


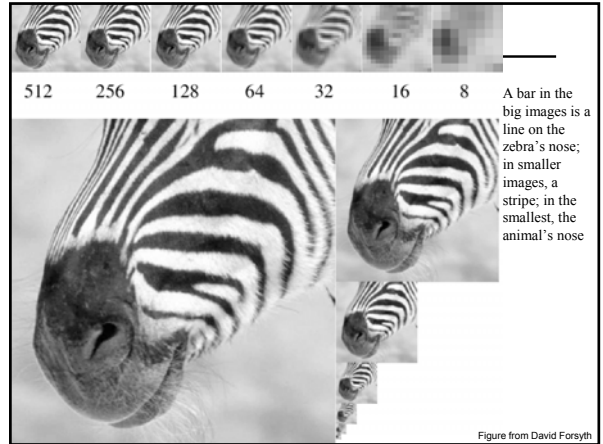
Image Pyramids

Idea: Represent $N \times N$ image as a "pyramid" of $1 \times 1, 2 \times 2, 4 \times 4, \dots, 2^k \times 2^k$ images (assuming $N=2^k$)



Known as a **Gaussian Pyramid** [Burt and Adelson, 1983]

- In computer graphics, a *mip map* [Williams, 1983]
- A precursor to *wavelet transform*



A bar in the big images is a line on the zebra's nose; in smaller images, a stripe; in the smallest, the animal's nose

Figure from David Forsyth

What are they good for?

Improve Search

- Search over translations
 - Classic coarse-to-fine strategy
- Search over scale
 - Template matching
 - E.g. find a face at different scales


Precomputation

- Need to access image at different blur levels
- Useful for texture mapping at different resolutions (called mip-mapping)

Image Processing


- Editing frequency bands separately
- E.g. image blending...

Example application: CMU face detector



From: <http://www.ius.cs.cmu.edu/IUS/har2/har/www/CMU-CS-95-158R/>

Gaussian pyramid construction



Repeat

- Filter
- Subsample

Until minimum resolution reached

- can specify desired number of levels (e.g., 3-level pyramid)


The whole pyramid is only $4/3$ the size of the original image!

Image sub-sampling



Throw away every other row and column to create a $1/2$ size image - called *image sub-sampling*

Image sub-sampling

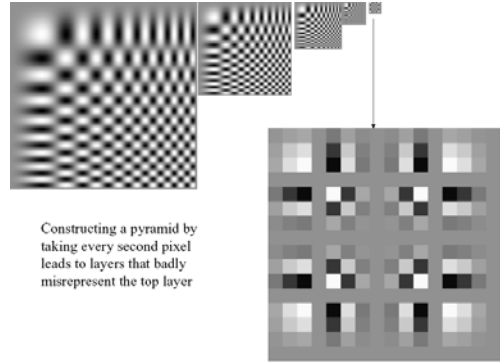


1/2

1/4 (2x zoom)


1/8 (4x zoom)

Why does this look so bad?



Constructing a pyramid by taking every second pixel leads to layers that badly misrepresent the top layer

Sampling




Good sampling:
-Sample often or,
-Sample wisely

Bad sampling:
-see aliasing in action!


Really bad in video

Imagine a spoked wheel moving to the right (rotating clockwise). Mark wheel with dot so we can see what's happening. If camera shutter is only open for a fraction of a frame time (frame time = 1/30 sec. for video, 1/24 sec. for film):




Without dot, wheel appears to be rotating slowly backwards! (counterclockwise)

Alias: n., an assumed name




Picket fence receding into the distance will produce aliasing...

WHY?



Not enough samples

Input signal:
Matlab output:




`x = 0:0.05:5; imagesc(sin((2.*x).^x))`


Smoothing as low-pass filtering

- The message of the FT is that high frequencies lead to trouble with sampling.
- A filter whose FT is a box is bad, because the filter kernel has infinite support
- Common solution: use a Gaussian
 - multiplying FT by Gaussian is equivalent to convolving image with Gaussian.
- Solution: suppress high frequencies before sampling
 - multiply the FT of the signal with something that suppresses high frequencies
 - or convolve with a low-pass filter


Sampling without smoothing. Top row shows the images, sampled at every second pixel to get the next; bottom row shows the magnitude spectrum of these images.



