Multi-resolution image editing

Image Blending



Feathering







Assign an alpha value to each pixel near the seam to make it less visible.

Effect of Window Size



α



Ghosting is most apparent here

Effect of Window Size







seam most visible here

α

0

Image Blending



Application #1: Pyramid Blending Algorithm

Input: Source images A, B & binary matte M Output: Blended image S

Α







Input: Source images A, B & binary matte M Output: Blended image S

Α



Pad to make size $(2^N + 1) \times (2^N + 1)$





Input: Source images A, B & binary matte M Output: Blended image S

Α







Input: Source images A, B & binary matte M Output: Blended image S







3. Compute M's Gaussian pyramid: Mg₀,..., Mg_N 2. Compute B's Laplacian pyramid: BL₀,..., BL_{N-1}, Bg_N

4. Compute the Laplacian pyramid, SL_0 , SL_1 ,..., SL_{N-1} , Sg_N by applying the matting equation with matte Mg_l .

 $SL_{l}(i,j) = Mg_{l}(i,j)AL_{l}(i,j) + (1 - Mg_{l}(i,j))BL_{l}(i,j)$



4. Compute the Laplacian pyramid, SL_0 , SL_1 ,..., SL_{N-1} , Sg_N by applying the matting equation with matte Mg_l .

 $SL_{l}(i,j) = Mg_{l}(i,j)AL_{l}(i,j) + (1 - Mg_{l}(i,j))BL_{l}(i,j)$



The algorithm effectively uses a different alpha matte for each level of detail.

↑ detail \Rightarrow ↓ feathering window ↓ detail \Rightarrow ↑ feathering window



5. Compute level 0 of the Gaussian pyramid of S from SL_0 , $SL_1, \ldots SL_{N-1}$, Sg_N .






Pyramid Blending Algorithm



Blending Mis-Matched Photos (still looks OK)



Merging Mis-Matched Photos (no blend)



Blending without Using Pyramid



Pyramid Blending





Horror Photo



© prof. dmartin