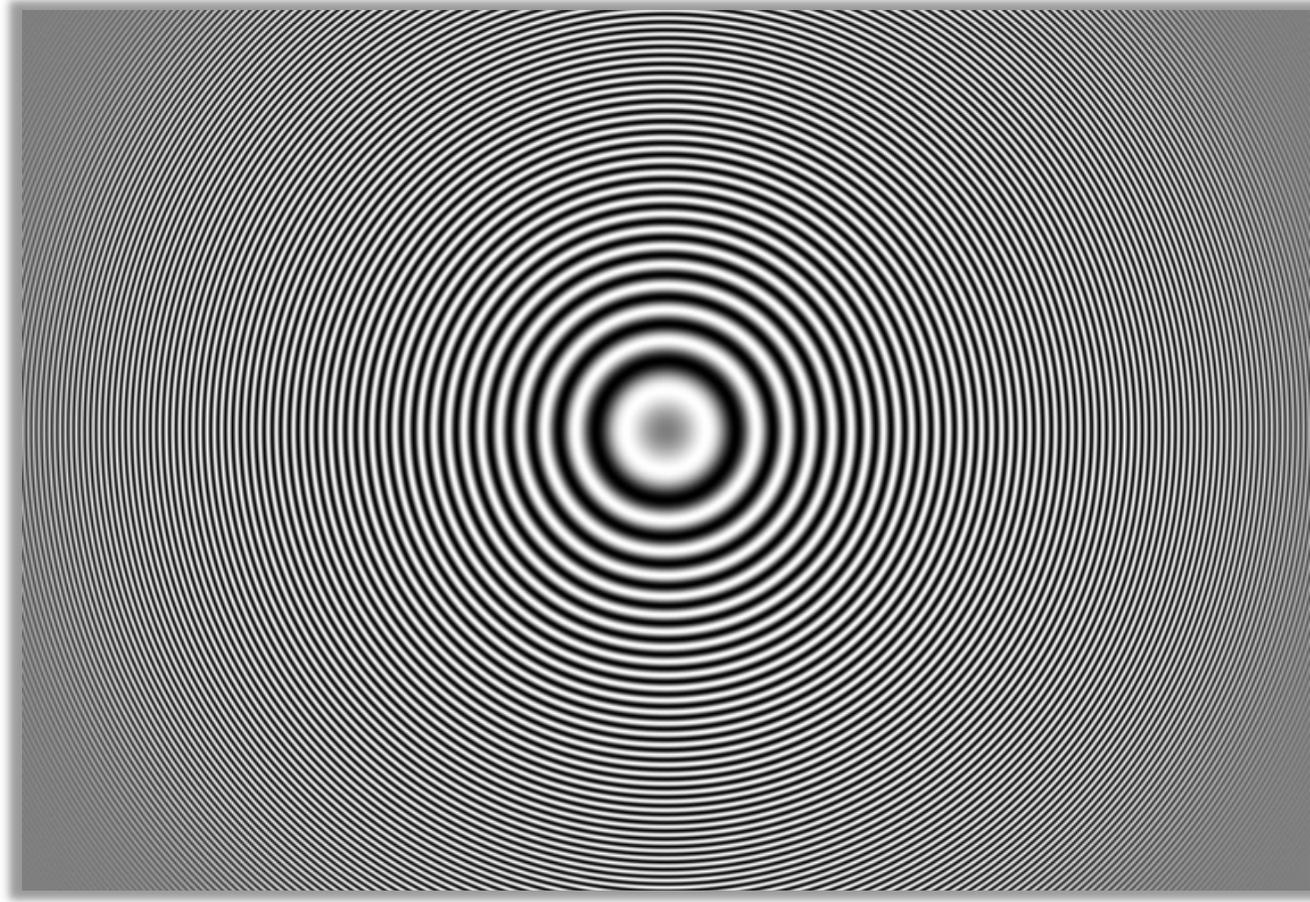


# Image Pyramids

Aliasing, Anti-Aliasing, Interpolation



CSC420

David Lindell

University of Toronto

[cs.toronto.edu/~lindell/teaching/420](http://cs.toronto.edu/~lindell/teaching/420)

Slide credit: Babak Taati ← Ahmed Ashraf ← Sanja Fidler

# Logistics

- HW1 is due Friday!
- Turn in on Markus
- Post on Piazza, 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



image



Template(filter)

# Finding Waldo

- He comes closer but our filter doesn't know that
- How can we find Waldo?



image



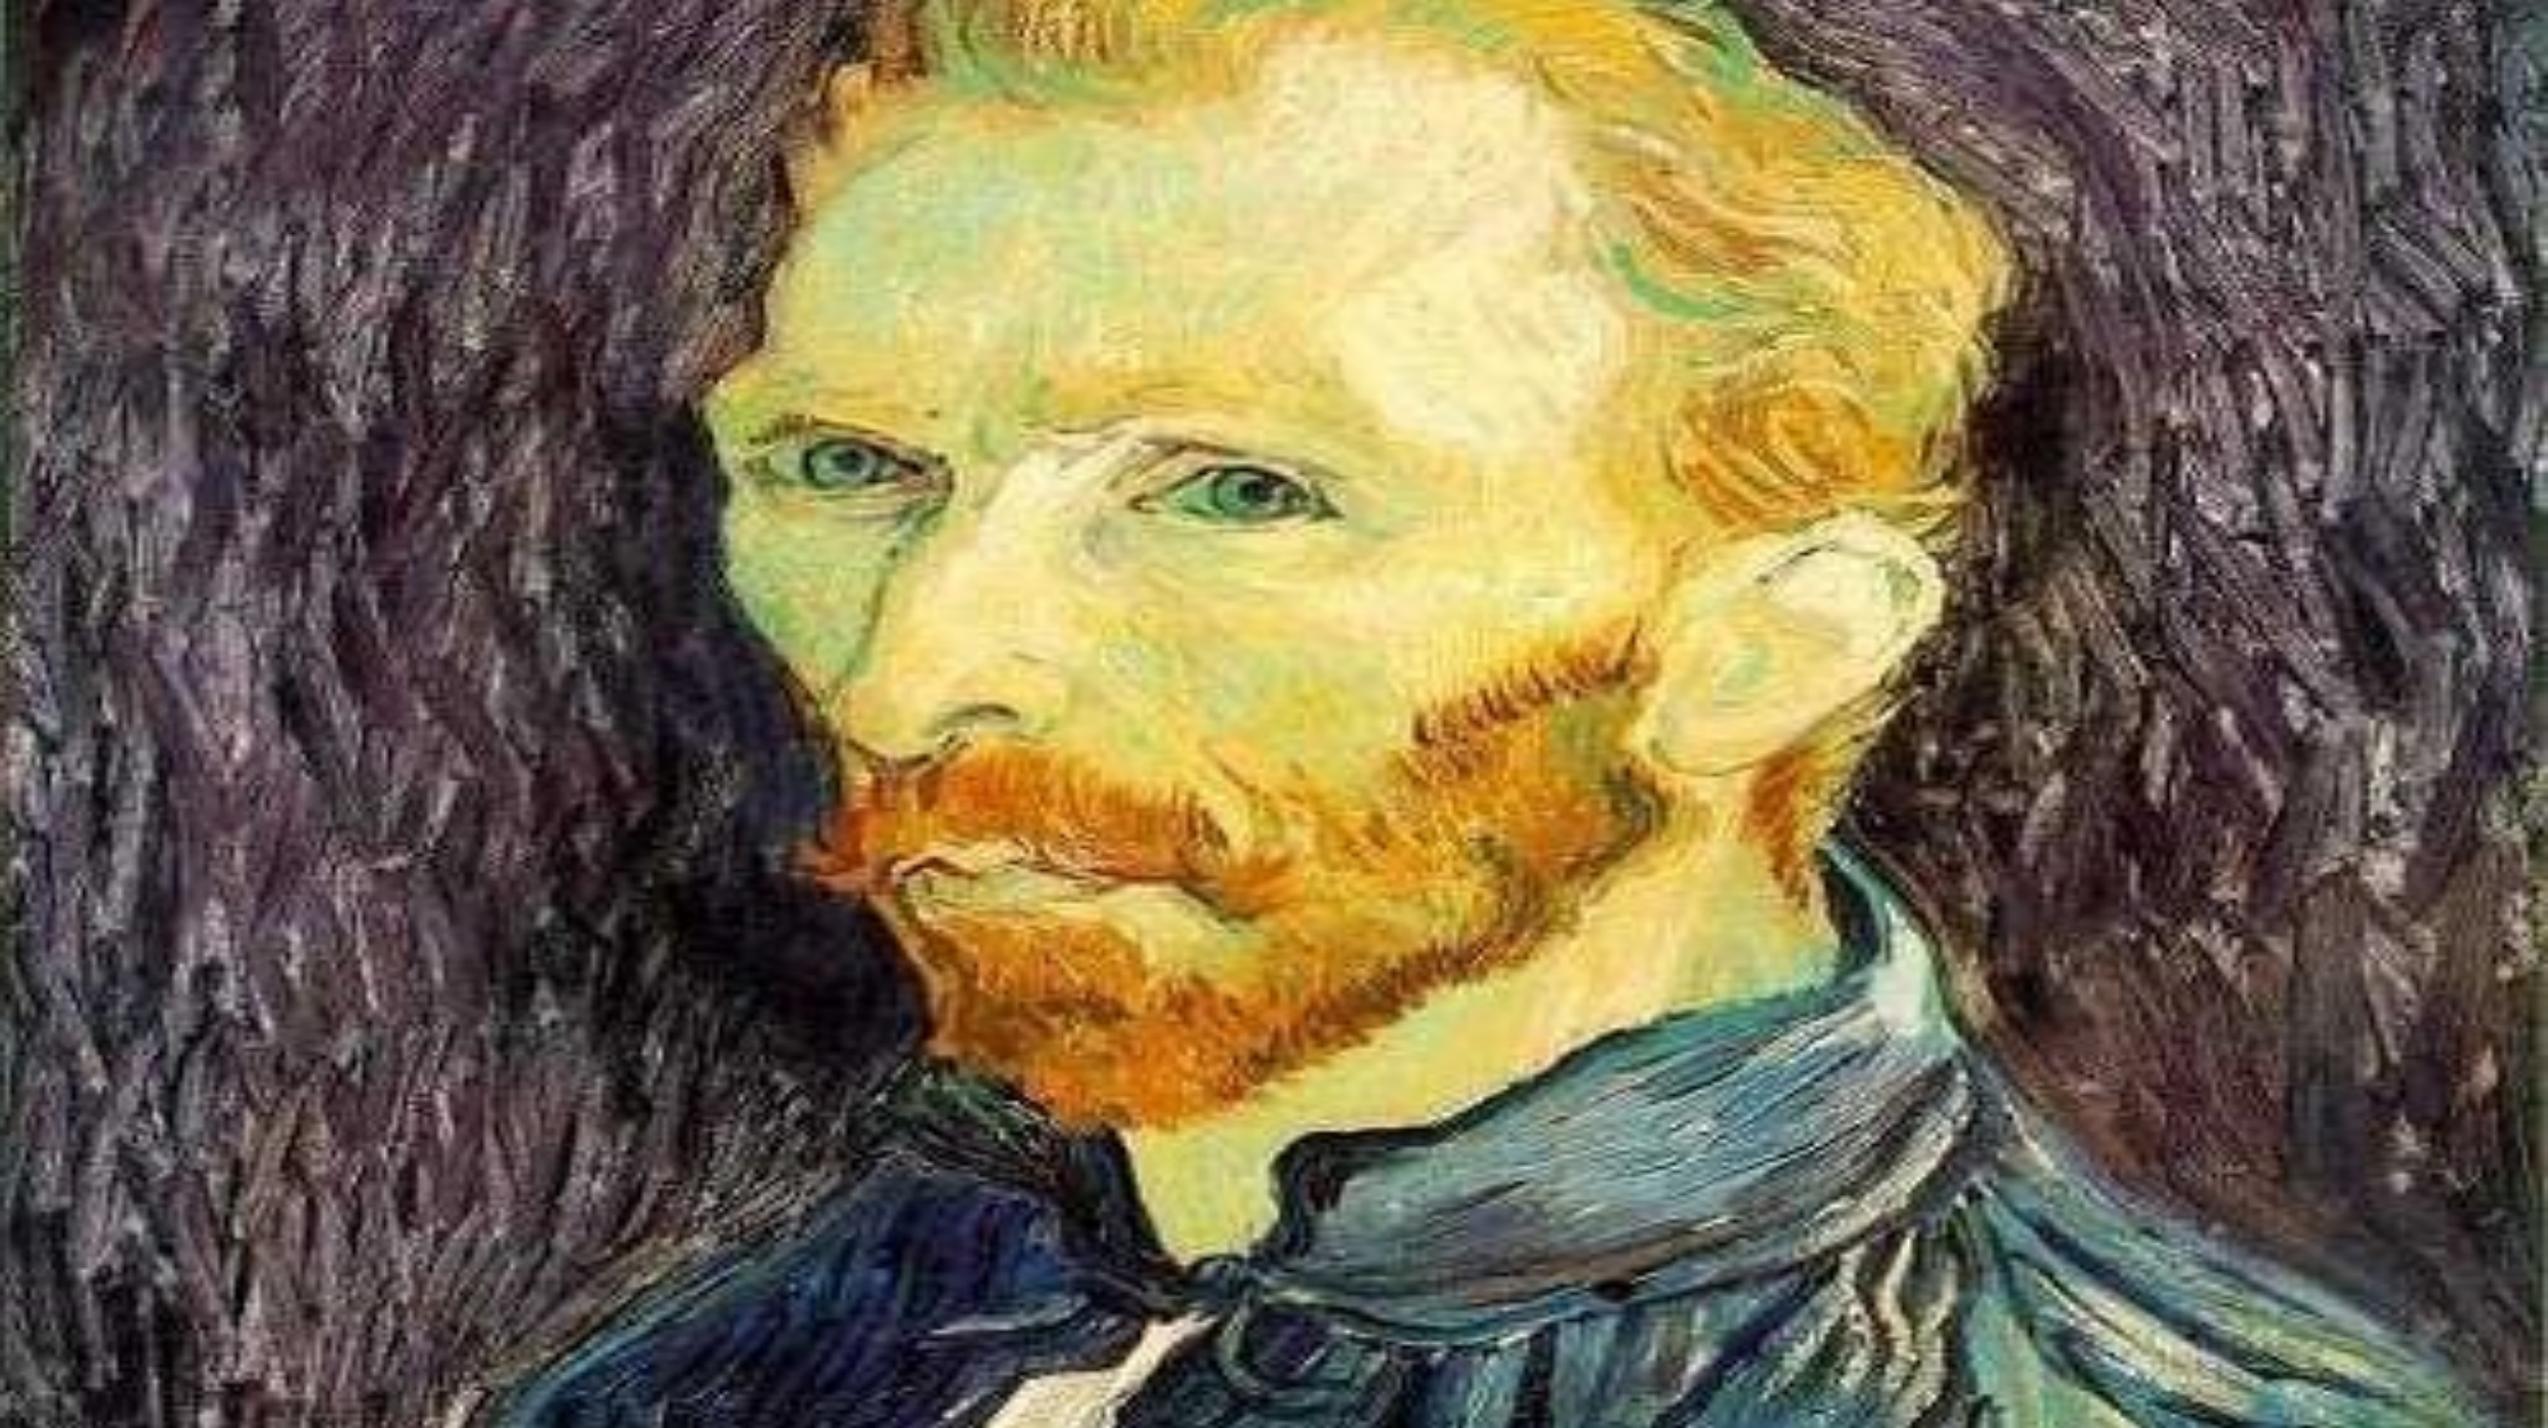
Template(filter)

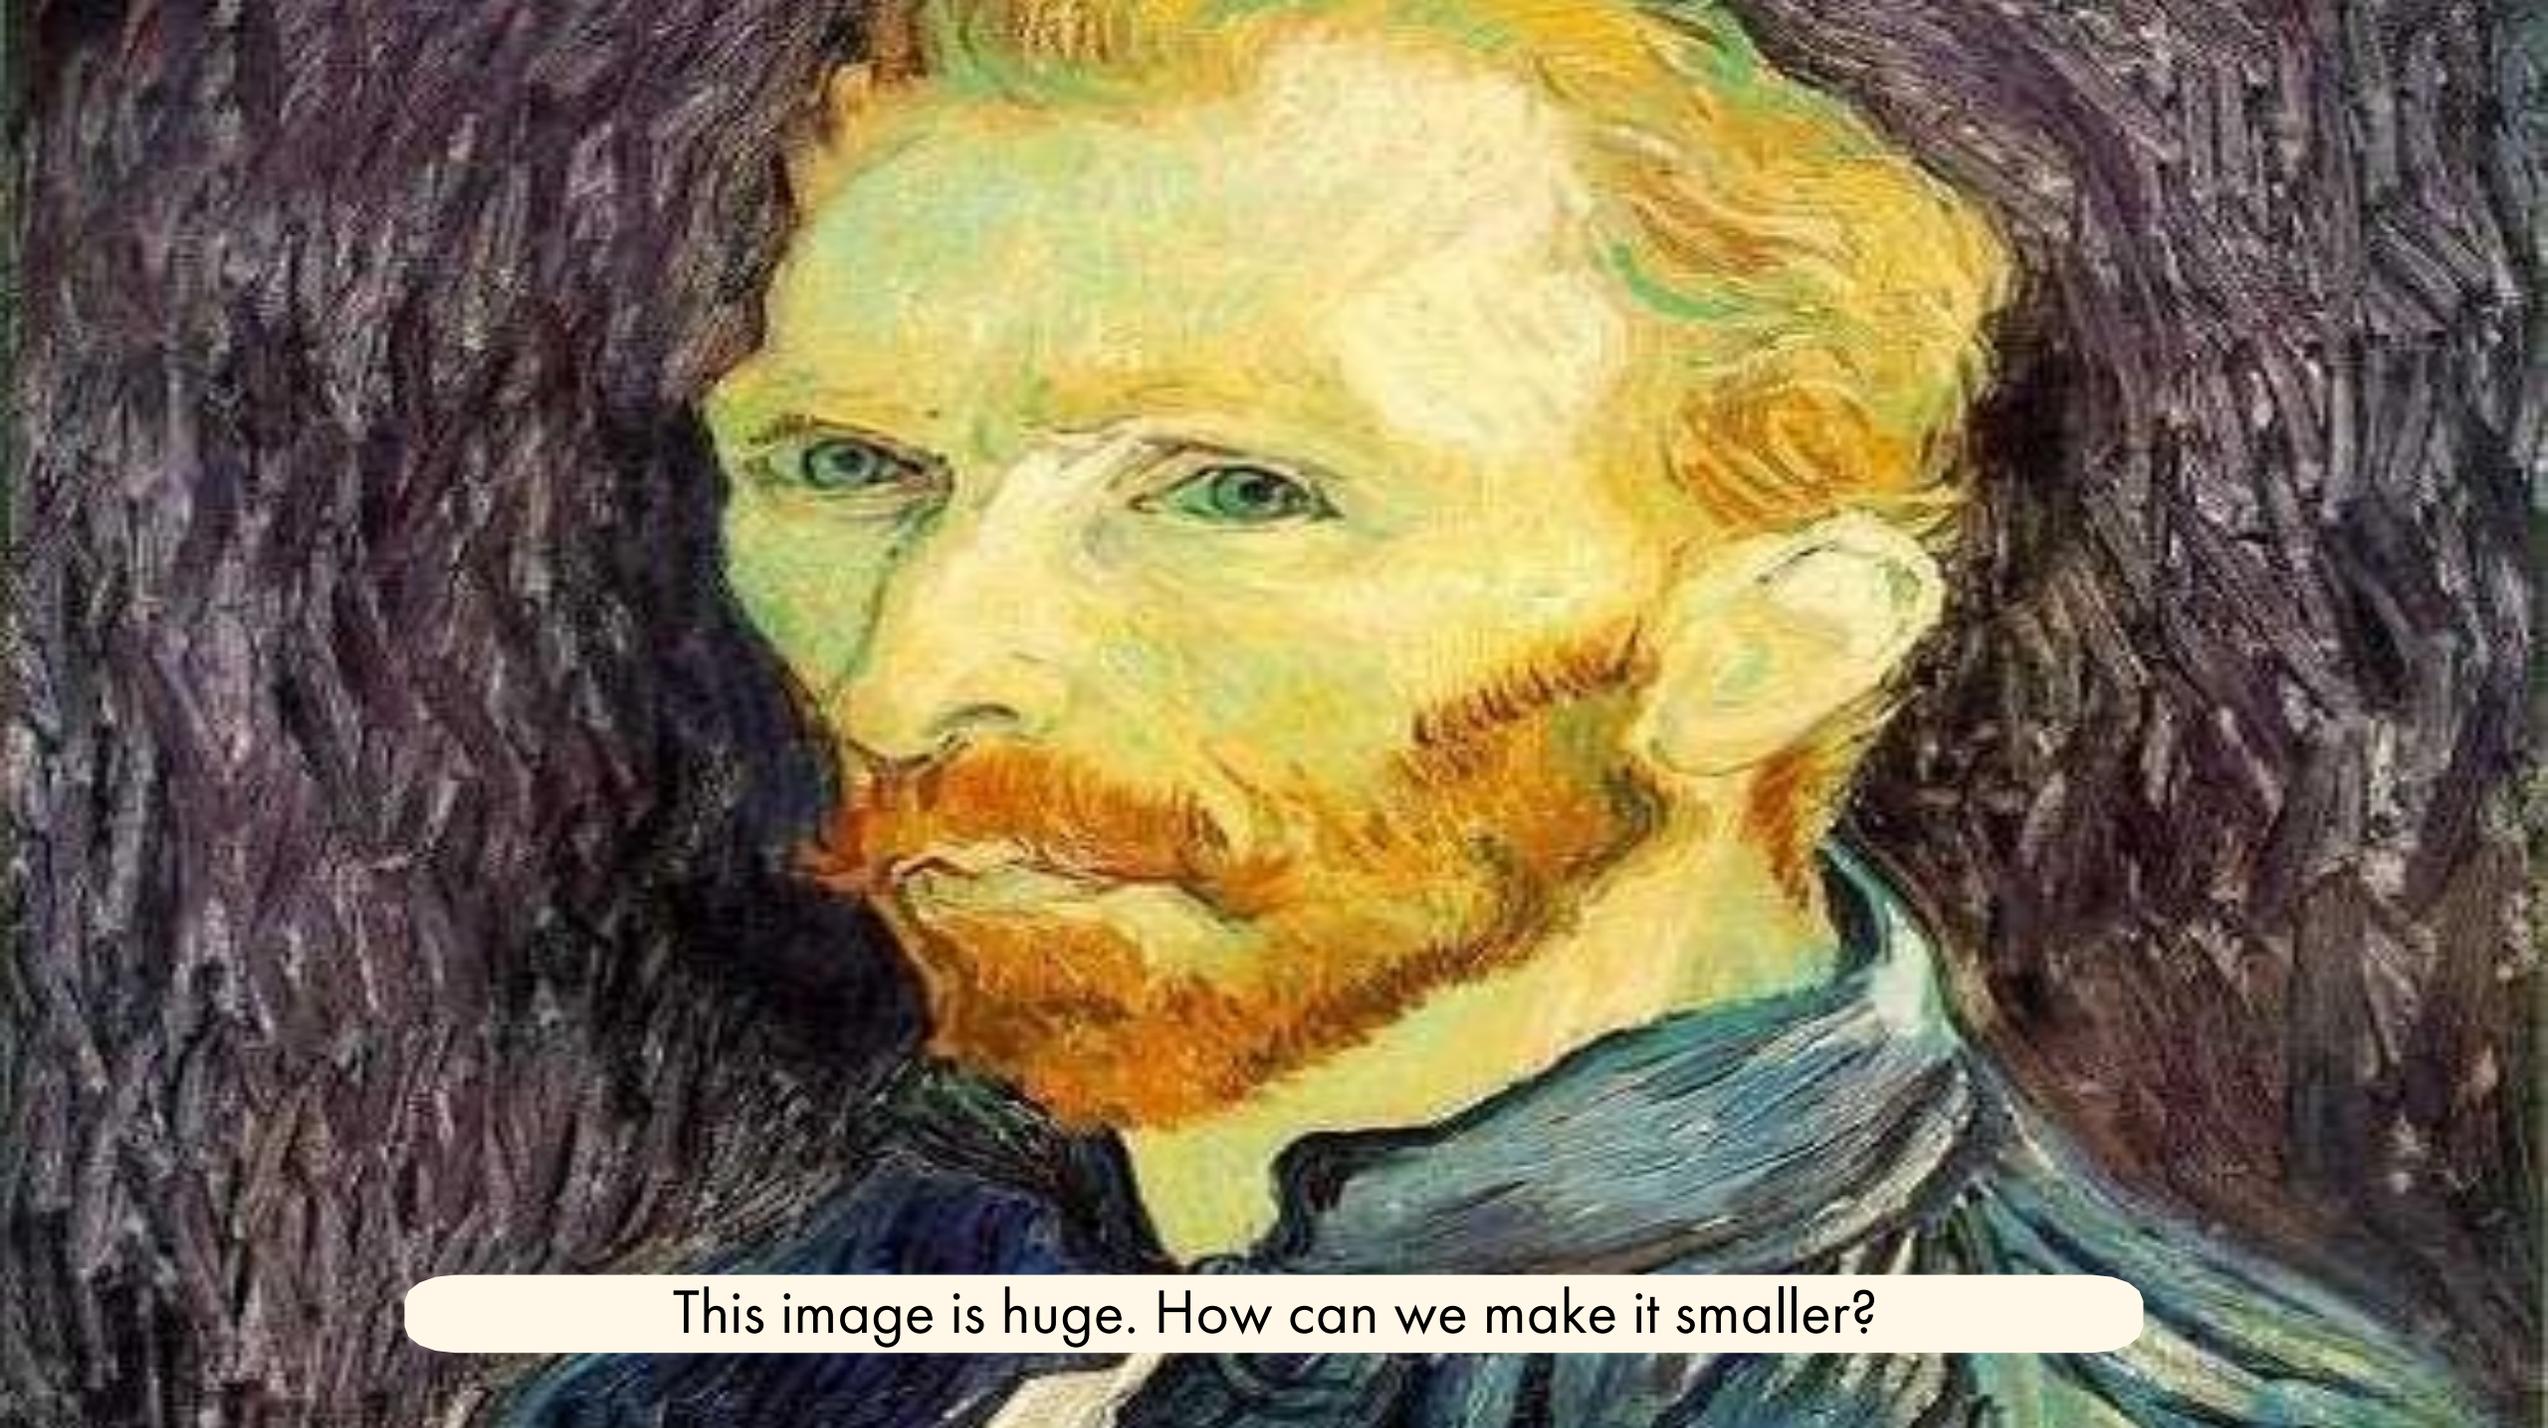
# Idea: Re-size Image

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



Template(filter)

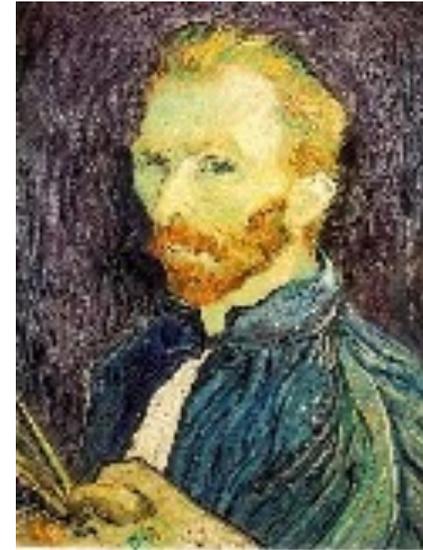
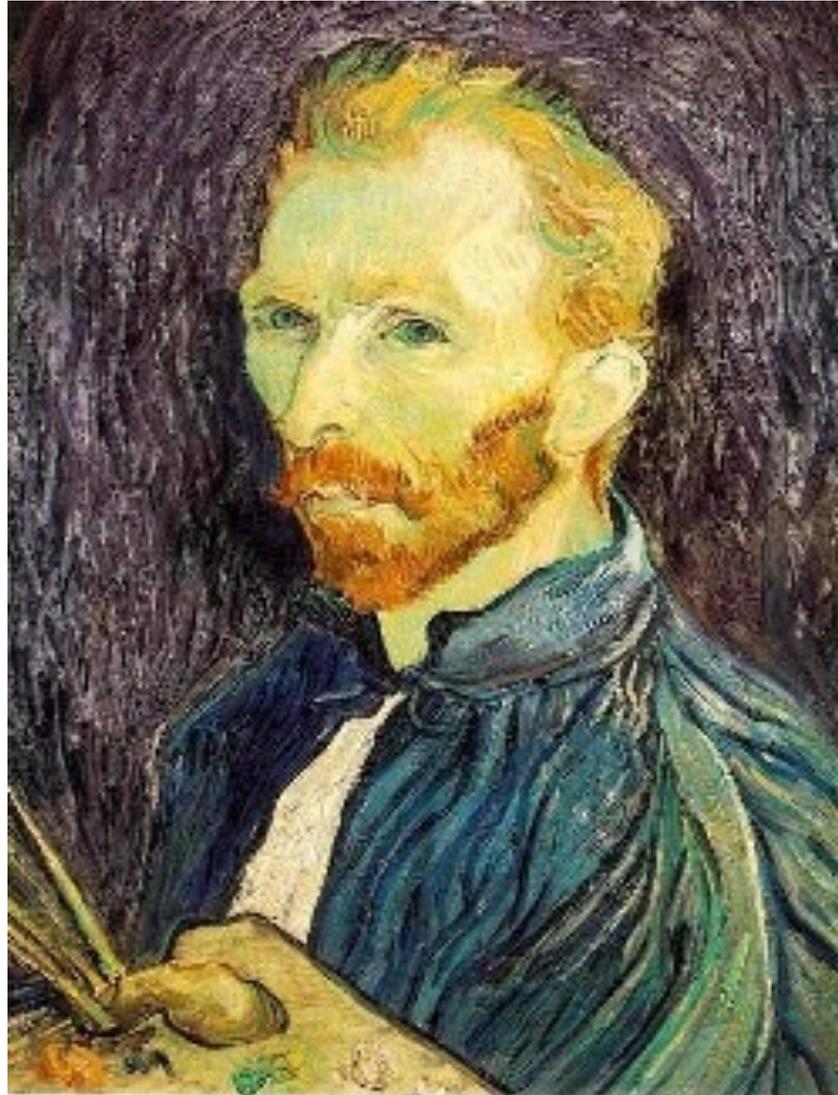




This image is huge. How can we make it smaller?

# Image Sub-Sampling

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



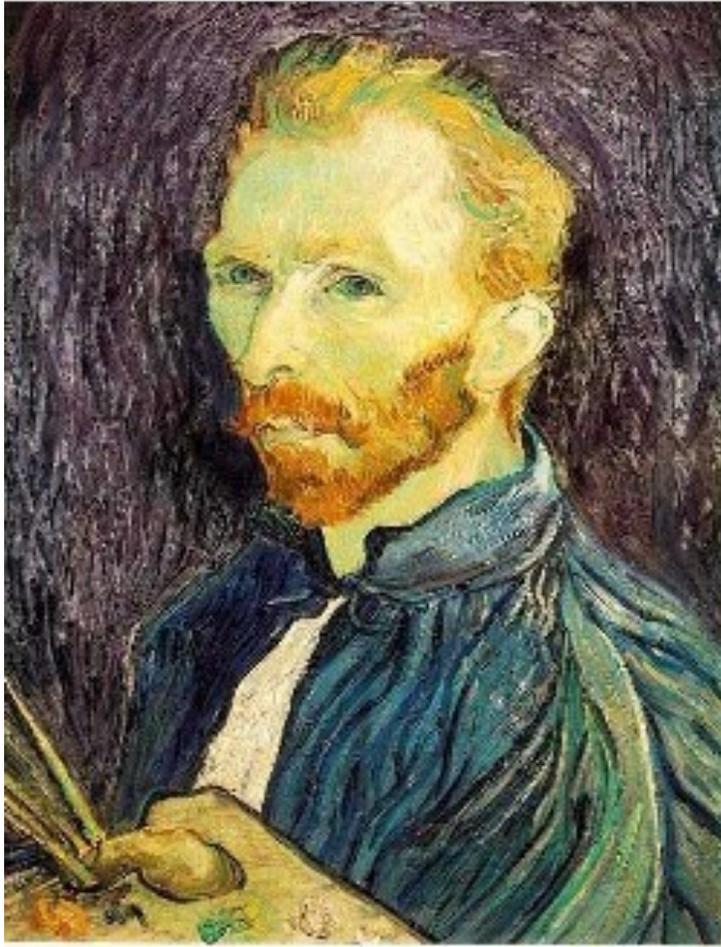
$1/4$



$1/8$

# Image Sub-Sampling

- Why does this look so cruffy?



$1/2$



$1/4$  (2x Zoom)



$1/8$  (4x Zoom)

[Source: S. Seitz]

# Even worse for synthetic images

- I want to resize my image by factor 2
- And I take every other column and every other row (1st, 3rd, 5th, etc)

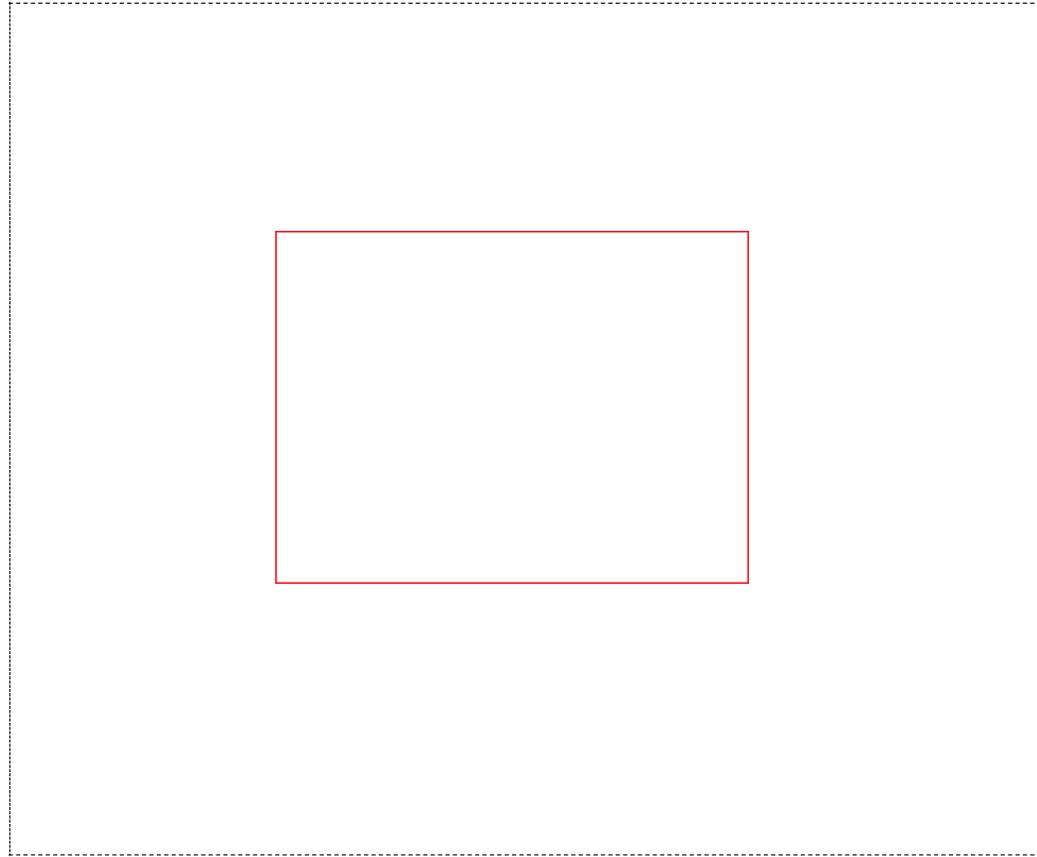
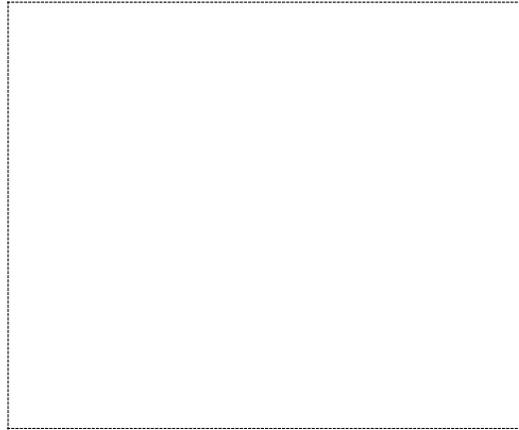


Figure: Dashed line denotes the border of the image (it's not part of the image)

# Even worse for synthetic images

- 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)

# Even worse for synthetic images

- 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)



# Even worse for synthetic images

- 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)

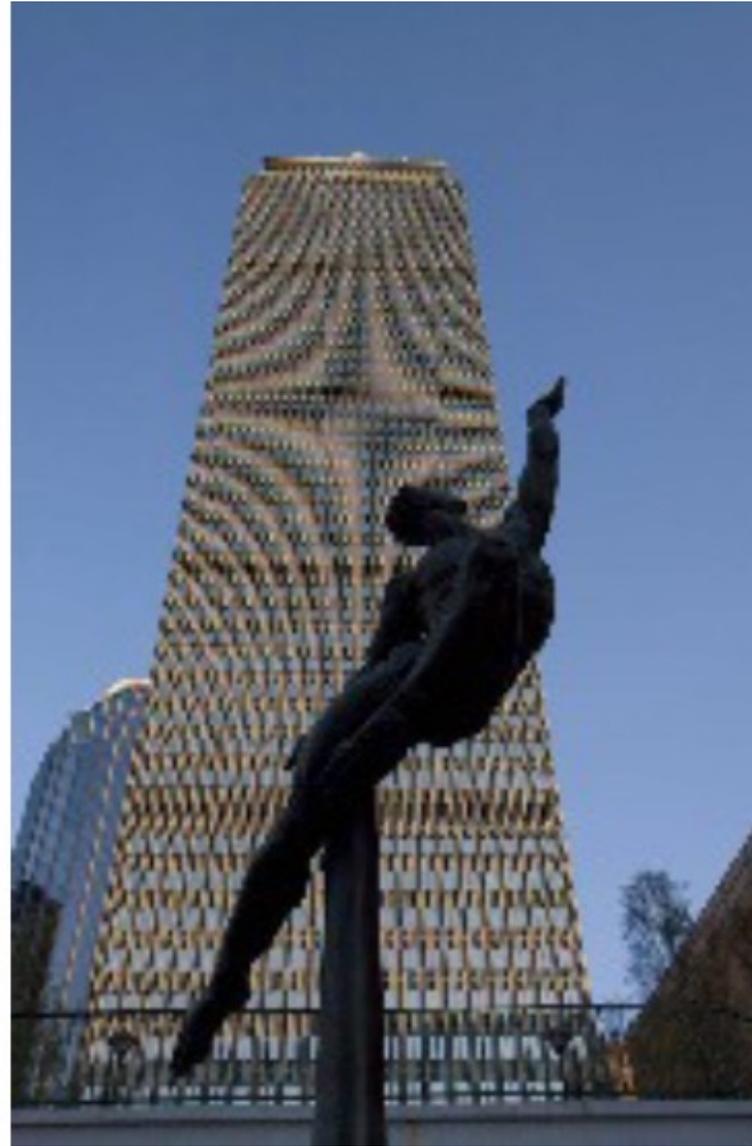
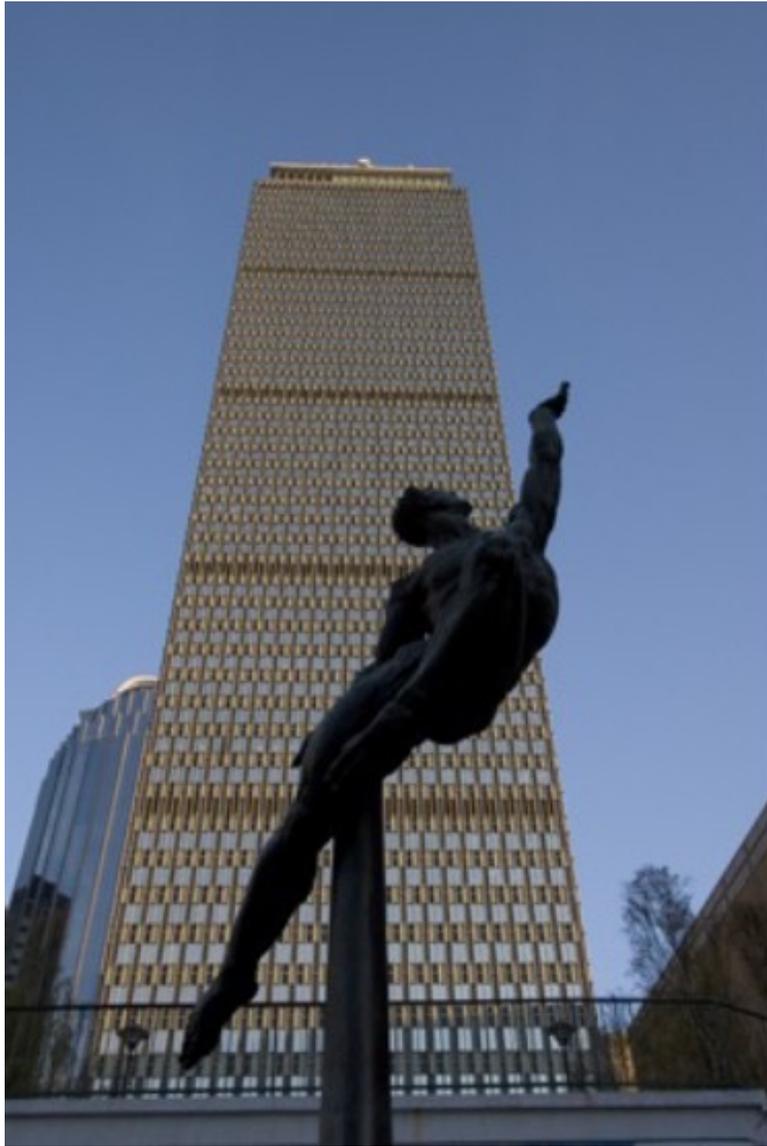


# Even worse for synthetic images

- 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!



