Image Pyramids Aliasing, Anti-Aliasing, Interpolation



CSC420 David Lindell University of Toronto <u>cs.toronto.edu/~lindell/teaching/420</u> Slide credit: Babak Taati ←Ahmed Ashraf ←Sanja Fidler



Logistics

- •HW1 is due Friday!
- •Turn in on Markup
- Post on Ed Discussion, go to TA office hours if you need help

Finding Waldo

•Let's revisit the problem of finding Waldo

•This time he is on the road





Template(filter)

image

Finding Waldo

He comes closer but our filter doesn't know thatHow can we find Waldo?





Template(filter)

image

Idea: Re-size Image

• Re-scale the image multiple times! Do correlation on every size!





Template(filter)



This image is huge. How can we make it smaller?

Image Sub-Sampling

[Source: S. Seitz]

• Idea: Throw away every other row and column to create a 1/2 size image





1/4



Image Sub-Sampling

•Why does this look so crufty?



[Source: S. Seitz]

•I want to resize my image by factor 2

•And I take every other column and every other row (1st, 3rd, 5th, etc)



Figure: Dashed line denotes the border of the image (it's not part of the image)

•I want to resize my image by factor 2

•And I take every other column and every other row (1st, 3rd, 5th, etc)

•Where is the rectangle!



Figure: Dashed line denotes the border of the image (it's not part of the image)

•What's in the image?

Now I want to resize my image by half in the width direction
And I take every other column (1st, 3rd, 5th, etc)



•What's in the image?

• Now I want to resize my image by half in the width direction

•And I take every other column (1st, 3rd, 5th, etc)



- •What's in the image?
- Now I want to resize my image by half in the width direction
- •And I take every other column (1st, 3rd, 5th, etc)
- •Where is the chicken!



Image Sub-Sampling



[Source: F. Durand]

•What's happening?



[Source: L. Zhang]

Aliasing

•Occurs when your sampling rate is not high enough to capture the amount of detail in your image



•To do sampling right, need to understand the structure of your signal/image

[Source: R. Urtasun]

Aliasing

•Occurs when your sampling rate is not high enough to capture the amount of detail in your image



To do sampling right, need to understand the structure of your signal/image
The minimum sampling rate is called the Nyquist rate

[Source: R. Urtasun]

Aliasing

•Occurs when your sampling rate is not high enough to capture the amount of detail in your image



To do sampling right, need to understand the structure of your signal/image
The minimum sampling rate is called the Nyquist rate

[Source: R. Urtasun]

Examples of Aliasing: Temporal Aliasing

• wagon wheel effect

youtube.com/watch?v=jHS9JGkEOmA

Examples of Aliasing: Temporal Aliasing



https://www.youtube.com/watch?v=yr3ngmRuGUc

Examples of Aliasing: Temporal Aliasing



https://www.youtube.com/shorts/eTW0rNgMcKk

Mr. Nyquist

- Harry Nyquist says that one should look at the frequencies of the signal.
- •One should find the highest frequency (via Fourier Transform)
- •To sample properly you need to sample with at least twice that frequency
- For those interested:

http://en.wikipedia.org/wiki/Nyquist%E2%80% 93Shannon_sampling_theorem







discrete sampled signal





Sample frequency: ? Signal frequency: ?



Sample frequency: 20 Hz Signal: 4 Hz





Sample frequency: 20 Hz Signal: 8 Hz



Sample frequency: 20 Hz Signal: 12 Hz



Sample frequency: 20 Hz Signal: 16 Hz



Sample frequency: 20 Hz Signal: 20 Hz



Sample frequency: 20 Hz Signal: 24 Hz



Sample frequency: 20 Hz Signal: 28 Hz


Sampling exercise

Sample frequency: 20 Hz Signal: 32 Hz



2D example



Good Sampling

Bad Sampling

[Source: N. Snavely]

Going back to Downsampling ...

- •When downsampling by a factor of two, the original image has frequencies that are too high
- High frequencies are caused by sharp edges
- How can we fix this?

Going back to Downsampling ...

- •When downsampling by a factor of two, the original image has frequencies that are too high
- High frequencies are caused by sharp edges
- How can we fix this?

Gaussian pre-filtering

• Solution: Filter out the higher frequency data. Blur the image via Gaussian, then subsample. Very simple!