Sampling with smoothing. Top row shows the images. We get the next image by smoothing the image with a Gaussian with sigma 1 pixel, then sampling at every second pixel to get the next; bottom row shows the magnitude spectrum of these images.



Sampling with smoothing. Top row shows the images. We get the next image by smoothing the image with a Gaussian with sigma 1.4 pixels, then sampling at every second pixel to get the next; bottom row shows the magnitude spectrum of these images.



Gaussian pre-filtering



G 1/8

G 1/4

Gaussian 1/2

Solution: filter the image, *then* subsample

Subsampling with Gaussian pre-filtering



Gaussian 1/2

G 1/4

G 1/8

Solution: filter the image, *then* subsample

- Filter size should double for each 1/2 size reduction. Why?
- How can we speed this up?

Compare with...




1/2

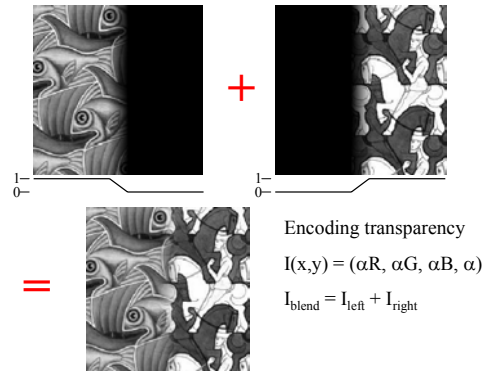
1/4 (2x zoom)

1/8 (4x zoom)


Image Blending




Feathering



Affect of Window Size



Affect of Window Size



Good Window Size



“Optimal” Window: smooth but not ghosted

What is the Optimal Window?

To avoid seams


- window \geq size of largest prominent feature

To avoid ghosting


- window $\leq 2 \times$ size of smallest prominent feature

Natural to cast this in the *Fourier domain*

- largest frequency $\leq 2 \times$ size of smallest frequency
- image frequency content should occupy one “octave” (power of two)



What if the Frequency Spread is Wide



Idea (Burt and Adelson)

- Compute $F_{\text{left}} = \text{FFT}(I_{\text{left}})$, $F_{\text{right}} = \text{FFT}(I_{\text{right}})$
- Decompose Fourier image into octaves (bands)
 - $F_{\text{left}} = F_{\text{left}}^1 + F_{\text{left}}^2 + \dots$
- Feather corresponding octaves F_{left}^i with F_{right}^i
 - Can compute inverse FFT and feather in spatial domain
- Sum feathered octave images in frequency domain

Better implemented in *spatial domain*

What does blurring take away?



original

What does blurring take away?



smoothed (5x5 Gaussian)

High-Pass filter




smoothed – original

Band-pass filtering

Gaussian Pyramid (low pass images)




Laplacian Pyramid

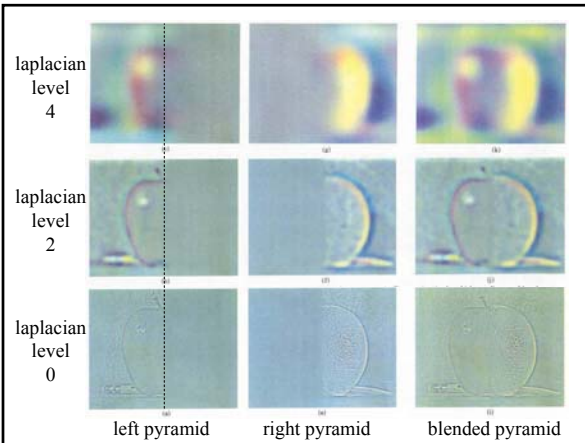



How can we reconstruct (collapse) this pyramid into the original image?

Pyramid Blending



Pyramid Blending



Laplacian Pyramid: Blending

General Approach:

1. Build Laplacian pyramids LA and LB from images A and B
2. Build a Gaussian pyramid GR from selected region R
3. Form a combined pyramid LS from LA and LB using nodes of GR as weights:
 - $LS(i,j) = GR(i,j) * LA(i,j) + (1-GR(i,j)) * LB(i,j)$
4. Collapse the LS pyramid to get the final blended image

Blending Regions