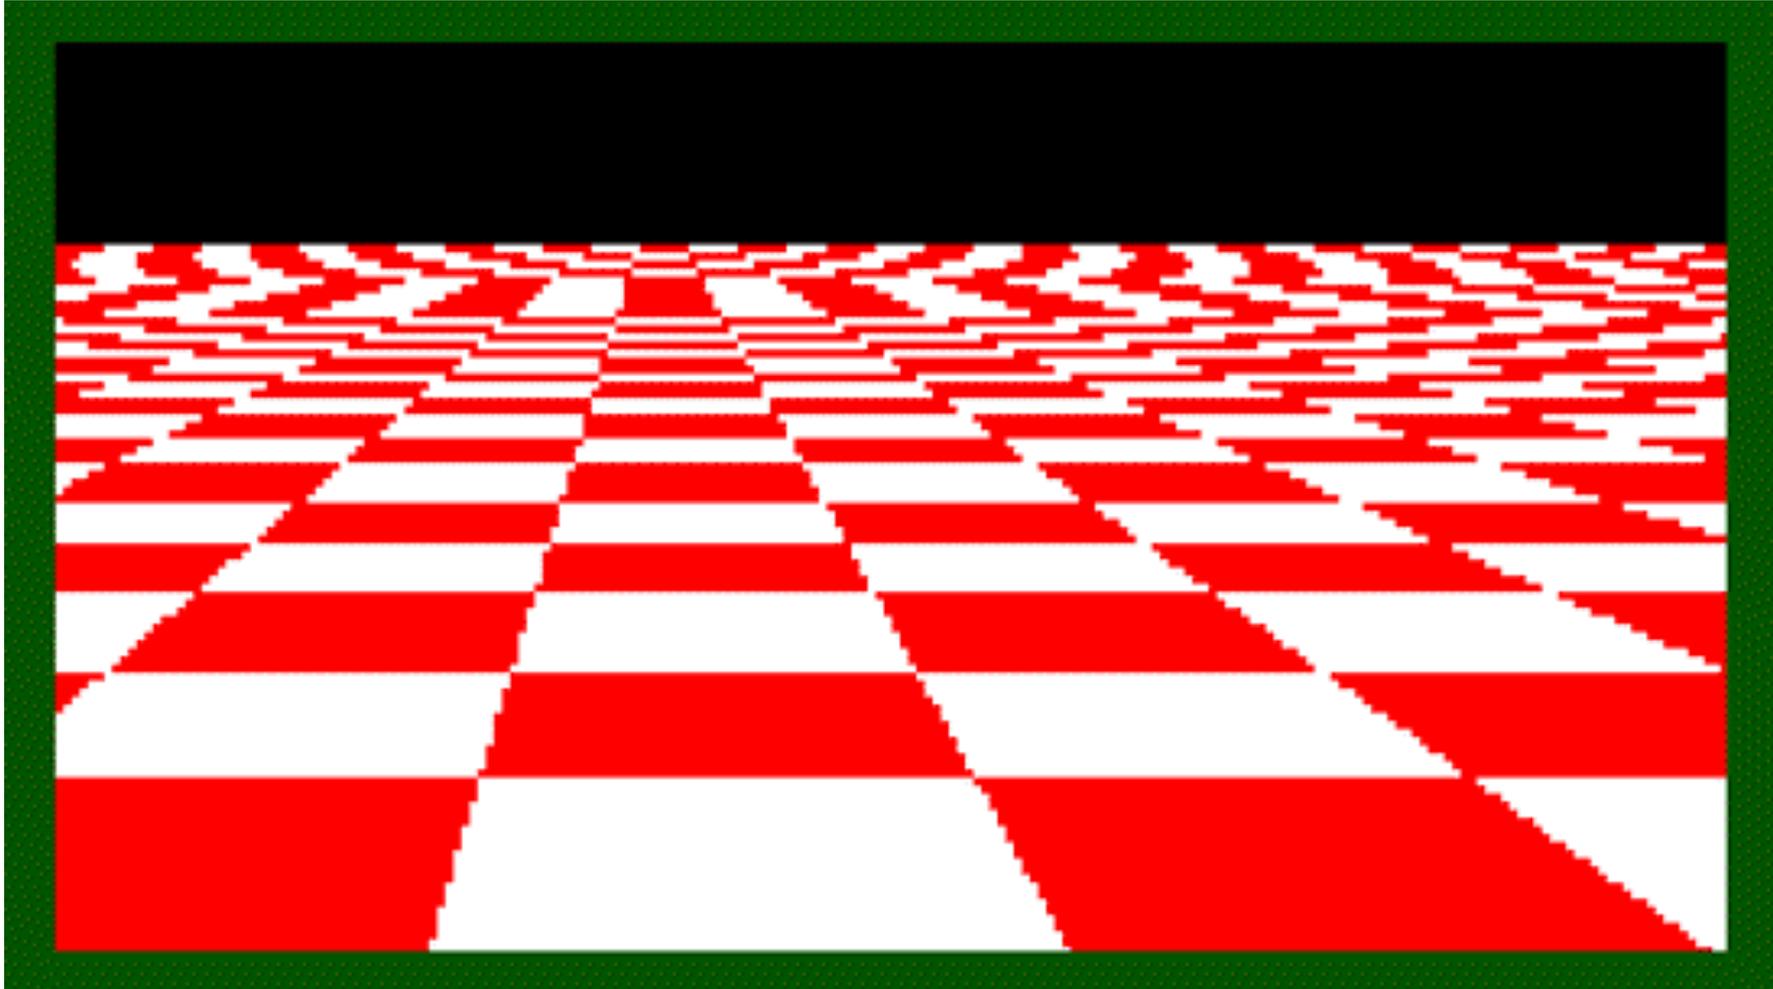
# Image Sub-Sampling



[Source: F. Durand]

# Even worse for synthetic images

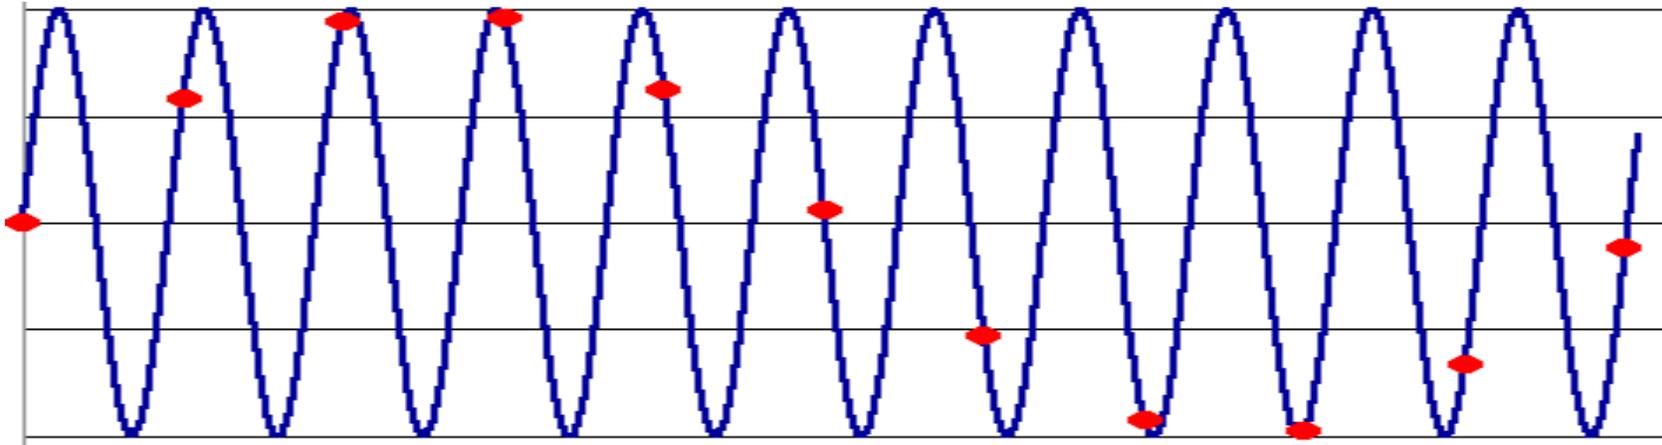
- What's happening?



[Source: L. Zhang]

# Aliasing

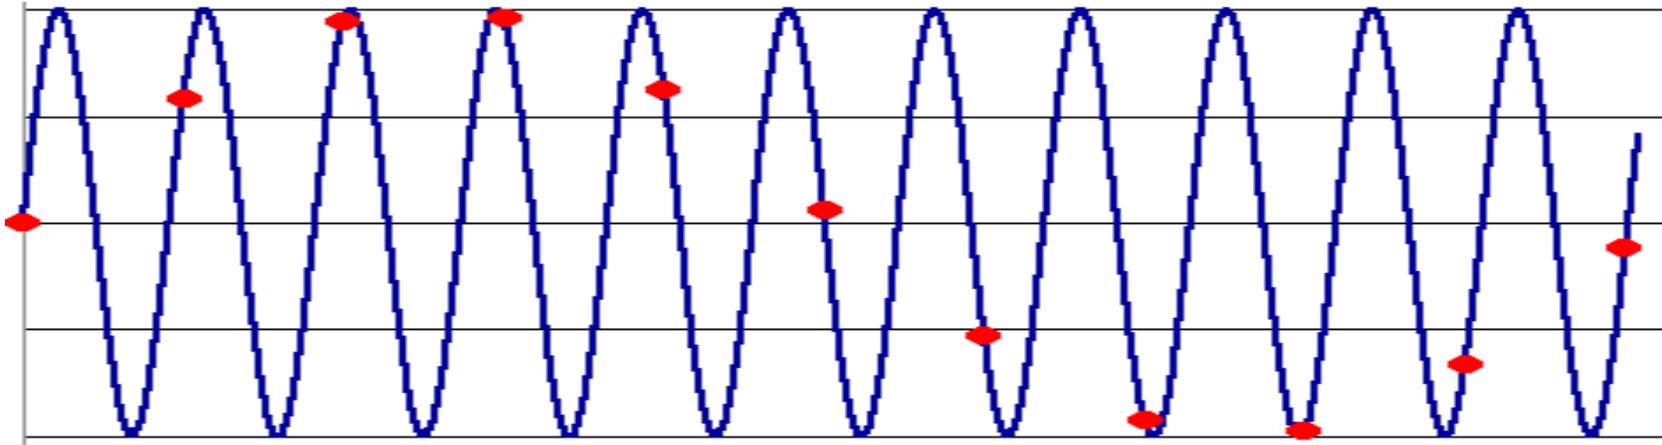
- Occurs when your sampling rate is not high enough to capture the amount of detail in your image



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

# Aliasing

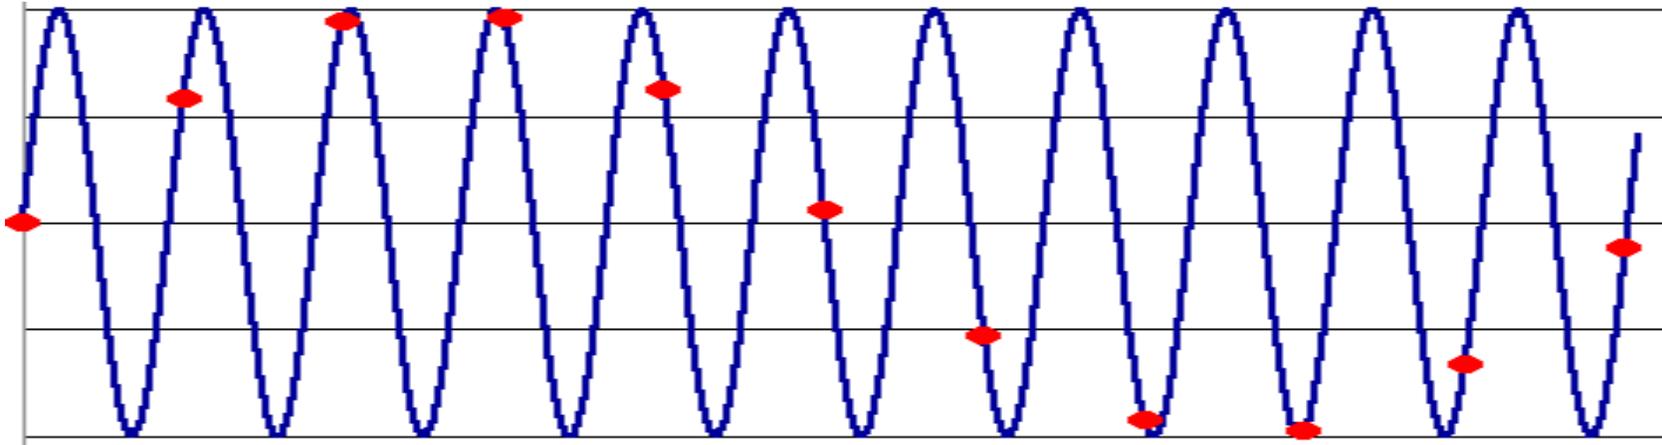
- Occurs when your sampling rate is not high enough to capture the amount of detail in your image



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

# 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

# Examples of Aliasing: Temporal Aliasing

- wagon wheel effect

[youtube.com/watch?v=jHS9JGkEOmA](https://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>



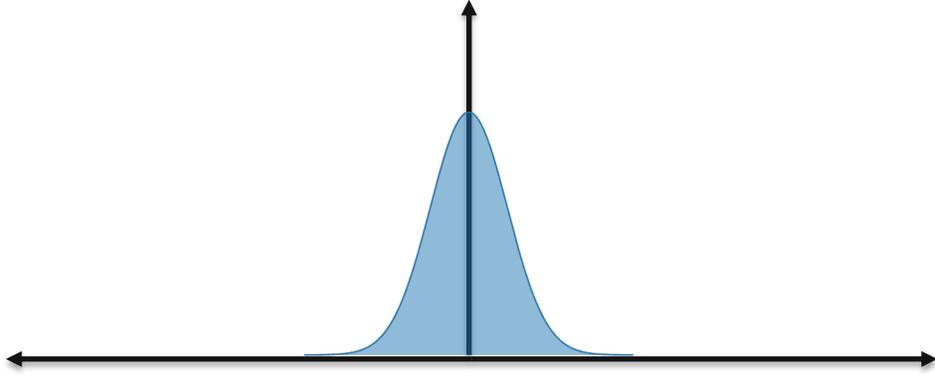
# Mr. Nyquist

- 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%99s\\_sampling\\_theorem](http://en.wikipedia.org/wiki/Nyquist%E2%80%99s_sampling_theorem)



# Sampling

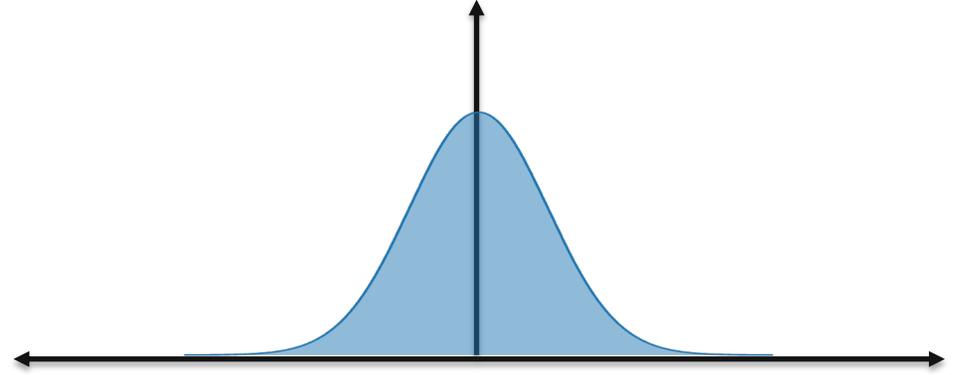
Primal Domain



$\mathcal{F}$

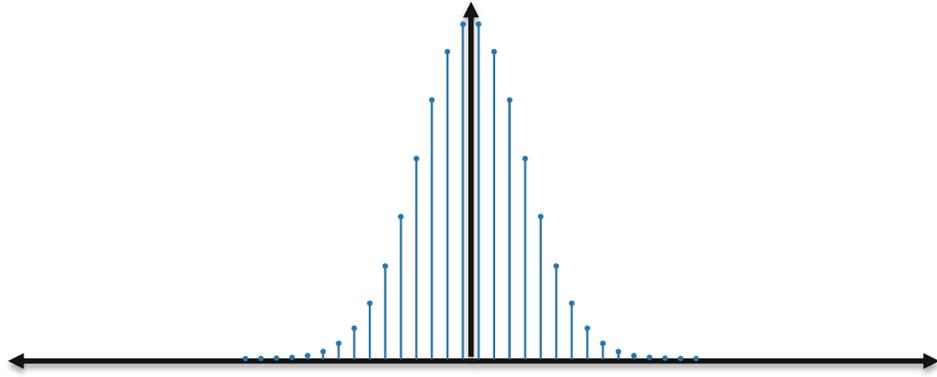


Fourier Domain



# Sampling

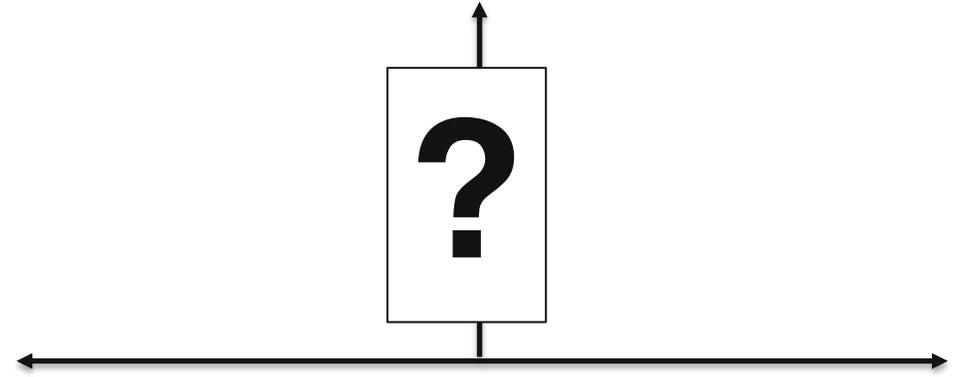
Primal Domain



discrete sampled signal

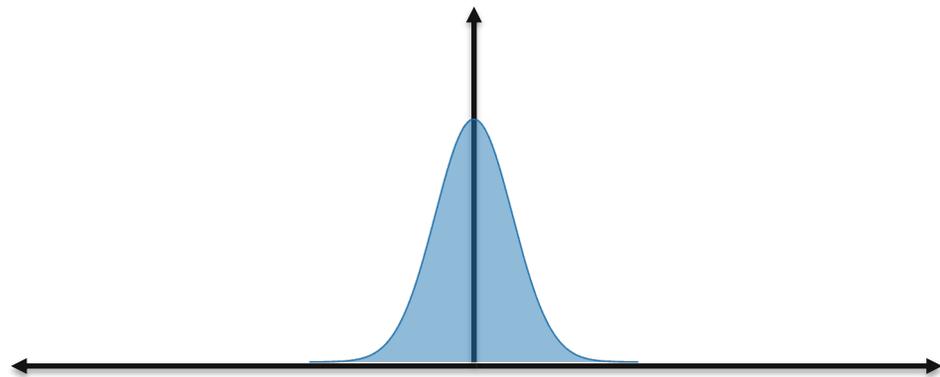
$\mathcal{F}$

Fourier Domain



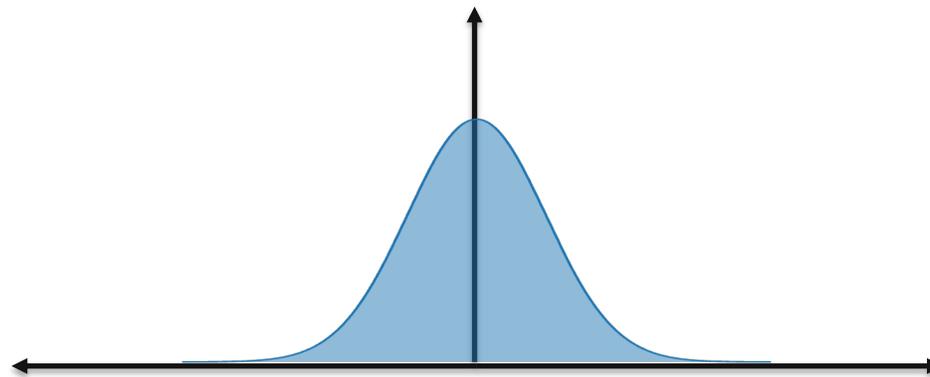
# Sampling

Primal Domain



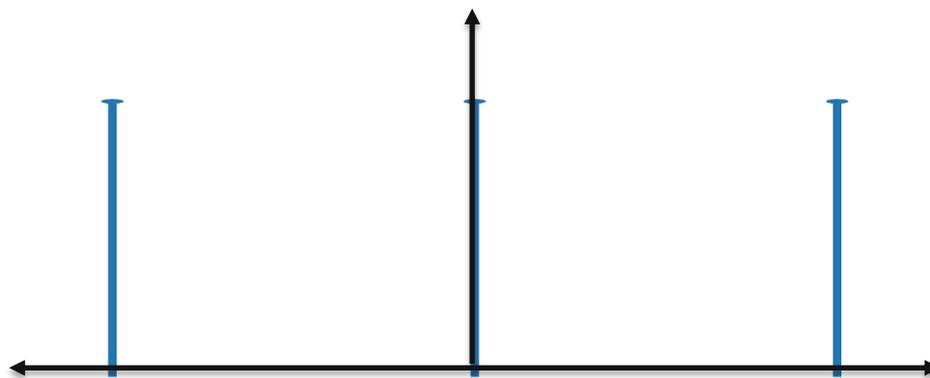
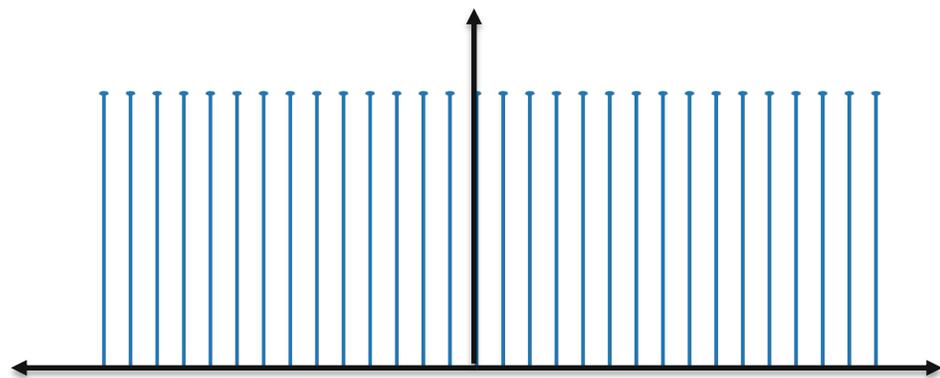
$\mathcal{F}$

Fourier Domain



 Sampling operator

$*$

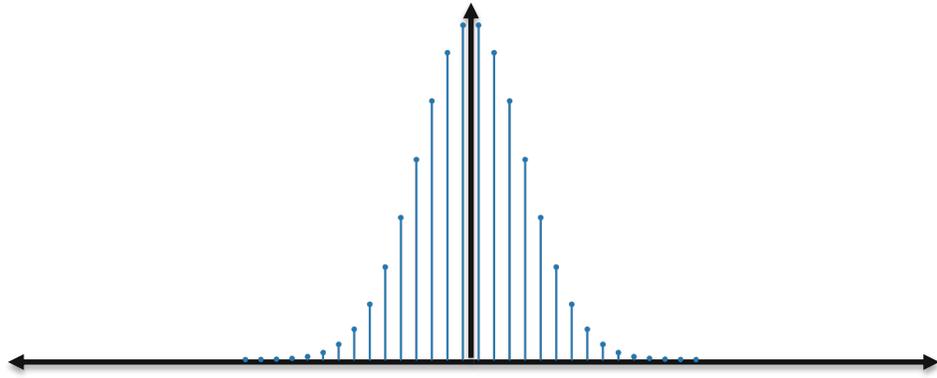


Sample rate of  $f_s$

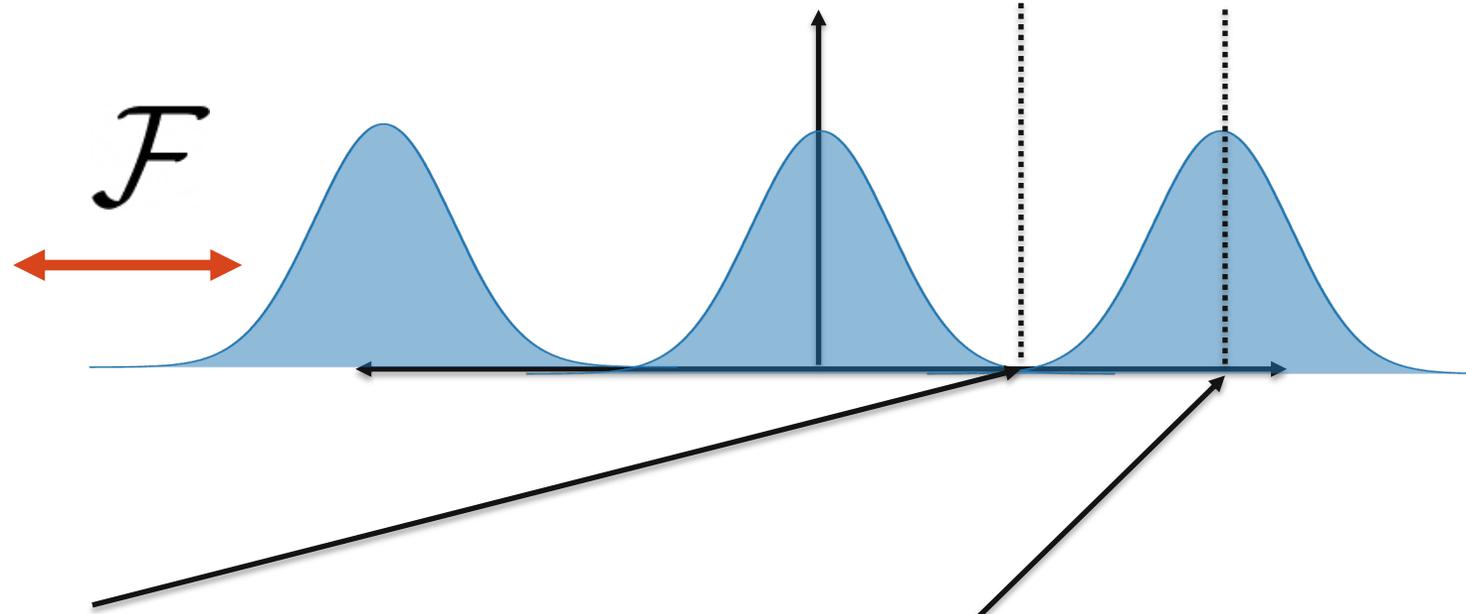
Shifted copies at  $f_s$

# Sampling

Primal Domain



Fourier Domain



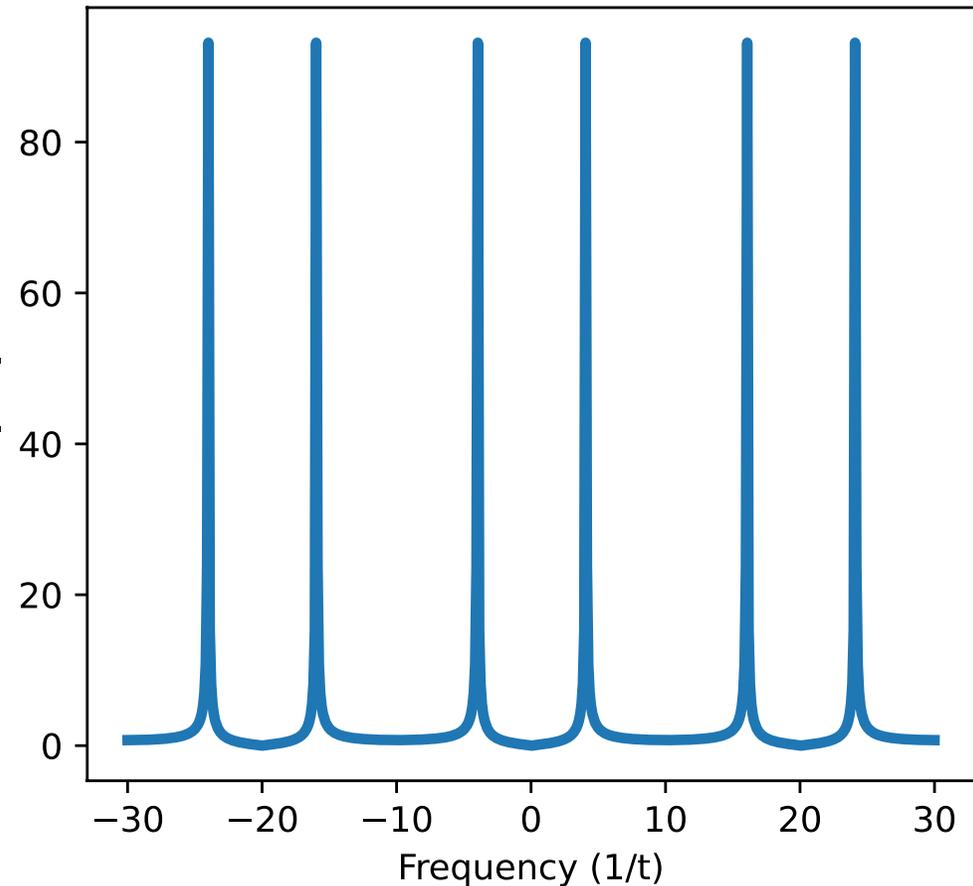
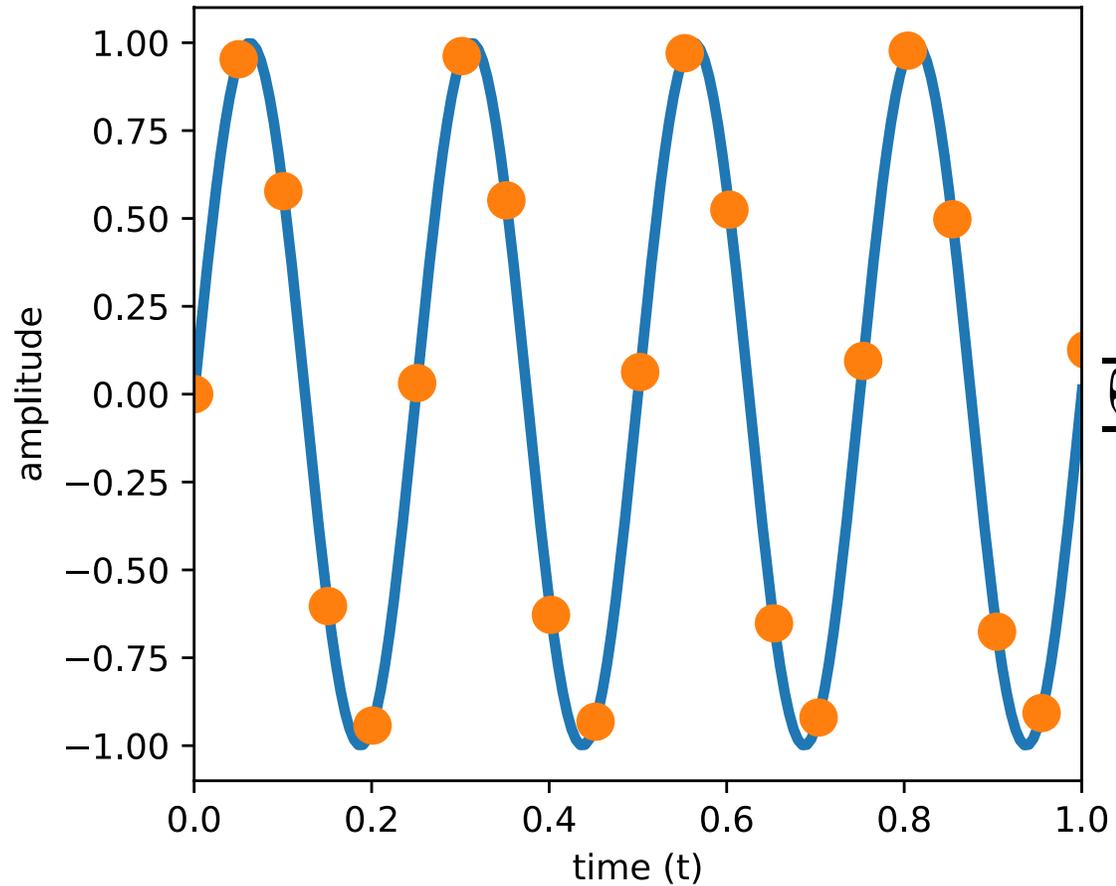
Highest frequency

Sample rate should be twice the highest frequency to avoid aliasing!

# Sampling exercise

Sample frequency: ?

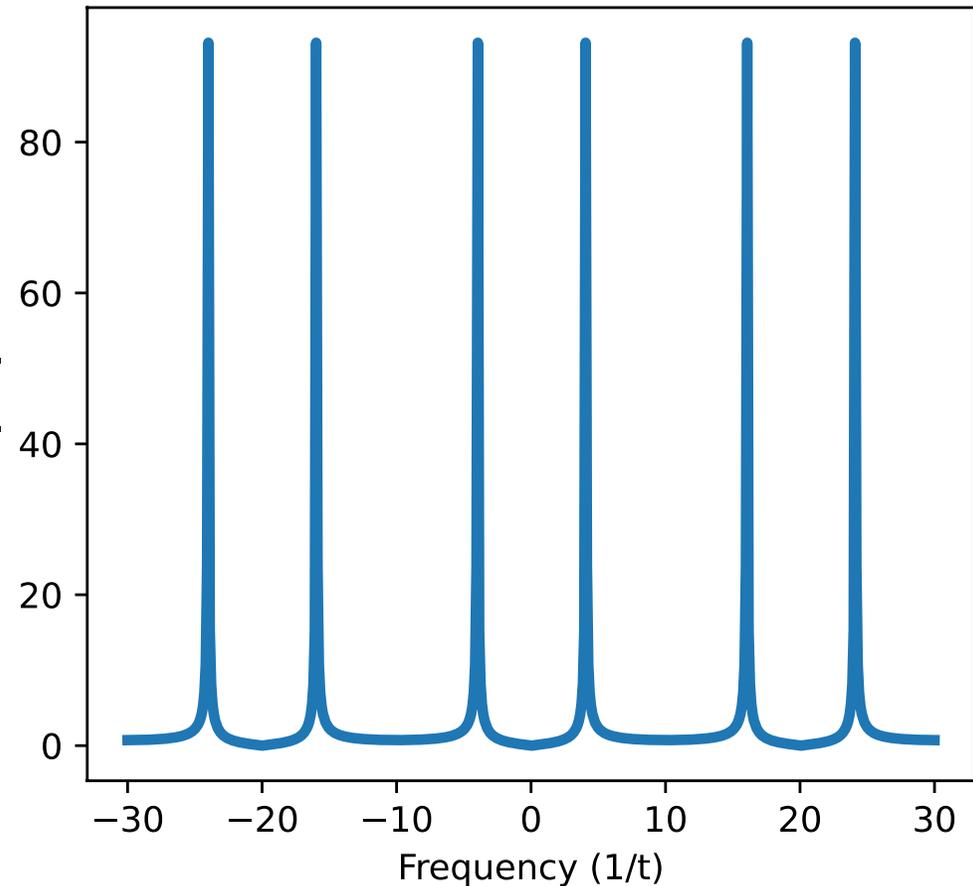
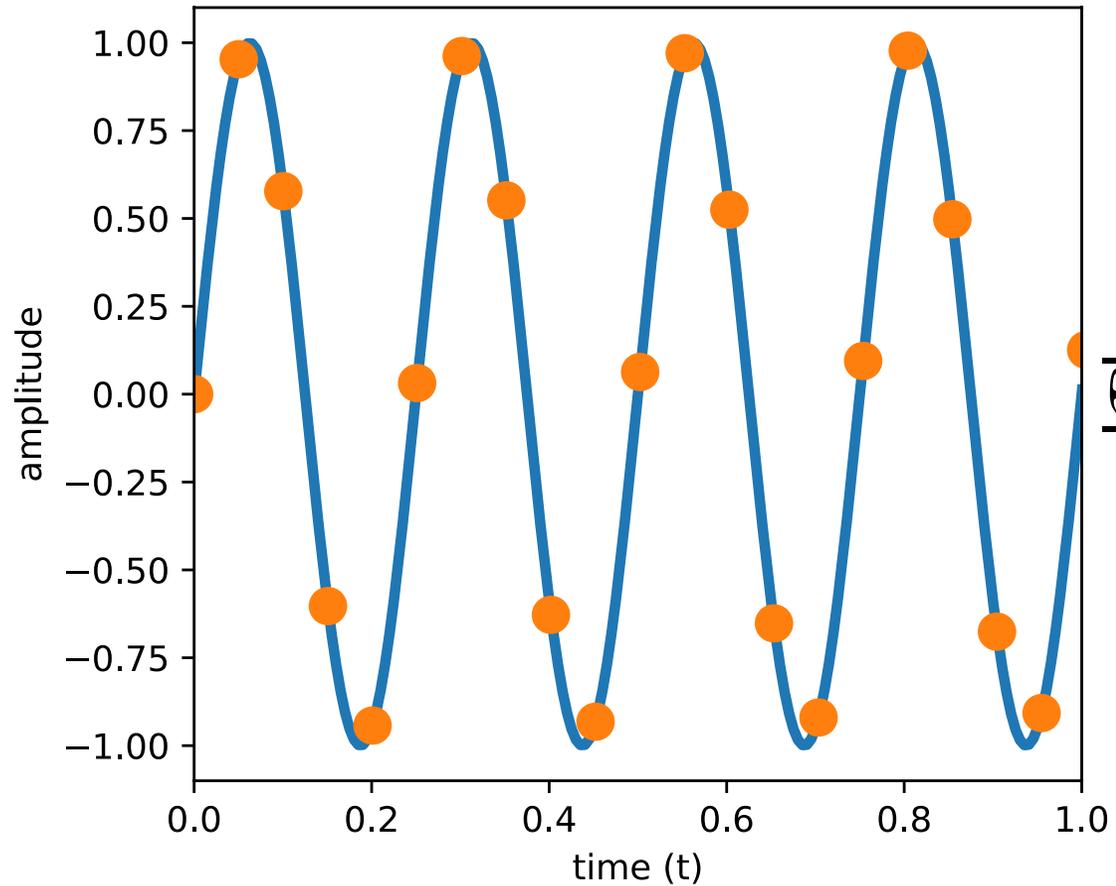
Signal frequency: ?



# Sampling exercise

Sample frequency: 20 Hz

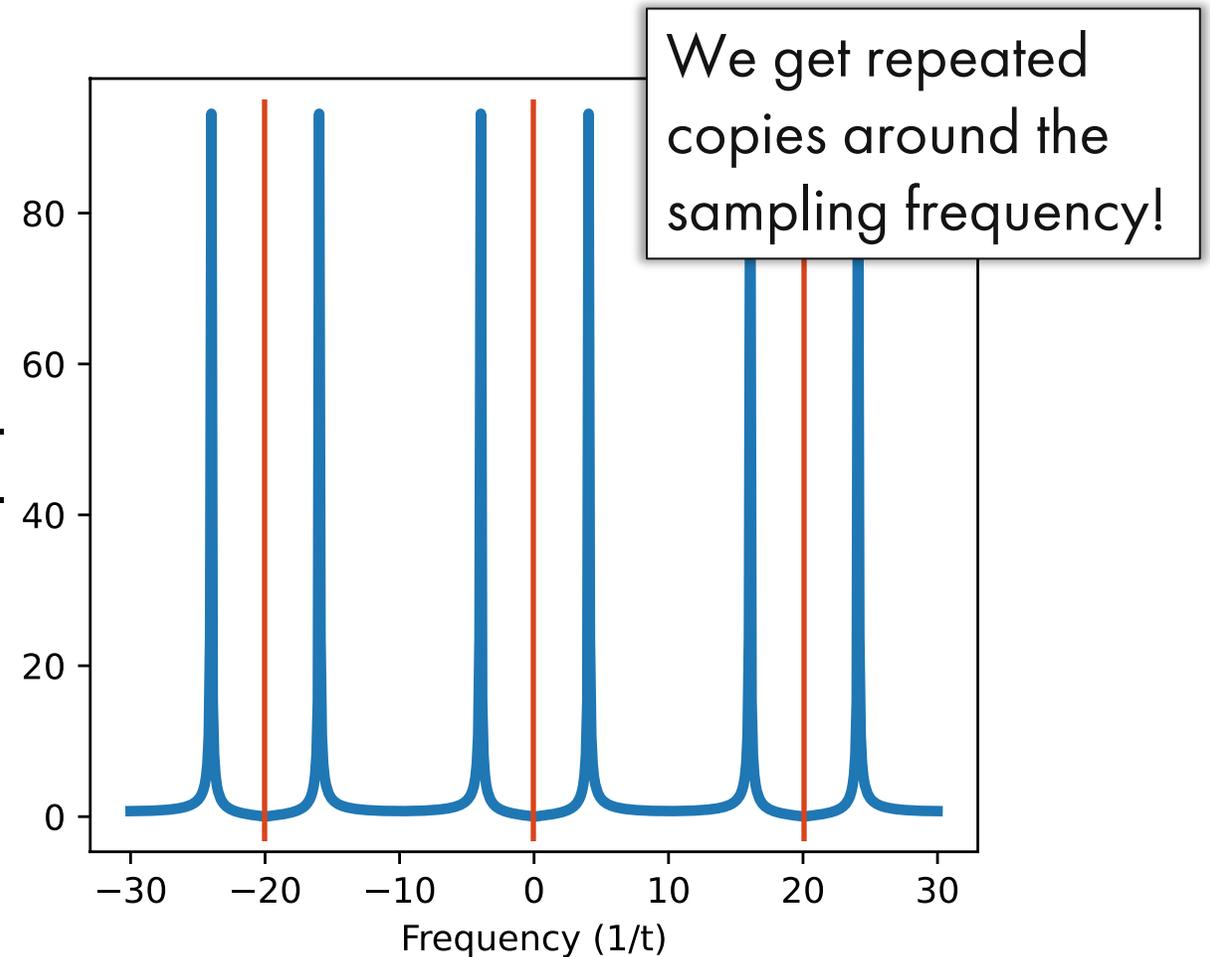
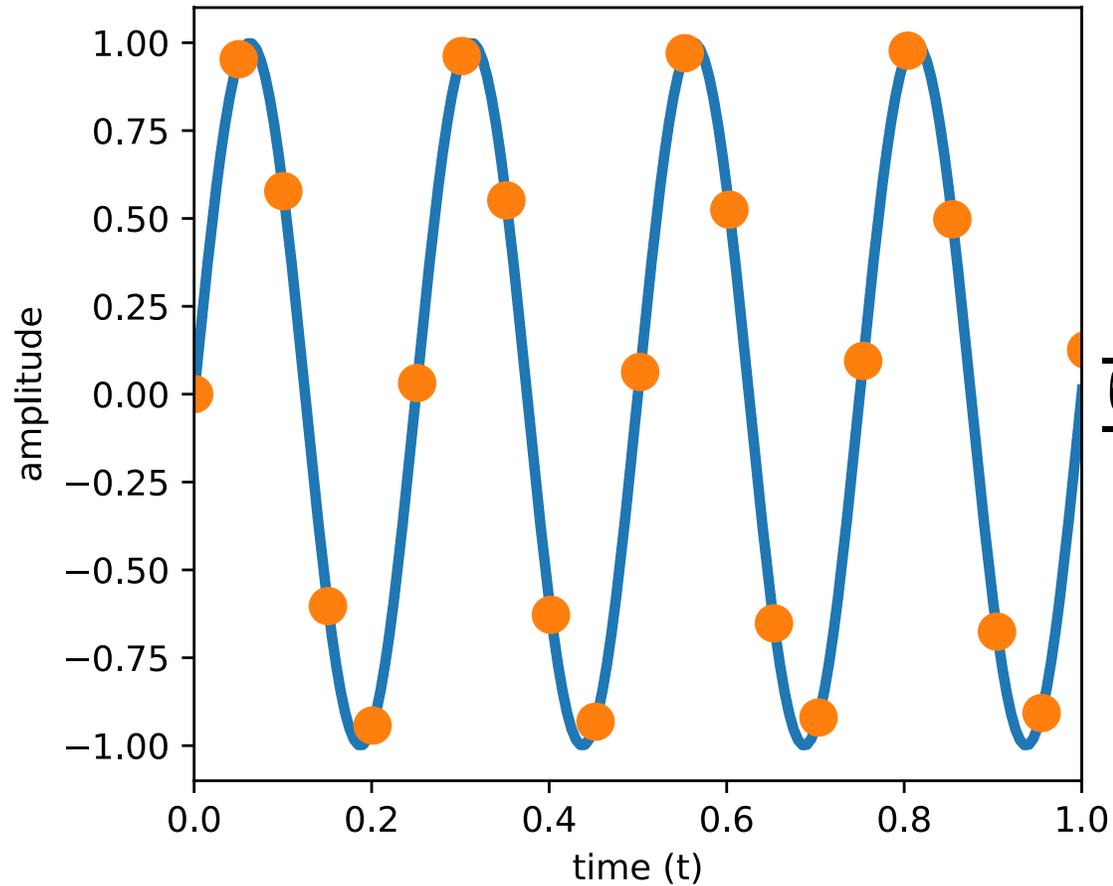
Signal: 4 Hz