21



[Source: N. Snavely]

Subsampling with Gaussian pre-filtering



Gaussian 1/2 [Source: S. Seitz]

G 1/4



Subsampling with Gaussian pre-filtering



¼ (2x Zoom)

1/8 (4x Zoom)

1/2 [Source: S. Seitz]

Where is the Rectangle?

• My image



Figure: Dashed line denotes the border of the image (it's not part of the image)

Where is the Rectangle?

•My image

•Let's blur



Figure: Dashed line denotes the border of the image (it's not part of the image)

Where is the Rectangle?

•My image

- •Let's blur
- •And now take every other row and column



Figure: Dashed line denotes the border of the image (it's not part of the image)

Where is the Chicken?

•My image



Where is the Chicken?

- •My image
- •Let's blur



Where is the Chicken?

- •My image
- •Let's blur
- •And now take every other column



Why does this work?

•What does blurring do in the frequency domain?

• How does that fix the aliasing problem?



What happens if we subsample in the primal domain?





What happens if we subsample in the primal domain?

- Shifted copies start to overlap! High frequencies *alias* into lower frequencies



- Shifted copies start to overlap! High frequencies *alias* into lower frequencies
- To solve: first low-pass filter



What happens if we subsample in the primal domain?

(what determines this?)

- Shifted copies start to overlap! High frequencies *alias* into lower frequencies
- To solve: first low-pass filter
- Then no aliasing after downsampling!

Filtering – Low-pass Filter

- low-pass filter: convolution in primal domain b = x * c
- convolution kernel *c* is also known as point spread function (PSF)



Filtering – Low-pass Filter

• low-pass filter: multiplication in frequency domain $F\{b\} = F\{x\} \cdot F\{c\}$



What's the picture in 2D?

Periodic spectral copies in the frequency domain



Periodic primal domain signal



What's the picture in 2D?

Periodic spectral copies in the frequency domain



Periodic primal domain signal



After low-pass filtering

What's the picture in 2D?

Periodic spectral copies in the frequency domain



Periodic primal domain signal



After sub-sampling (no aliasing!)

Image Downsampling (& Upsampling)

• "anti-aliasing" \rightarrow before re-sampling, apply appropriate filter!

• how much filtering? Shannon-Nyquist sampling theorem:





Parmar et al. 2021

Gaussian Pyramids [Burt and Adelson, 1983]

•A sequence of images created with Gaussian blurring and downsampling is called a Gaussian Pyramid

•In computer graphics, a mip map [Williams, 1983]



How much space does a Gaussian pyramid take compared to original image?

Gaussian Pyramids [Burt and Adelson, 1983]

•A sequence of images created with Gaussian blurring and downsampling is called a Gaussian Pyramid

•In computer graphics, a mip map [Williams, 1983]





$$\sum_{n=1}^{\infty} \frac{1}{4^n} = \frac{1/4}{1-1/4} = \frac{1}{3}$$

How much space does a Gaussian pyramid take compared to original image?

Example of Gaussian Pyramid



[Source: N. Snavely]

Image Up-Sampling

•This image is too small, how can we make it 10 times as big?



[Source: N. Snavely, R. Urtasun]

Image Up-Sampling

•This image is too small, how can we make it 10 times as big?



• Simplest approach: repeat each row and column 10 times



[Source: N. Snavely, R. Urtasun]



• It is a discrete point-sampling of a continuous function

• If we could somehow reconstruct the original function, any new image could be generated, at any resolution and scale



• It is a discrete point-sampling of a continuous function

• If we could somehow reconstruct the original function, any new image could be generated, at any resolution and scale



• It is a discrete point-sampling of a continuous function

• If we could somehow reconstruct the original function, any new image could be generated, at any resolution and scale

Interpolation



•What if we don't know f?

Interpolation



•What if we don't know f?

• Guess an approximation: for example nearest-neighbor
Interpolation



•What if we don't know f?

- Guess an approximation: for example nearest-neighbor
- Guess an approximation: for example linear

[Source: N. Snavely, S. Seitz]

Interpolation



•What if we don't know f?

- Guess an approximation: for example nearest-neighbor
- Guess an approximation: for example linear
- More complex approximations: cubic, B-splines

[Source: N. Snavely, S. Seitz]

Interpolation



•What if we don't know f?

