


Resizing and resampling

# Aliasing


- Images are made up of high frequency and low frequency components
- High frequency components: pixel-to-pixel details
- Low frequency components: high-level structure
- What subsampling should do: remove pixel-to-pixel details, keep high-level structure
- What naïve subsampling does: converts pixel-to-pixel details to new coarse structures → problem

Aliasing




Aliasing

Aliasing artifacts



# Image sub-sampling



1/2



1/4 (2x zoom)



1/16 (4x zoom)

Why does this look so cruffy? Aliasing!

# Why does aliasing happen?

- Consider sampling every  $P$  pixels


- $$B_{\frac{N}{P}+k}(Pn) = e^{\frac{i2\pi(\frac{N}{P}+k)Pn}{N}} = e^{i2\pi n} e^{\frac{i2\pi kPn}{N}} = e^{\frac{i2\pi kPn}{N}} = B_k(Pn)$$

- In fact 
$$B_{\frac{mN}{P}+k}(Pn) = B_k(Pn)$$

- The high frequency component  $B_{\frac{N}{P}+k}$  gets aliased as the low frequency component  $B_k$
- This is because  $B_{\frac{N}{P}+k}$  computes an additional cycle between two samples and ends up in the same place as  $B_k$


# Why does aliasing happen?

Blue gets aliased as orange



# Why does aliasing happen?


Blue gets aliased as orange





# Why does aliasing happen?

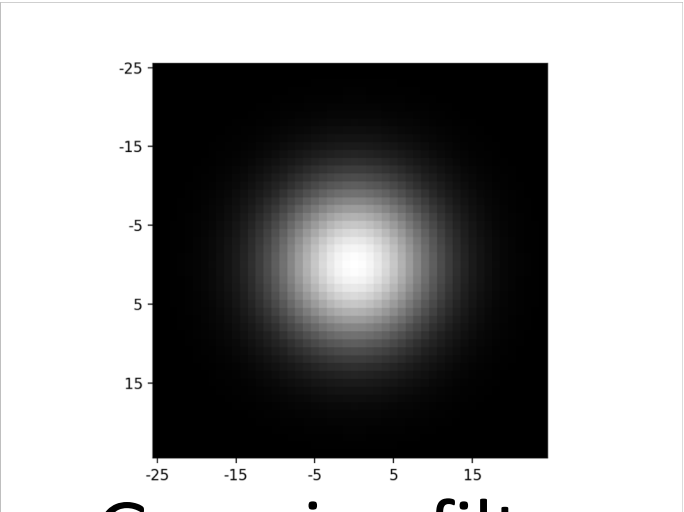

Blue gets aliased as orange



# How to avoid aliasing

- To recover a sinusoid, need to sample at least twice per cycle
- For a general image, need to sample at least twice the rate of the highest frequency component
- **Nyquist sampling theorem:**  $2\nu_{max} < \nu_{sample}$
- To subsample, *remove high frequency components*
- To remove high frequency components, *blur the image with a Gaussian*



# Image



Gaussian filters

Zeros out  
high  
frequencies

# Fourier transform



Keeps low  
frequencies




Fourier transform

# Gaussian pre-filtering


- Solution: filter the image, *then* subsample



*Gaussian pyramid*




...



# Gaussian pyramids

## [Burt and Adelson, 1983]


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  $N = 2^k$ )





- In computer graphics, a *mip map* [Williams, 1983]

Gaussian Pyramids have all sorts of applications in computer vision

# Gaussian pyramids - Searching over scales




# Gaussian pyramids - Searching over scales






# The Gaussian Pyramid





# Gaussian pyramid and stack





## Memory Usage

- Each color is a separate pyramid
- 3 pyramids fit into  $2W \times 2H$  image



# What about upsampling?

- Simple solution: Fill rest of the pixels with zeros
- Obviously wrong. How can we do better?




# Upsampling

- Need to *interpolate* intermediate pixels. What is the best way to interpolate?
- Recall: before subsampling, we removed high frequencies
- Key idea: upsampled image should not have high frequencies either
- Gaussian blur again!