# Sampling exercise

Sample frequency: 20 Hz

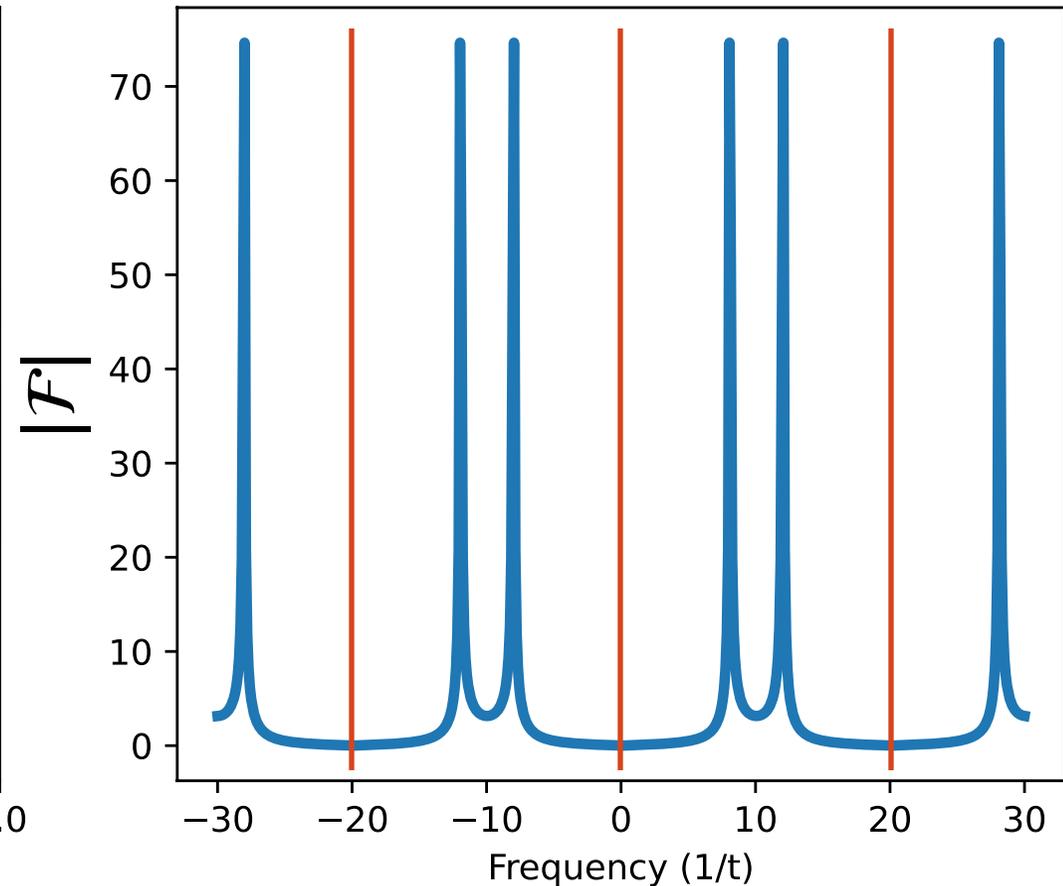
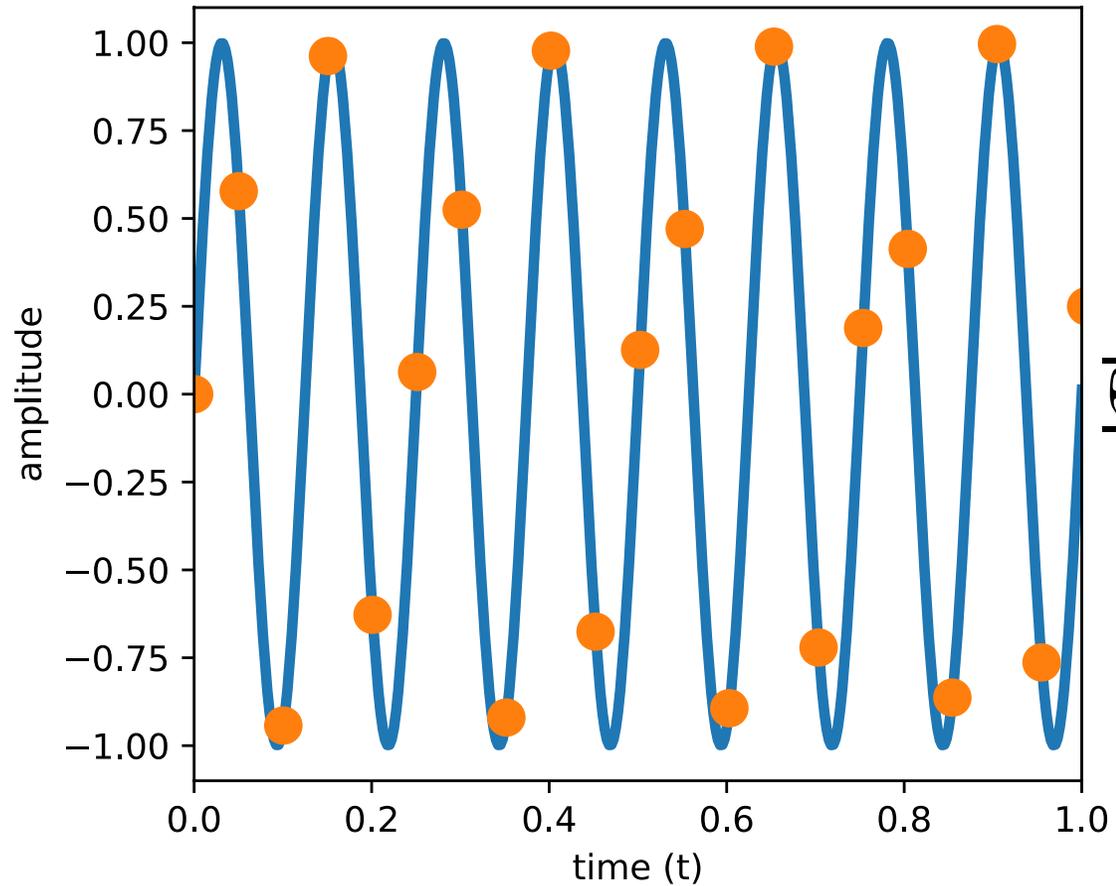
Signal: 4 Hz



# Sampling exercise

Sample frequency: 20 Hz

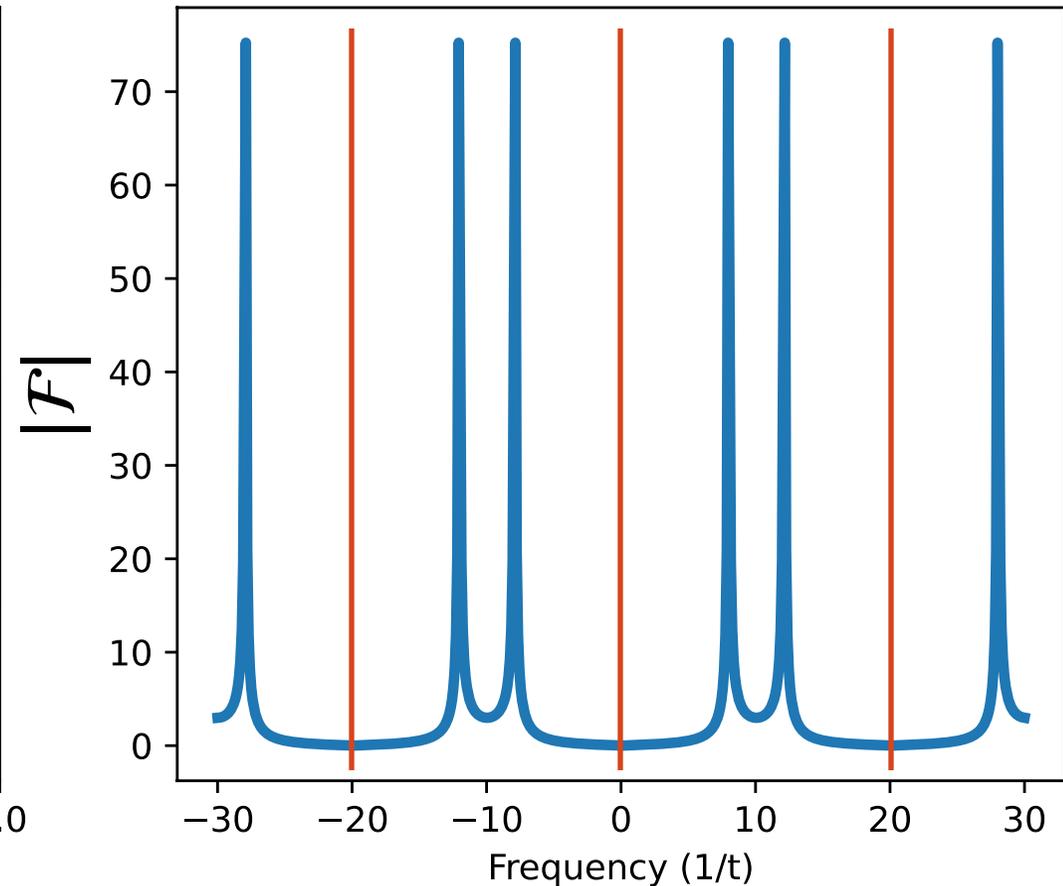
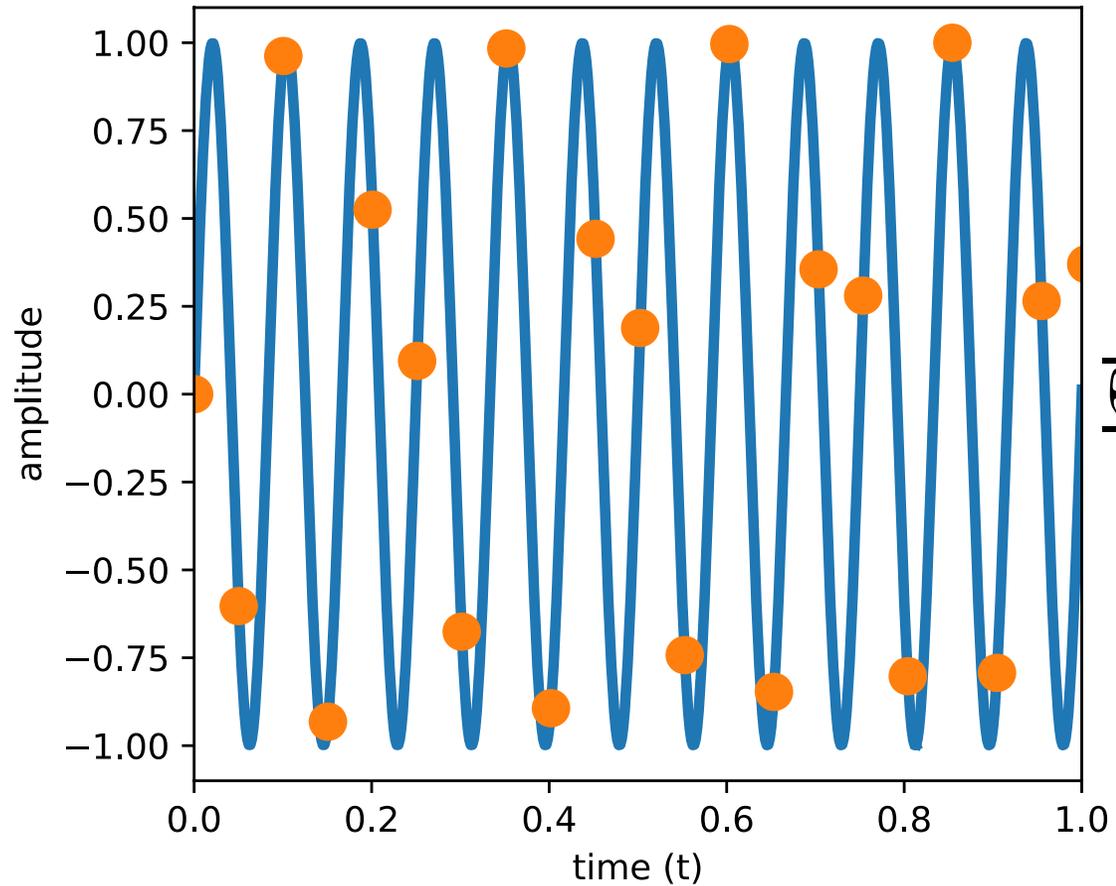
Signal: 8 Hz



# Sampling exercise

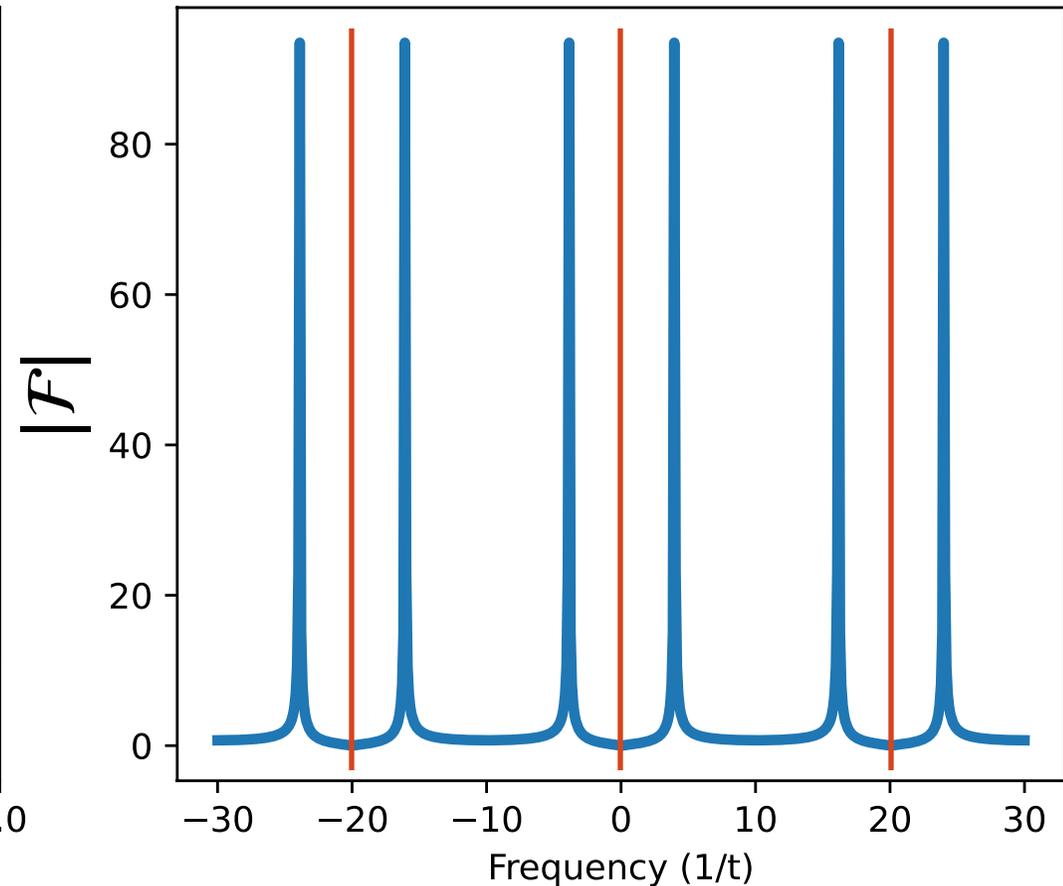
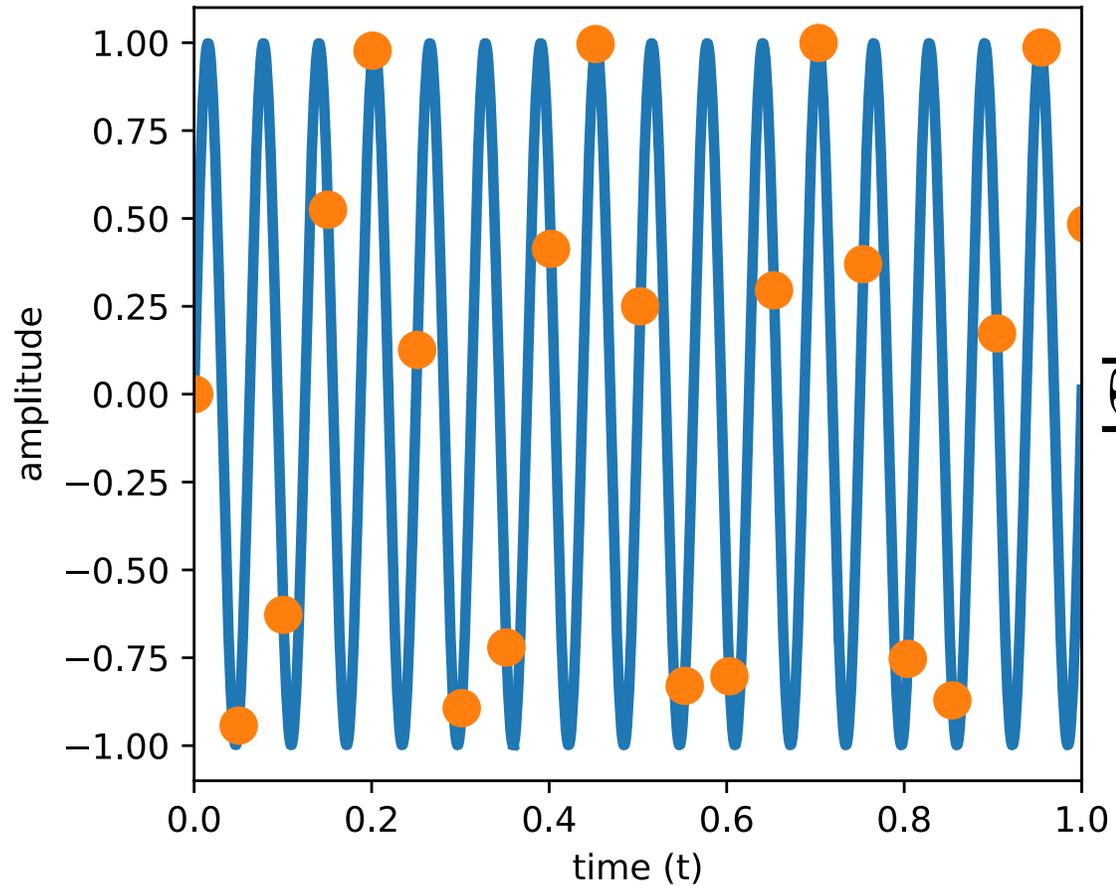
Sample frequency: 20 Hz

Signal: 12 Hz



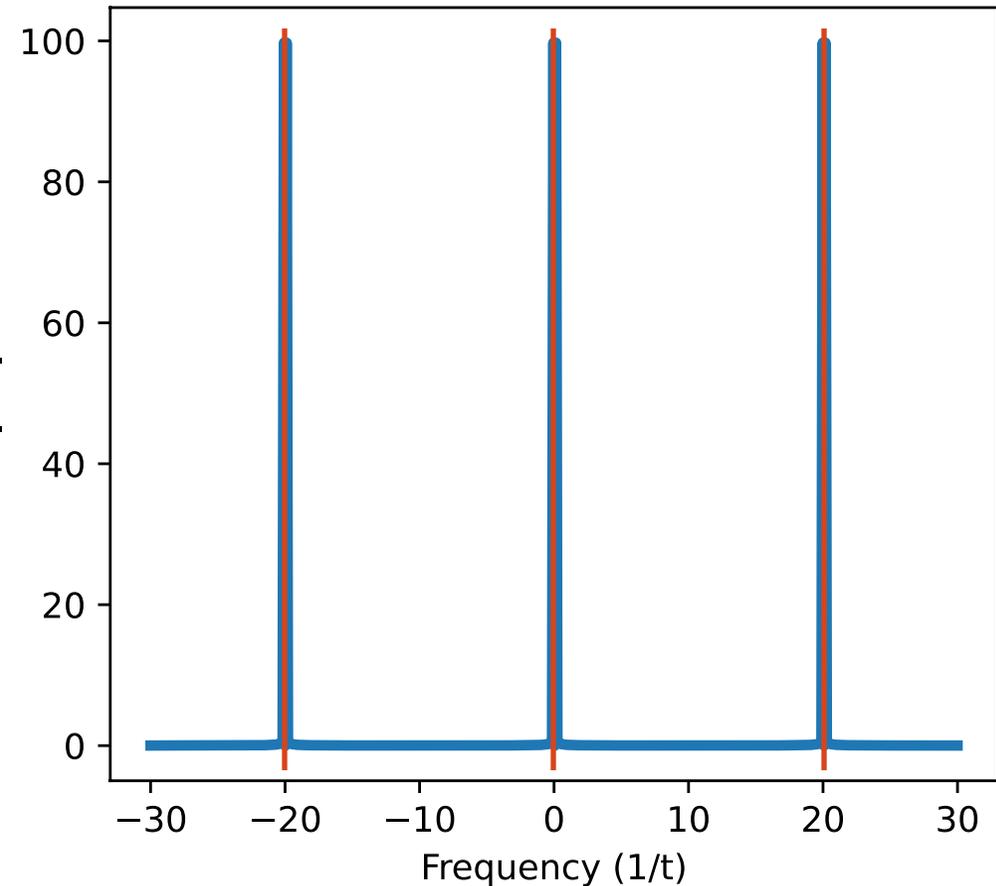
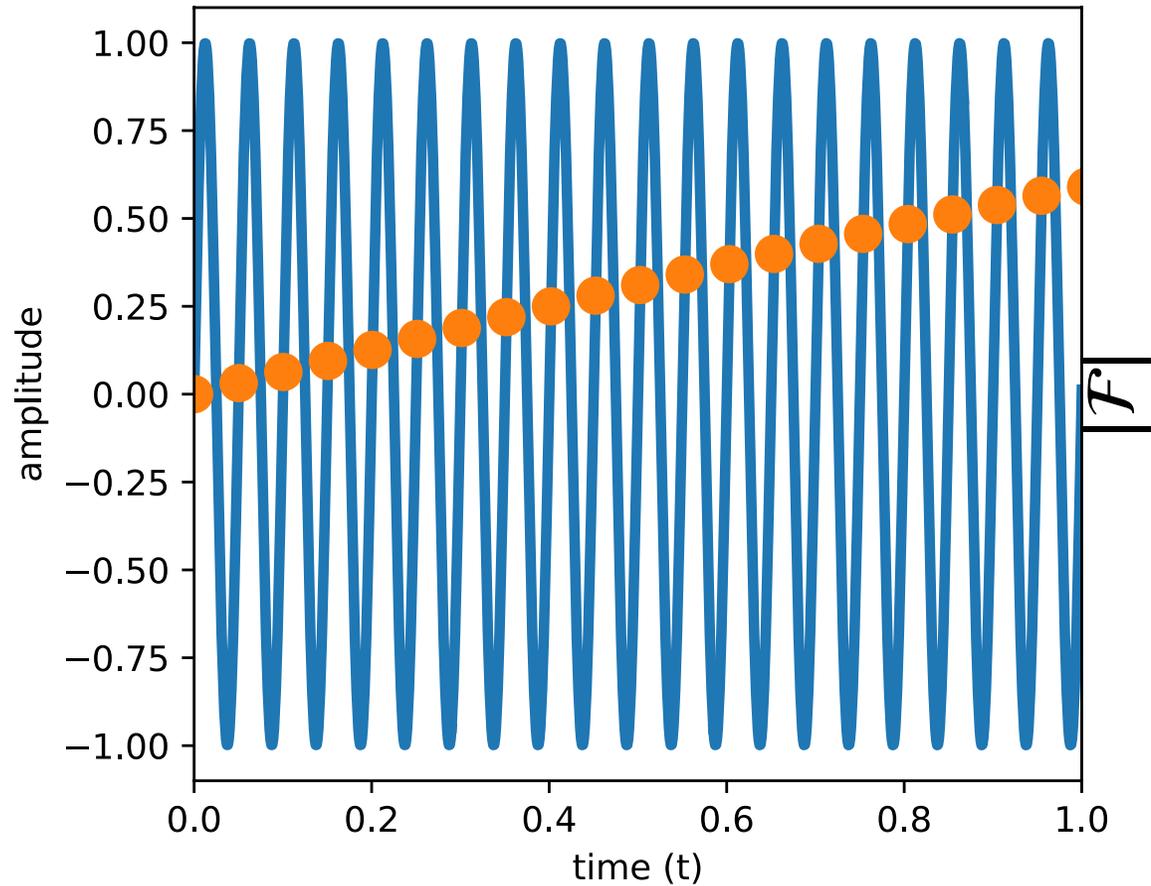
# Sampling exercise

Sample frequency: 20 Hz  
Signal: 16 Hz



# Sampling exercise

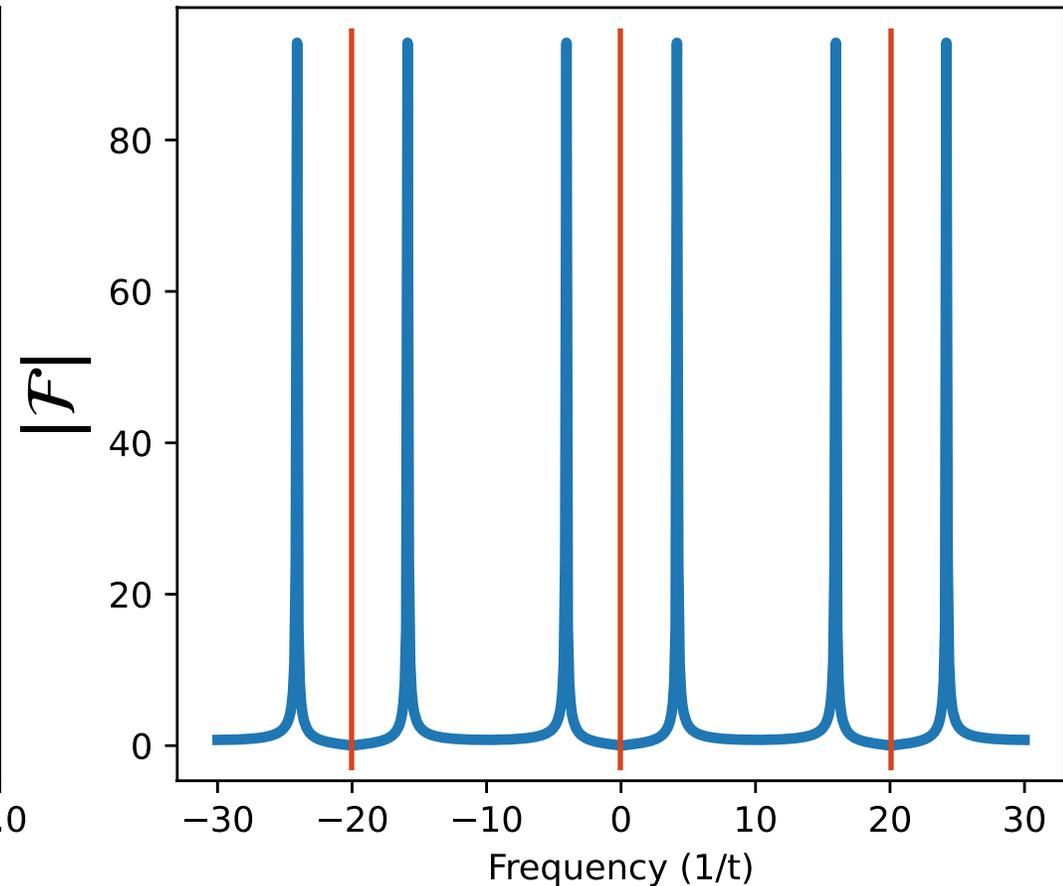
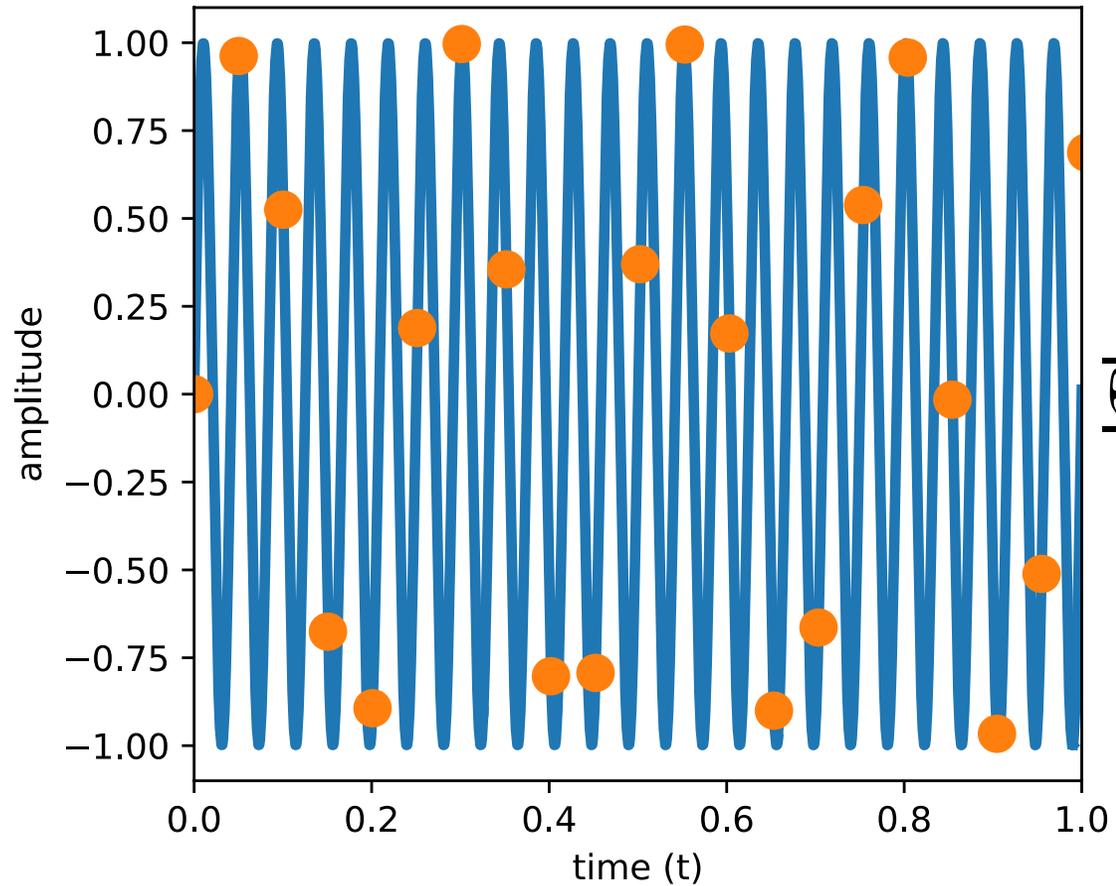
Sample frequency: 20 Hz  
Signal: 20 Hz



# Sampling exercise

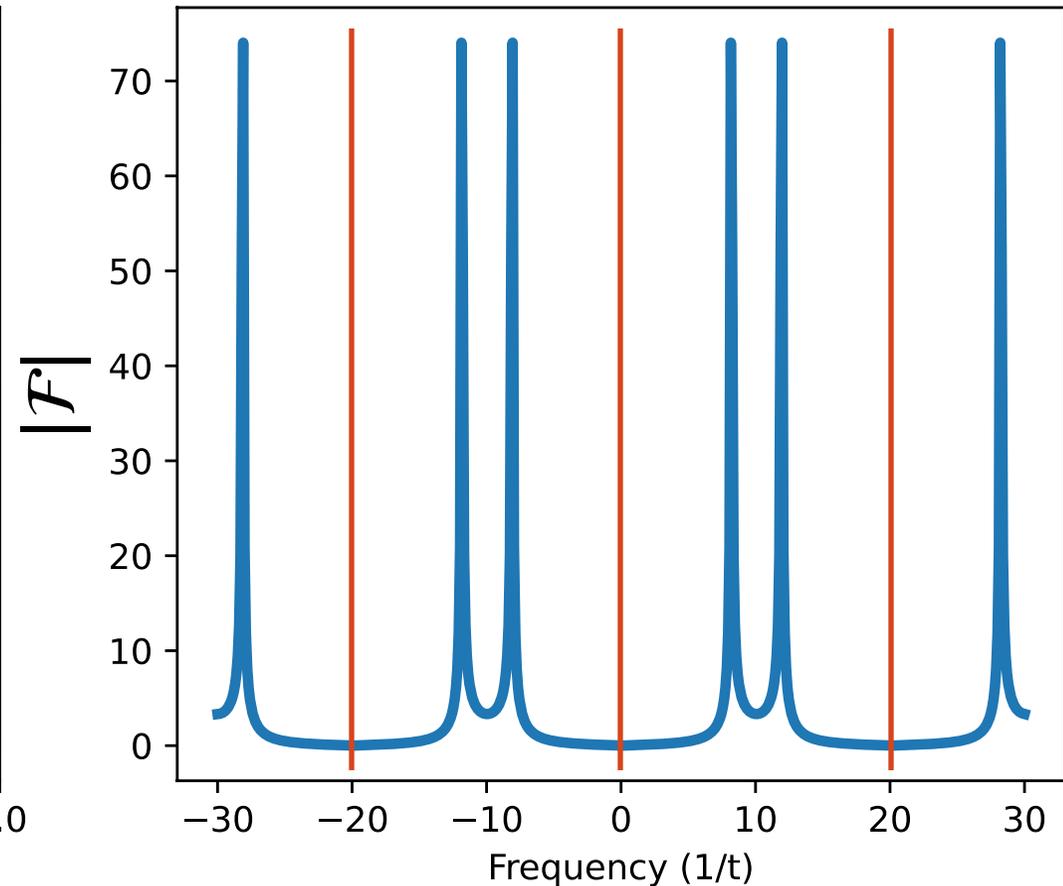
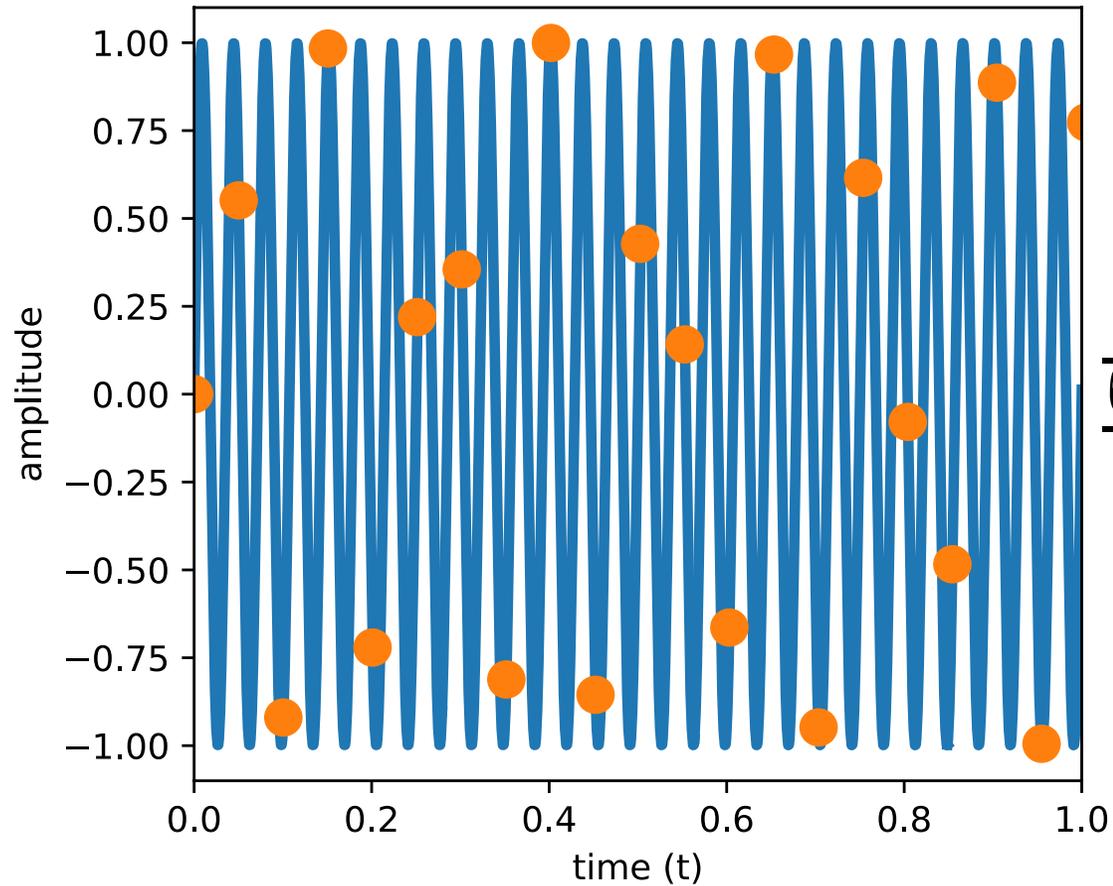
Sample frequency: 20 Hz

Signal: 24 Hz



# Sampling exercise

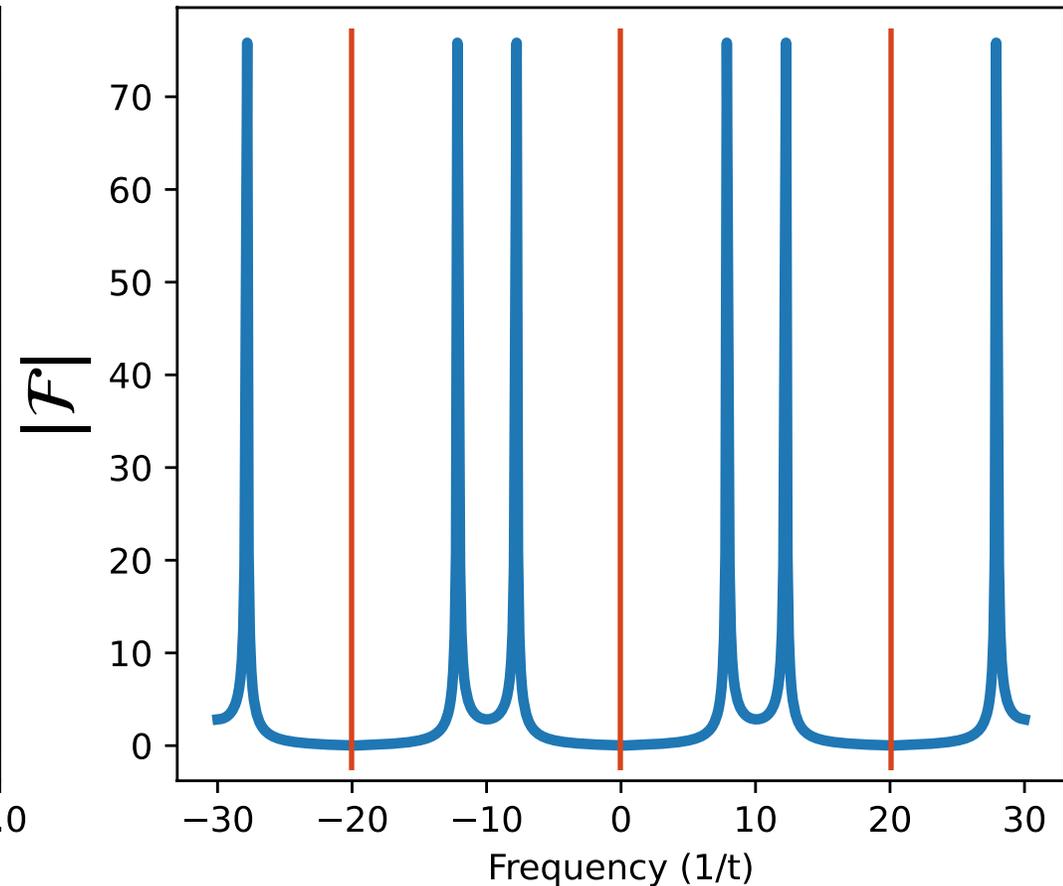
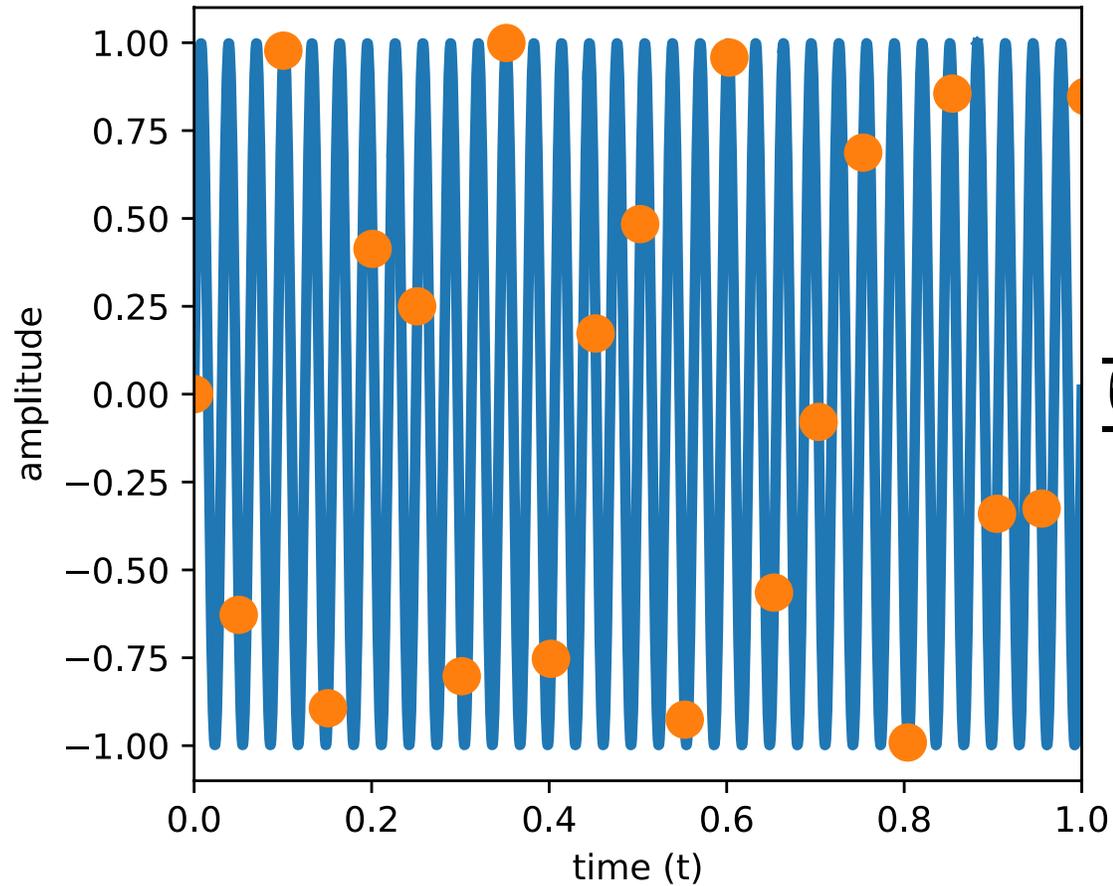
Sample frequency: 20 Hz  
Signal: 28 Hz



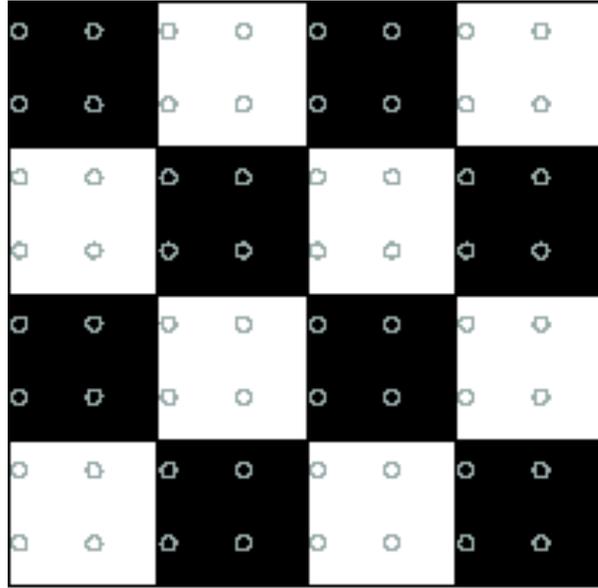
# Sampling exercise

Sample frequency: 20 Hz

Signal: 32 Hz

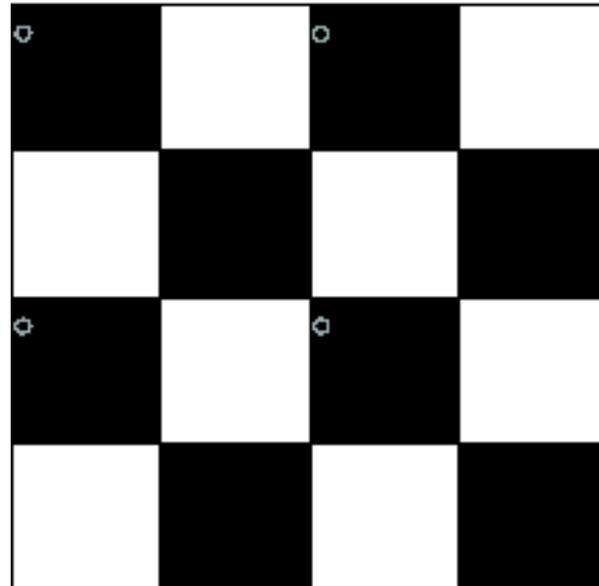
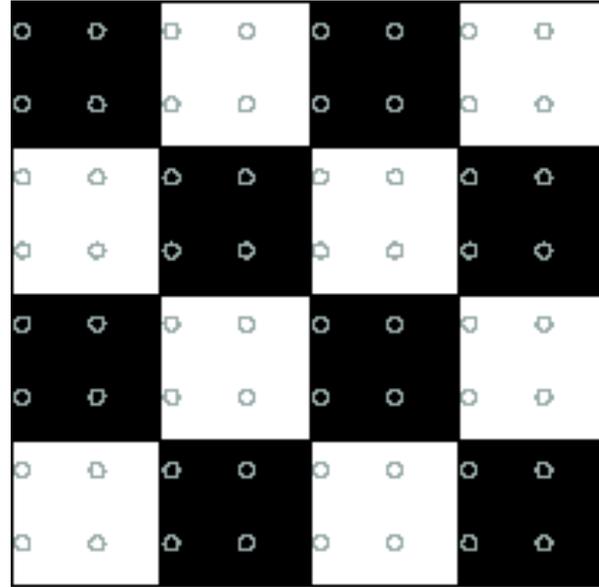


# 2D example



above or  
below  
Nyquist?

