



# Fourier transforms and rescaling

# Fourier transforms


Image








Fourier transform



Image




Fourier transform




# Low-pass filtering




\*



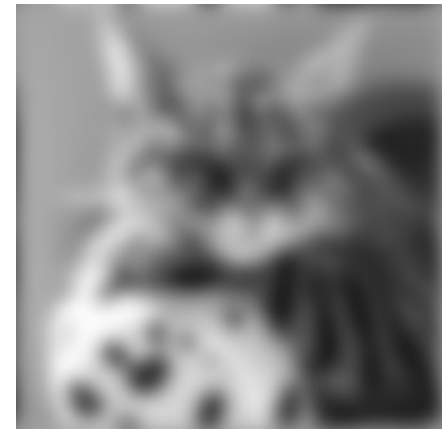
Fourier




Fourier



=



\*




=



# High-pass filtering



-




=



-



=




# Band-pass filtering




-





=




-





=




-




=




# Hybrid images (PA1)



# Hybrid images (PA1)




# Hybrid images (PA1)





Resizing and resampling

Let's enhance!




Louis Daguerre, 1838

# Let's enhance!

- When is enhancement possible?
- How can we model what happens when we upsample or downsample an image?
- Resizing up or down very common operation
  - Searching across scales
  - applications have different memory/quality tradeoffs


# What is a (digital) image?

- True image is a function from  $\mathbb{R}^2$  to  $\mathbb{R}$
- Digital image is a sample from it
- 1D example:



- To enhance, we need to recover the original signal and sample again

# Undersampling




# Undersampling

- What if we “missed” things between the samples?
- Simple example: undersampling a sine wave
  - unsurprising result: information is lost
  - surprising result: indistinguishable from lower frequency
  - also was always indistinguishable from higher frequencies
  - *aliasing*: signals “traveling in disguise” as other frequencies

# Aliasing

- When sampling is not adequate, impossible to distinguish between low and high frequency signal



# Aliasing in time






# Aliasing in time




# Image Scaling

What happens if we naively upsample?



Source: S. Seitz

# Image sub-sampling




1/4



1/16

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)




1/16 (4x zoom)

Why does this look so cruffy? Aliasing!

# Image sub-sampling




# Point sampling in action




# How many samples do we need?

- 1 sample per time period is too less:



# How many samples do we need?

- 2 samples per time-period is enough

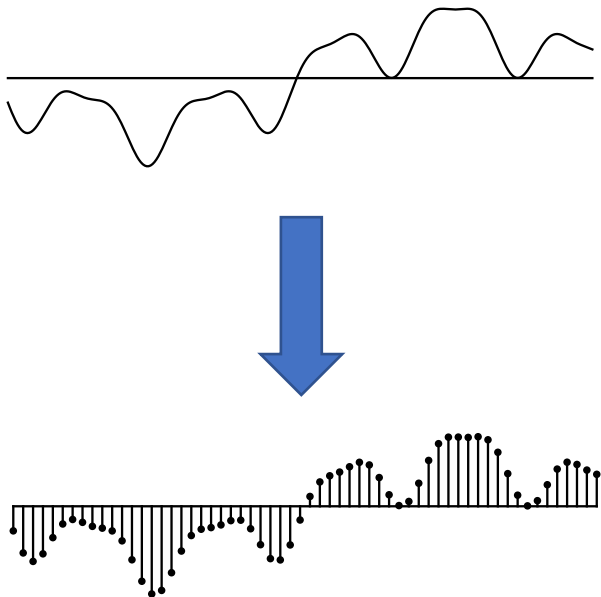


- Nyquist sampling theorem: Need to sample at least 2 times the frequency
- General signals? Need to sample at least 2 times the maximum frequency




# Nyquist sampling: why?

Spatial domain




Sampling = Keep values at  $t = kT_s$ , make everything else 0

Frequency domain




Sampling = Make frequency domain periodic with period  $\nu = 1/T_s$  by making copies