# Upsampling

- Step 1: upsample and fill with 0s
- Step 2: Gaussian blur to interpolate
- Step 3: Scale correction
  - Gaussian blur is just weighted average
  - But we just introduced a bunch of zeros ==> need to scale up the resulting image

# Laplacian pyramid



# Laplacian pyramid



$$L_4 = G_4 =$$

$$L_3 = G_3 - \text{expand}(G_4) =$$

$$L_2 = G_2 - \text{expand}(G_3) =$$

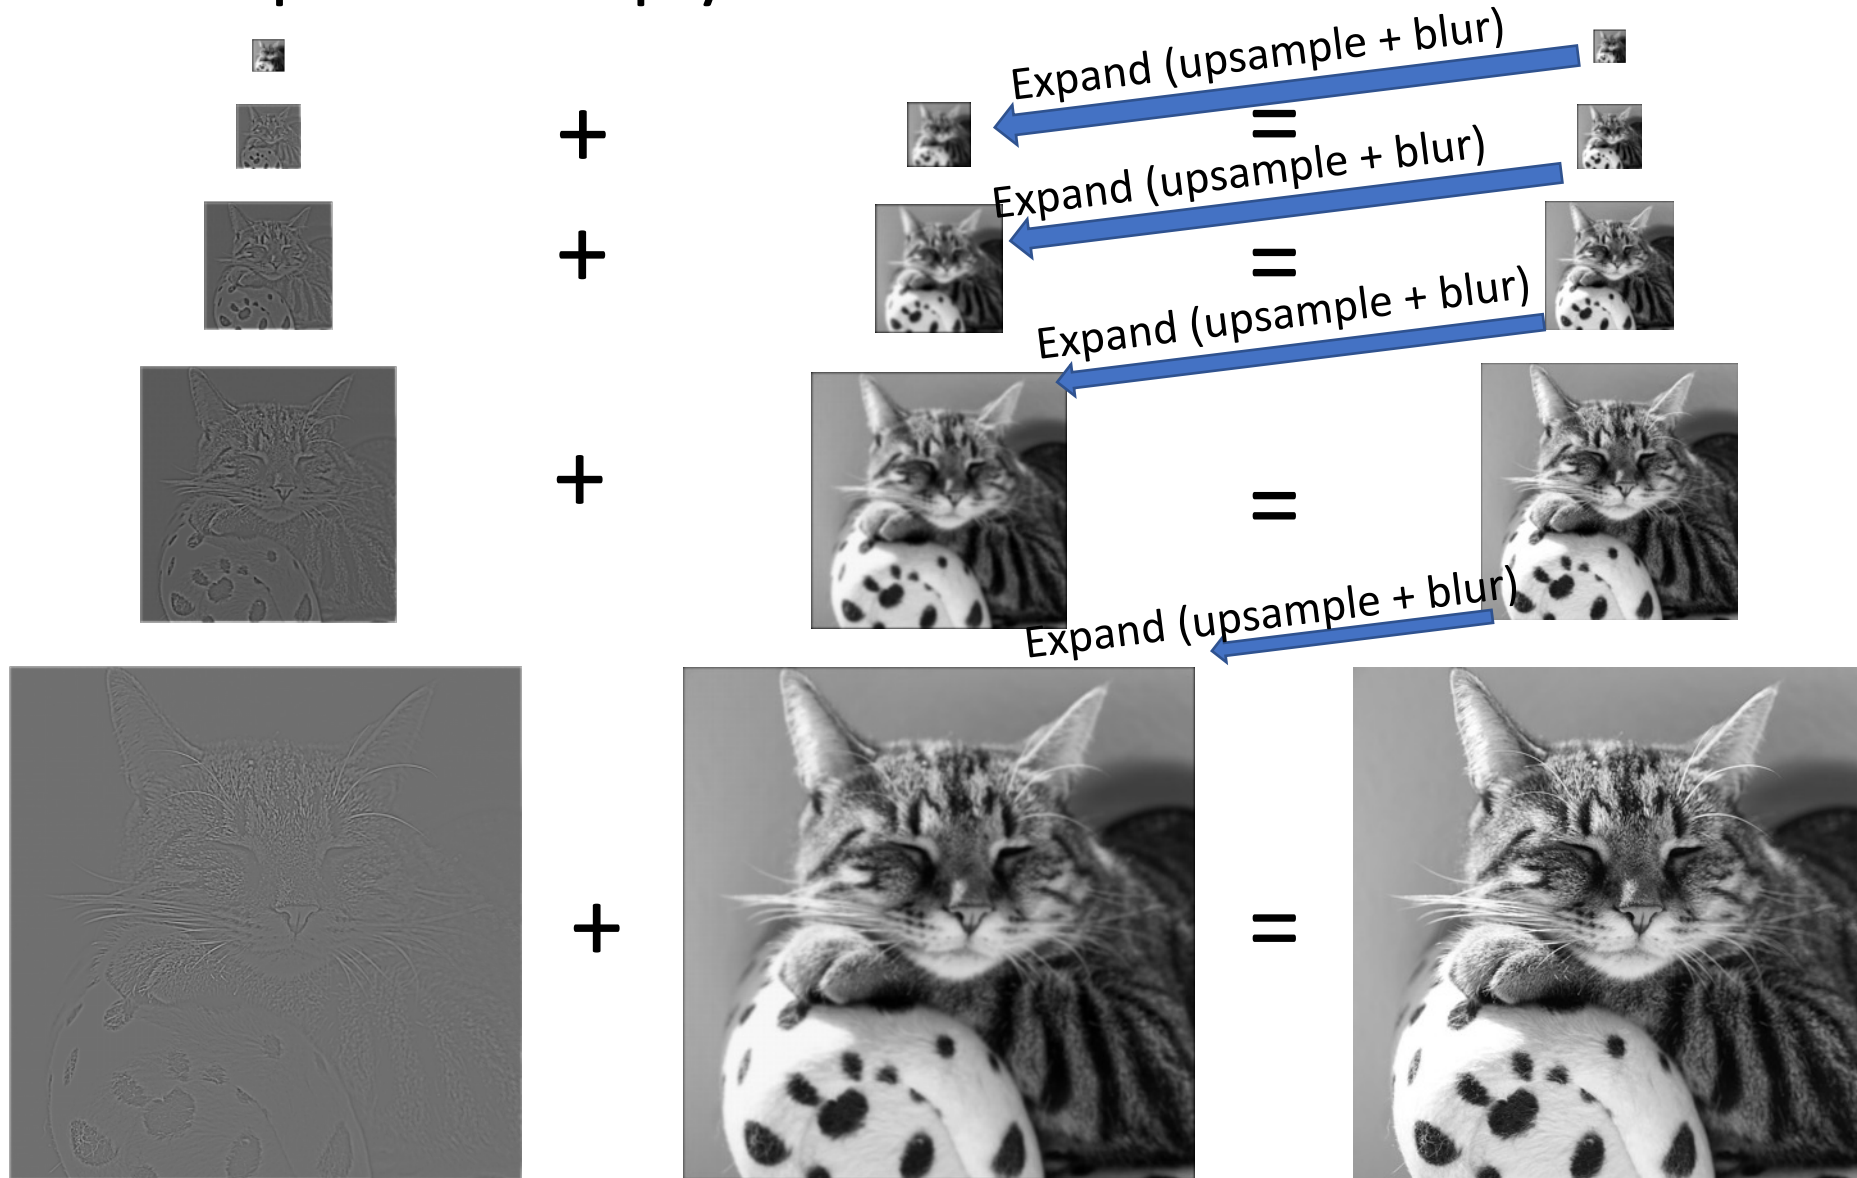
$$L_1 = G_1 - \text{expand}(G_2) =$$

$$L_0 = G_0 - \text{expand}(G_1) =$$






# Reconstructing the image from a Laplacian pyramid



# Laplacian pyramid



512

256


128

64

32

16

8



# Low-pass and high-pass filtering

- Convoluting with a Gaussian = remove high frequencies
- “Low-pass” filtering: low frequencies “pass” through filter, high frequencies don’t
- Identity – Low-pass filtered image = “High-pass filtering”

# Hybrid images (PA1)

- From afar, images look tiny, we only see low frequencies
- Up close, we see high frequencies
- Low frequencies of one image + high frequencies of another