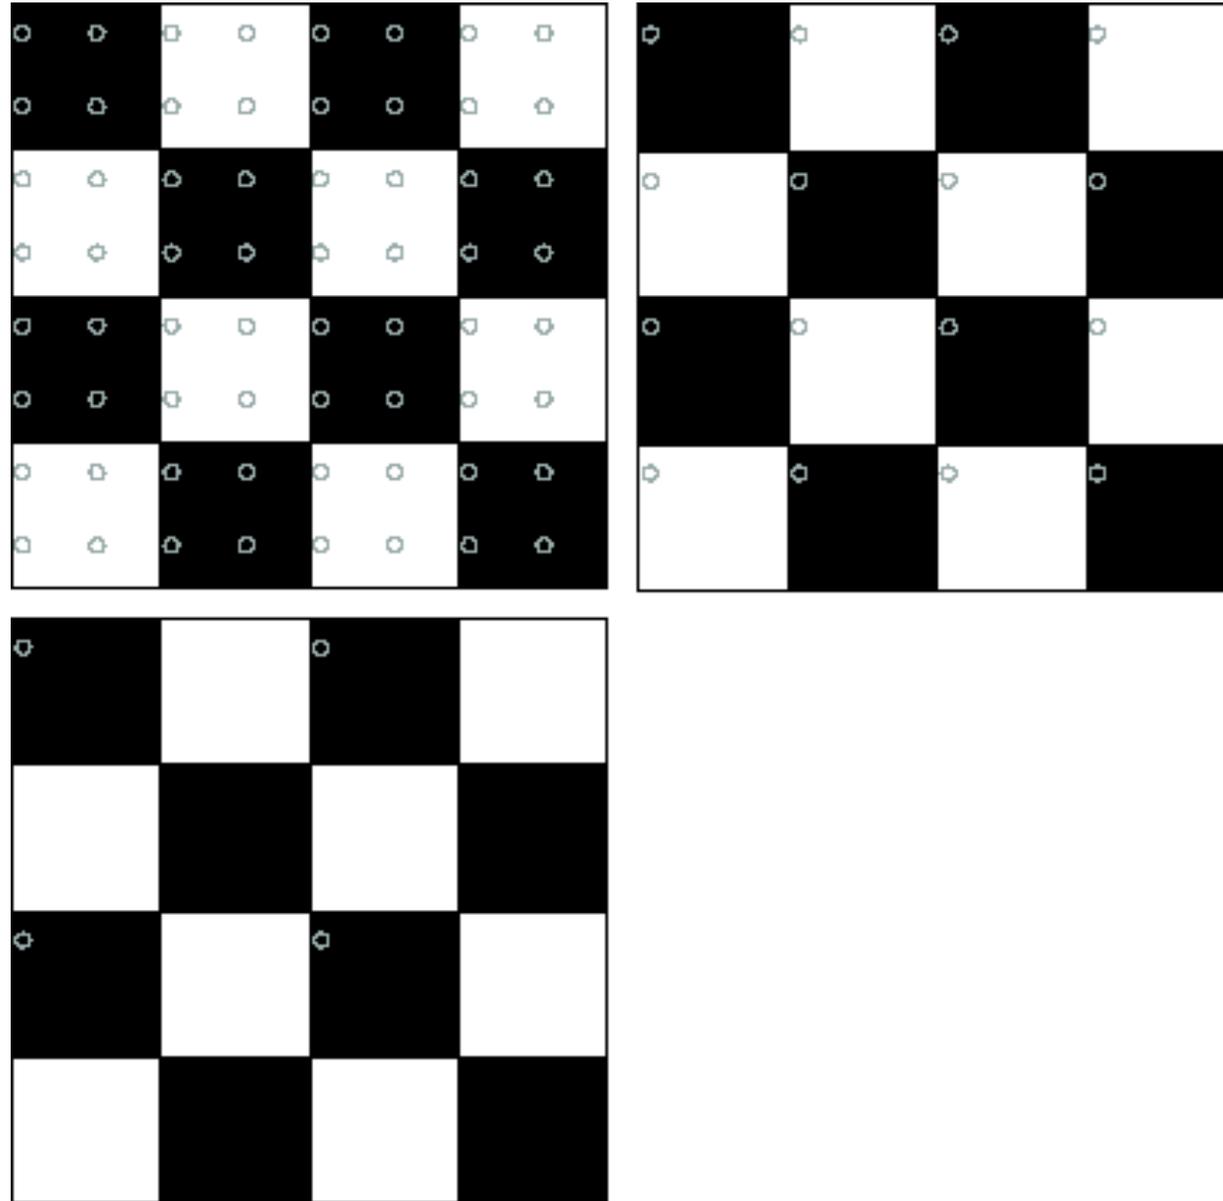
# 2D example



above or  
below  
Nyquist?

[Source: N. Snavely]

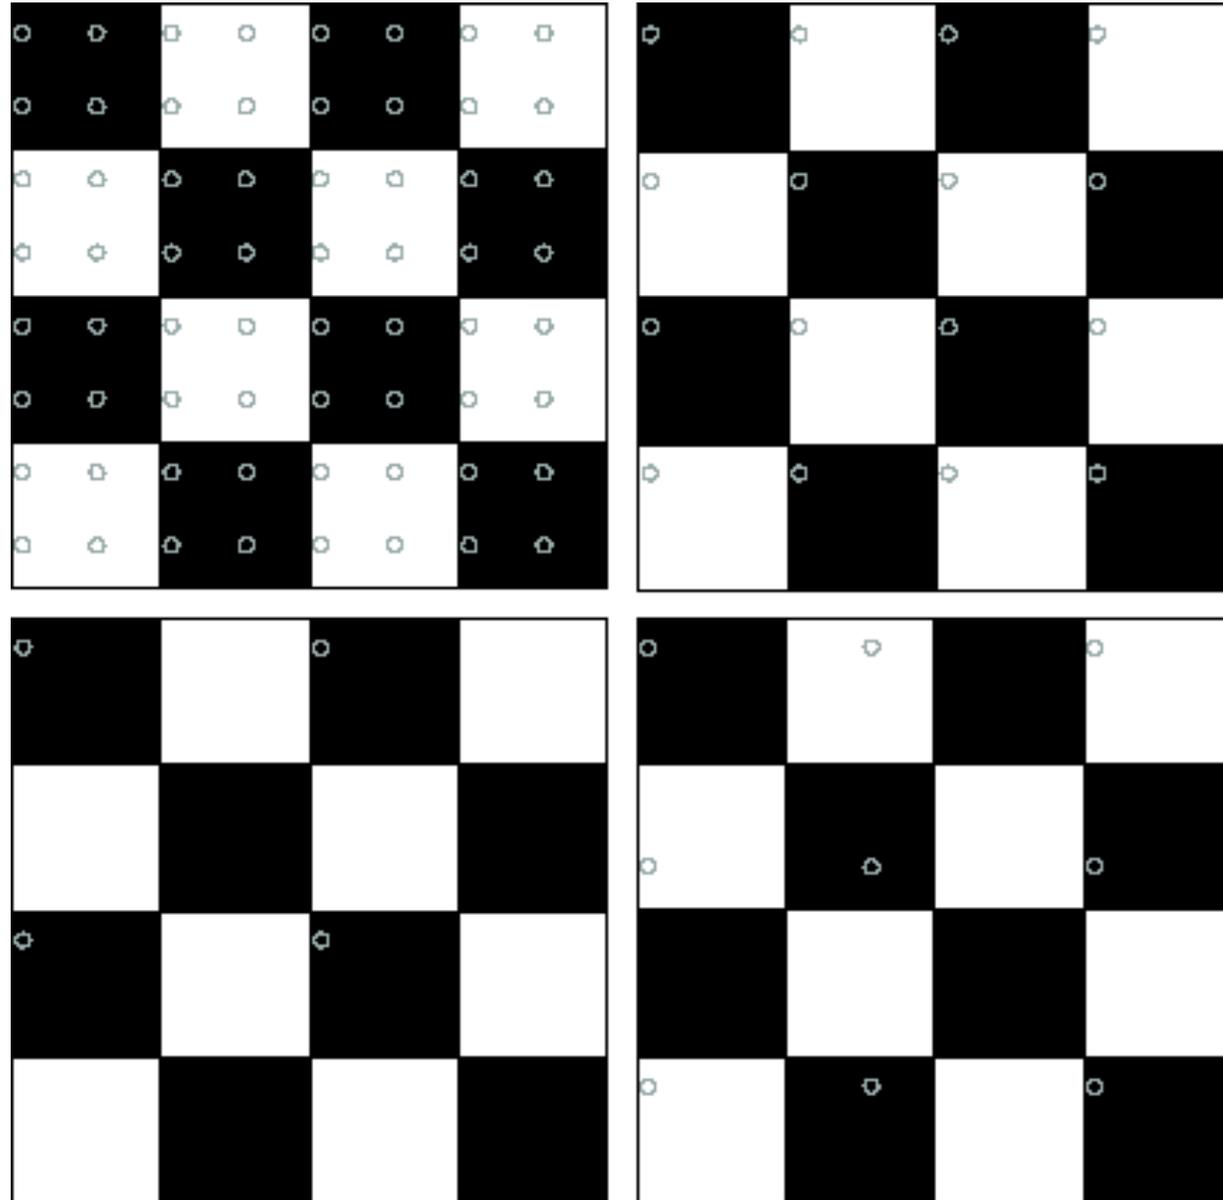
# 2D example



above or  
below  
Nyquist?

[Source: N. Snavely]

# 2D example



above or  
below  
Nyquist?

[Source: N. Snavely]

# Going back to Downsampling ...

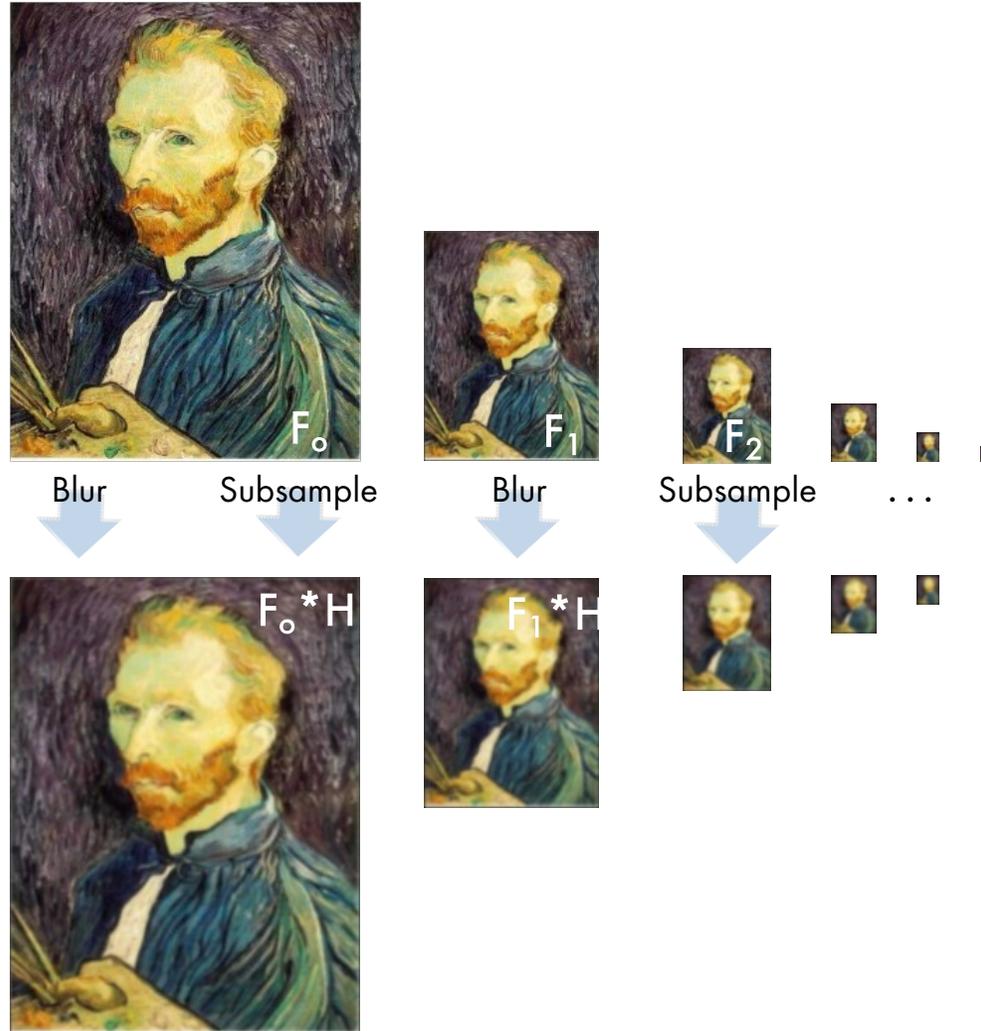
- When downsampling by a factor of two, the original image has frequencies that are too high
- High frequencies are caused by sharp edges
- How can we fix this?

# 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?

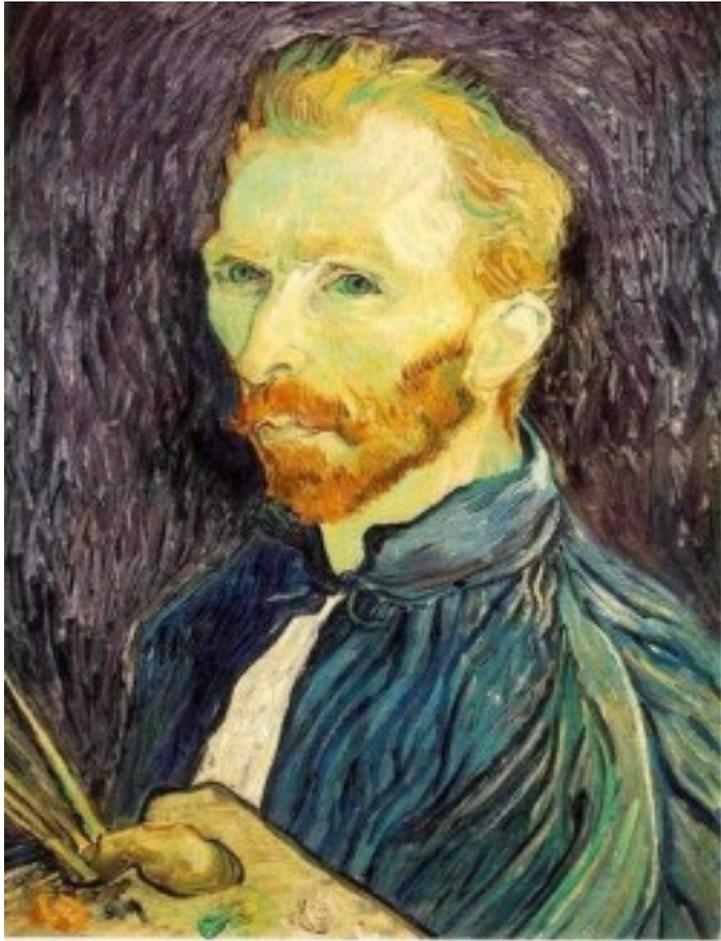
# Gaussian pre-filtering

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

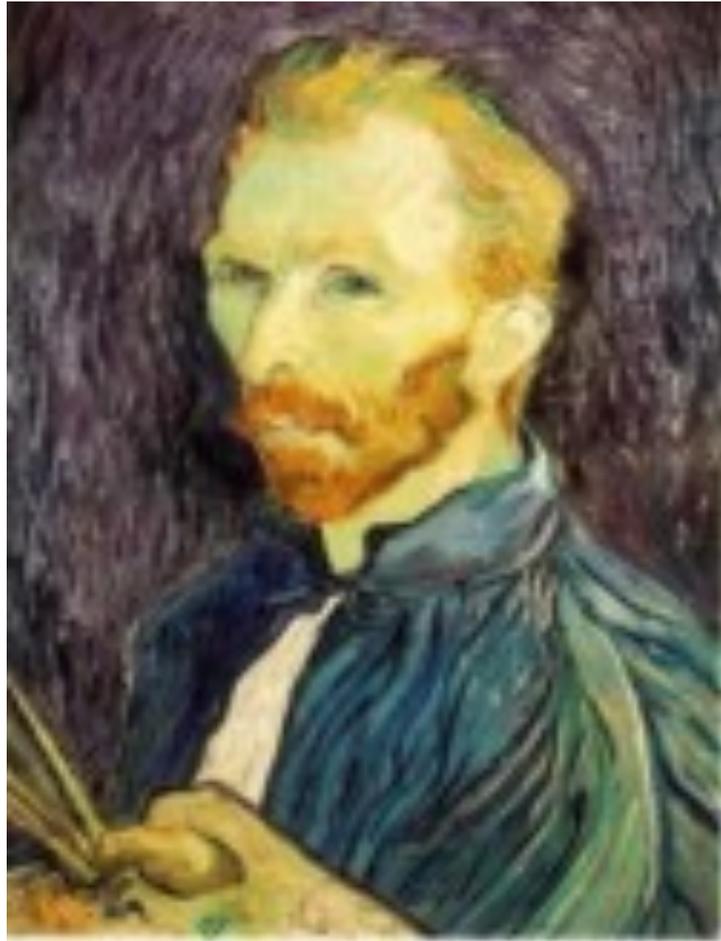


[Source: N. Snavely]

# Subsampling with Gaussian pre-filtering



Gaussian  $1/2$



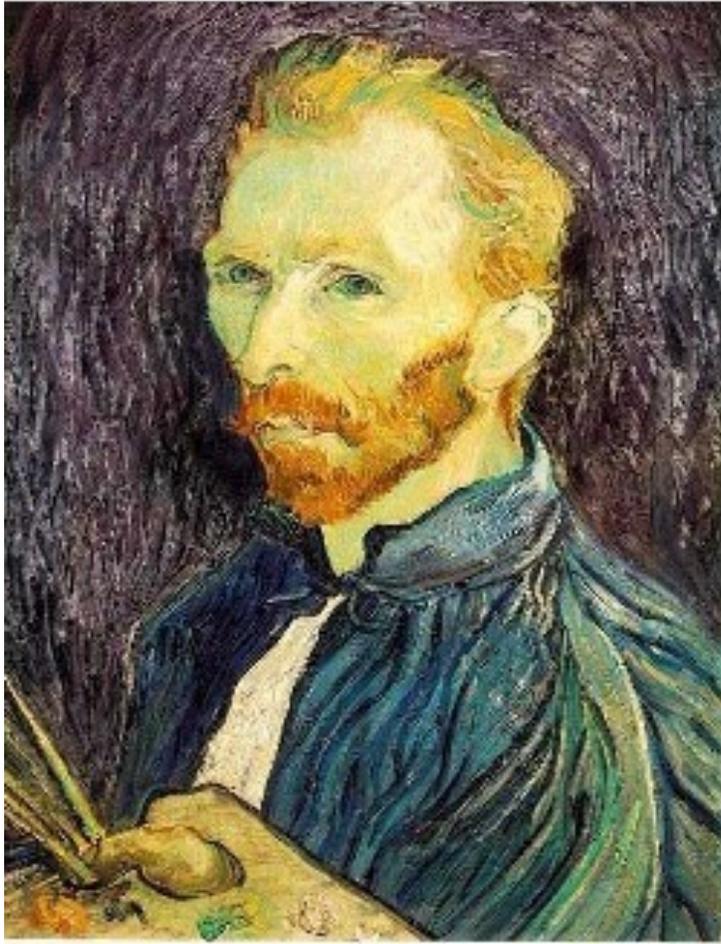
G  $1/4$



G  $1/8$

[Source: S. Seitz]

# Subsampling with Gaussian pre-filtering



$1/2$



$1/4$  (2x Zoom)



$1/8$  (4x Zoom)

[Source: S. Seitz]

# Where is the Rectangle?

- My image

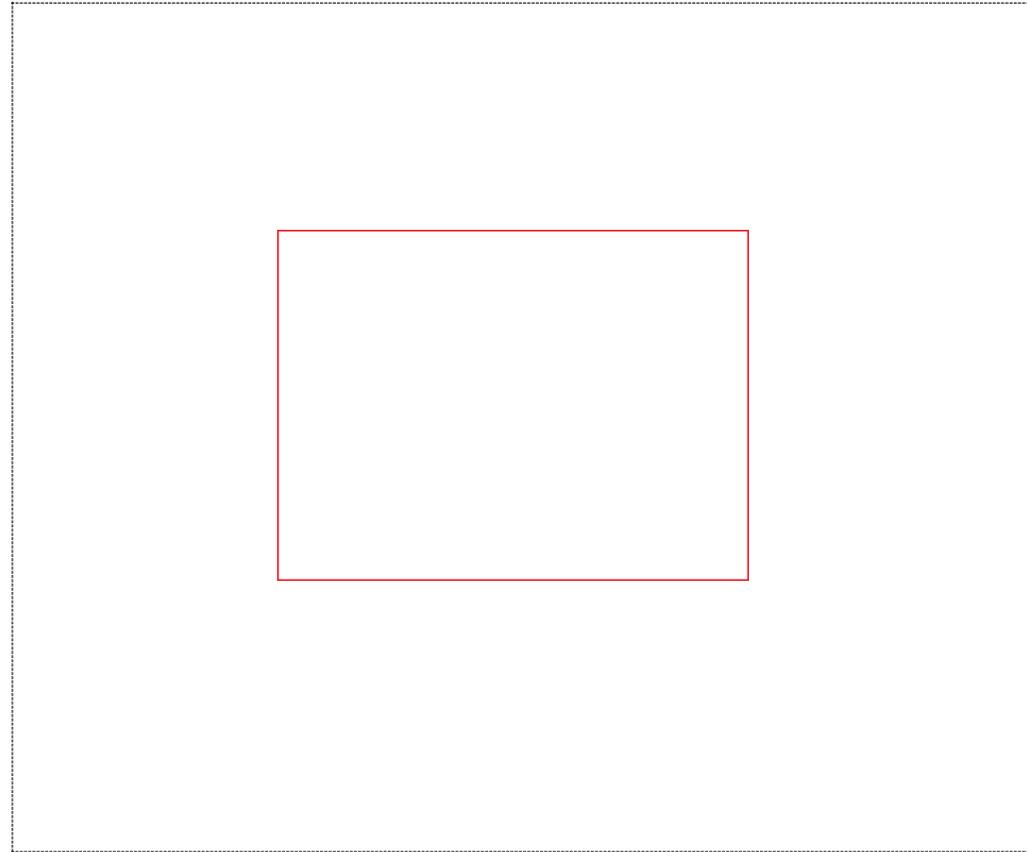


Figure: Dashed line denotes the border of the image (it's not part of the image)

# Where is the Rectangle?

- My image
- Let's blur

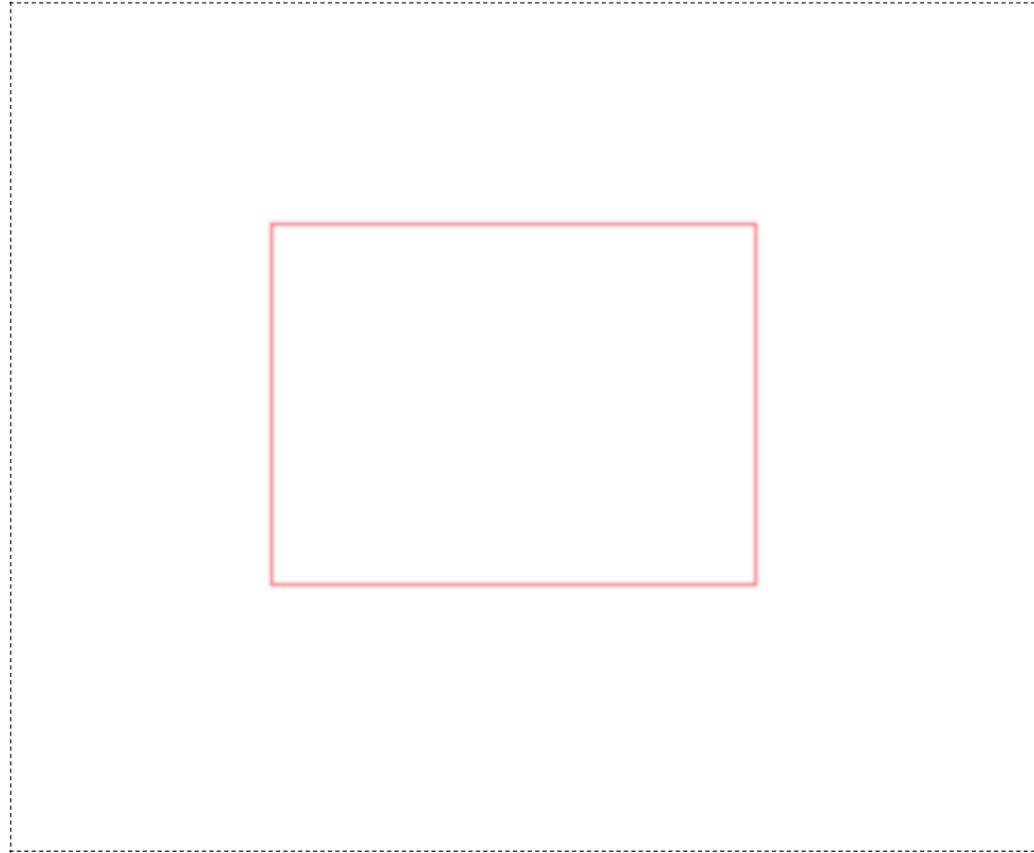
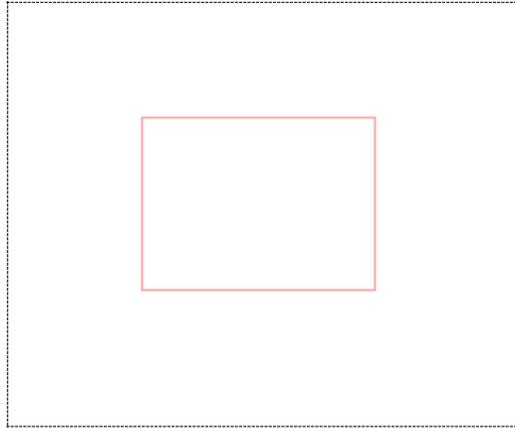


Figure: Dashed line denotes the border of the image (it's not part of the image)

# 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

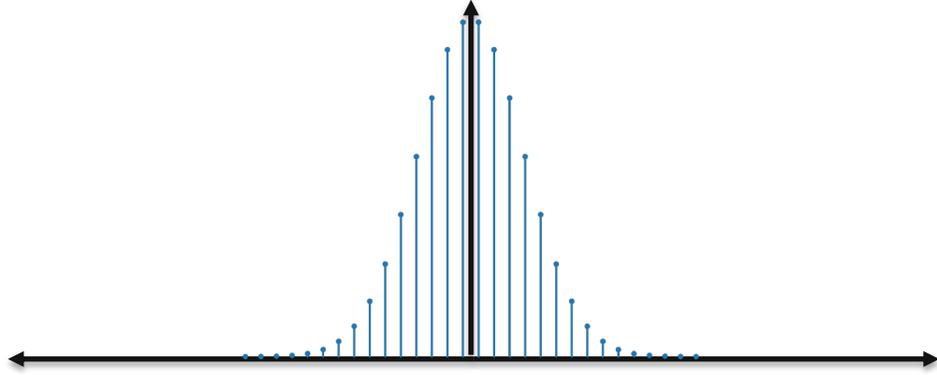


# Why does this work?

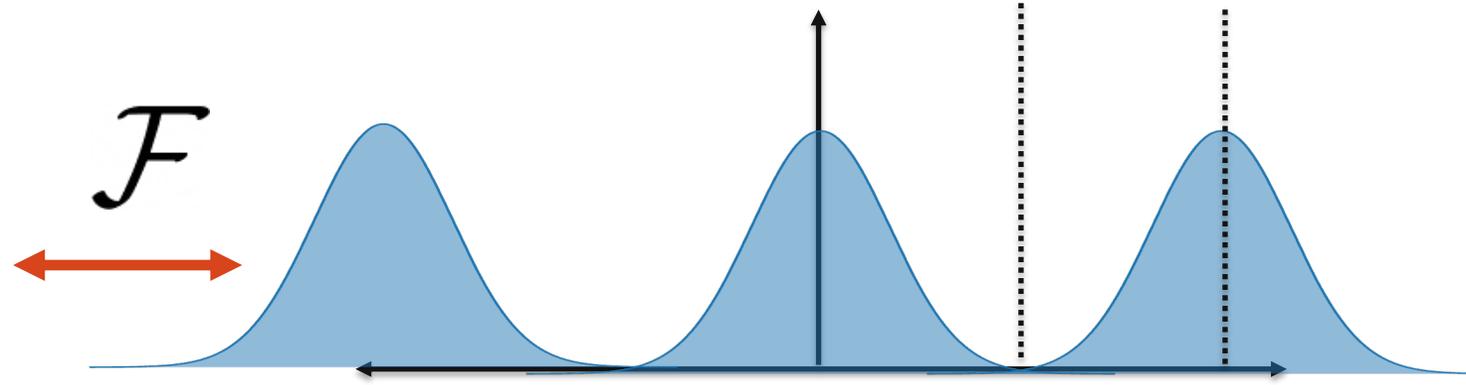
- What does blurring do in the frequency domain?
- How does that fix the aliasing problem?

# Sampling

Primal Domain



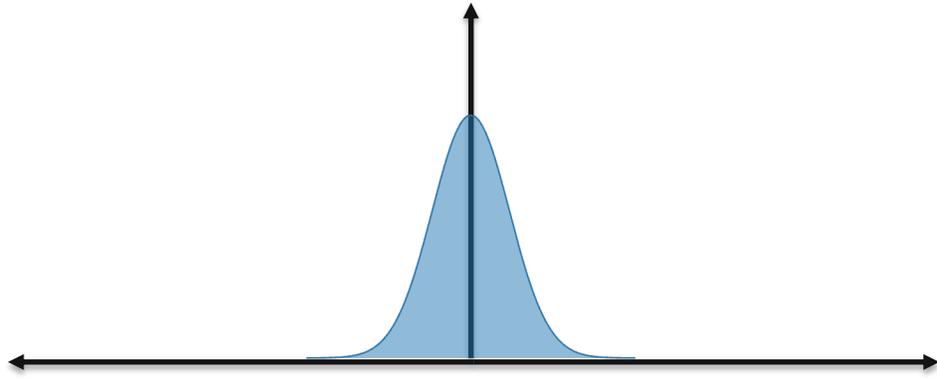
Fourier Domain



What happens if we subsample in the primal domain?

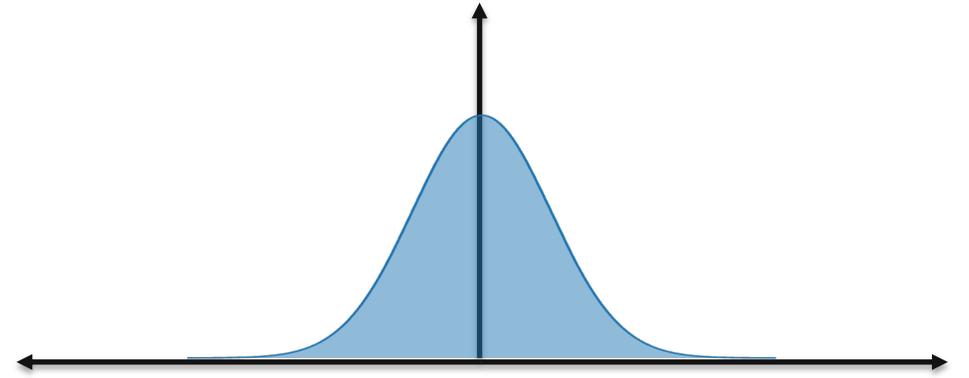
# Sampling

Primal Domain



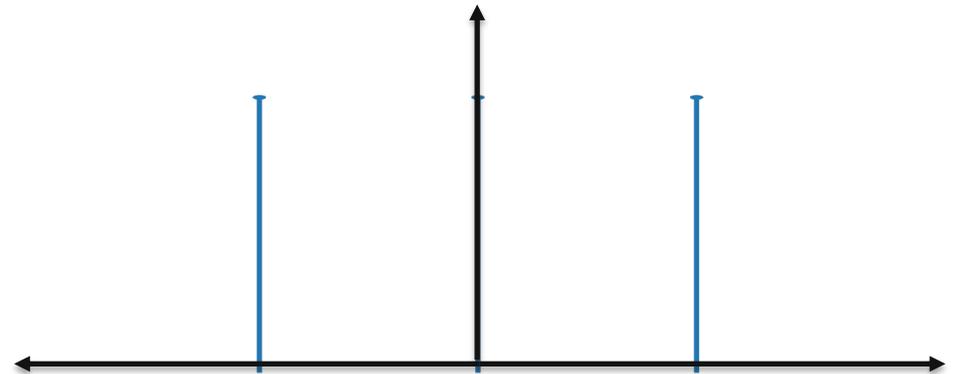
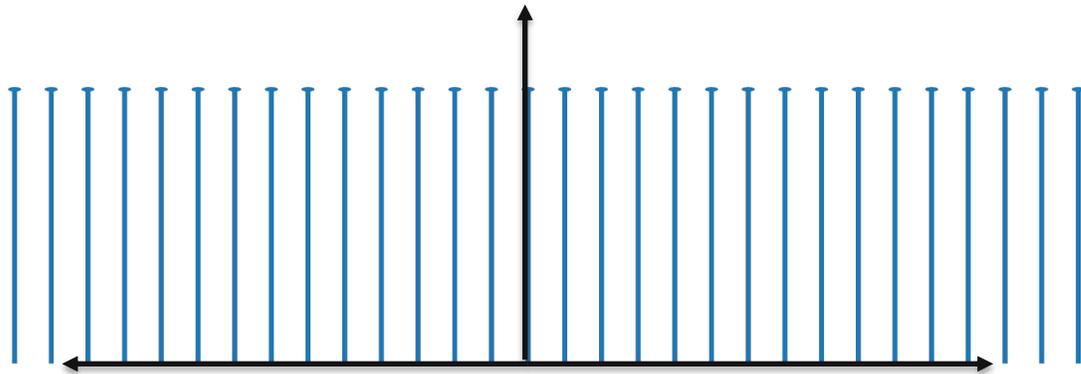
$\mathcal{F}$

Fourier Domain



 Sampling operator

$*$

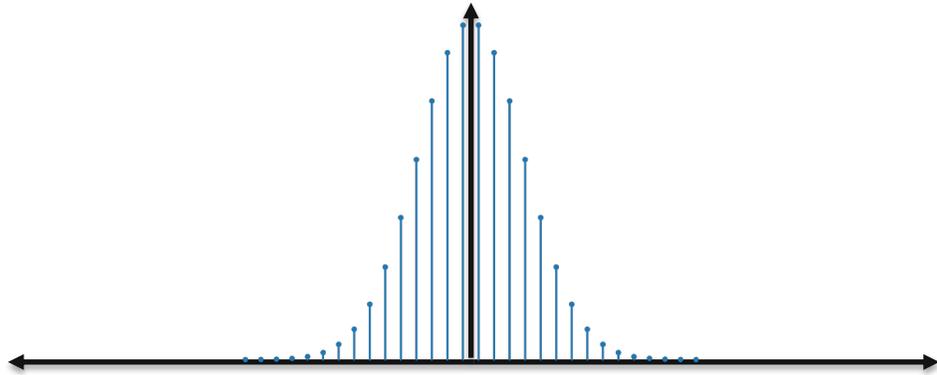


Sample rate of  $f_s$

Shifted copies at  $f_s$

# Sampling

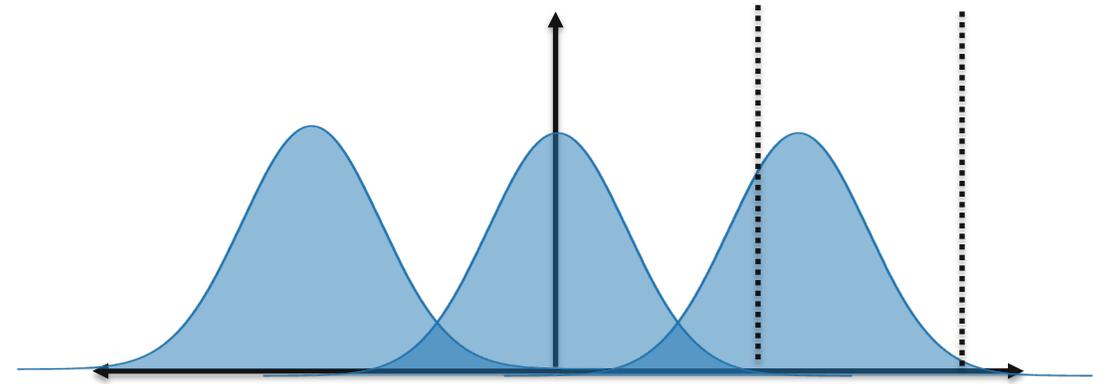
Primal Domain



$\mathcal{F}$



Fourier Domain

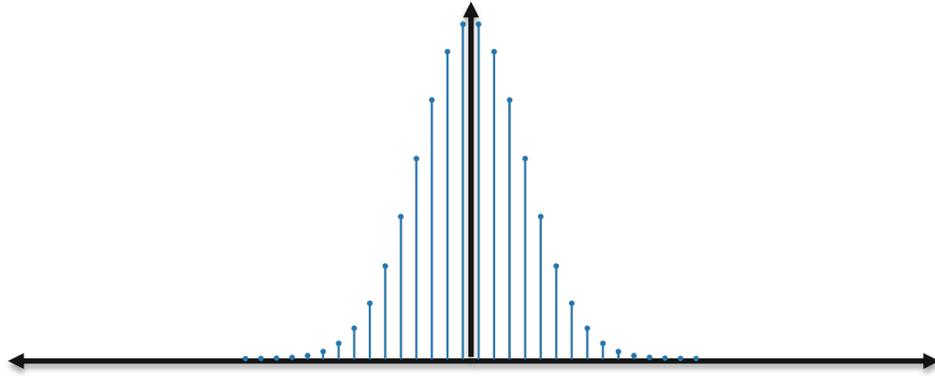


What happens if we subsample in the primal domain?

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

# Sampling

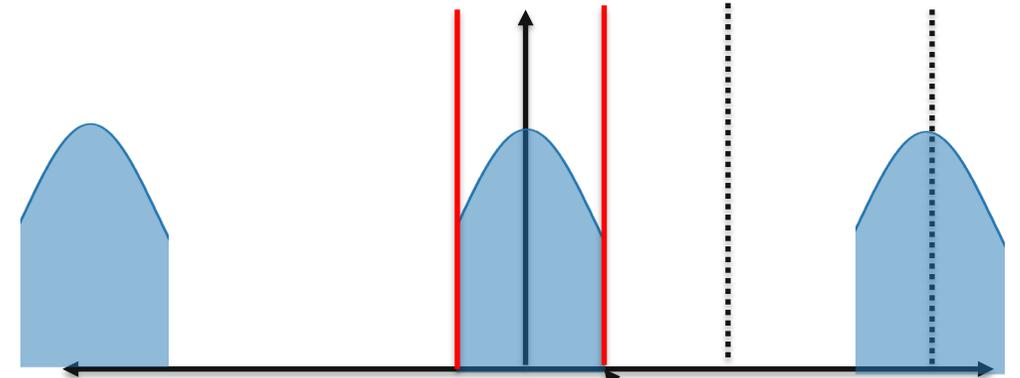
Primal Domain



$\mathcal{F}$



Fourier Domain



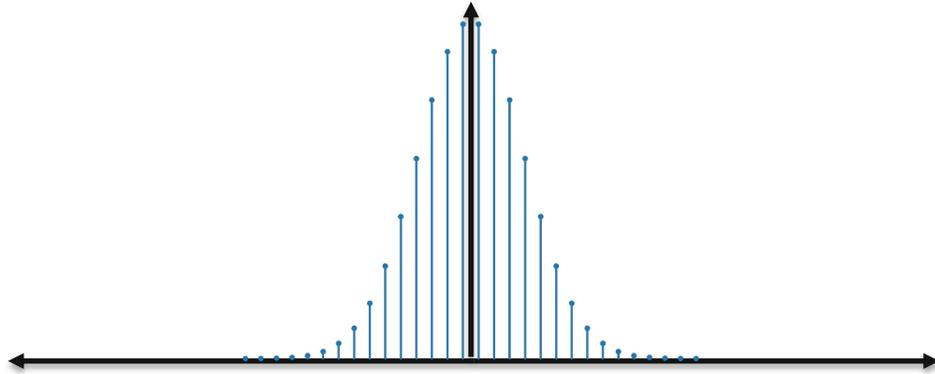
Filter cutoff frequency  
(what determines this?)

What happens if we subsample in the primal domain?

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

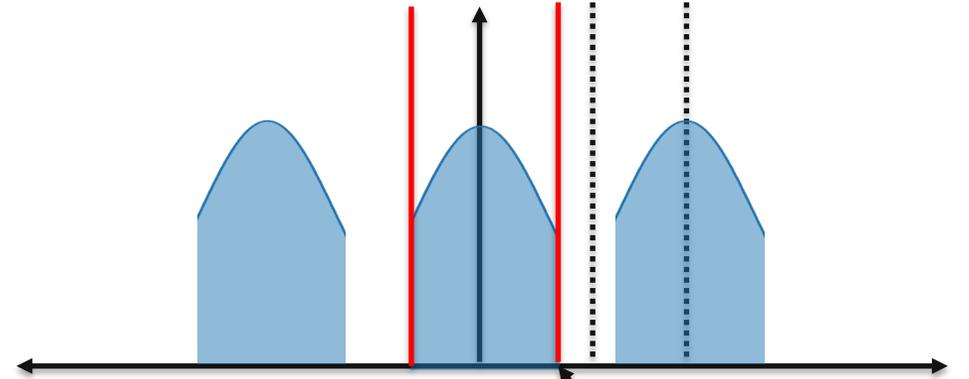
# Sampling

Primal Domain



$\mathcal{F}$

Fourier Domain



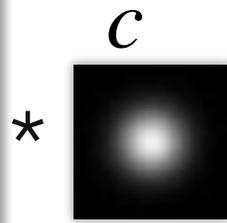
Filter cutoff frequency  
(what determines this?)

What happens if we subsample in the primal domain?