# Nyquist sampling: why?



$$\frac{1}{T_s} - \nu_{\max} > \nu_{\max} \Rightarrow \frac{1}{T_s} > 2\nu_{\max}$$


# Nyquist sampling: why?



# Aliasing and downsampling


- Nyquist says must sample at at least twice maximum frequency
- When downsampling by a factor of two
  - Original image has frequencies that are too high
- How can we fix this?
- Eliminate them before sampling!
  - Convert to frequency space
  - Multiply with low-pass filter

# Eliminating High Frequencies




# Process

- Can we do this in spatial domain?
  - Yes!
- **Multiplication** in frequency domain = **convolution** in spatial domain
- **Box filter** in frequency domain = **sinc** in spatial domain
- **Multiplication** with **box filter** in frequency domain = **convolution** with **sinc filter** in spatial domain




# Reconstruction from samples

Spatial domain




Sampling = Keep values at  $t = kT_s$ , make everything else 0

Frequency domain



Sampling = Make frequency domain periodic with period  $\nu = 1/T_s$  by making copies

# Reconstruction from samples



$$F_{recons}(\nu) = F_{sampled}(\nu)B(\nu)$$

Box filter in frequency space




# Reconstruction from samples

$$F_{recons}(\nu) = F_{sampled}(\nu)B(\nu)$$

- Multiplication in frequency domain = convolution in spatial domain
- Box filter in frequency domain = sinc filter in spatial domain
- Convolve sampled signal with sinc filter to reconstruct

# Reconstruction from samples

- "Sampled signal" is non-zero at sample points and 0 everywhere else
  - i.e., has holes



# Recap: subsampling and reconstruction


## Subsampling

1. Convolve with sinc filter to eliminate high frequencies
2. Sample by picking only values at sample points

## Reconstruction

1. Start with sampled signal (0 at non-sample points)
2. Convolve with sinc to reconstruct



# Sinc is annoying




# Sinc and Gaussian

- Sinc is annoying: infinite spatial extent
- Use Gaussian instead!


Sinc/box



Gaussian



Spatial domain



Frequency domain

# Subsampling images

- Step 1: Convolve with Gaussian to eliminate high frequencies
- Step 2: Drop unneeded pixels




Subsampling without removing high frequencies



Subsampling after removing high frequencies

# Subsampling images correctly



Gaussian 1/2



G 1/4



G 1/8

- 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



# Compare with...



1/2



1/4 (2x zoom)



1/8 (4x zoom)

# Upsampling images



Step 1: blow up to  
original size with 0's  
in between



# Upsampling images



Step 2: Convolve with  
Gaussian



# Take-away

- Subsampling causes aliasing
  - High frequencies masquerading as low frequencies
- Remove low frequencies by blurring!
  - Ideal: sinc
  - Common: Gaussian
- When upsampling, reconstruct missing values by convolution
  - Ideal: sinc
  - Common: Gaussian



# So... can we enhance?

- Nyquist theorem limits frequencies we can reconstruct from subsampled image
- Can only reconstruct max sampling frequency/2
- Sorry CSI!