- Guess an approximation: for example nearest-neighbor
- Guess an approximation: for example linear
- More complex approximations: cubic, B-splines
- But more isn't always better!

[Source: N. Snavely, S. Seitz]

Linear Interpolation



• Linear interpolation from our discretized F :

$$G(x) = \frac{x_2 - x}{x_2 - x_1} F(x_1) + \frac{x - x_1}{x_2 - x_1} F(x_2)$$



•Let's make this signal triple length



Make a vector G with d times the size of F

G

•Let's make this signal triple length (d=3)



G

F(1

G

F(n)F(n)F(3)F(2)F(3)F(2)F(1)if $\frac{i}{d}$ integer: G(i) = F(i/d)otherwise: $G(i) = \frac{x_2 - x}{x_2 - x_1}F(x_1) + \frac{x - x_1}{x_2 - x_1}F(x_2)$ x = i/dwhere $x_1 = \lfloor i/d \rfloor$ •Let's make this signal triple length $x_2 = \lceil i/d \rceil$ • If i/d is an integer, just copy from the signal

•Otherwise use the interpolation formula

Linear Interpolation via Convolution F[x]

x

 $\overline{x_1}$

1

• Linear interpolation:

$$G(x) = \frac{x_2 - x}{x_2 - x_1} F(x_1) + \frac{x - x_1}{x_2 - x_1} F(x_2)$$

 $\overline{x_2}$

4

x

5

Linear Interpolation via Convolution



•Linear interpolation:

$$G(x) = \frac{x_2 - x}{x_2 - x_1} F(x_1) + \frac{x - x_1}{x_2 - x_1} F(x_2)$$

• In terms of discrete points, what if we want to add one sample between two samples? Make a longer array G' (initialized to zeros and F at alternating indices)

$$G_{interpolated}(x_i) = \frac{1}{2}G'^{(x_{i-1})} + G'(x_i) + \frac{1}{2}G'(x_{i+1})$$

Linear Interpolation via Convolution



•Linear interpolation:

$$G(x) = \frac{x_2 - x}{x_2 - x_1} F(x_1) + \frac{x - x_1}{x_2 - x_1} F(x_2)$$

• In terms of discrete points, what if we want to add one sample between two samples? Make a longer array G' (initialized to zeros and F at alternating indices)

$$G_{interpolated}(x_i) = \frac{1}{2}G'^{(x_{i-1})} + G'(x_i) + \frac{1}{2}G'(x_{i+1})$$
$$G_{interpolated}(x_i) = \left[\frac{1}{2}1\frac{1}{2}\right] * G'$$

F(1) $F(2)$) $F(3)$	F(n)
-------------	----------	------

•Let's make this signal triple length

G'



•Let's make this signal triple length (d=3)



G'

*

- •Let's make this signal triple length (d=3)
- •What should be my "reconstruction" filter h (such that G = h * G')?



- •Let's make this signal triple length (d=3)
- •What should be my "reconstruction" filter h (such that G = h * G')?
- $h = [0, \frac{1/d}{d}, \dots, \frac{d-1/d}{d}, 1, \frac{d-1/d}{d}, \dots, \frac{1/d}{d}, 0]$, where d my upsampling factor

$$h = \left[0, \frac{1}{3}, \frac{2}{3}, 1, \frac{2}{3}, \frac{1}{3}, 0\right] \qquad \qquad * \ G'$$



•Let's make this signal triple length (d=3)

•What should be my "reconstruction" filter h (such that G = h * G')?

• $h = [0, \frac{1/d}{d}, \dots, \frac{d-1/d}{d}, 1, \frac{d-1/d}{d}, \dots, \frac{1/d}{d}, 0]$, where d my upsampling factor

$$h = \left[0, \frac{1}{3}, \frac{2}{3}, 1, \frac{2}{3}, \frac{1}{3}, 0\right] \qquad \qquad * \ G'$$



•Let's make this signal triple length (d=3)

•What should be my "reconstruction" filter h (such that G = h * G')? • $h = [0, \frac{1/d}{2}, \dots, \frac{d-1/d}{2}, 1, \frac{d-1/d}{2}, \dots, \frac{1/d}{2}, 0]$, where d my upsampling factor

$$h = \left[0, \frac{1}{3}, \frac{2}{3}, 1, \frac{2}{3}, \frac{1}{3}, 0\right] \qquad \qquad * G'$$



•Let's make this signal triple length (d=3)

•What should be my "reconstruction" filter h (such that G = h * G')?