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

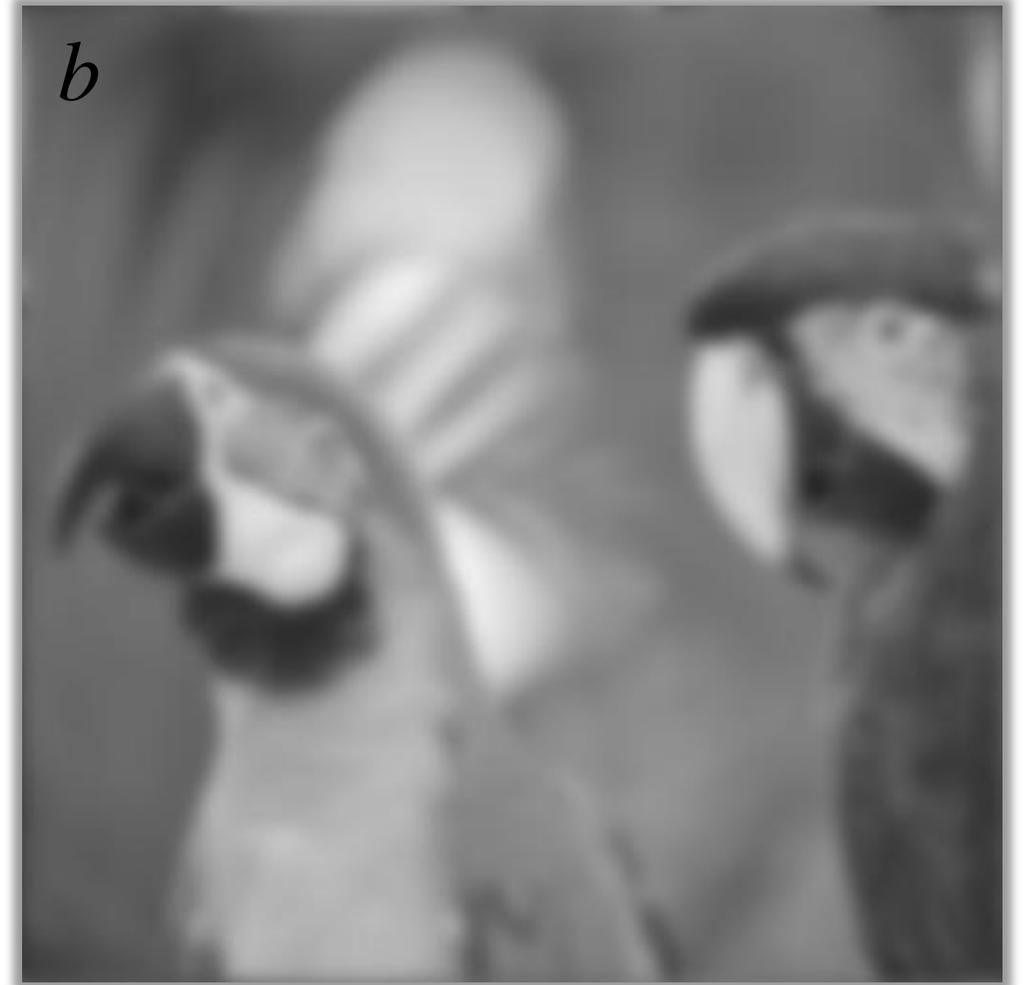
# Filtering – Low-pass Filter

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



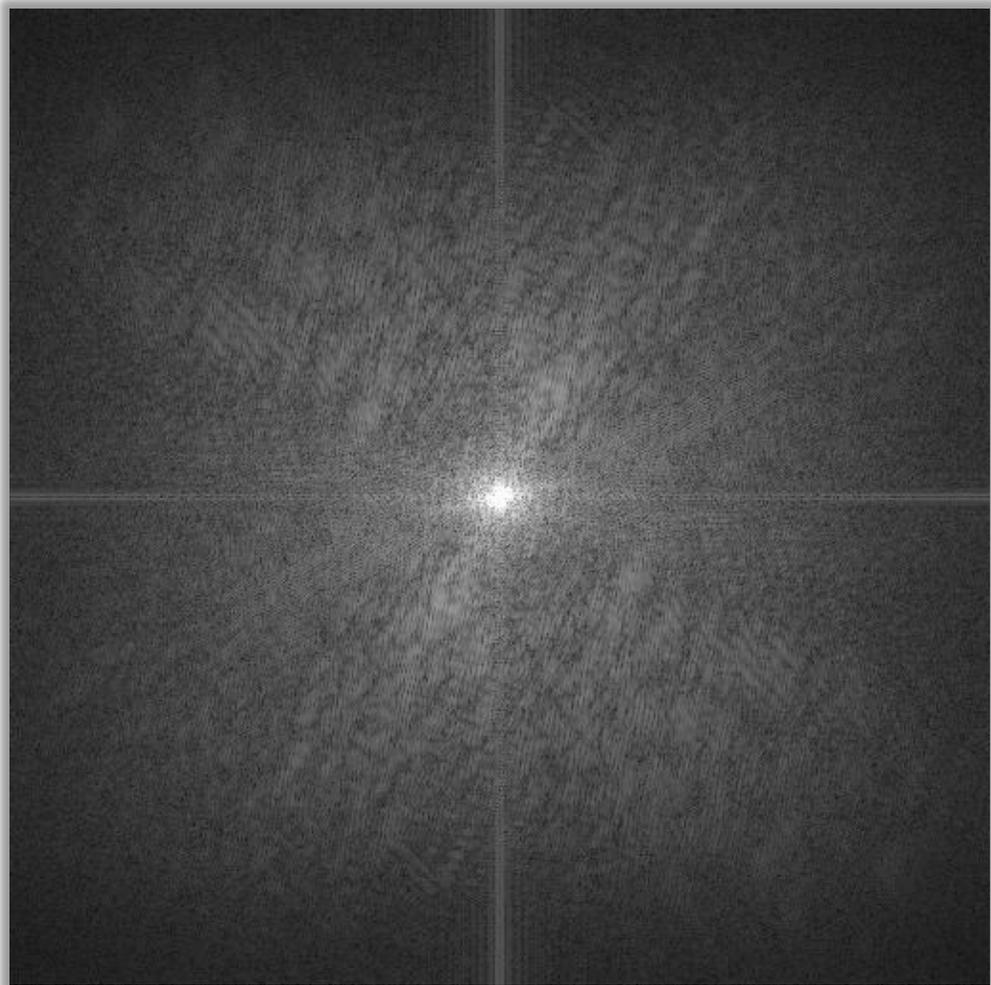
\*

=

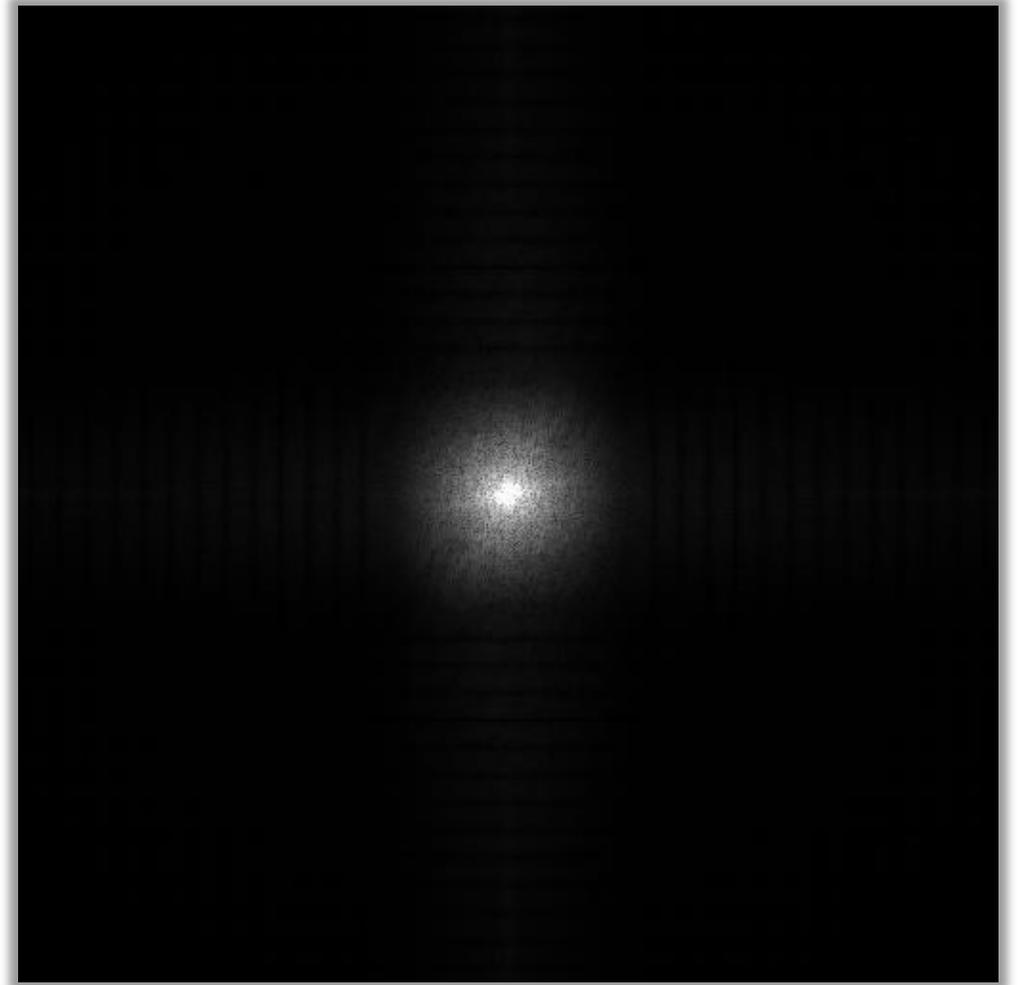


# Filtering – Low-pass Filter

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

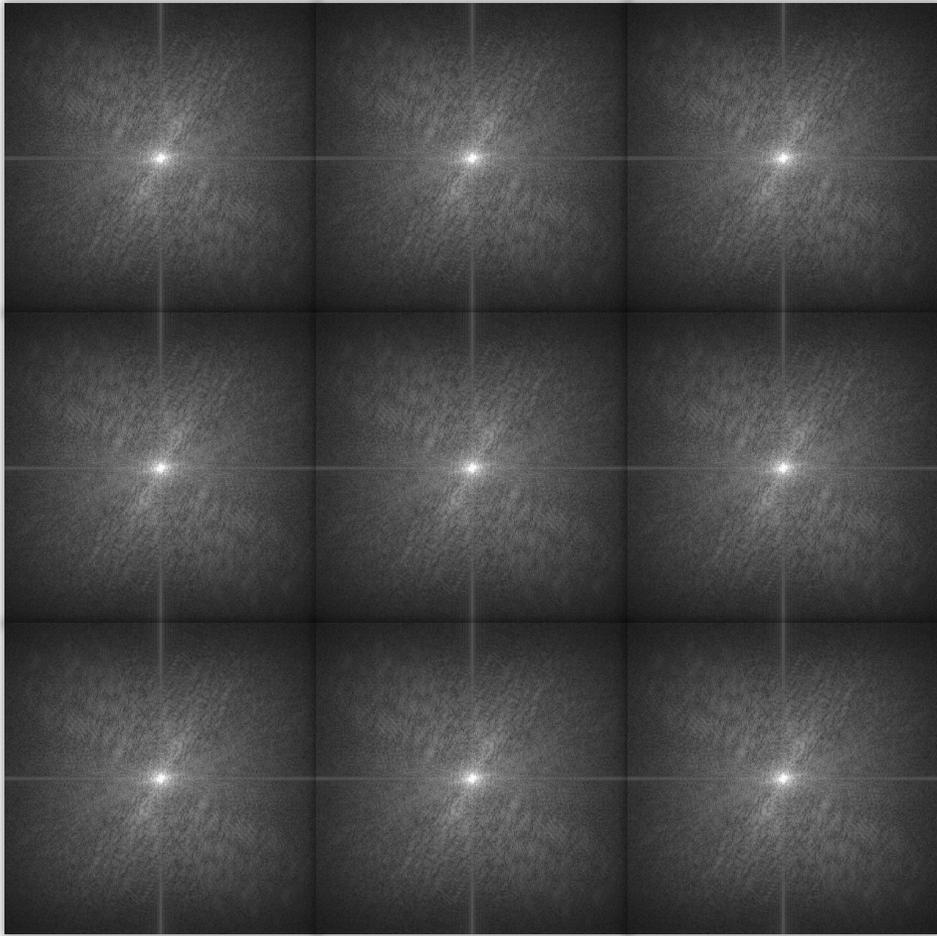


=



# What's the picture in 2D?

Periodic spectral copies in the frequency domain

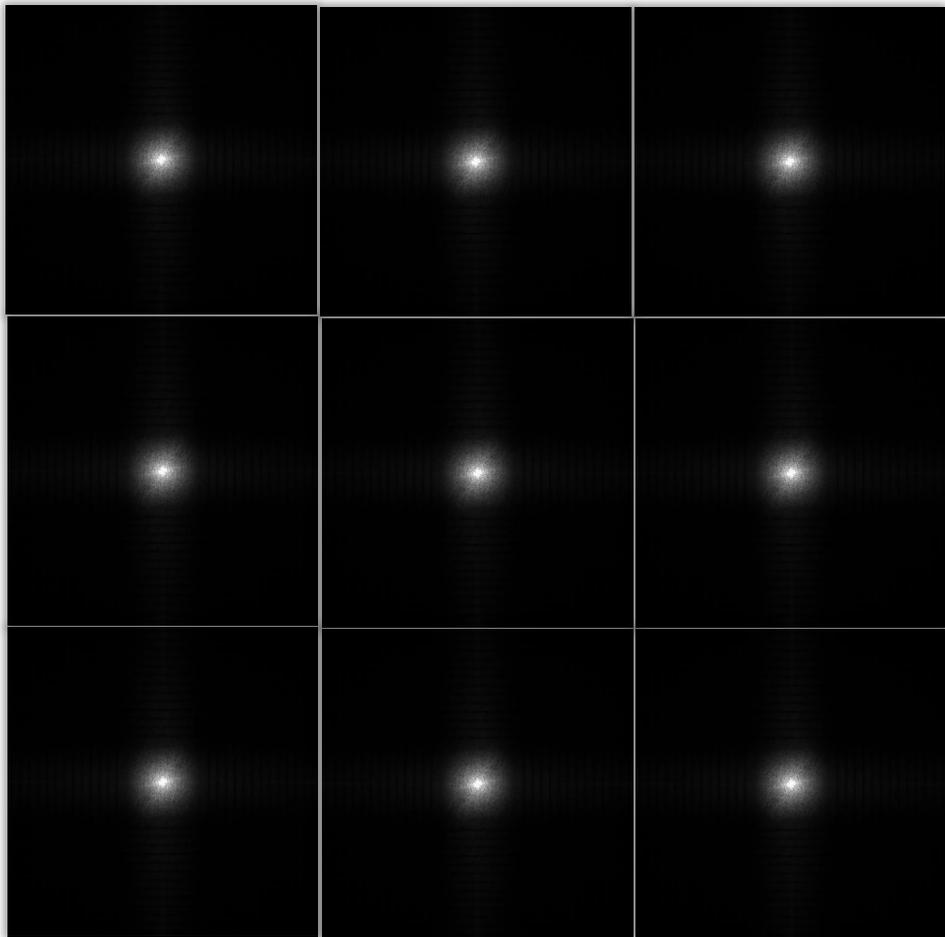


Periodic primal domain signal

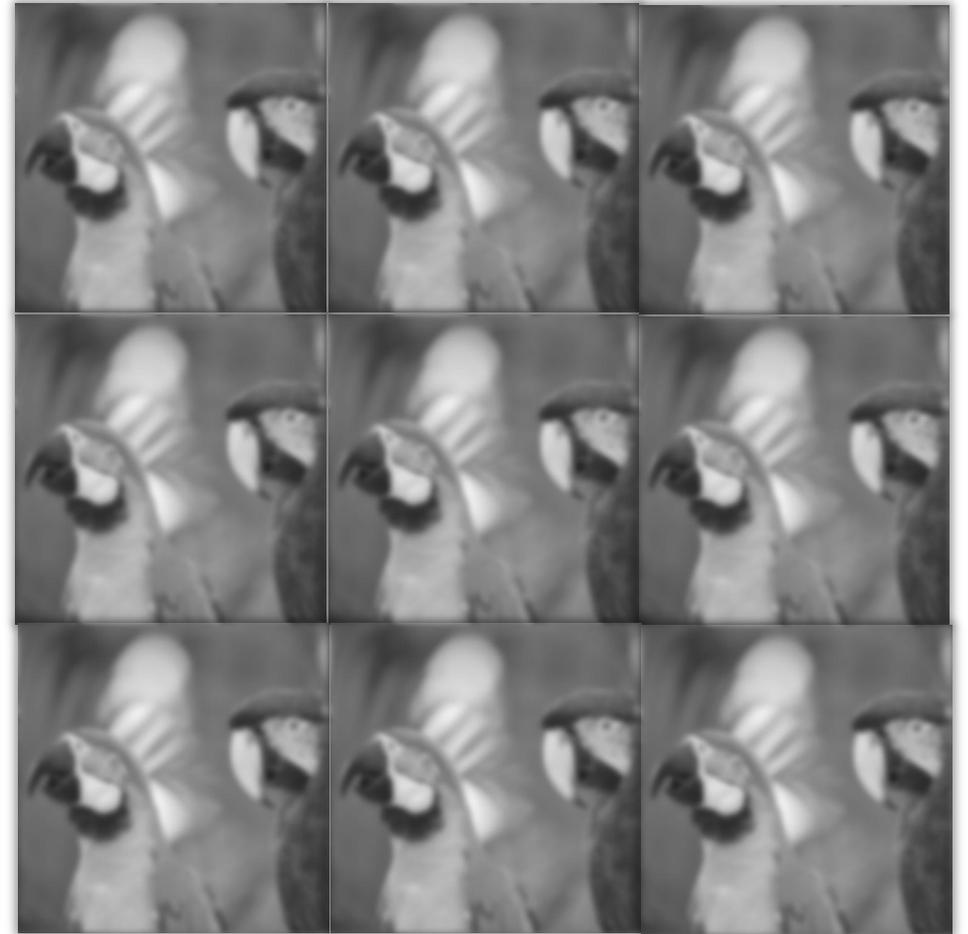


# What's the picture in 2D?

Periodic spectral copies in the frequency domain



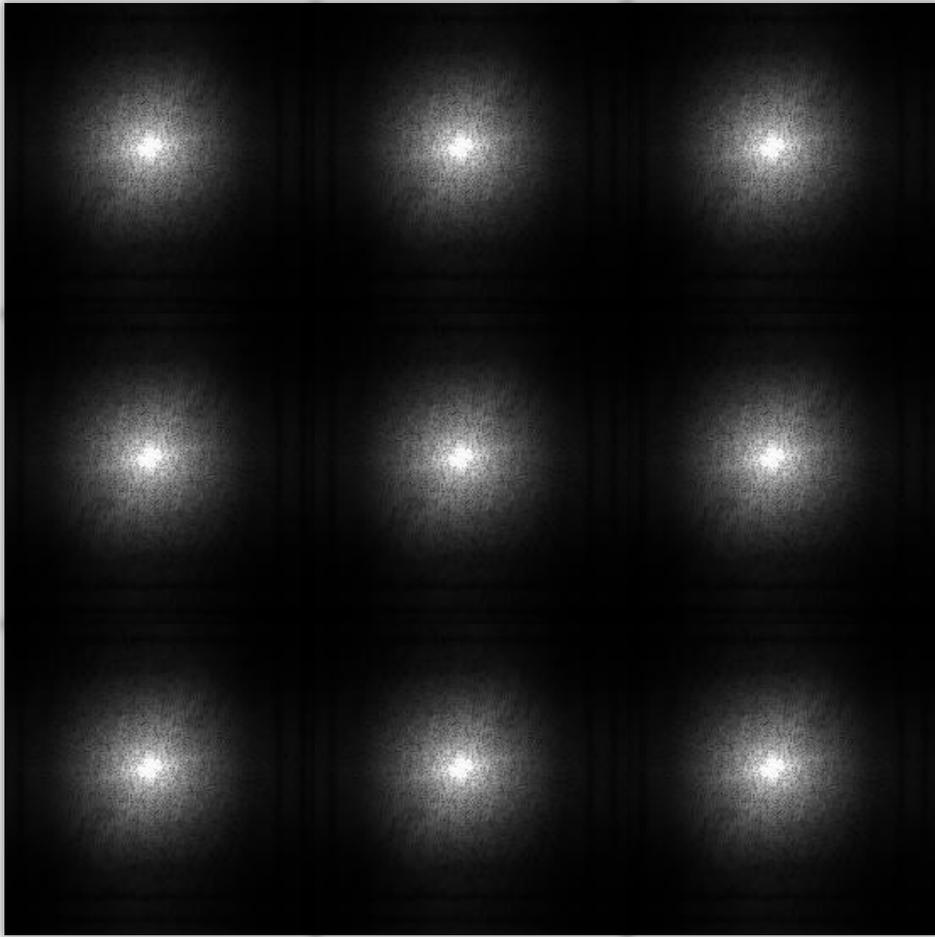
Periodic primal domain signal



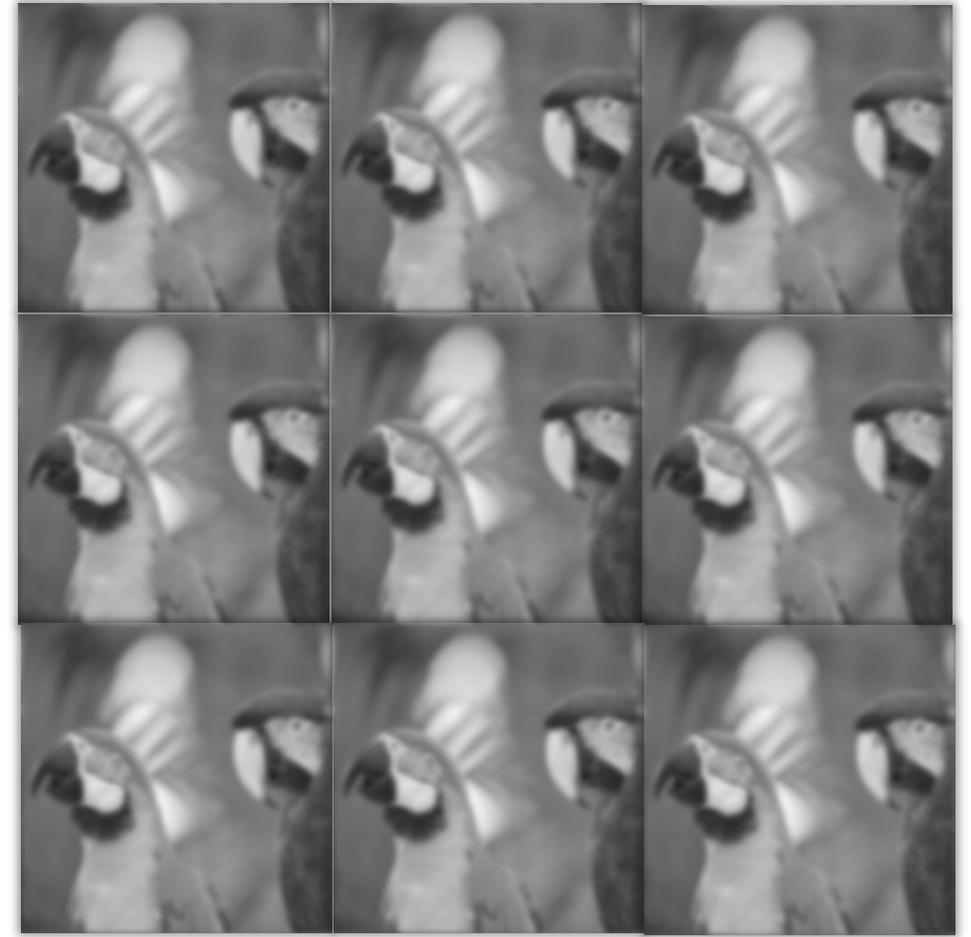
After low-pass filtering

# What's the picture in 2D?

Periodic spectral copies in the frequency domain



Periodic primal domain signal

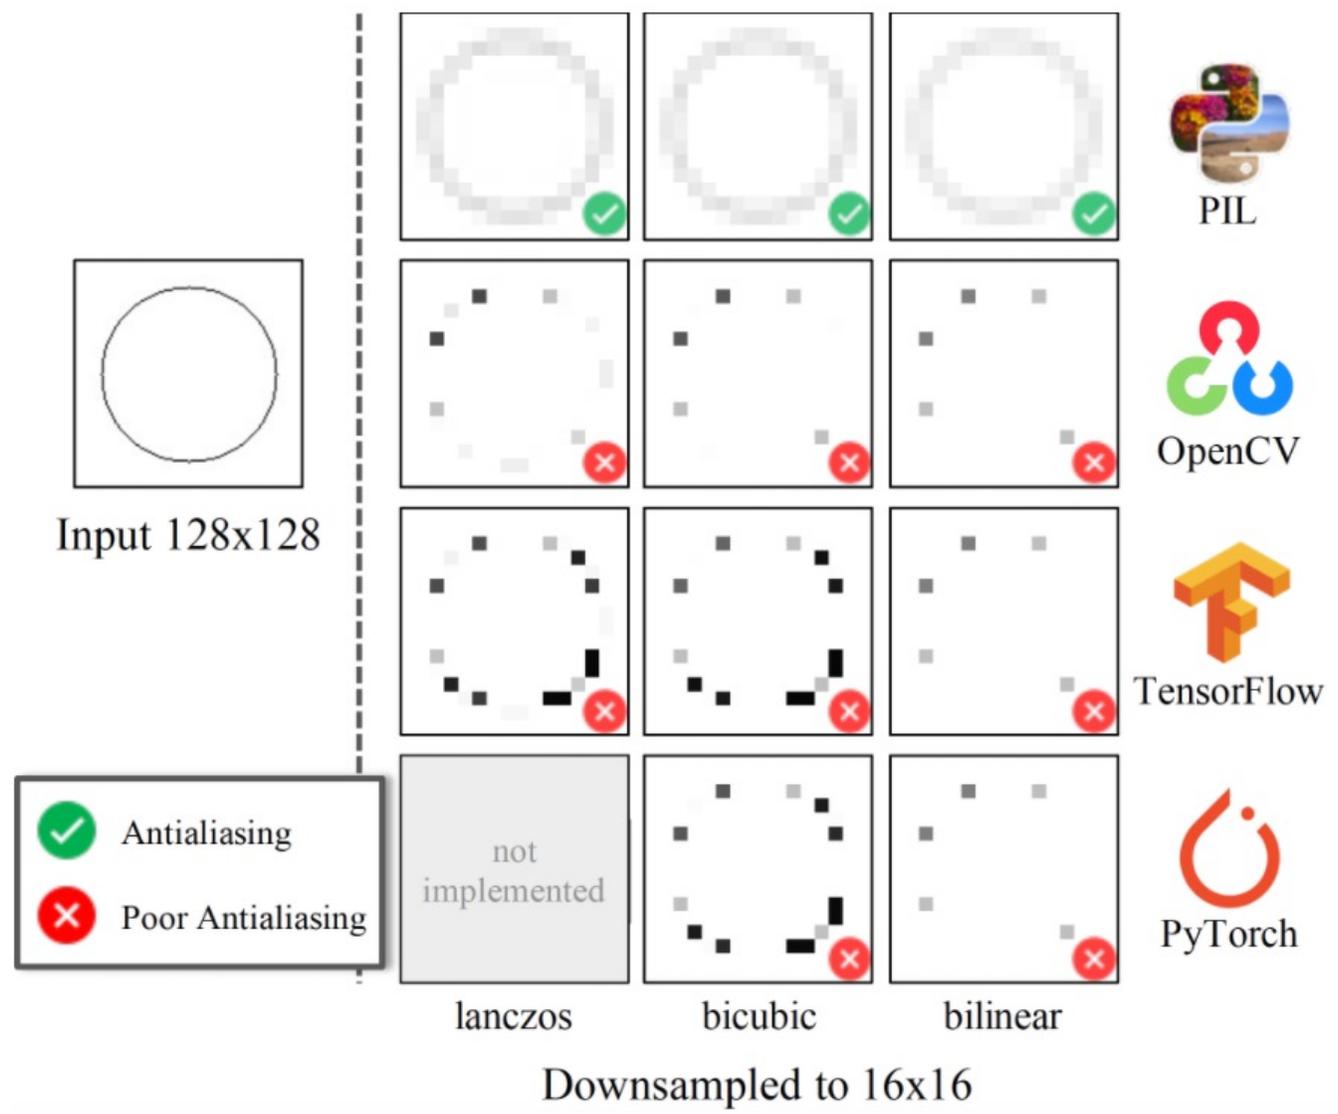


After sub-sampling (no aliasing!)

# Image Downsampling (& Upsampling)

- “anti-aliasing” → **before** re-sampling, apply appropriate filter!
- how much filtering? Shannon-Nyquist sampling theorem:

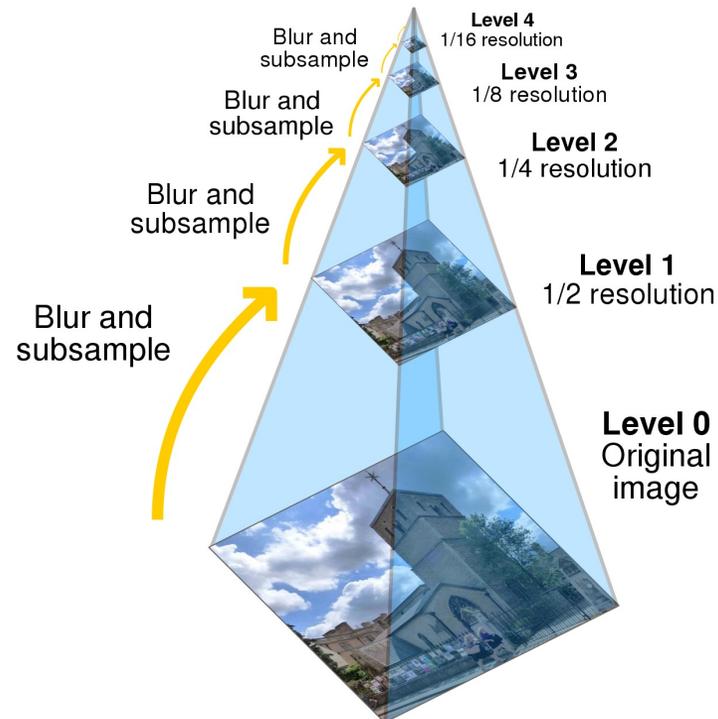
$$f_s \geq 2 f_{\max}$$



# 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]

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  $N=2^k$ )



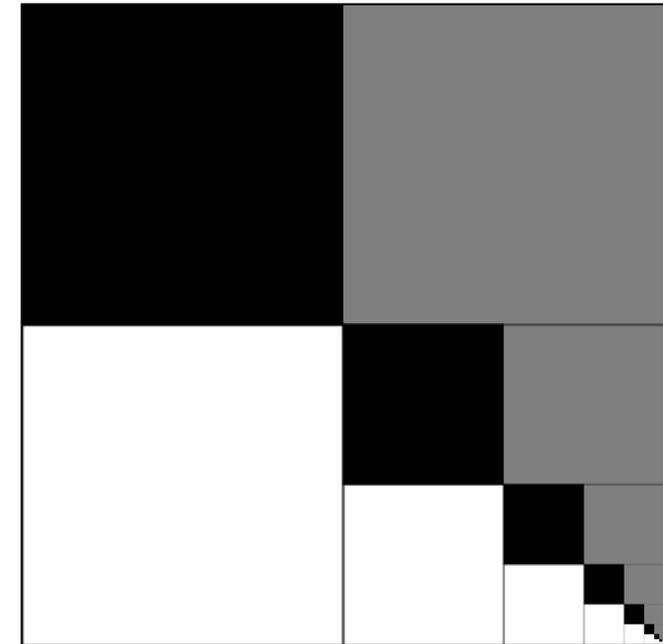
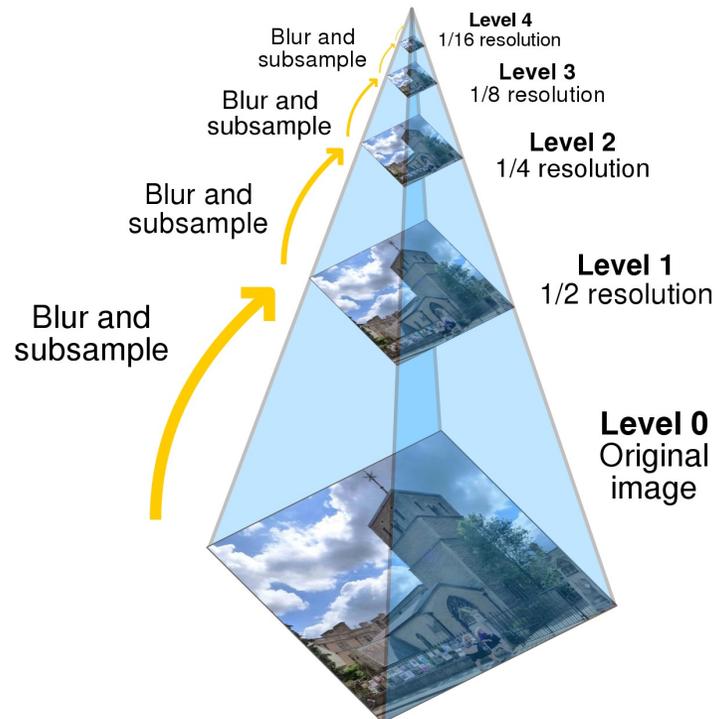
[Source: J. SEITZ]

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

# Gaussian Pyramids [Burt and Adelson, 1983]

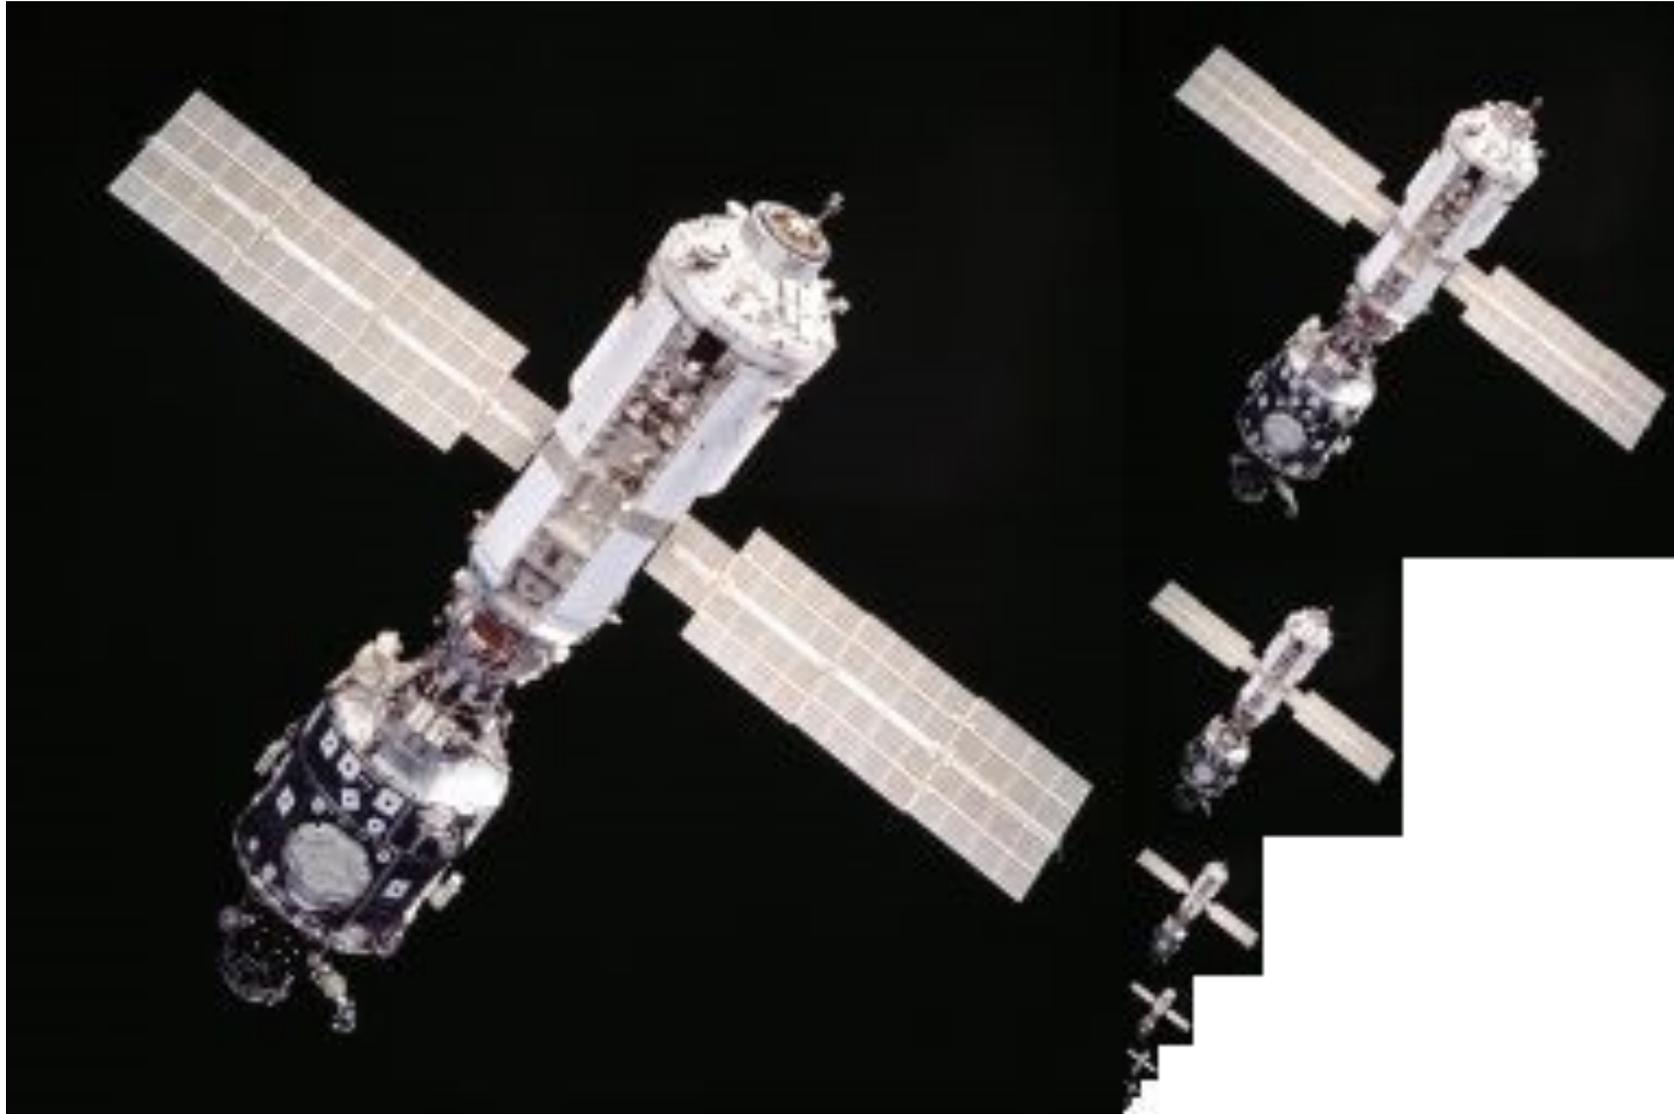
- A sequence of images created with Gaussian blurring and downsampling is called a Gaussian Pyramid
- In computer graphics, a mip map [Williams, 1983]

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  $N = 2^k$ )



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?



# Image Up-Sampling

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



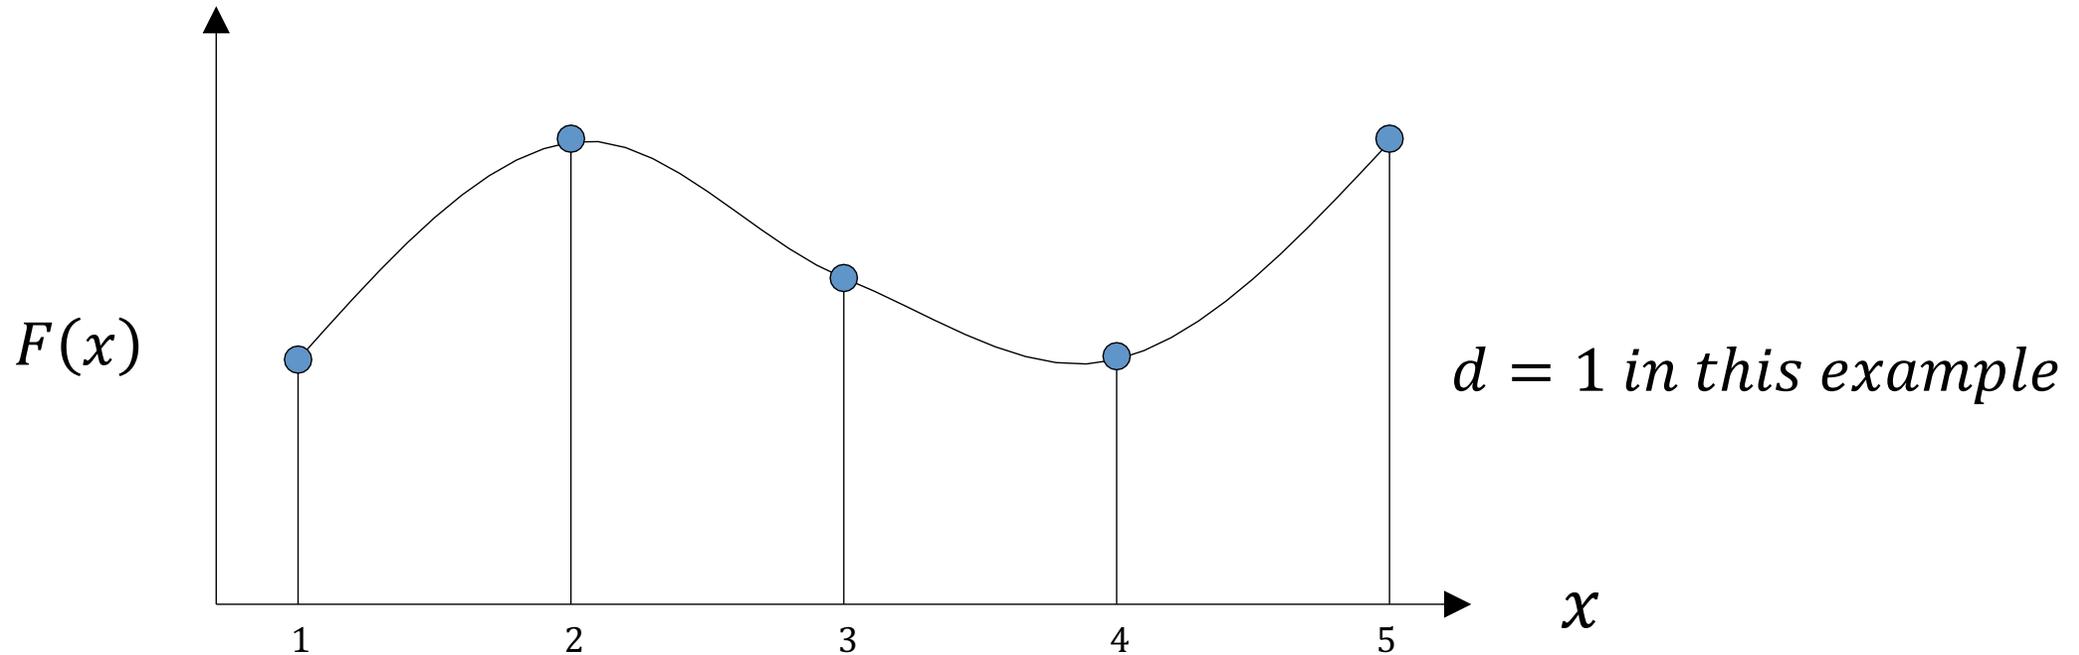
- Simplest approach: repeat each row and column 10 times



[Source: N. Snavely, R. Urtasun]

# Interpolation

# Interpolation

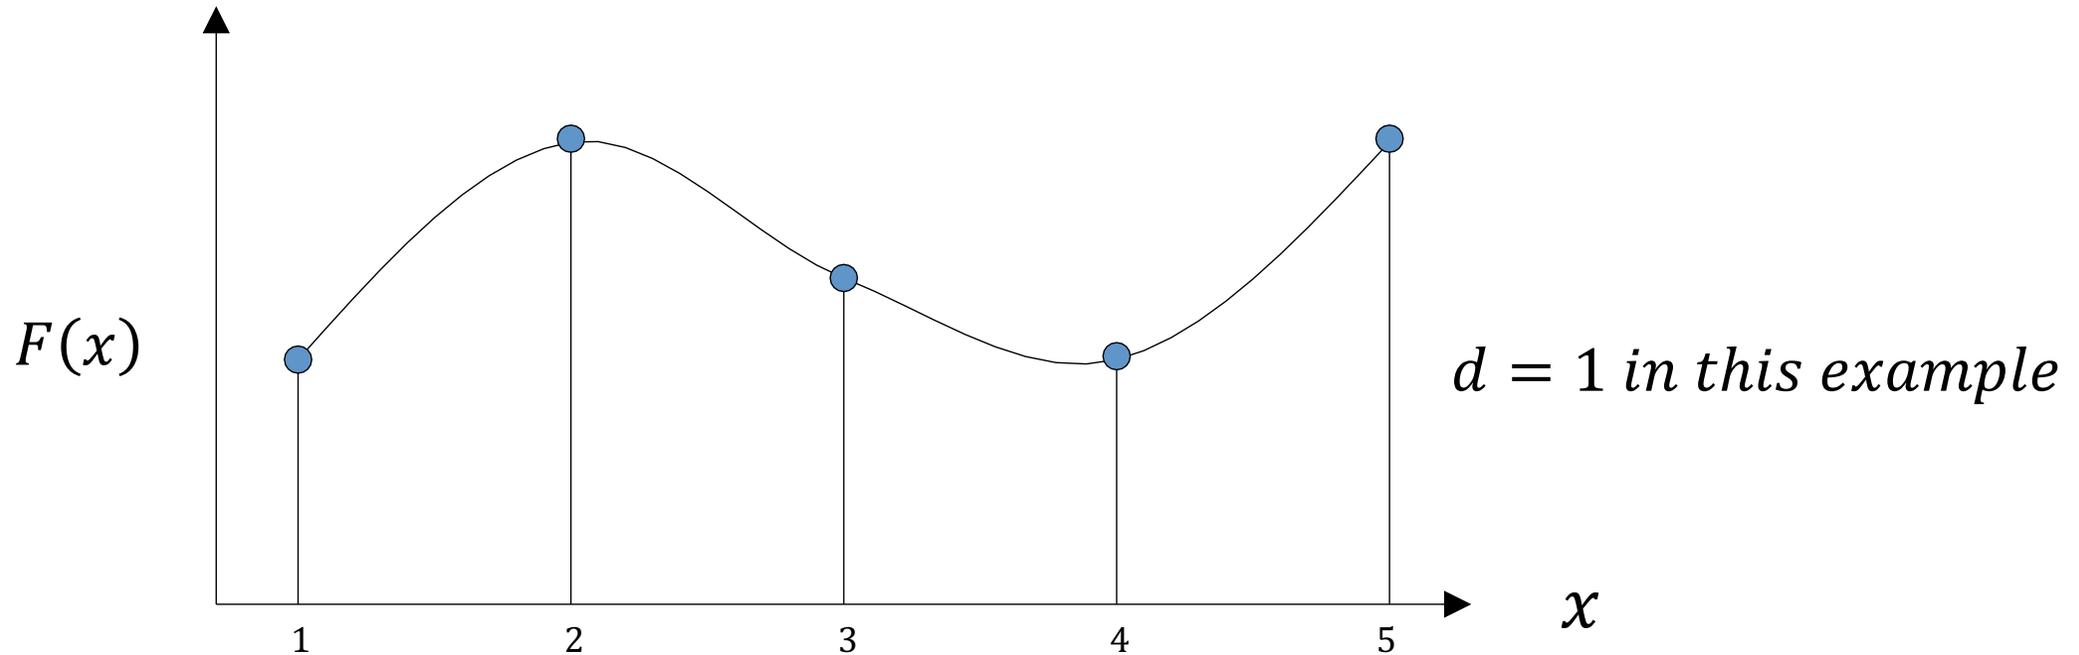


- Recall how digital image is formed

$$F[x, y] = \text{quantize} \left\{ f \left( \frac{x}{d}, \frac{y}{d} \right) \right\}$$