Pyramids

# 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



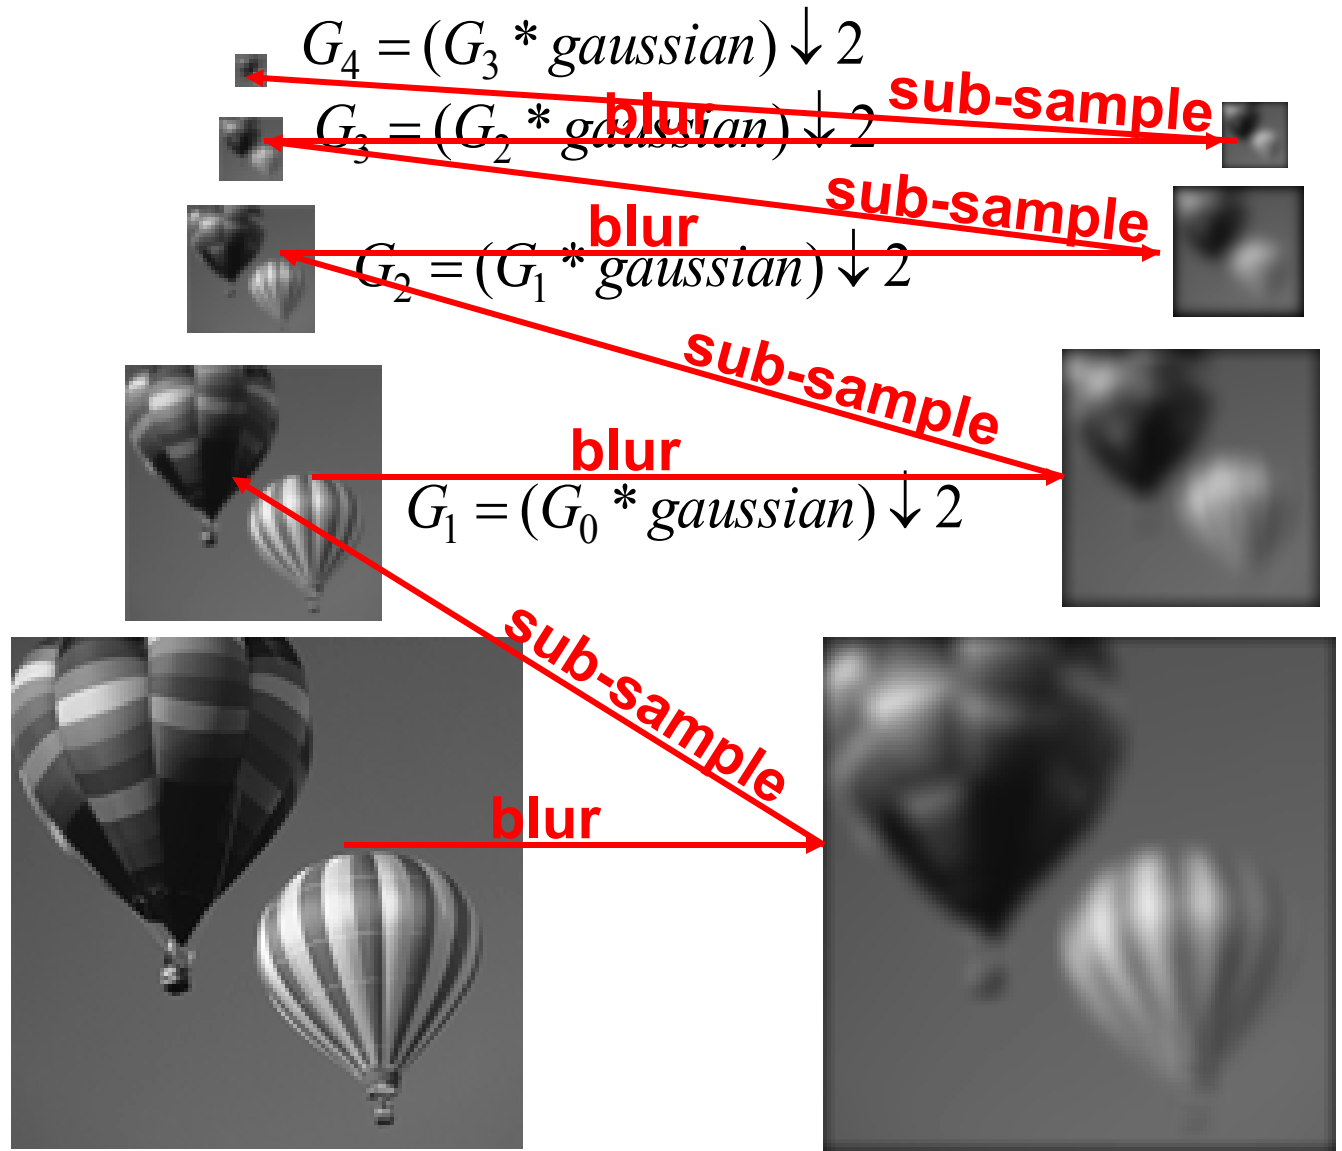
# The Gaussian Pyramid

Low resolution




High resolution

$G_0 = \text{Image}$