• $h = [0, \frac{1/d}{d}, \dots, \frac{d-1/d}{d}, 1, \frac{d-1/d}{d}, \dots, \frac{1/d}{d}, 0]$, where d my upsampling factor

Interpolation via Convolution (1D)



"Ideal" reconstruction



Image Interpolation (2D)



•Let's make this image triple size

•Copy image in every third pixel. What about the remaining pixels in G?

Image Interpolation (2D)



- •Let's make this image triple size
- •Copy image in every third pixel. What about the remaining pixels in G?
- How shall we compute this value?



- •Let's make this image triple size
- •Copy image in every third pixel. What about the remaining pixels in G?
- Better: bilinear interpolation; linear interpolation in x, then y, resulting in a quadratic interpolation.
 Check out details: <u>http://en.wikipedia.org/wiki/Bilinear_interpolation</u>



- •Let's make this image triple size
- •Copy image in every third pixel. What about the remaining pixels in G?

•What does the 2D version of this hat function look like?

h(x)

h(x,y)



Performs Linear Interpolation

(tent function) **Performs** bilinear Interpolation

•What does the 2D version of this hat function look like?



h(x,y)

Performs Linear Interpolation



bilinear Interpolation

•And filter for nearest neighbor interpolation?

•What does the 2D version of this hat function look like?



h(x,y)

Performs Linear Interpolation

•And filter for nearest neighbor interpolation?



•What does the 2D version of this hat function look like?



h(x,y)



Performs Linear Interpolation

bilinear Interpolation

• Better filters give better resampled images: Bicubic is a common choice

Cubic interpolation

•Uses third degree polynomial

• Formula: (interpolating between [0, 1] for points at -1, 0, 1, 2)

$$f(p_0, p_1, p_2, p_3, x) = \left(-\frac{1}{2}p_0 + \frac{3}{2}p_1 - \frac{3}{2}p_2 + \frac{1}{2}p_3\right)x^3$$

$$+ \left(p_0 - \frac{5}{2}p_1 + 2p_2 - \frac{1}{2}p_3\right)x^2 + \left(-\frac{1}{2}p_0 + \frac{1}{2}p_2\right)x + p_1 \quad (2)$$



Smoother interpolation between points (can constrain both values and derivatives at each point)

https://www.paulinternet.nl/?page=bicubic

•What happens to the global minimum/maximum of discrete samples after cubic interpolation?

- •What happens to the global minimum/maximum of discrete samples after cubic interpolation?
- •Nearest neighbor



https://www.paulinternet.nl/?page=bicubic

•What happens to the global minimum/maximum of discrete samples after cubic interpolation?

• Bilinear interpolation



https://www.paulinternet.nl/?page=bicubic

- •What happens to the global minimum/maximum of discrete samples after cubic interpolation?
- Bicubic interpolation



Image Interpolation via Convolution (2D)



•Let's make this image triple size: copy image values in every third pixel, place zeros everywhere else

Image Interpolation via Convolution (2D)



- •Let's make this image triple size: copy image values in every third pixel, place zeros everywhere else
- •Convolution with a reconstruction filter (e.g., bilinear) and you get the interpolated image

Image Interpolation

Original image











Bicubic interpolation

[Source: N. Snavely]

Summary – Stuff You Should Know

- •To down-scale an image: blur it with a small Gaussian (e.g., σ = 1.4) and downsample
- •What is aliasing in the Fourier domain? What is anti-aliasing?
- •To up-scale an image: interpolation (nearest neighbor, bilinear, bicubic, etc)
- •Gaussian pyramid: Blur with Gaussian filter, downsample result by factor 2, blur it with the Gaussian, downsample by 2...

OpenCV functions:

- imresize: with interpolation options INTER_LINEAR, INTER_CUBIC, ...
- pyrDown: Blurs an image and downsamples it (1/2 size)
- pyrUp: Upsamples an image and then blurs it ($\times 2$ size)
- img[::2,::2,:]: takes every second row and column