- 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

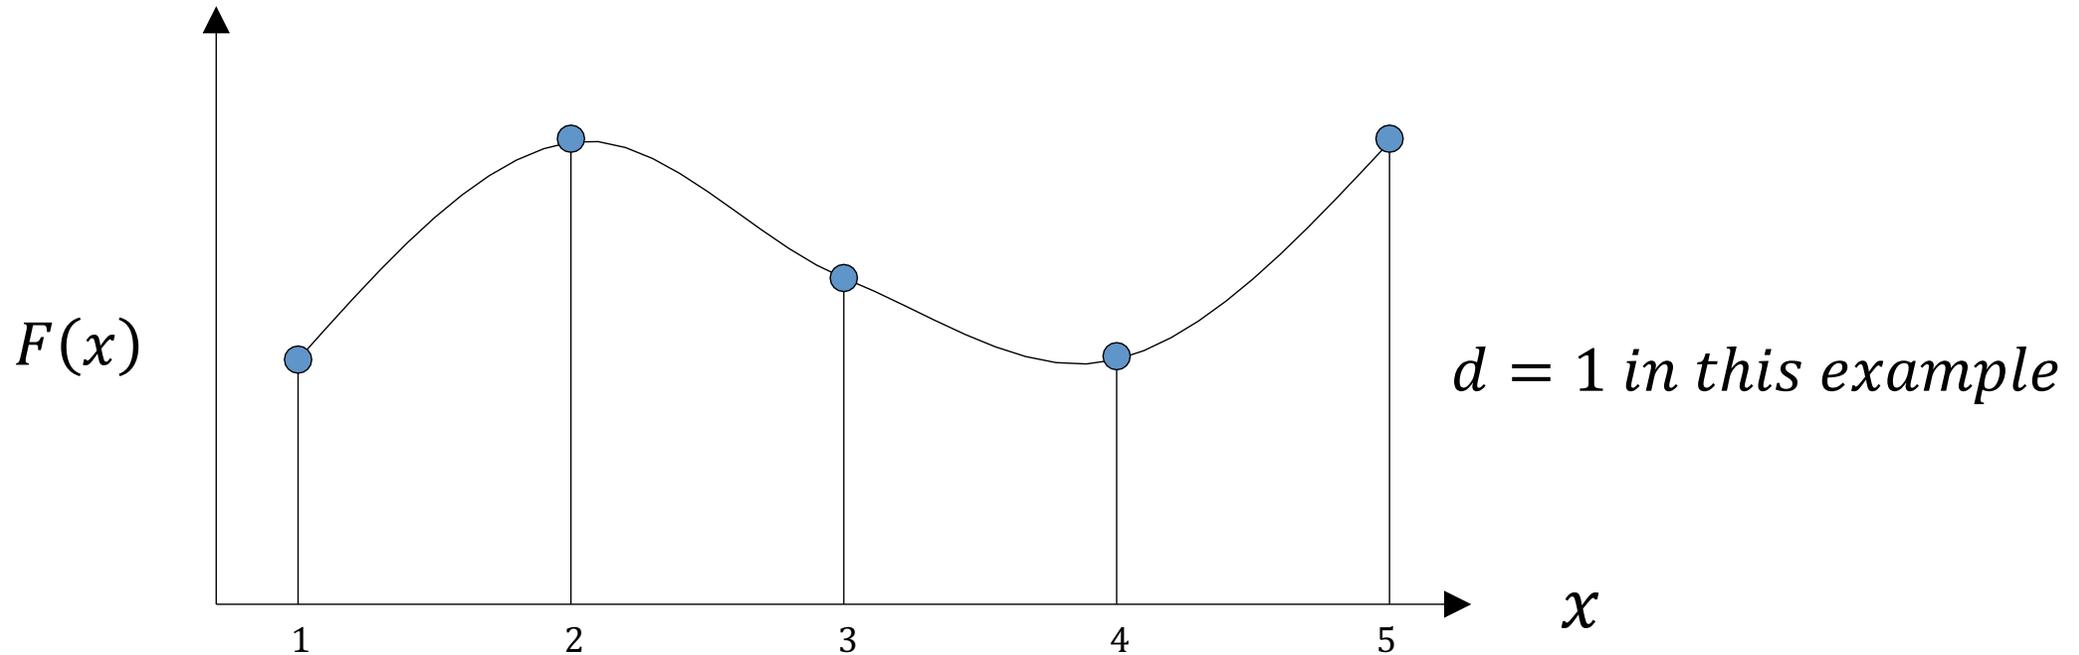


- Recall how digital image is formed

$$F[x, y] = \text{quantize} \left\{ f \left( \frac{x}{d}, \frac{y}{d} \right) \right\}$$

- 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

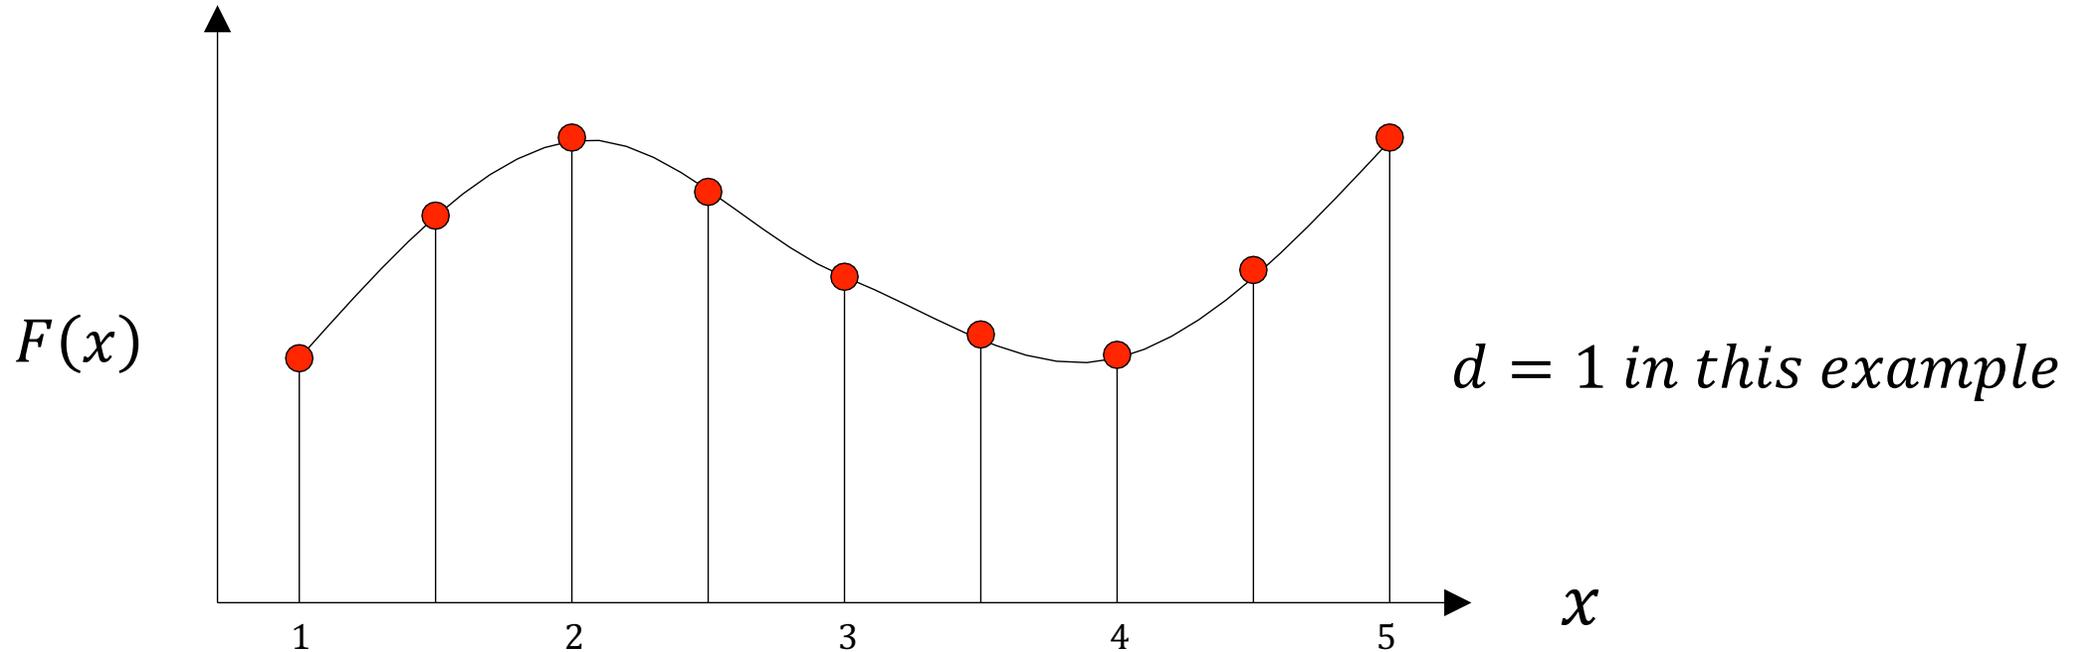


- Recall how digital image is formed

$$F[x, y] = \text{quantize} \left\{ f \left( \frac{x}{d}, \frac{y}{d} \right) \right\}$$

- 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

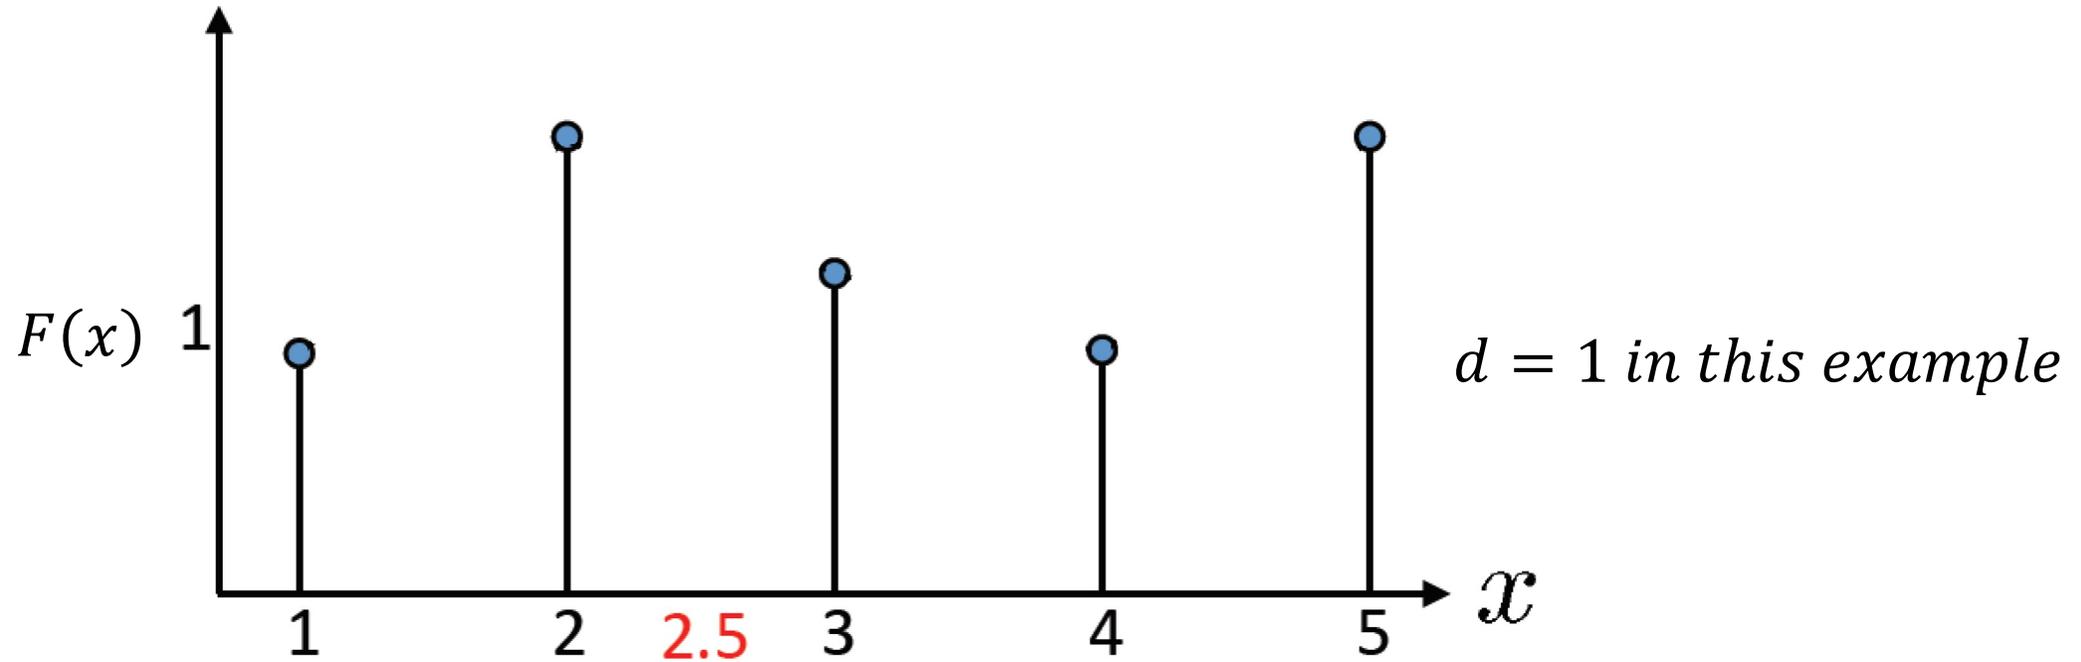


- Recall how digital image is formed

$$F[x, y] = \text{quantize} \left\{ f \left( \frac{x}{d}, \frac{y}{d} \right) \right\}$$

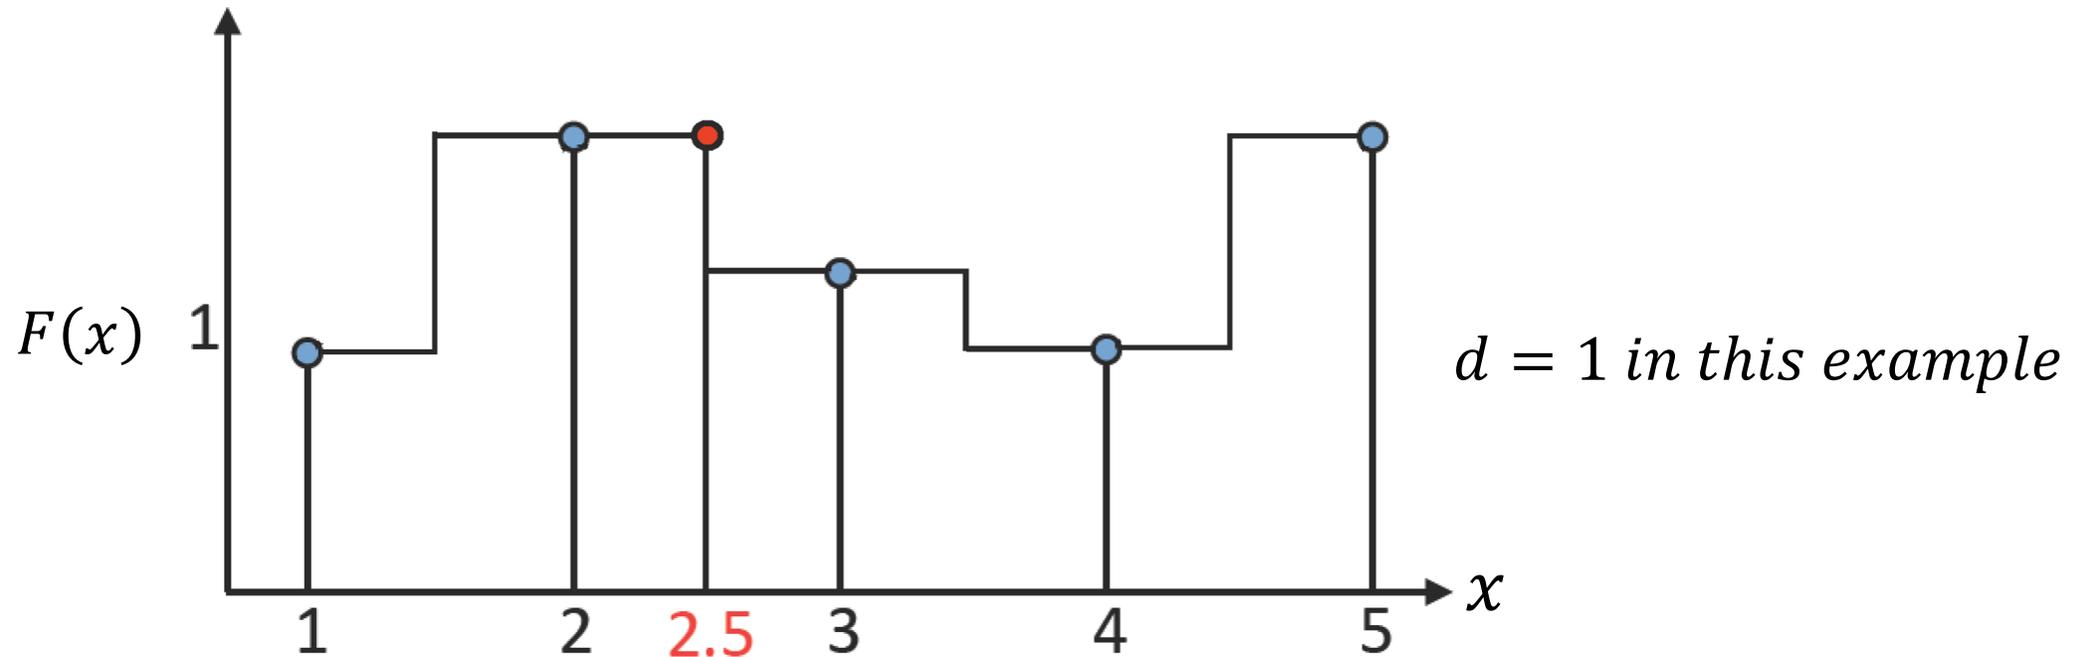
- It is a discrete point-sampling of a continuous function
- If we could somehow reconstruct the original function, any new image could be generated, at any resolution and scale

# Interpolation



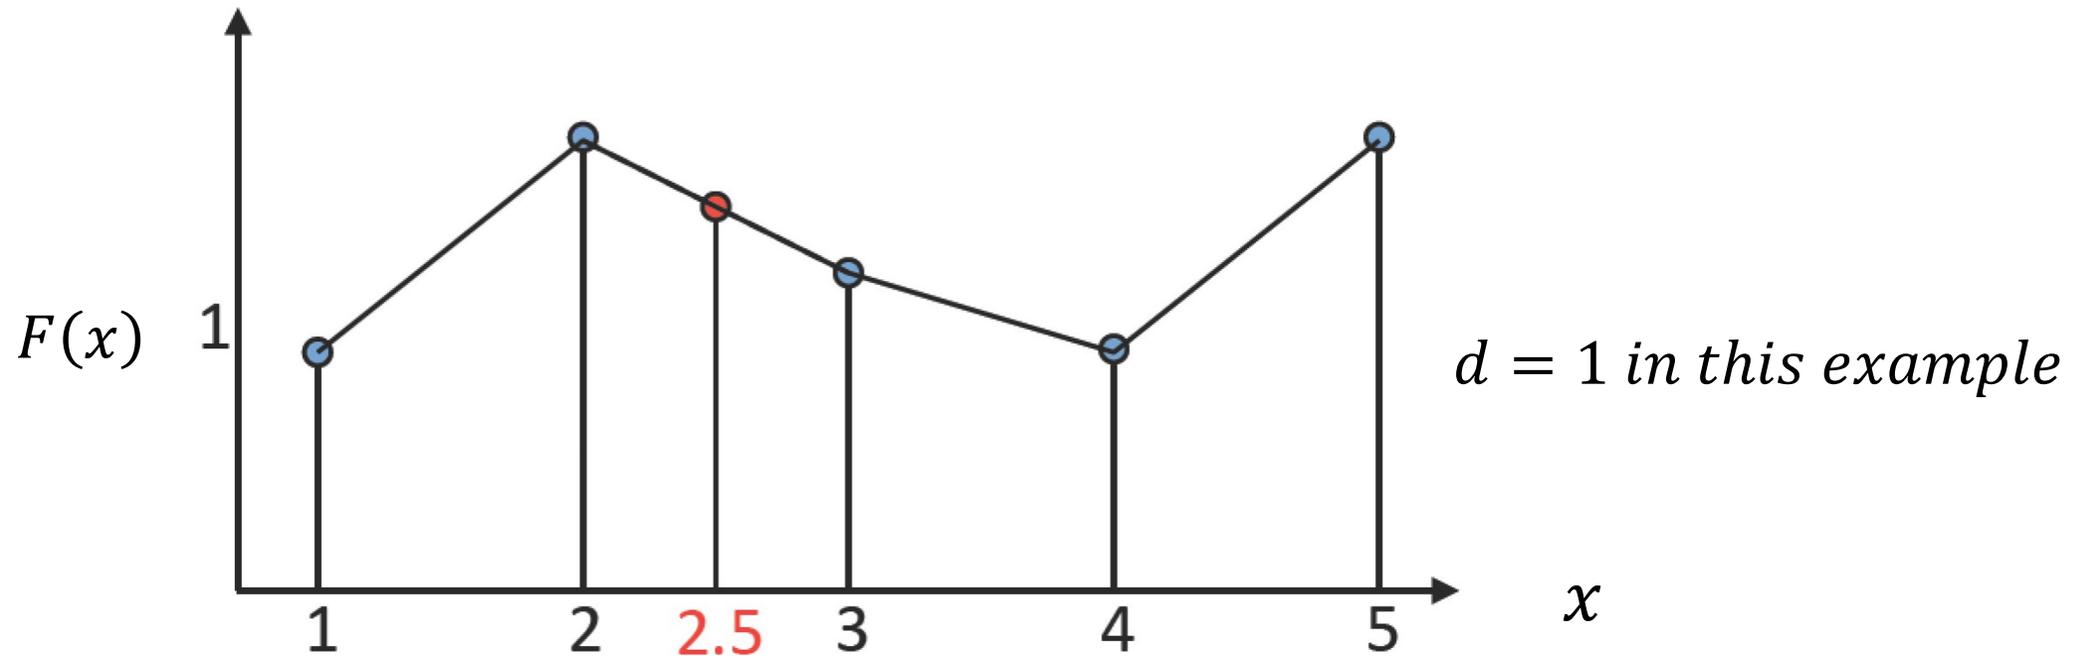
- What if we don't know  $f$ ?

# Interpolation



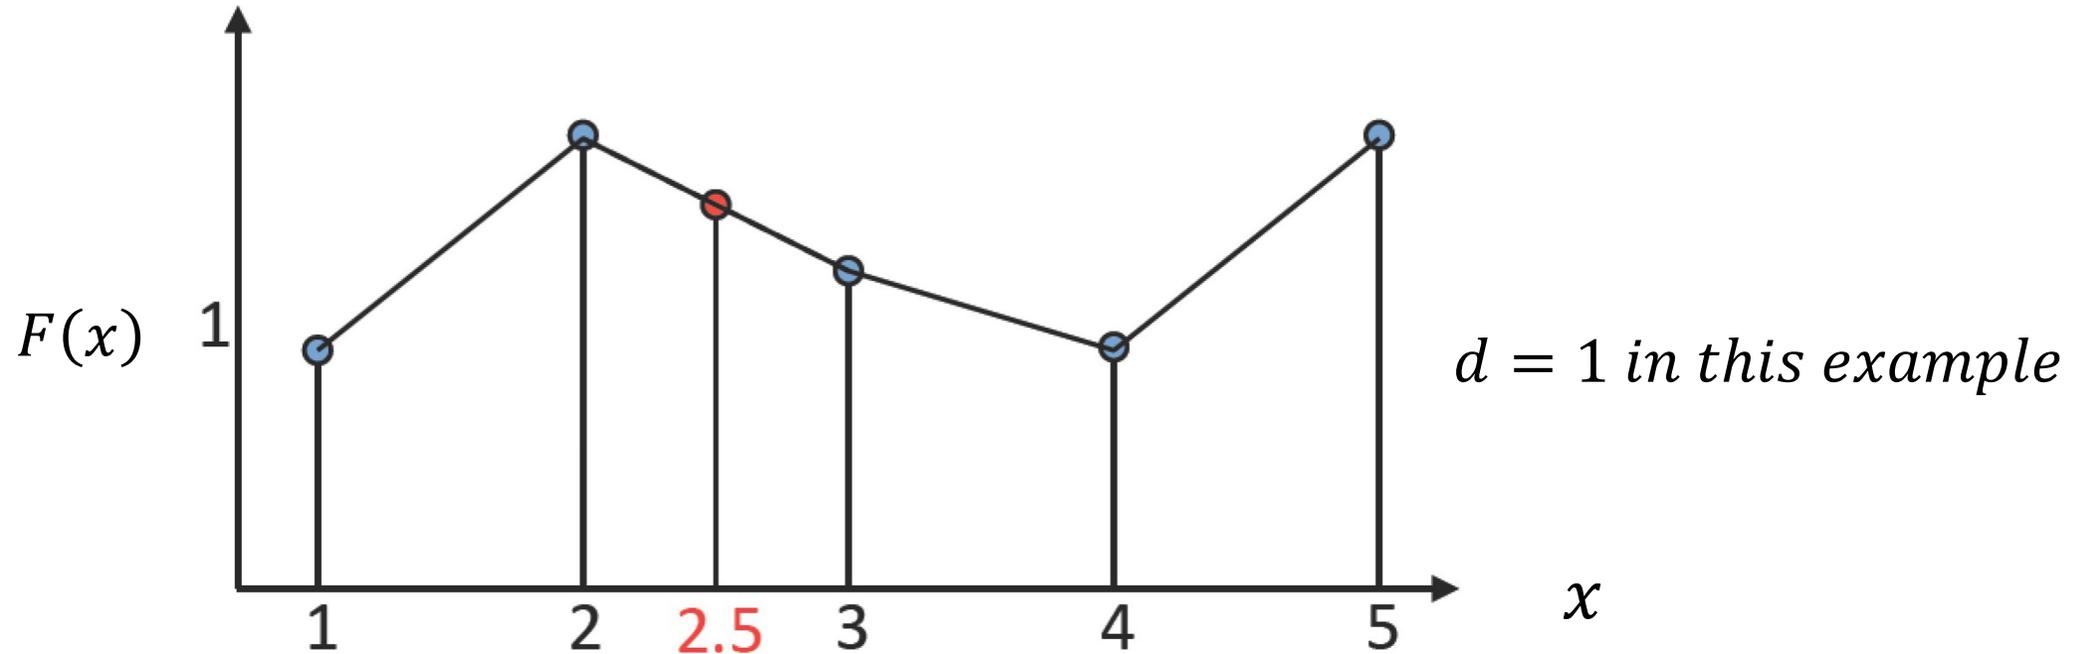
- What if we don't know  $f$ ?
  - Guess an approximation: for example nearest-neighbor

# Interpolation



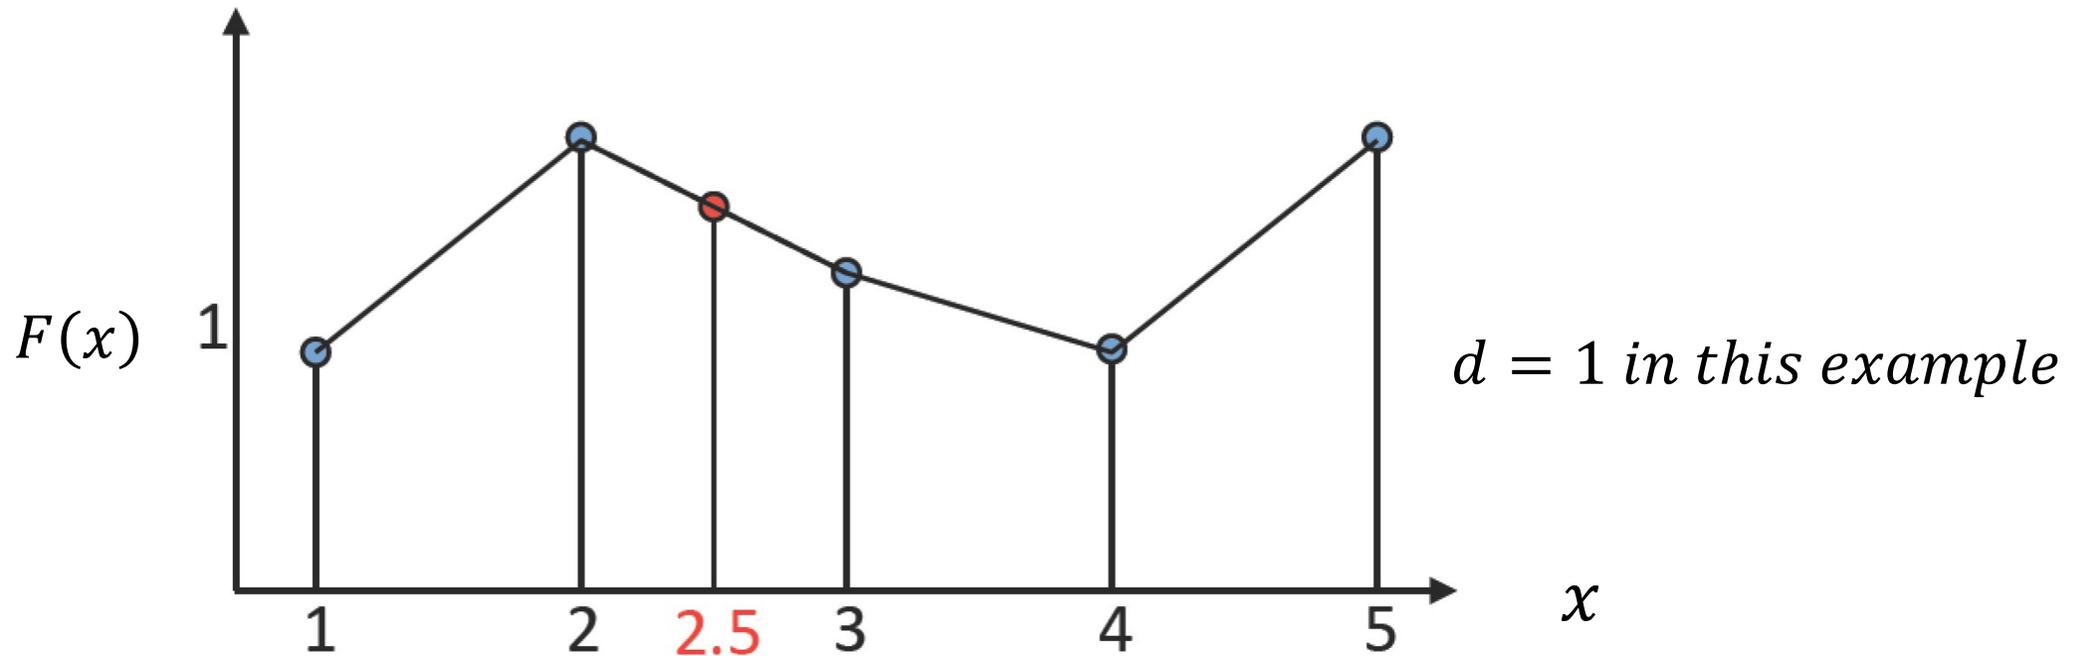
- What if we don't know  $f$ ?
  - Guess an approximation: for example nearest-neighbor
  - Guess an approximation: for example linear

# Interpolation



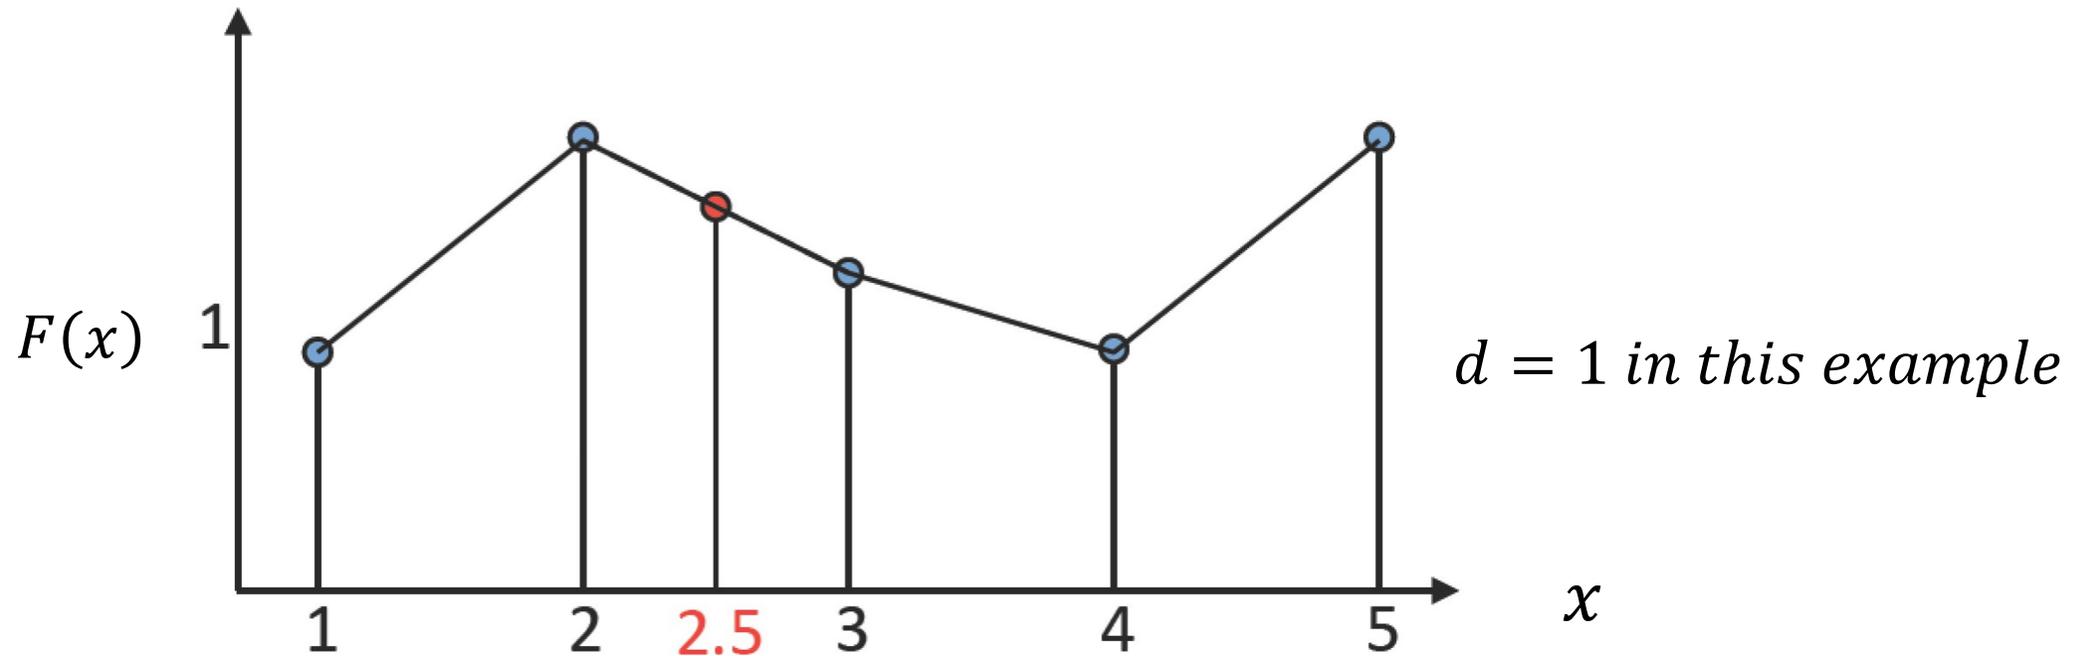
- What if we don't know  $f$ ?
  - Guess an approximation: for example nearest-neighbor
  - Guess an approximation: for example linear
  - More complex approximations: cubic, B-splines

# 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!

# Linear Interpolation



- Linear interpolation from our discretized  $F$  :

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

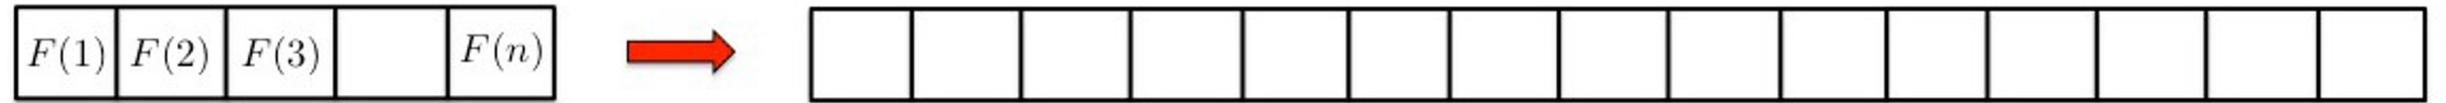
# Interpolation: 1 D Example

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

- Let's upsample by a factor of  $d=3$

# Interpolation: 1 D Example

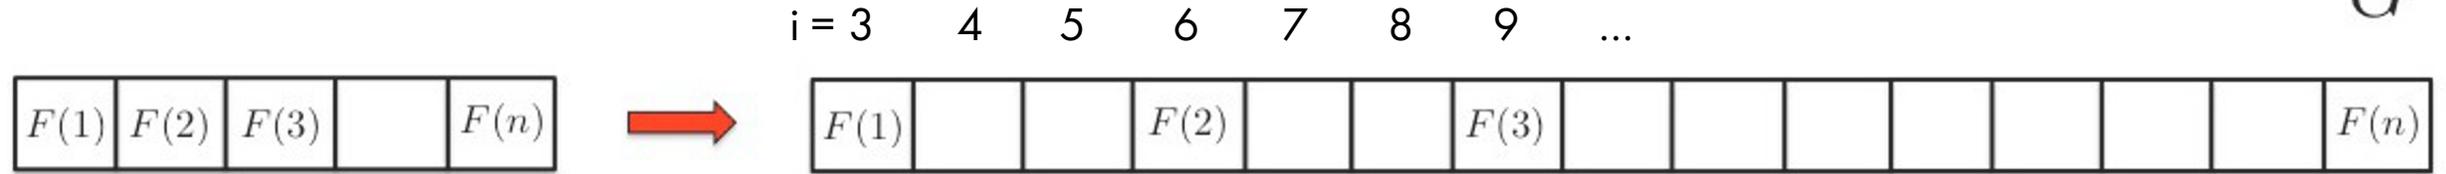
$G$



Make a vector  $G$  with  $d$  times the size of  $F$

- Let's upsample by a factor of  $d=3$

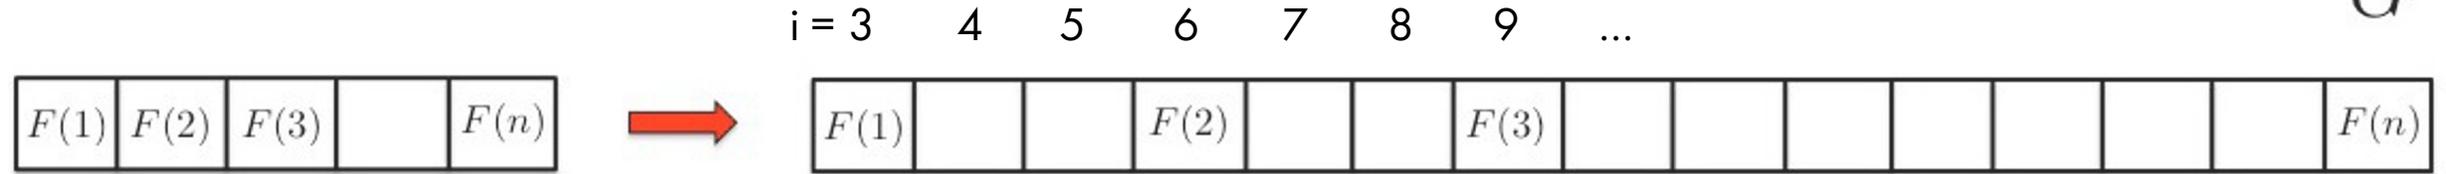
# Interpolation: 1 D Example



$$\text{if } \frac{i}{d} \text{ integer: } G(i) = F(i/d)$$

- Let's upsample by a factor of  $d=3$
- if  $i/d$  is integer, just copy over the value

# Interpolation: 1 D Example



$$\text{if } \frac{i}{d} \text{ integer: } G(i) = F(i/d)$$

$$\text{otherwise: } G(i) = \frac{x_2 - x}{x_2 - x_1} F(x_1) + \frac{x - x_1}{x_2 - x_1} F(x_2)$$

$$x = i/d$$

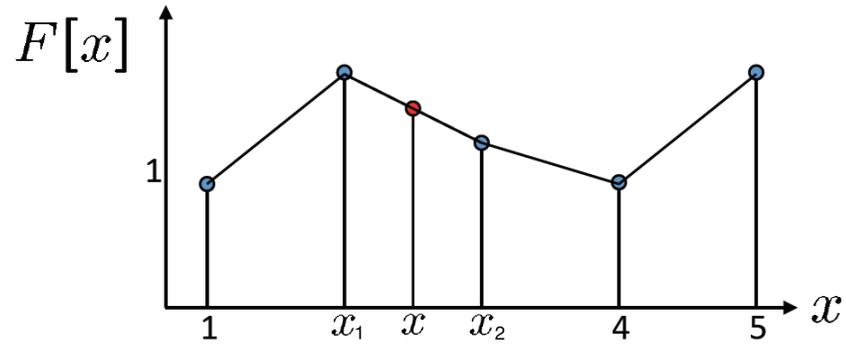
where

$$x_1 = \lfloor i/d \rfloor$$

$$x_2 = \lceil i/d \rceil$$

- Let's upsample by a factor of  $d=3$
- if  $i/d$  is integer, just copy over the value
- otherwise, use interpolation formula

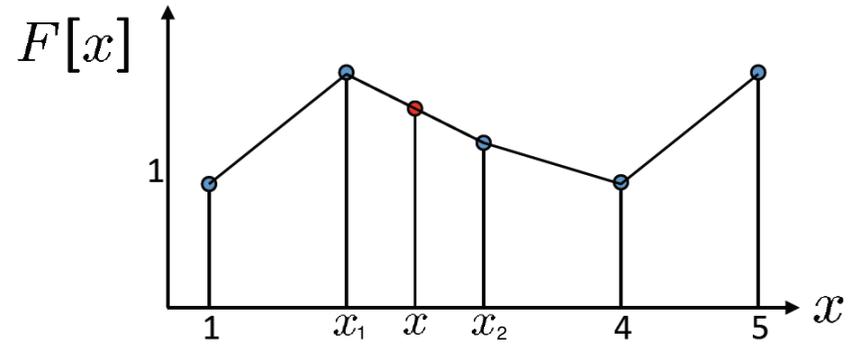
# Linear Interpolation via Convolution



- Linear interpolation:

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

# Linear Interpolation via Convolution

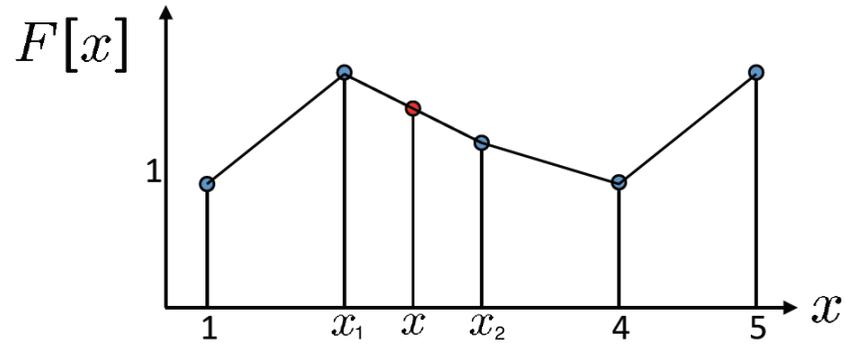


- Linear interpolation:

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

- how can we do this with a convolution?

# Linear Interpolation via Convolution

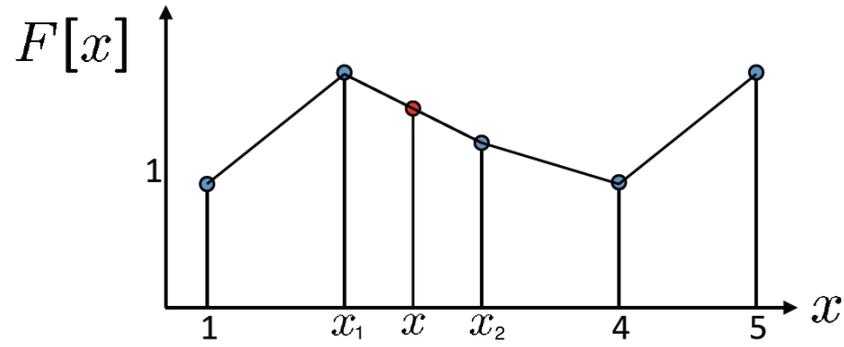


- Linear interpolation:

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

- how can we do this with a convolution?
- what should be the input?