# Gaussian pyramid and stack

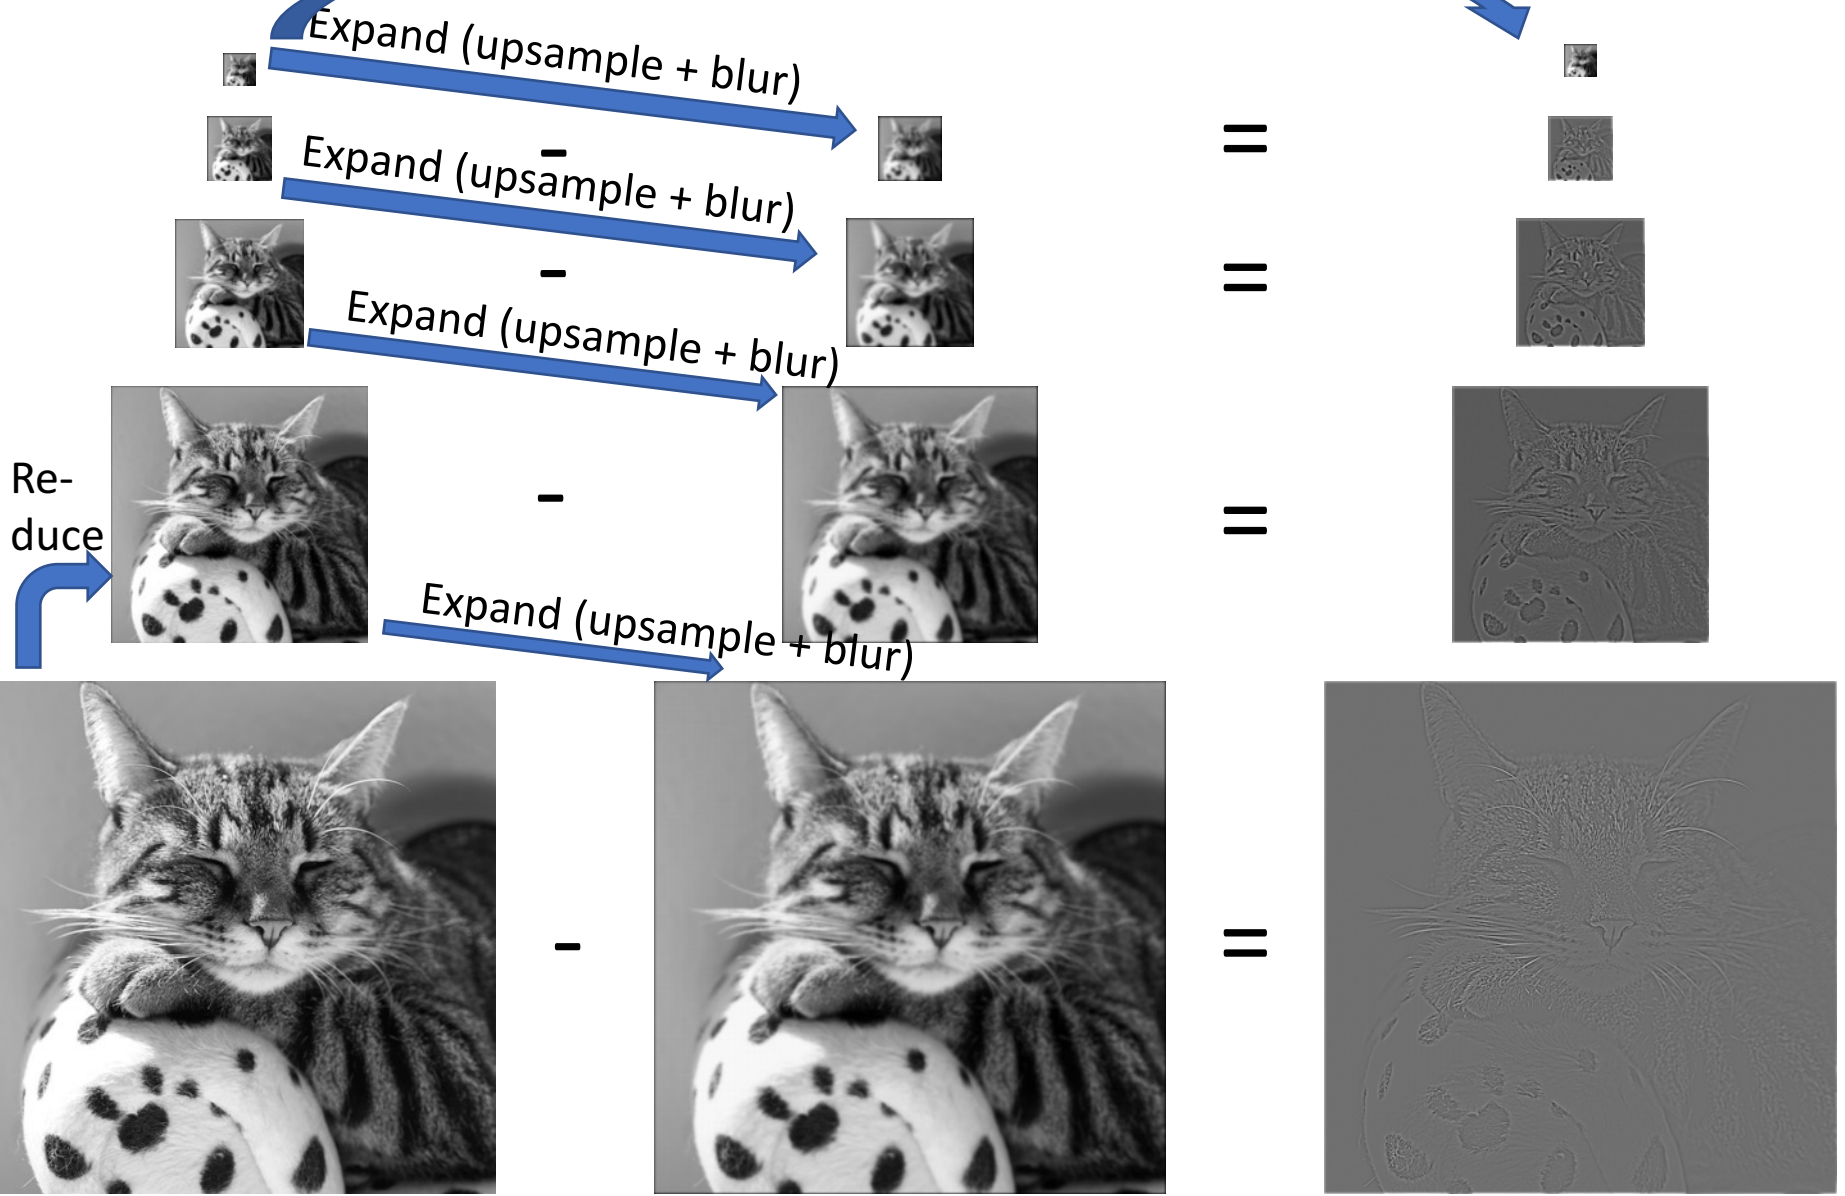


# Memory Usage

- What is the size of the pyramid?



# Laplacian pyramid



# Laplacian pyramid




$$L_4 = G_4 =$$

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

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

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

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







# Laplacian pyramid



512

256

128

64

32

16

8