# Linear Interpolation via Convolution

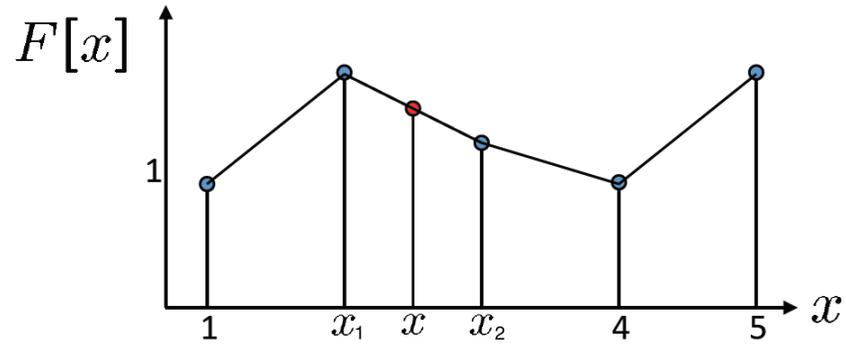


- Linear interpolation:

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

- how can we do this with a convolution?
- what should be the input? (original signal interleaved with zeroes)

# Linear Interpolation via Convolution

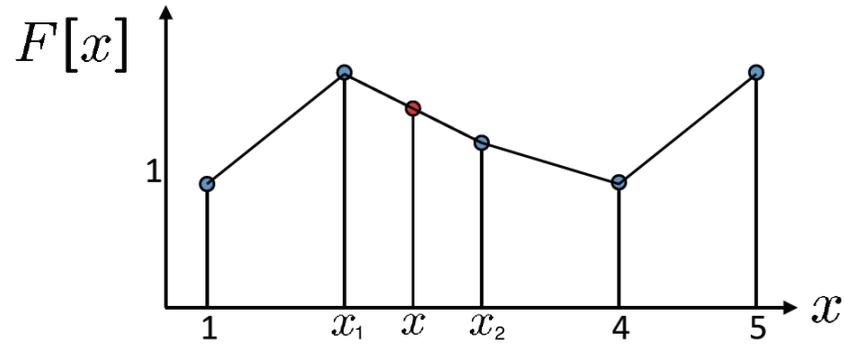


- Linear interpolation:

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

- how can we do this with a convolution?
- what should be the input? (original signal interleaved with zeroes)
- what should be the filter?

# Linear Interpolation via Convolution



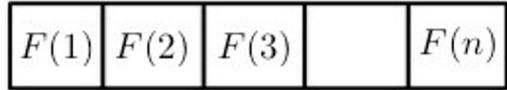
- Linear interpolation:

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

- how can we do this with a convolution?
- what should be the input? (original signal interleaved with zeroes)

- what should be the filter?  $G_{interpolated}(x_i) = \left[ \frac{1}{2}, 1, \frac{1}{2} \right] * G'$

# Interpolation via Convolution: 1 D Example

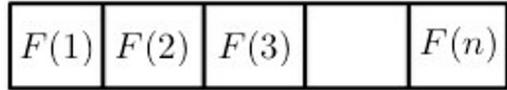


- Let's make this signal triple length (what will this look like?)

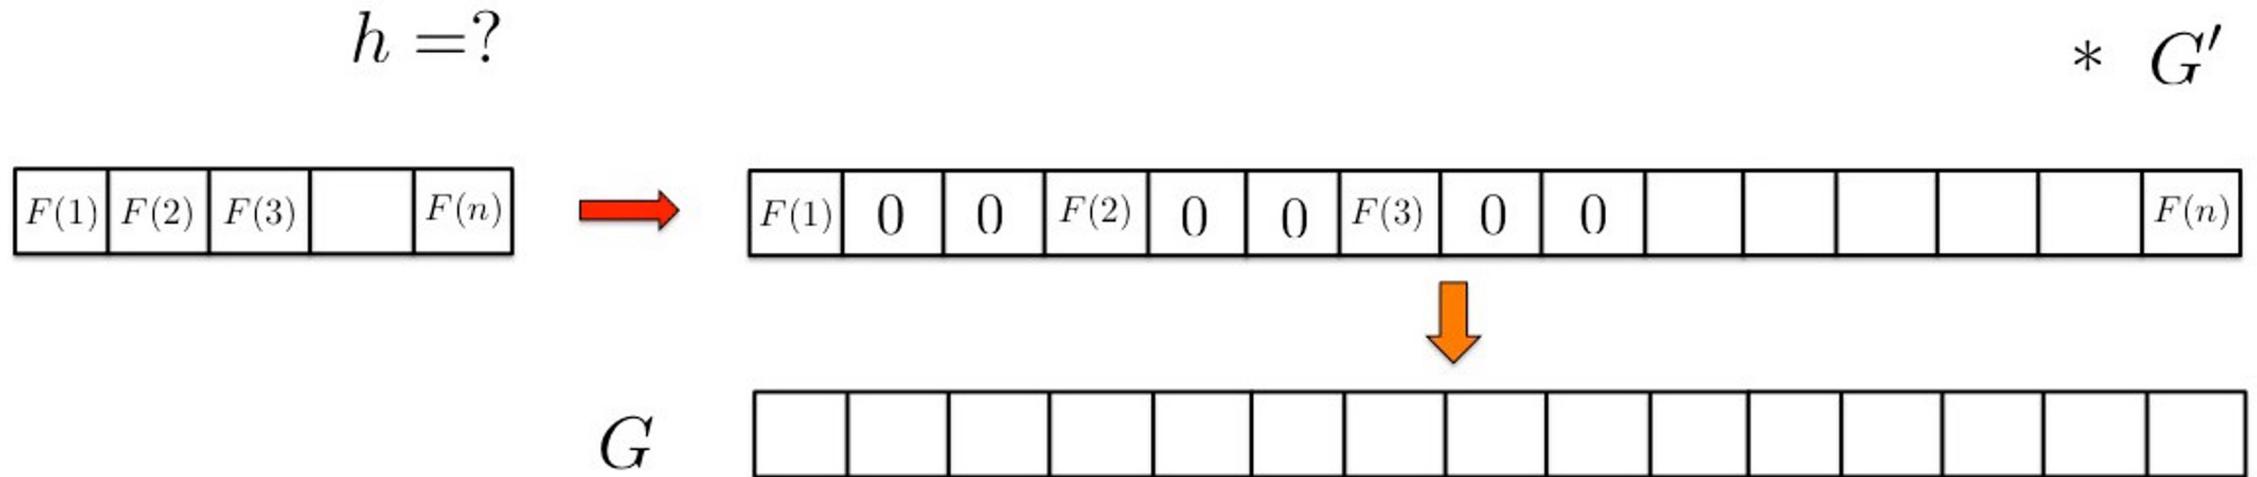
# Interpolation via Convolution: 1D Example

$h = ?$

$* G'$

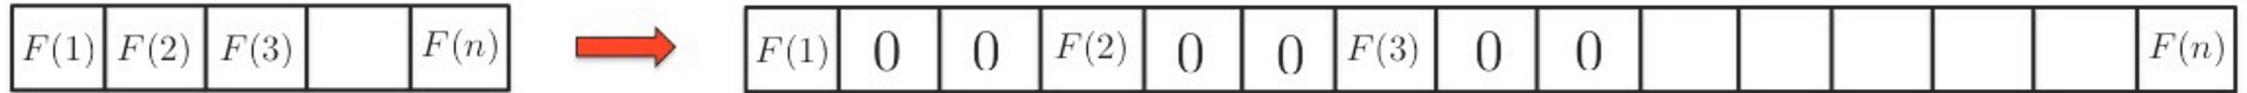


# Interpolation via Convolution: 1 D Example



- how does this get copied to the output?

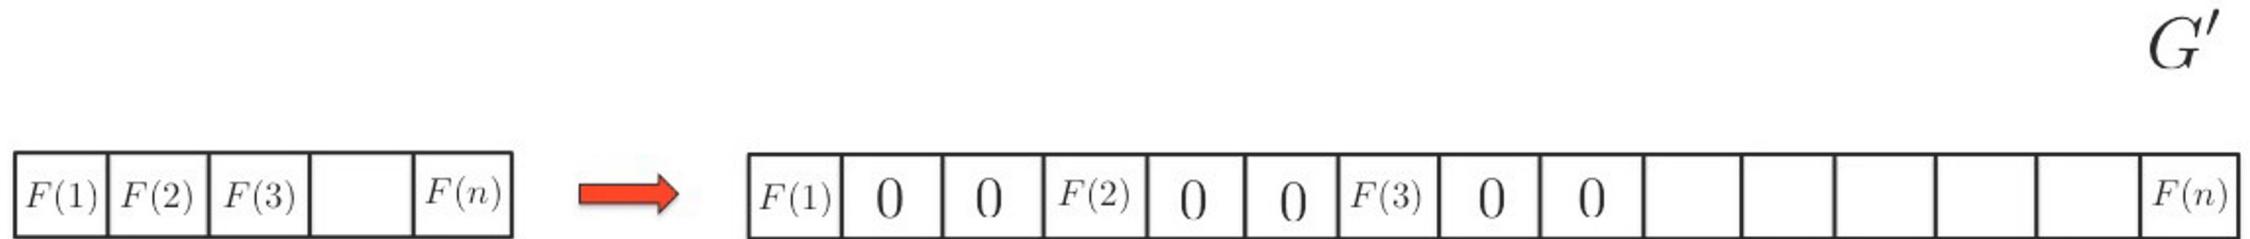
# Interpolation via Convolution: 1D Example

 $G'$ 

if  $\frac{i}{d}$  integer:  $G'(i) = F(i/d)$

otherwise: 0

# Interpolation via Convolution: 1D Example



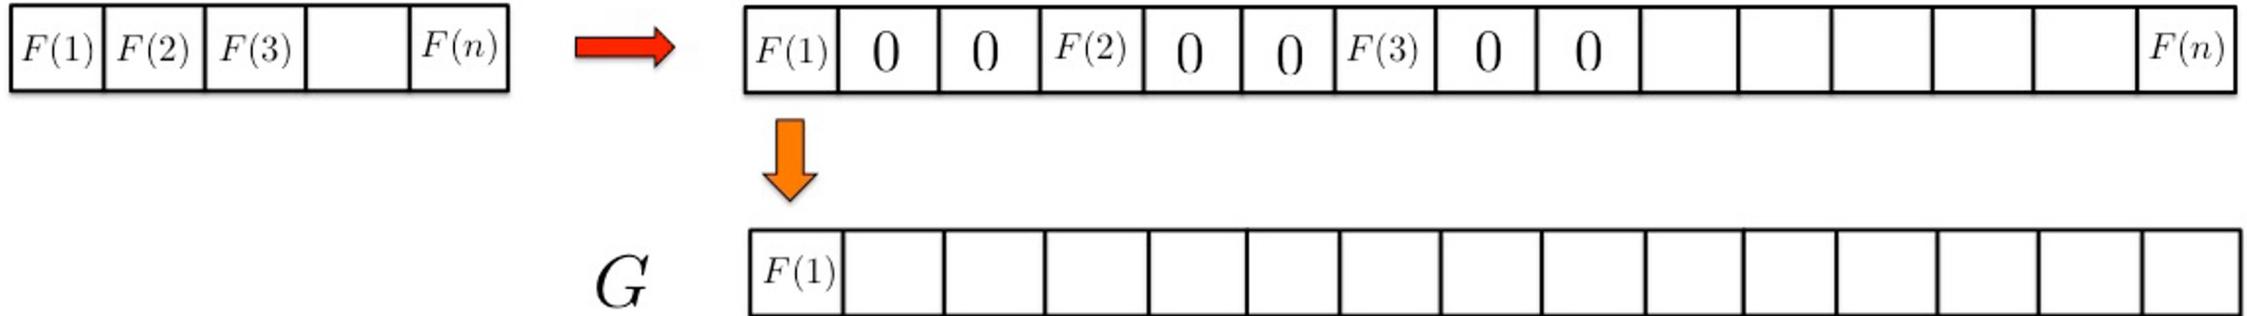
$$\text{if } \frac{i}{d} \text{ integer: } G'(i) = F(i/d)$$

otherwise: 0

- what does the convolution kernel look like?

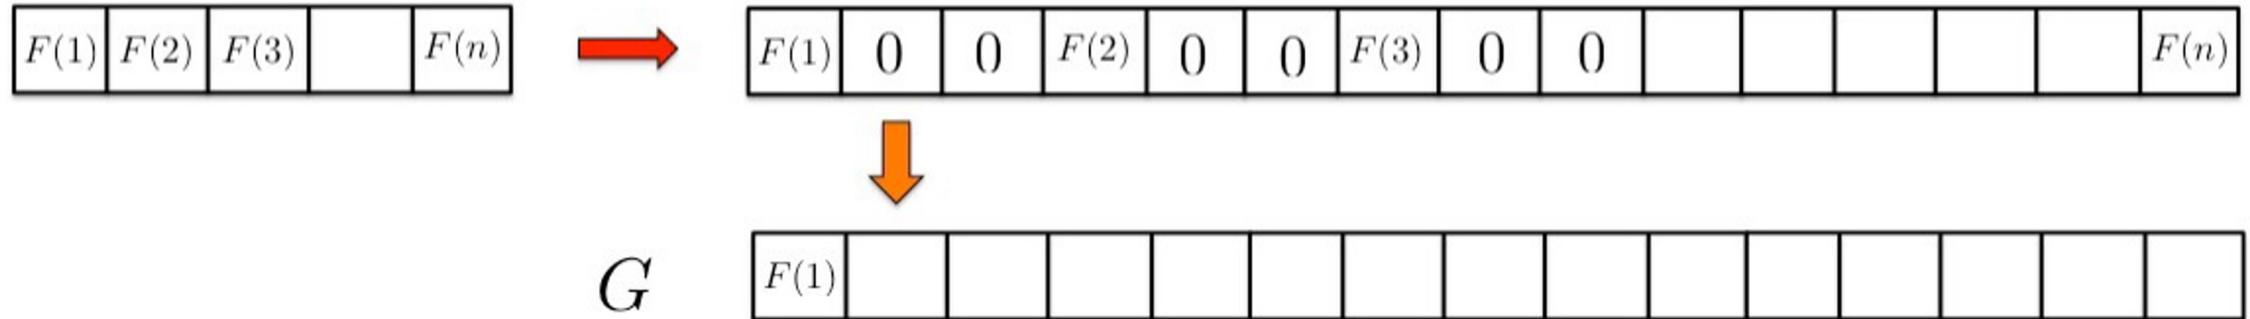
# Interpolation via Convolution: 1D Example

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



# Interpolation via Convolution: 1D Example

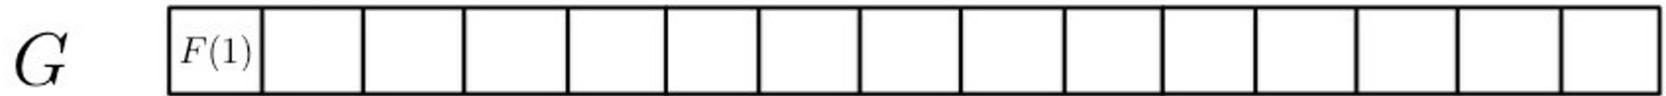
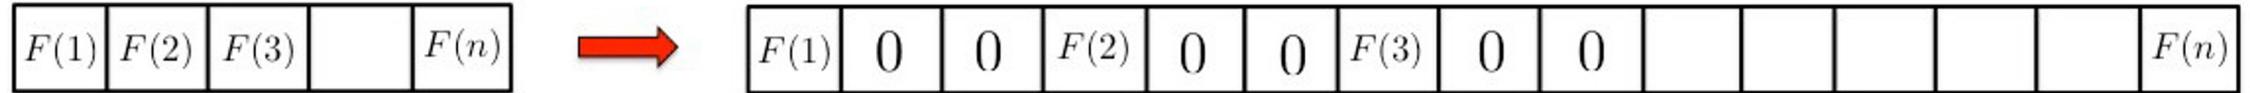
$$h = \left[ 0, \frac{1}{3}, \frac{2}{3}, 1, \frac{2}{3}, \frac{1}{3}, 0 \right] \quad * \quad G'$$



- how do we calculate this entry?

# Interpolation via Convolution: 1D Example

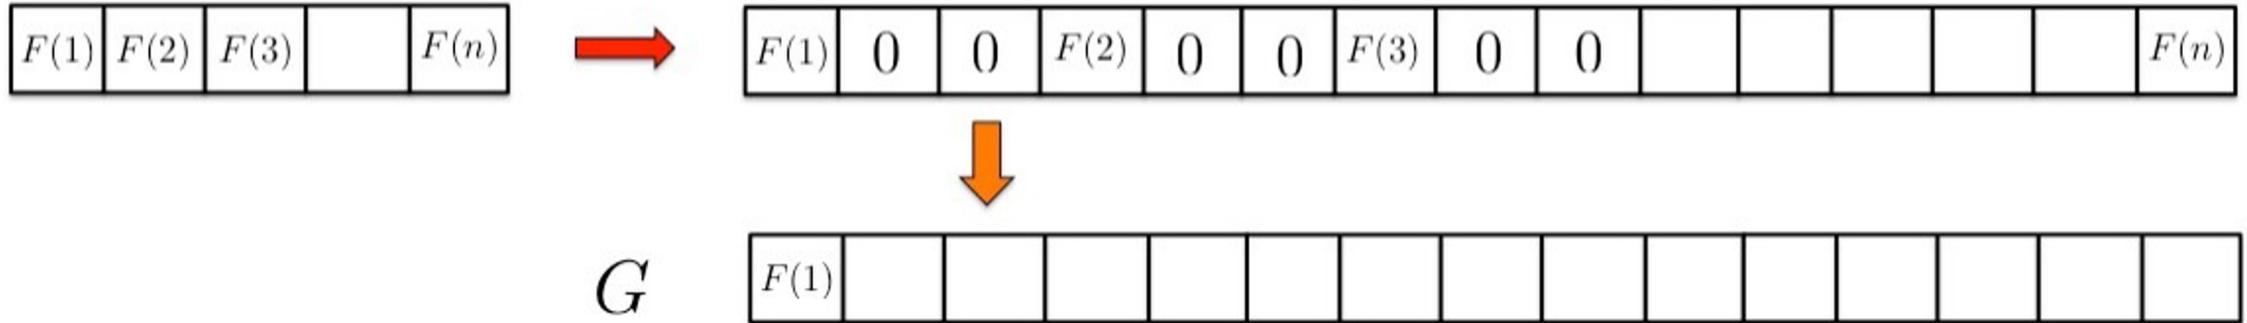
$$h = \left[ 0, \frac{1}{3}, \frac{2}{3}, 1, \frac{2}{3}, \frac{1}{3}, 0 \right] \quad * \quad G'$$



$$\frac{2}{3}F(1) + \frac{1}{3}F(2)$$

# Interpolation via Convolution: 1D Example

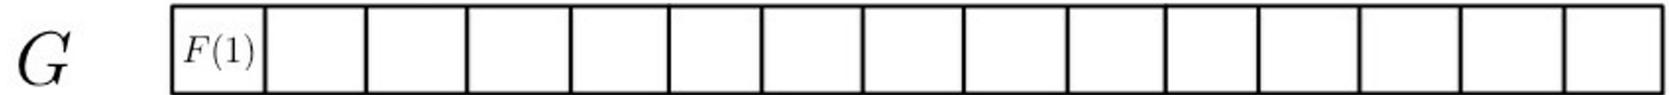
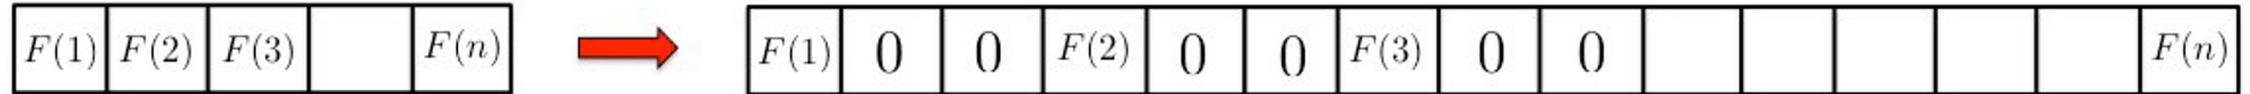
$$h = \left[ 0, \frac{1}{3}, \frac{2}{3}, 1, \frac{2}{3}, \frac{1}{3}, 0 \right] \quad * \quad G'$$



- how do we calculate this entry?

# Interpolation via Convolution: 1D Example

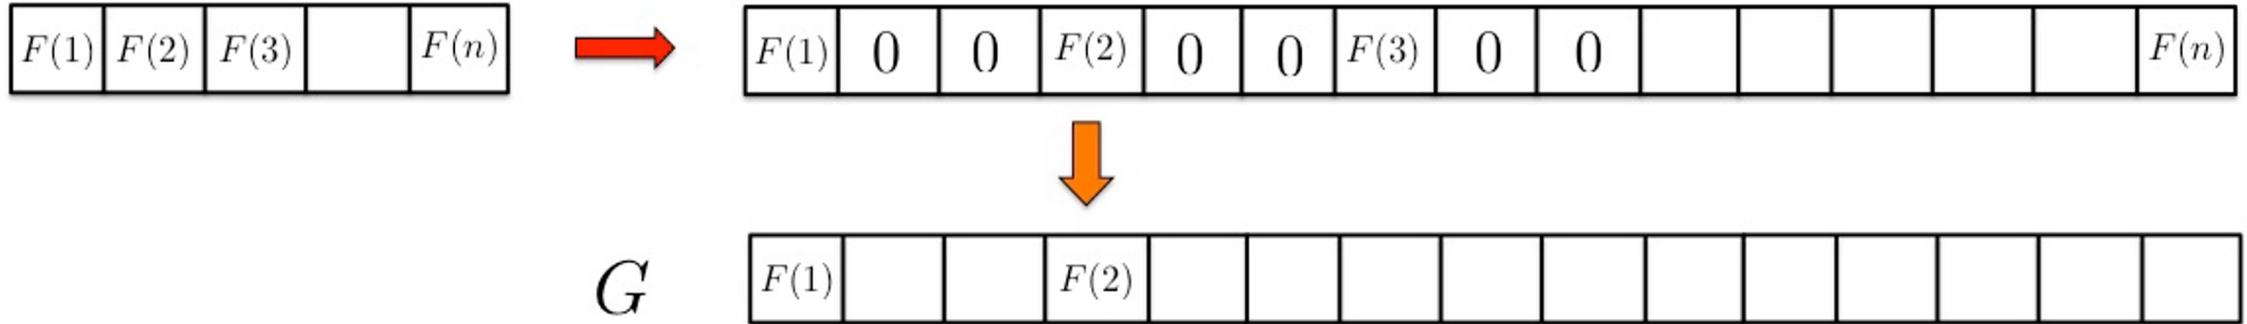
$$h = \left[ 0, \frac{1}{3}, \frac{2}{3}, 1, \frac{2}{3}, \frac{1}{3}, 0 \right] \quad * \quad G'$$



$$\frac{1}{3}F(1) + \frac{2}{3}F(2)$$

# Interpolation via Convolution: 1D Example

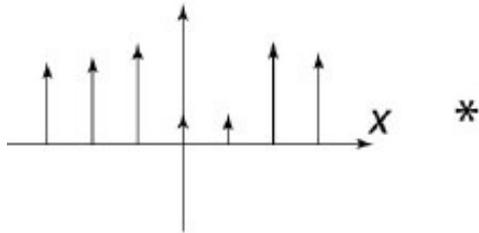
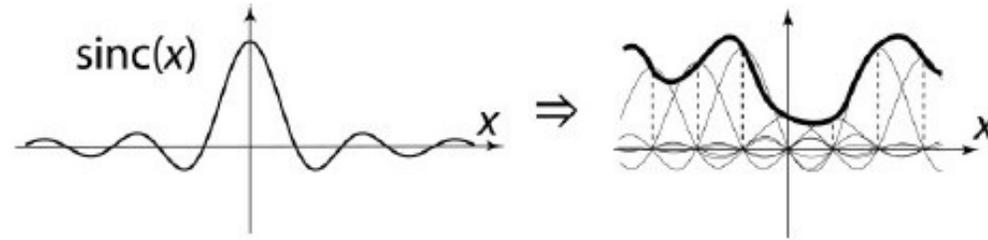
$$h = \left[0, \frac{1}{3}, \frac{2}{3}, 1, \frac{2}{3}, \frac{1}{3}, 0\right] \quad * \quad G'$$



- ... and so on and so forth

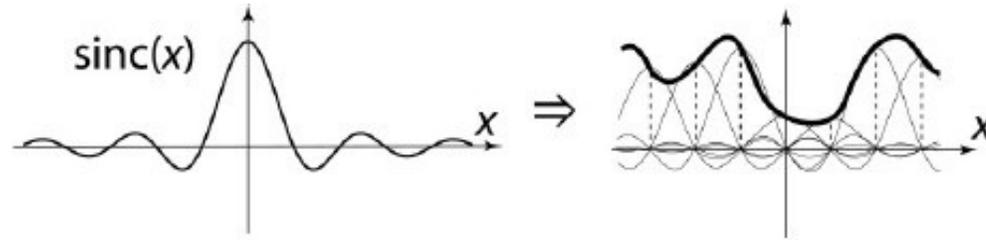
# Interpolation via Convolution (1D)

overview of  
interpolation  
kernels

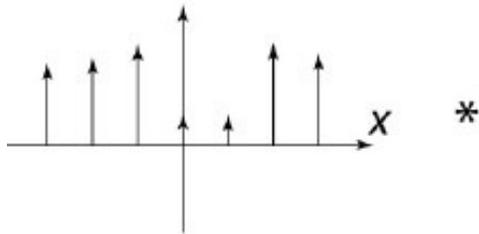


# Interpolation via Convolution (1D)

overview of  
interpolation  
kernels

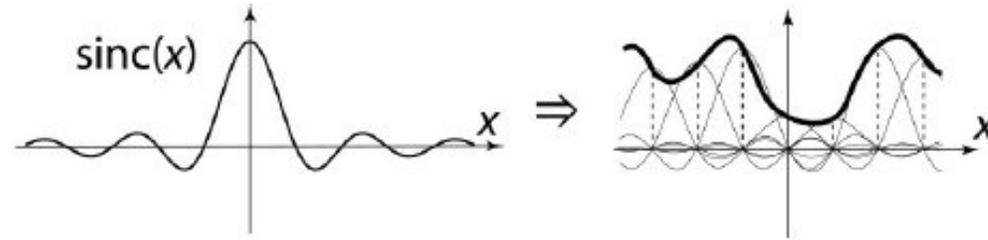
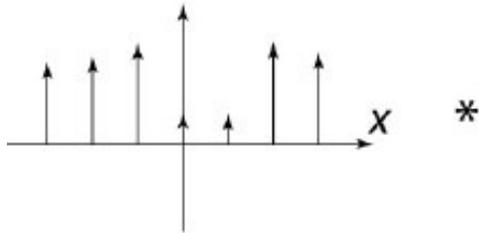


"Ideal" reconstruction

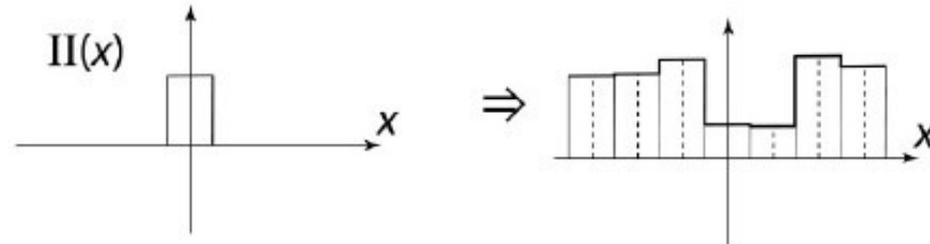


# Interpolation via Convolution (1D)

overview of  
interpolation  
kernels

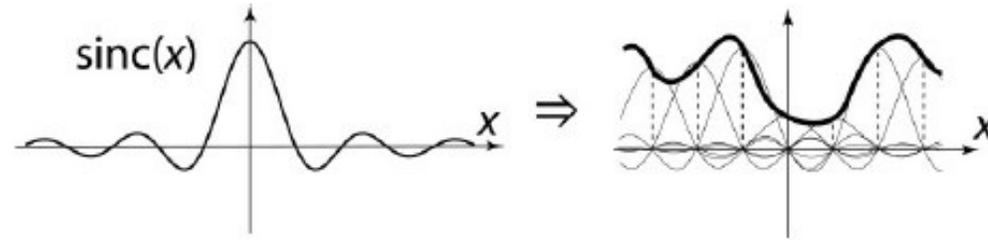
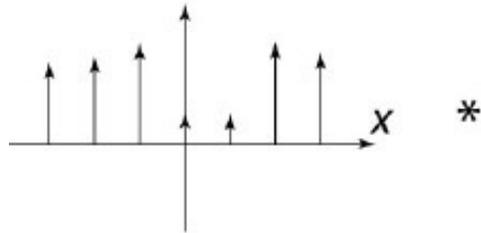


"Ideal" reconstruction

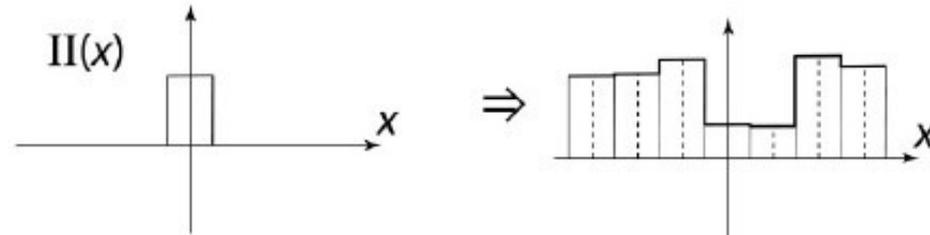


# Interpolation via Convolution (1D)

overview of  
interpolation  
kernels



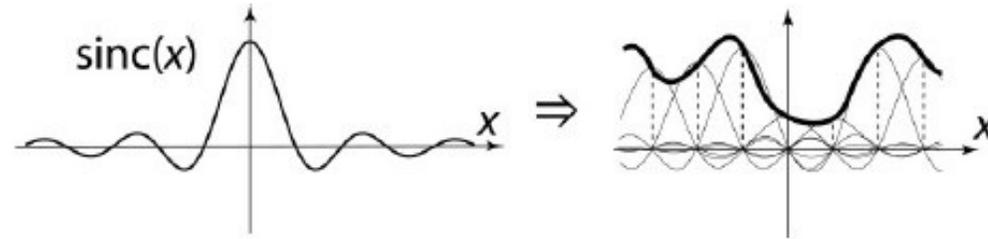
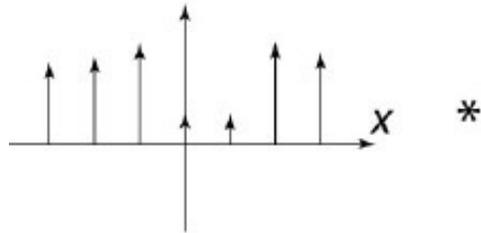
"Ideal" reconstruction



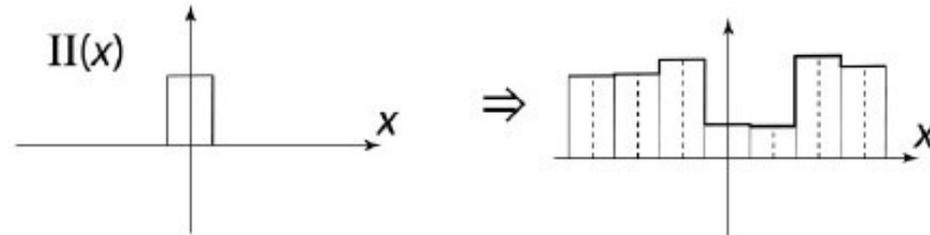
Nearest-Neighbor Interpolation

# Interpolation via Convolution (1D)

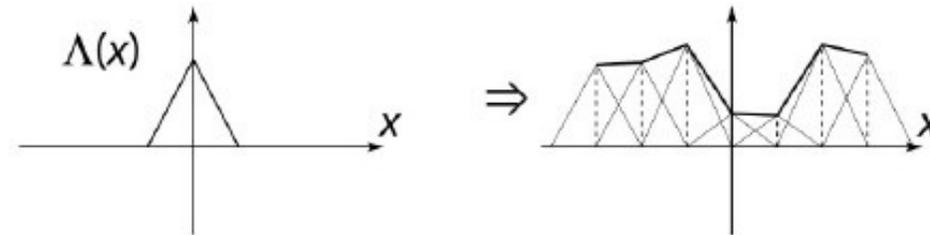
overview of  
interpolation  
kernels



"Ideal" reconstruction

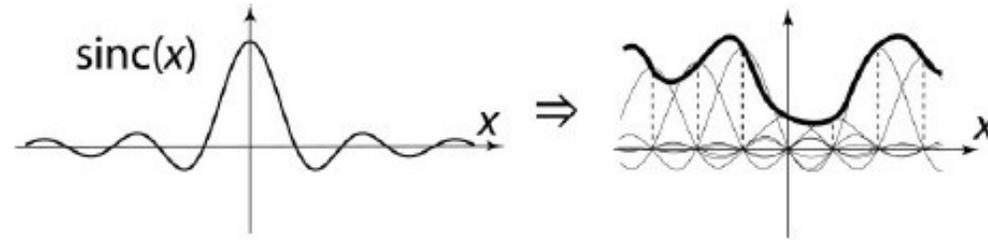
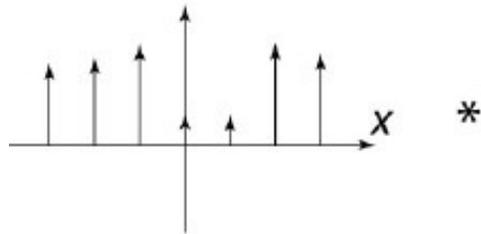


Nearest-Neighbor Interpolation

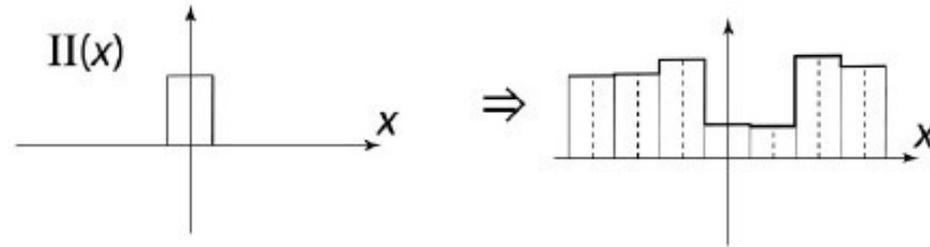


# Interpolation via Convolution (1D)

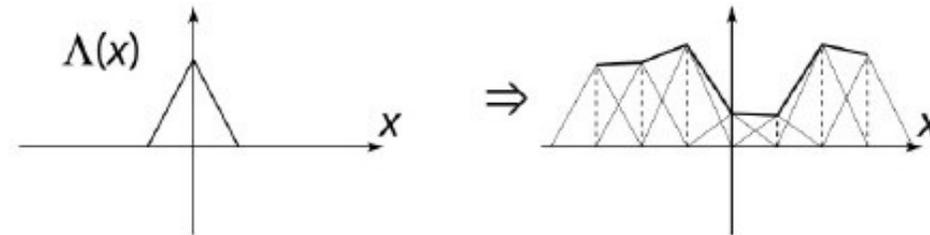
overview of  
interpolation  
kernels



"Ideal" reconstruction



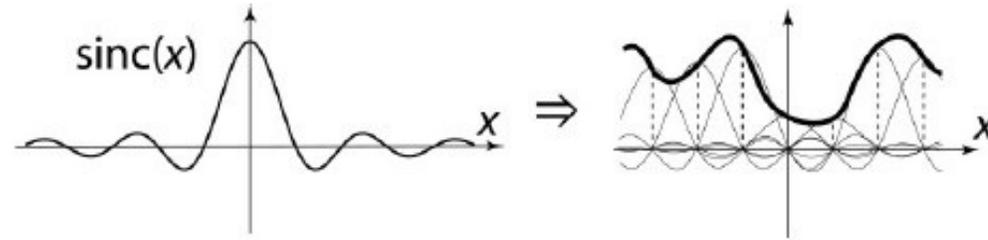
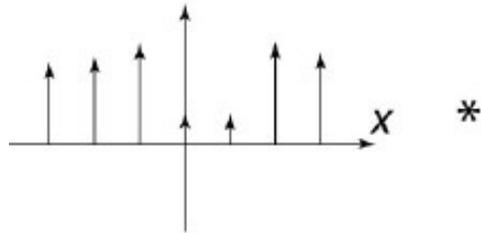
Nearest-Neighbor Interpolation



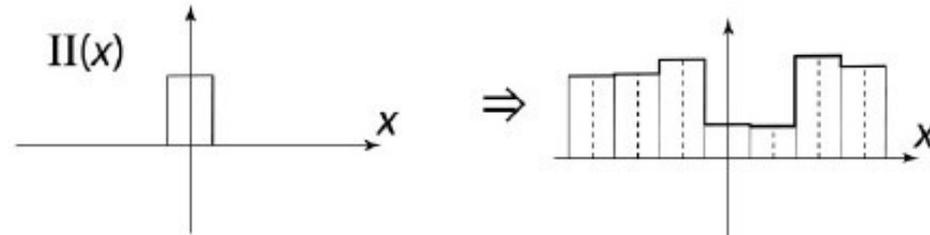
Linear Interpolation

# Interpolation via Convolution (1D)

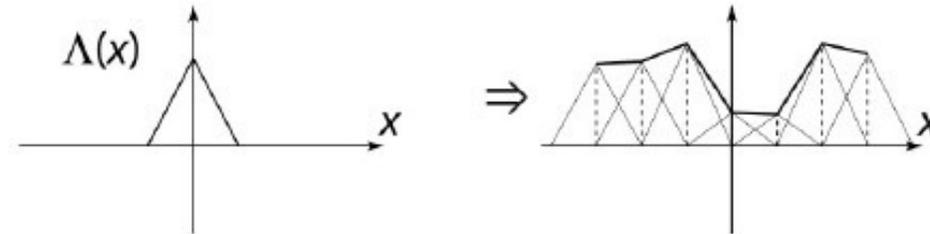
overview of  
interpolation  
kernels



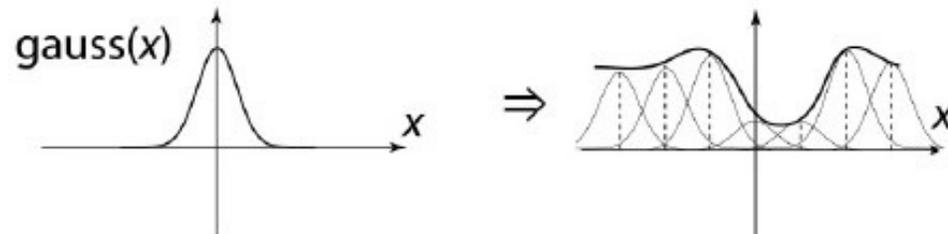
"Ideal" reconstruction



Nearest-Neighbor Interpolation

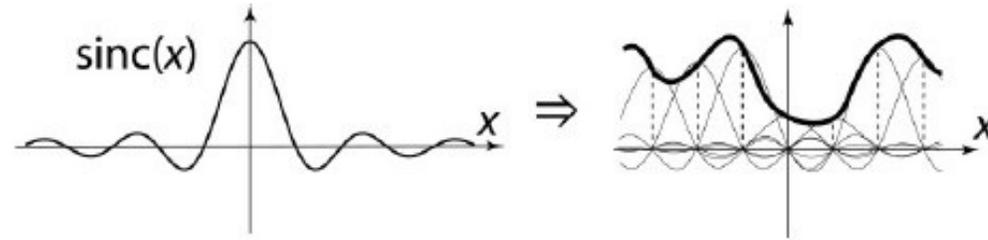
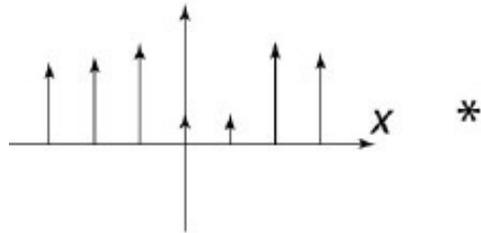


Linear Interpolation

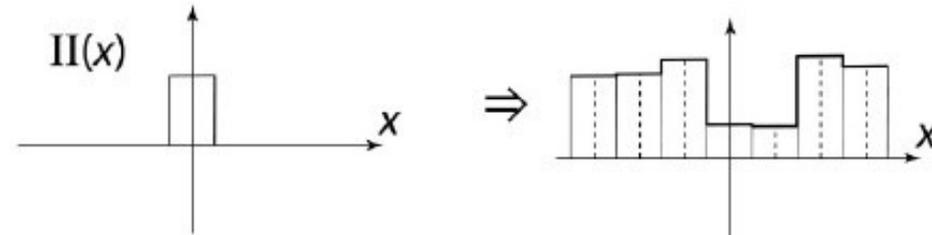


# Interpolation via Convolution (1D)

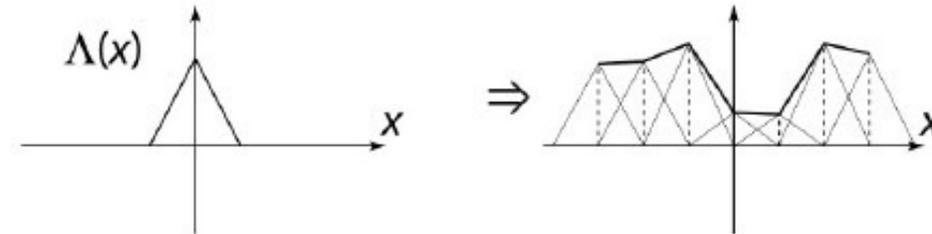
overview of  
interpolation  
kernels



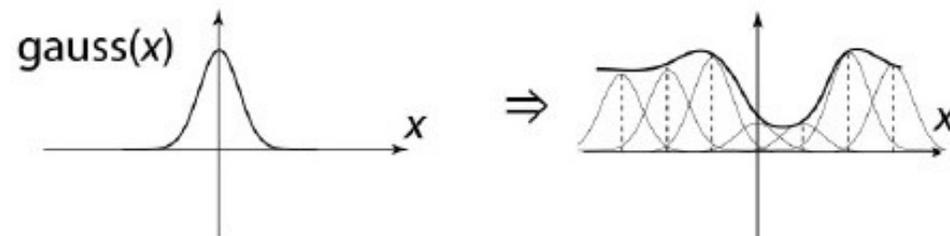
"Ideal" reconstruction



Nearest-Neighbor Interpolation



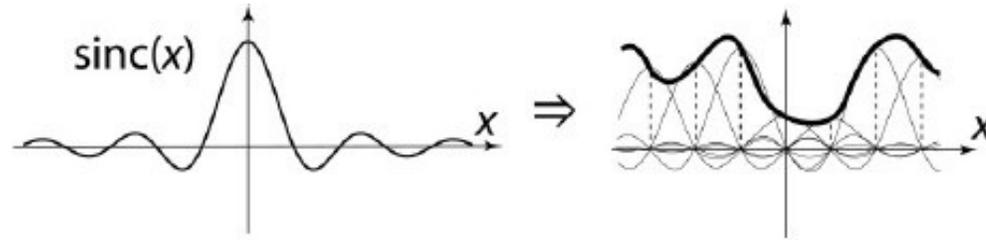
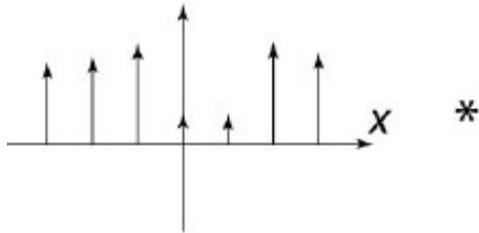
Linear Interpolation



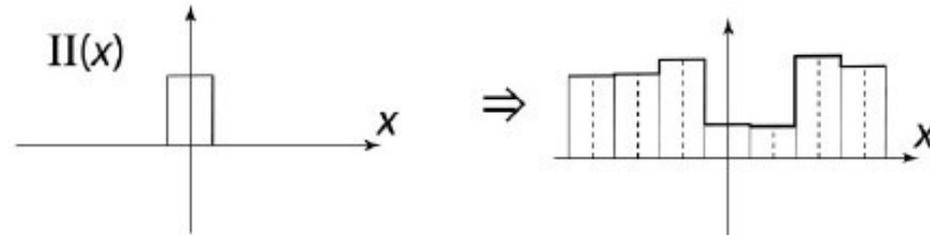
Gaussian reconstruction

# Interpolation via Convolution (1D)

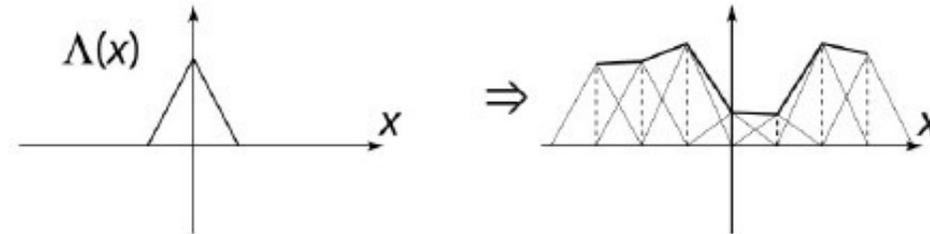
overview of  
interpolation  
kernels



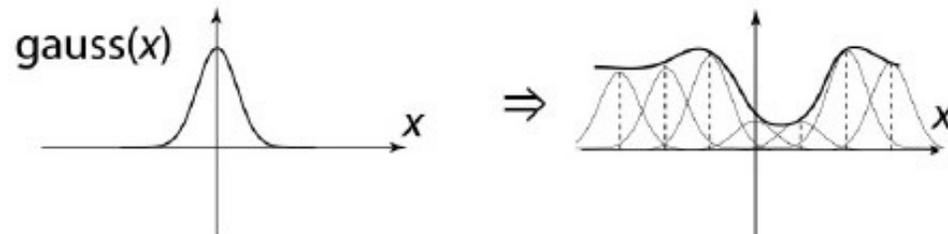
"Ideal" reconstruction



Nearest-Neighbor Interpolation



Linear Interpolation



Gaussian reconstruction

pros and cons of each?

# Image Interpolation (2D)

# Image Interpolation (2D)

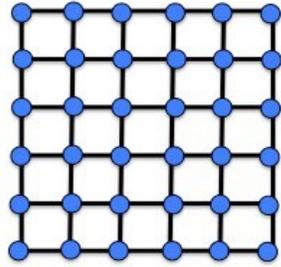
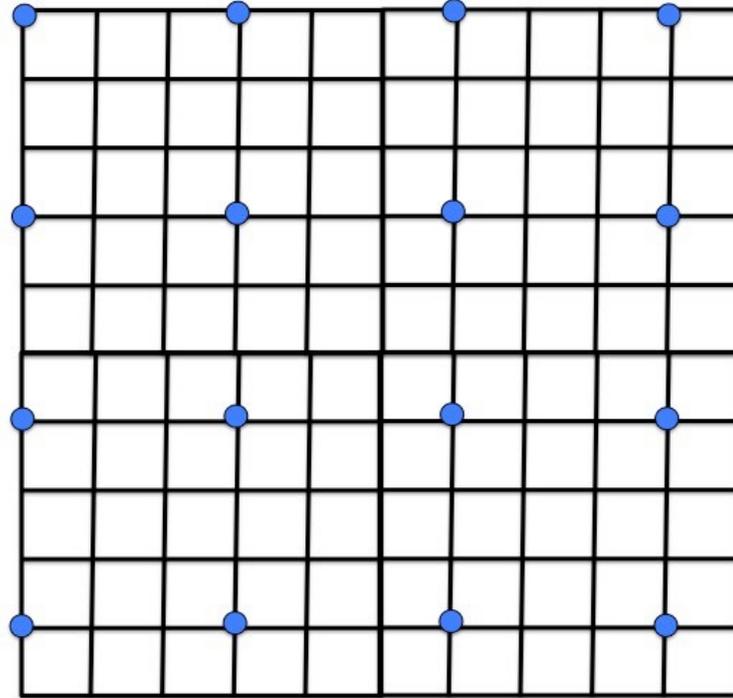


image  $I$



- Let's make this image triple size

# Image Interpolation (2D)

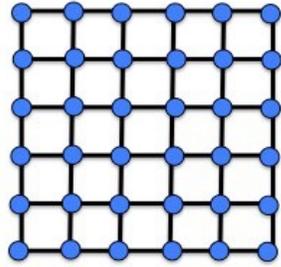
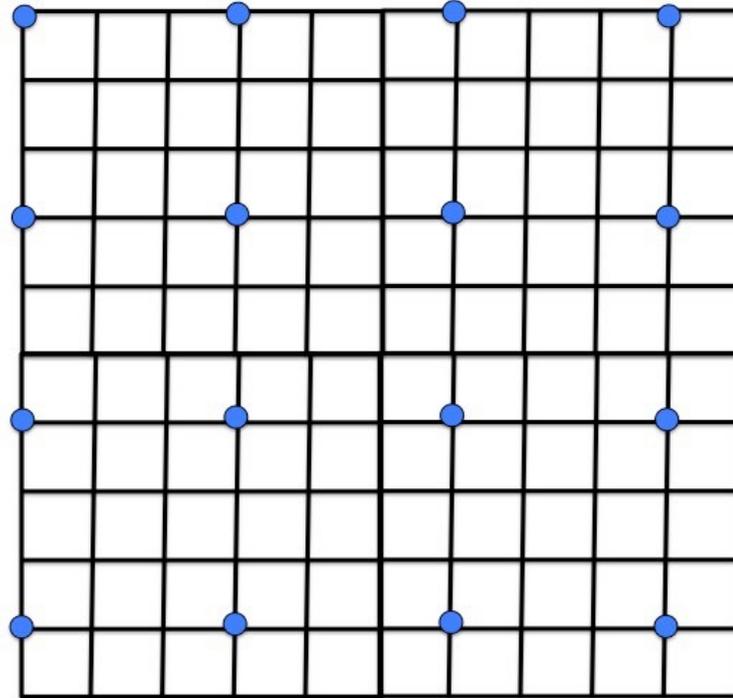
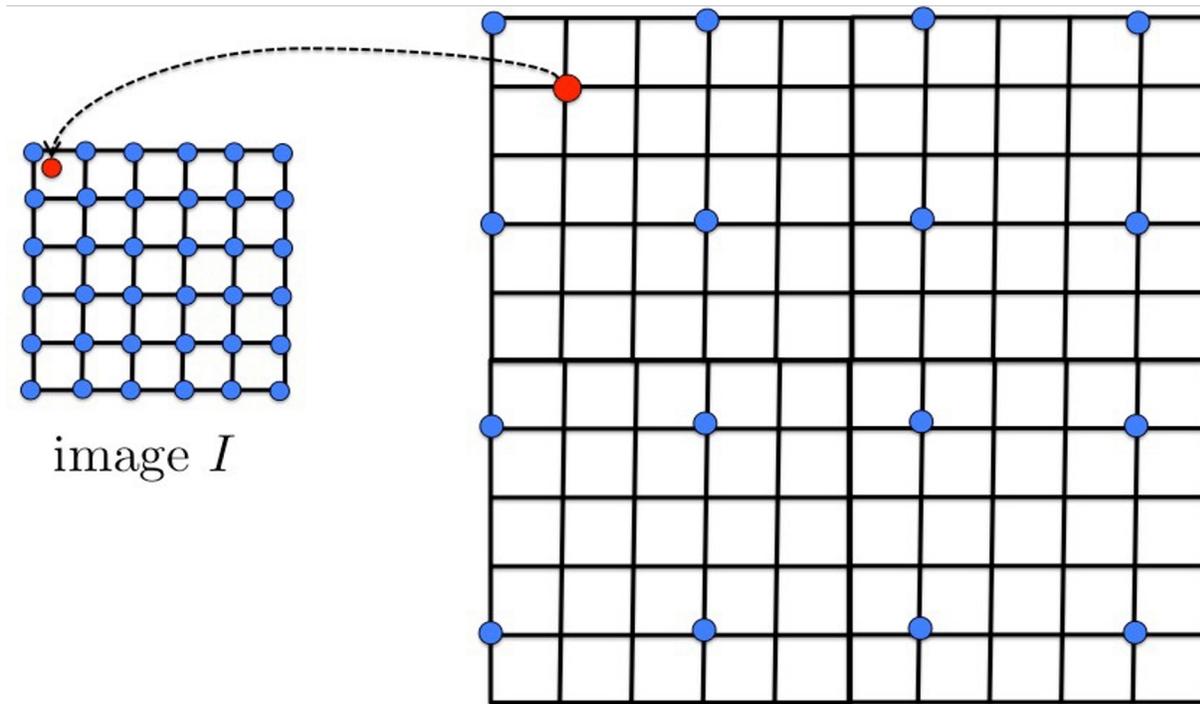


image  $I$



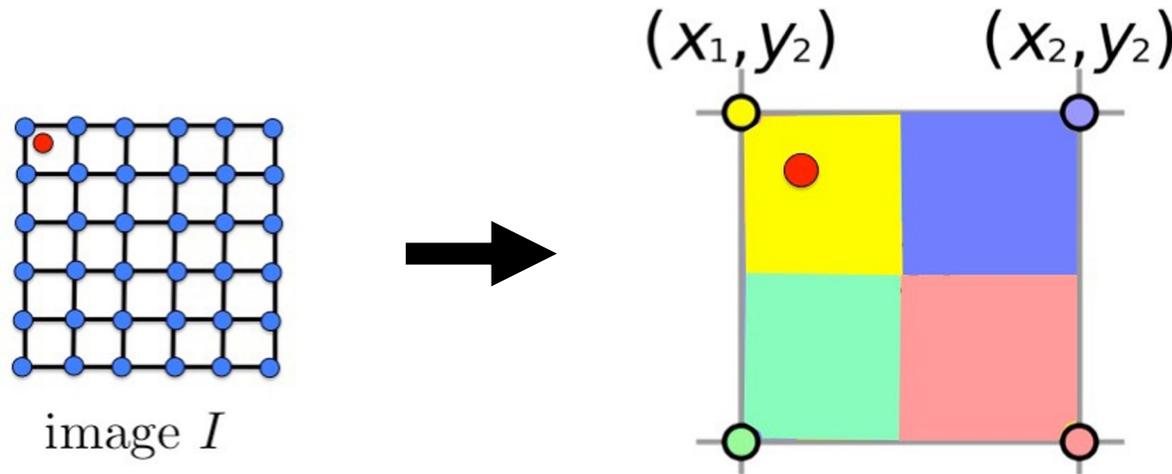
- Let's make this image triple size
- Copy image in every third pixel.

# Image Interpolation (2D)



- Let's make this image triple size
- Copy image in every third pixel.
- What about the remaining pixels in  $G$ ? (how would we compute this value?)

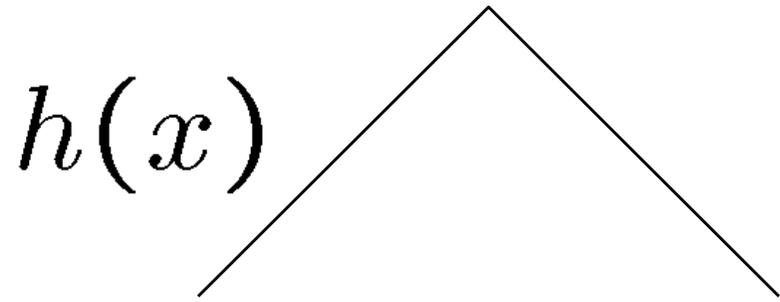
# Image Interpolation (2D)



- Let's make this image triple size
- Copy image in every third pixel.
- What about the remaining pixels in  $G$ ? (how would we compute this value?)
  - *bilinear interpolation (linear interpolation in  $x$  and  $y$ , resulting in quadratic interpolation)*
  - Check out details: [http://en.wikipedia.org/wiki/Bilinear\\_interpolation](http://en.wikipedia.org/wiki/Bilinear_interpolation)

# Reconstruction Filters

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

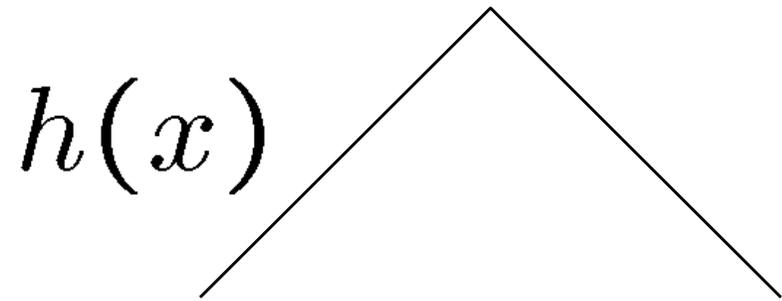


Performs Linear Interpolation

$h(x, y)$  ?

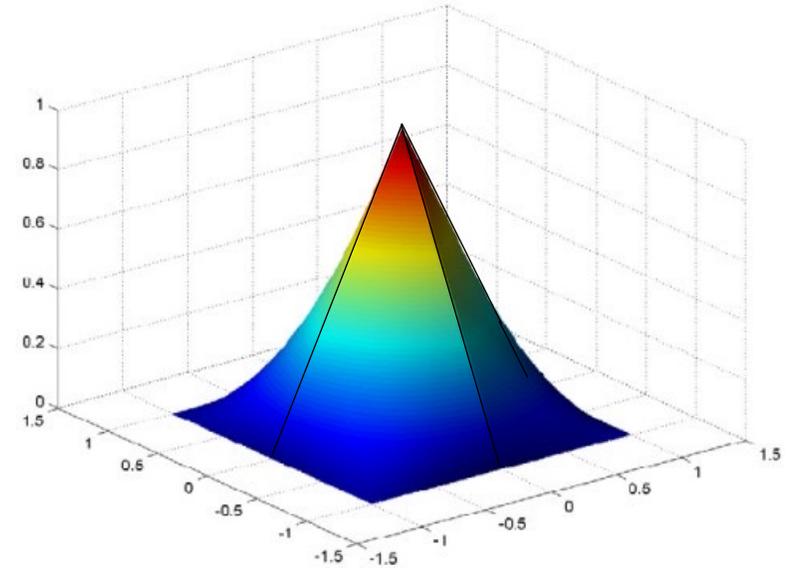
# Reconstruction Filters

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



Performs Linear Interpolation

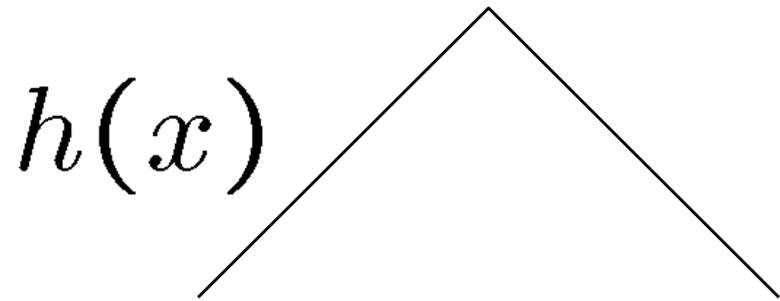
$h(x, y)$



(tent function) Performs  
bilinear Interpolation

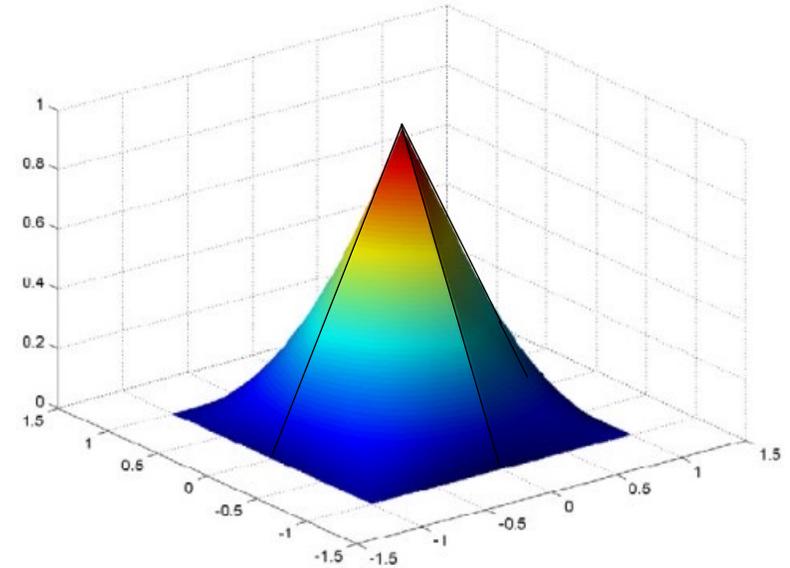
# Reconstruction Filters

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



Performs Linear Interpolation

$h(x, y)$

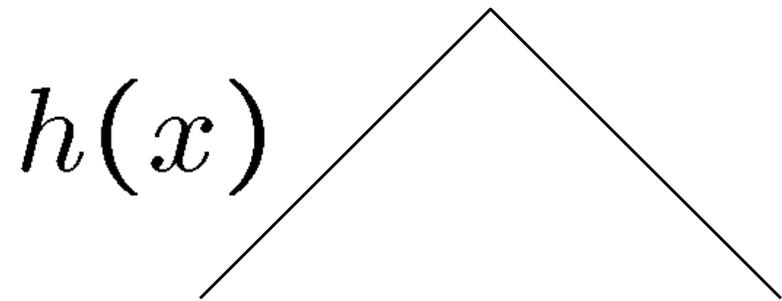


(tent function) Performs  
bilinear Interpolation

- And filter for nearest neighbor interpolation?

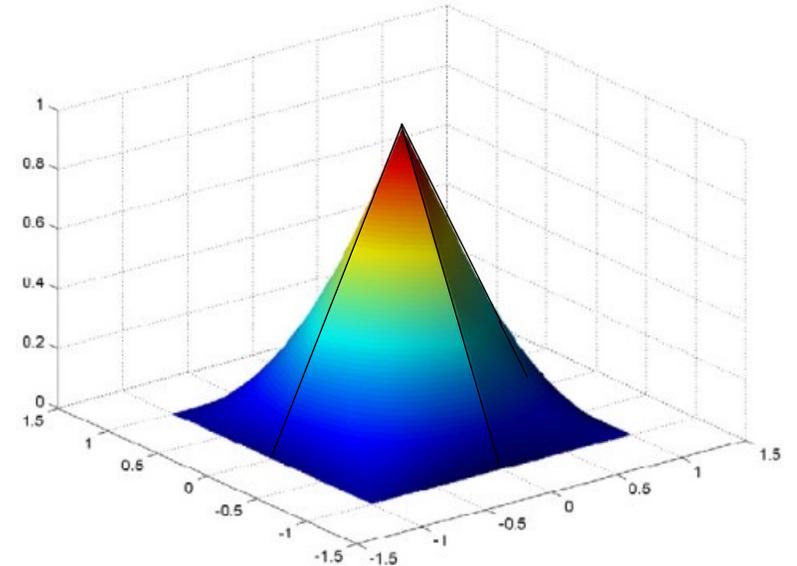
# Reconstruction Filters

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



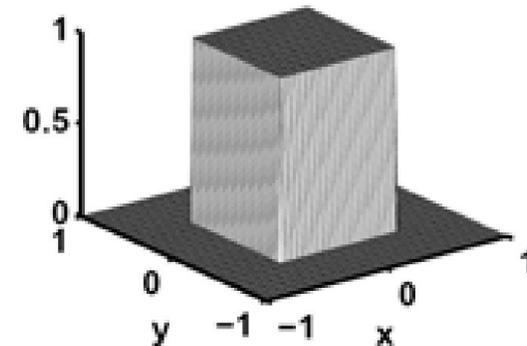
Performs Linear Interpolation

$h(x, y)$



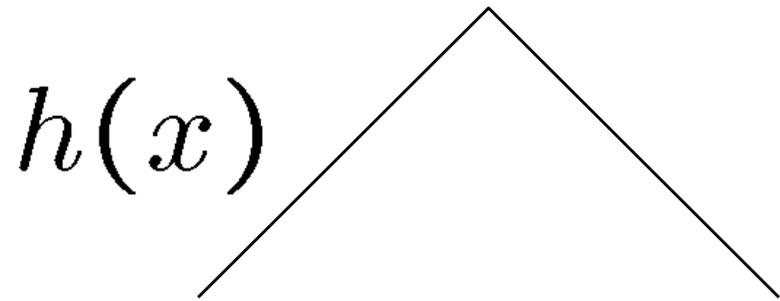
(tent function) Performs bilinear Interpolation

- And filter for nearest neighbor interpolation?



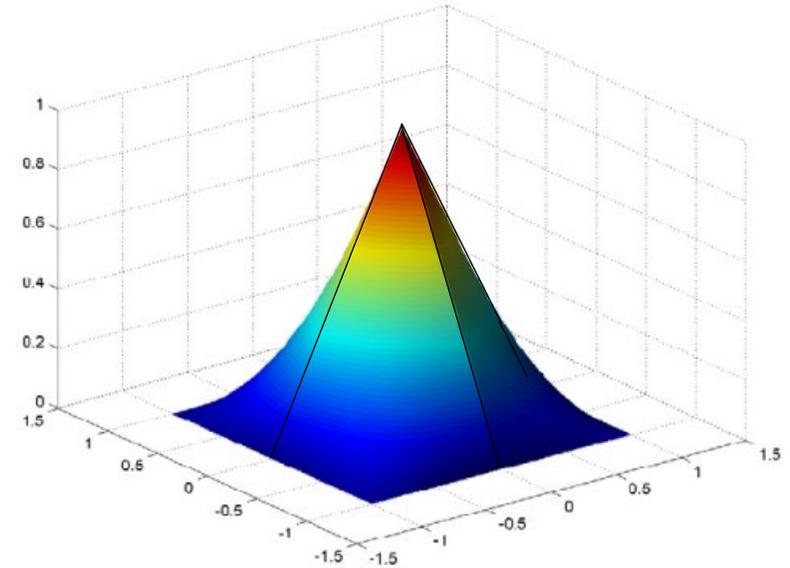
# Reconstruction Filters

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



Performs Linear Interpolation

$h(x, y)$



(tent function) Performs  
bilinear Interpolation

- higher order filters can give better resampled images: Bicubic is a common choice

# Image Interpolation via Convolution (2D)

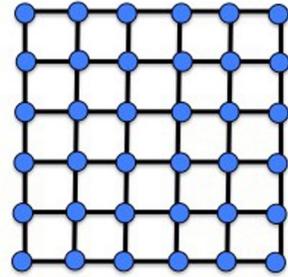
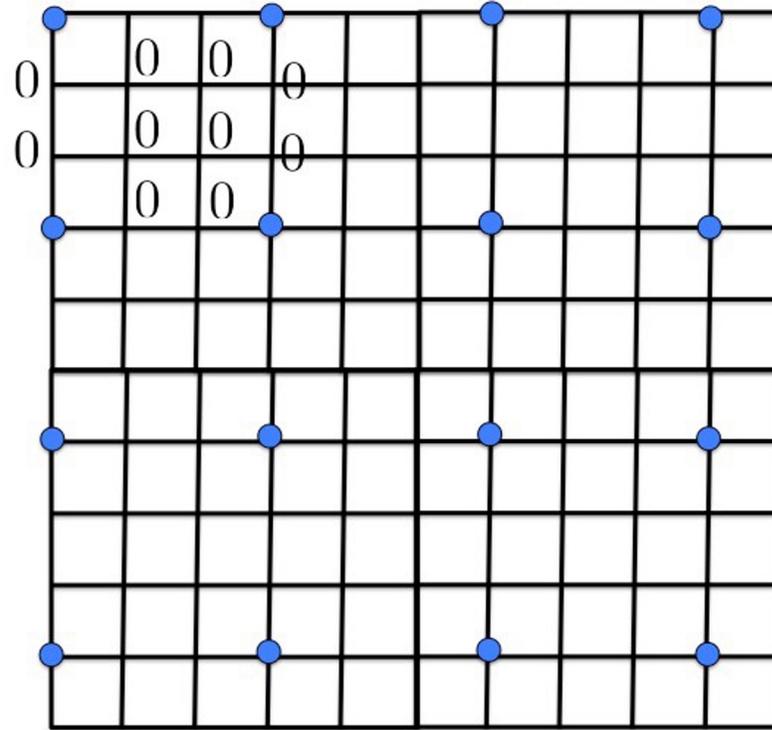


image  $I$



- Let's make this image triple size: copy image values in every third pixel, place zeros everywhere else
- what filter do we use?

# Image Interpolation via Convolution (2D)

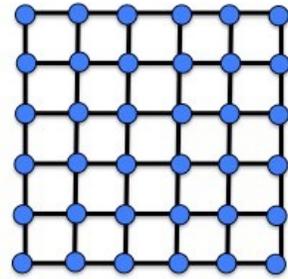
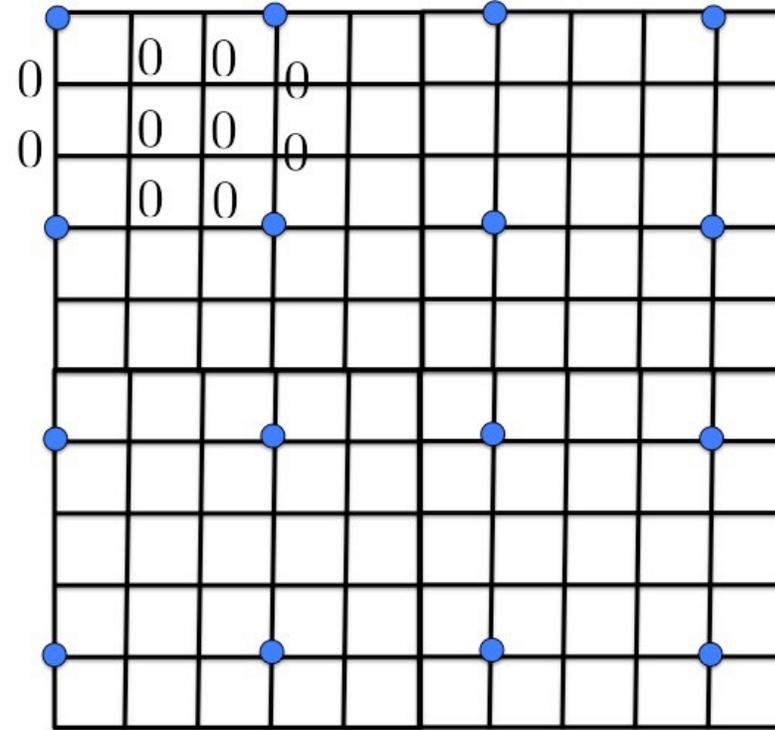
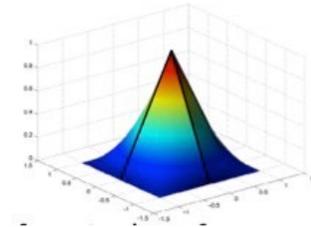


image  $I$



$* h(x, y)$



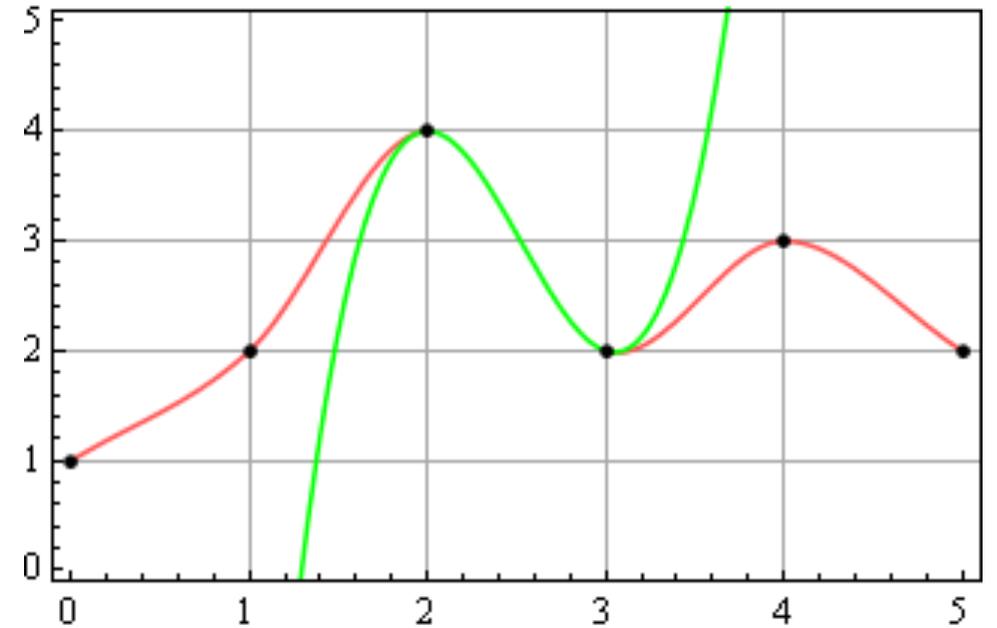
- Let's make this image triple size: copy image values in every third pixel, place zeros everywhere else
- Convolution with a reconstruction filter (e.g., bilinear) and you get the interpolated image

# Cubic interpolation

- Uses third degree polynomial

$$f(x) = ax^3 + bx^2 + cx + d$$

$$f'(x) = 3ax^2 + 2bx + c$$



# Cubic interpolation

- Uses third degree polynomial

$$f(x) = ax^3 + bx^2 + cx + d$$

$$f'(x) = 3ax^2 + 2bx + c$$

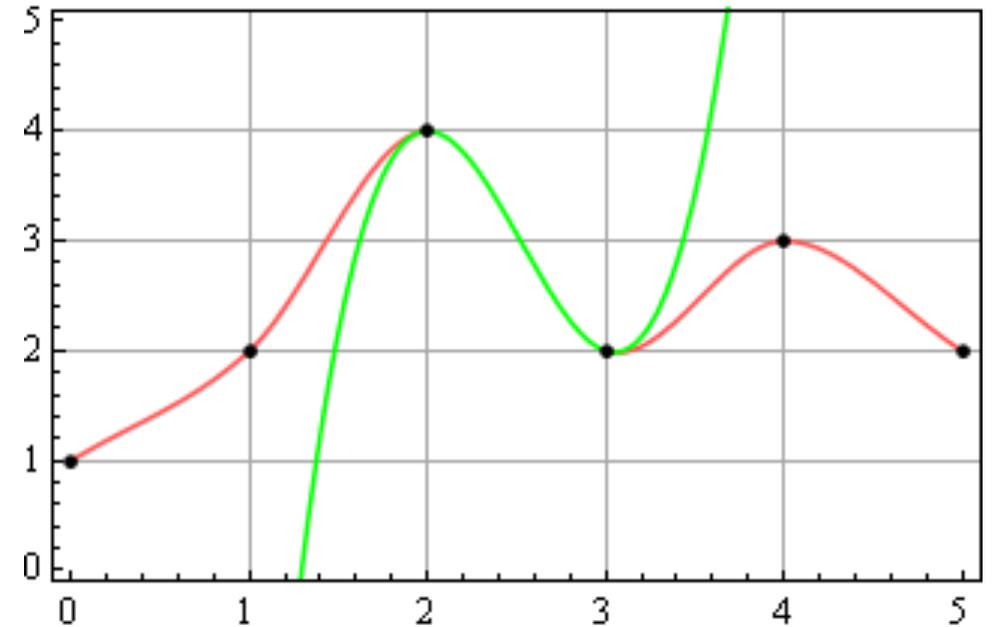
- plug in values at  $x=0$  and  $x=1$  (pick two points)

$$f(0) = d$$

$$f(1) = a + b + c + d$$

$$f'(0) = c$$

$$f'(1) = 3a + 2b + c$$



# Cubic interpolation

- Uses third degree polynomial

$$f(x) = ax^3 + bx^2 + cx + d$$

$$f'(x) = 3ax^2 + 2bx + c$$

- plug in values at  $x=0$  and  $x=1$  (pick two points)

$$f(0) = d$$

$$f(1) = a + b + c + d$$

$$f'(0) = c$$

$$f'(1) = 3a + 2b + c$$

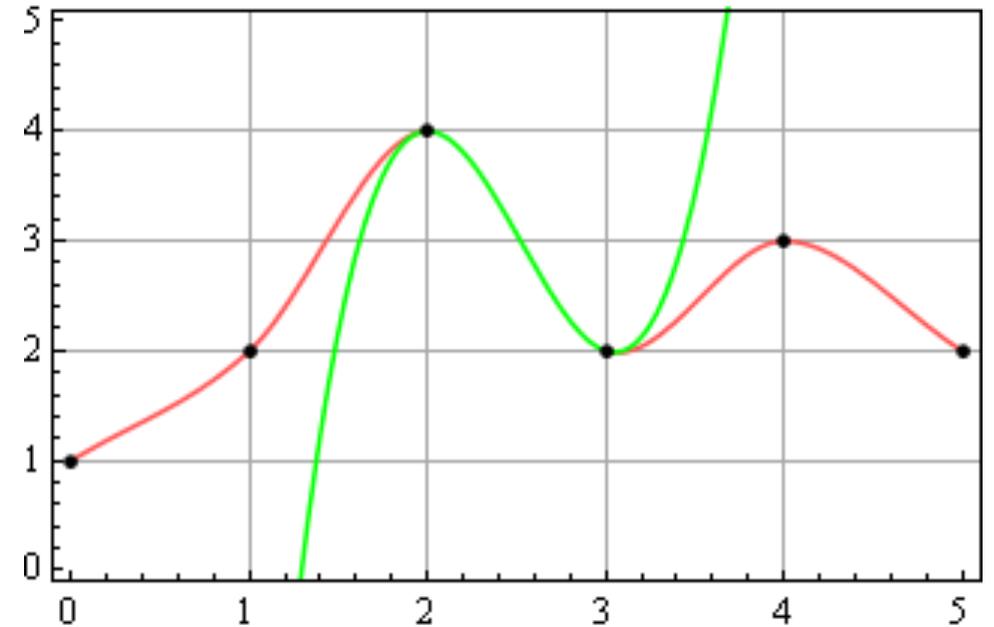
- rewrite the equations as

$$a = 2f(0) - 2f(1) + f'(0) + f'(1)$$

$$b = -3f(0) + 3f(1) - 2f'(0) - f'(1)$$

$$c = f'(0)$$

$$d = f(0)$$



# Cubic interpolation

- rewrite the equations as

$$a = 2f(0) - 2f(1) + f'(0) + f'(1)$$

$$b = -3f(0) + 3f(1) - 2f'(0) - f'(1)$$

$$c = f'(0)$$

$$d = f(0)$$

- plug in known values for  $f$

$$f(0) = p_0$$

$$f(1) = p_1$$

$$f'(0) = \frac{p_2 - p_0}{2}$$

$$f'(1) = \frac{p_3 - p_1}{2}$$

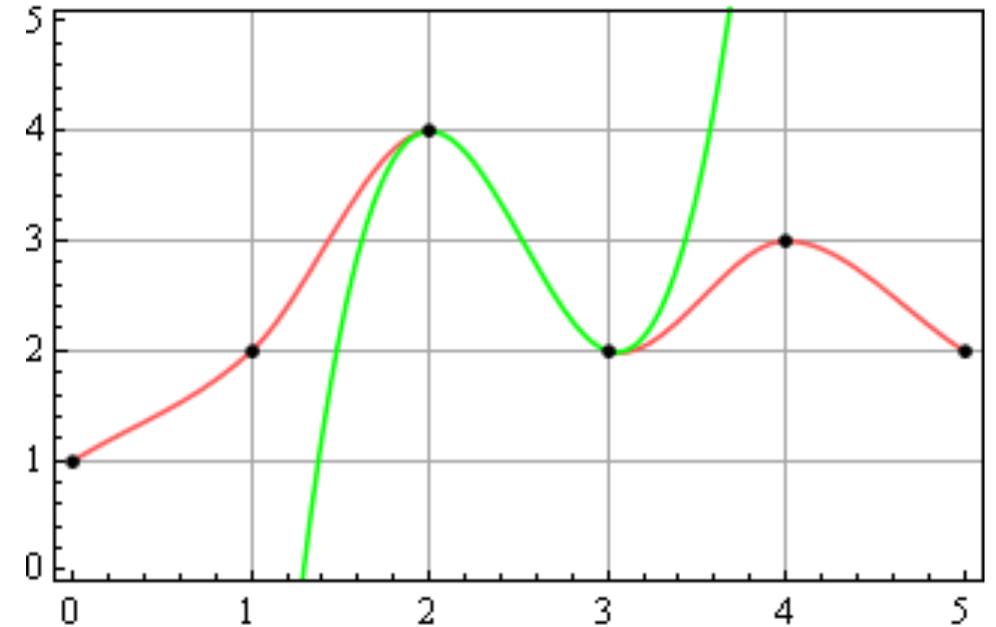
- plug back in!

$$a = -\frac{1}{2}p_0 + \frac{3}{2}p_1 - \frac{3}{2}p_2 + \frac{1}{2}p_3$$

$$c = -\frac{1}{2}p_0 + \frac{1}{2}p_2$$

$$b = p_0 - \frac{5}{2}p_1 + 2p_2 - \frac{1}{2}p_3$$

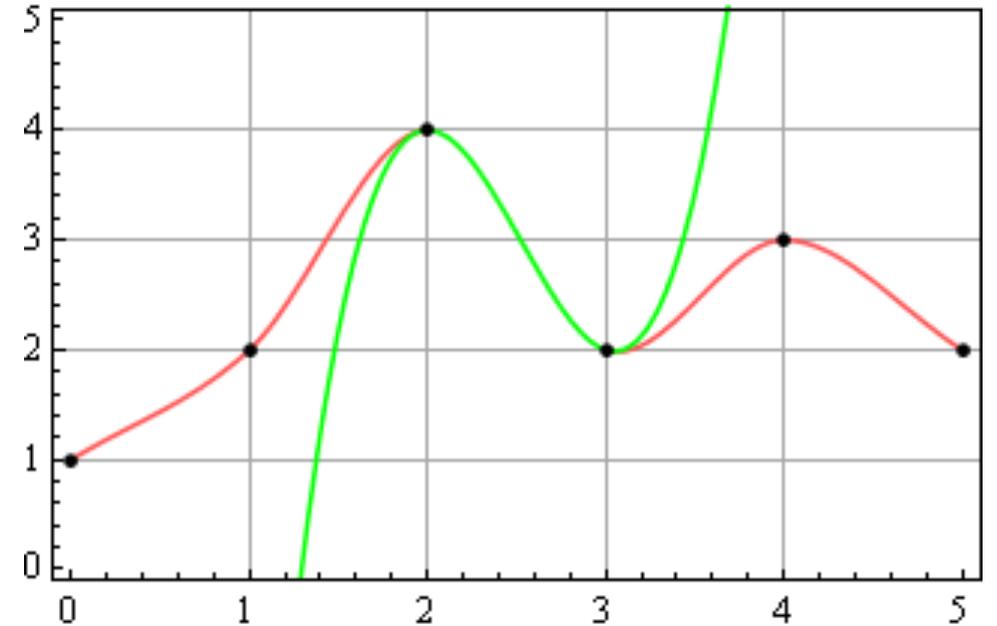
$$d = p_0$$



# Cubic interpolation

The whole formula looks like:

$$f(p_0, p_1, p_2, p_3, x) = \left( -\frac{1}{2}p_0 + \frac{3}{2}p_1 - \frac{3}{2}p_2 + \frac{1}{2}p_3 \right) x^3 \\ + \left( p_0 - \frac{5}{2}p_1 + 2p_2 - \frac{1}{2}p_3 \right) x^2 + \left( -\frac{1}{2}p_0 + \frac{1}{2}p_2 \right) x + p_1$$



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

# Bicubic interpolation

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

# Bicubic interpolation

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



Input image

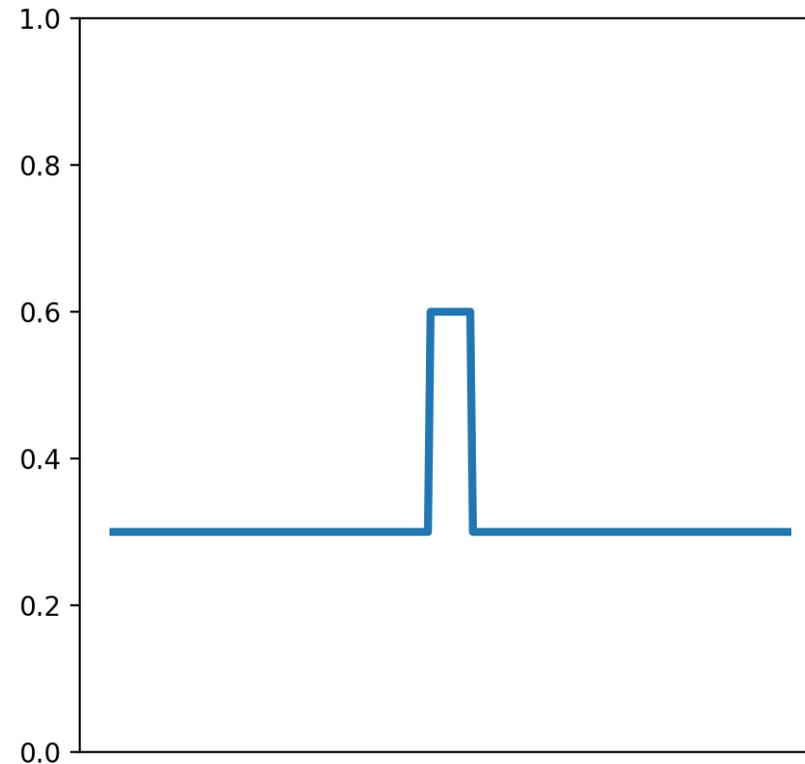


image slice

# Bicubic interpolation

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



Input image

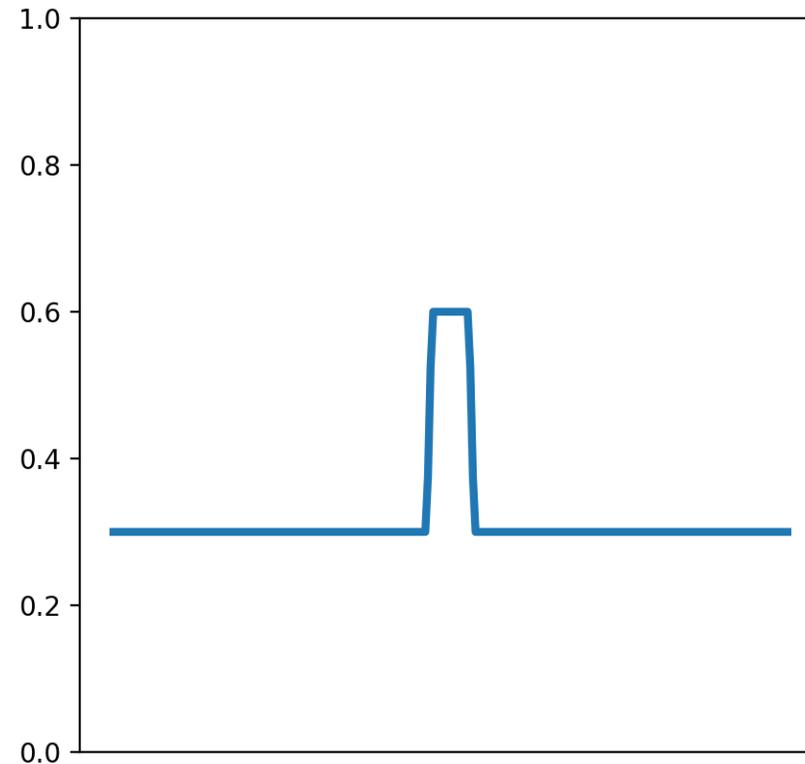
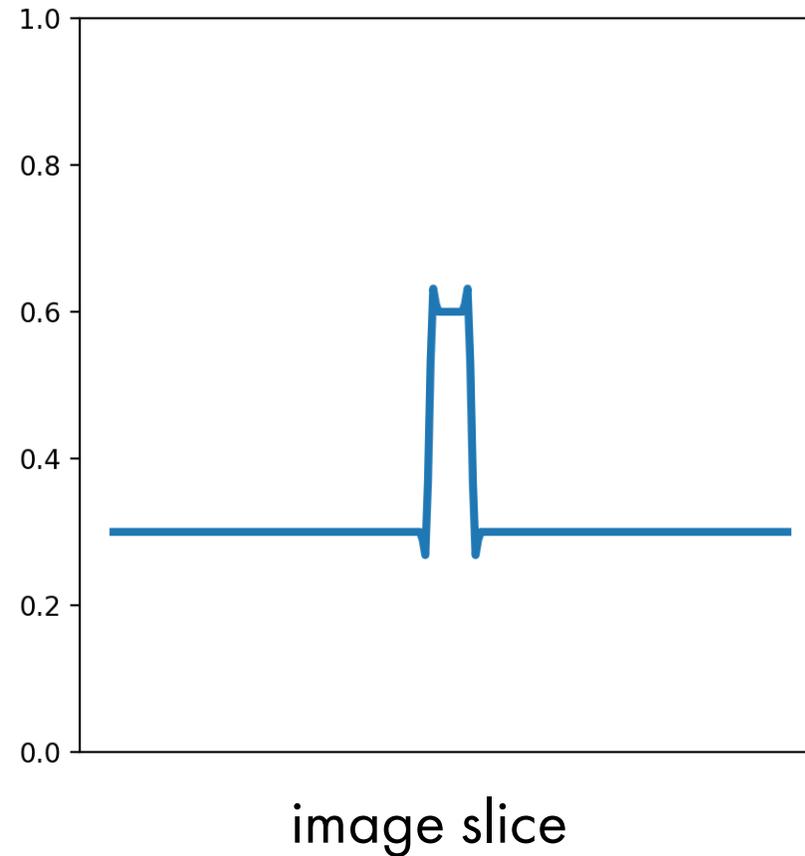
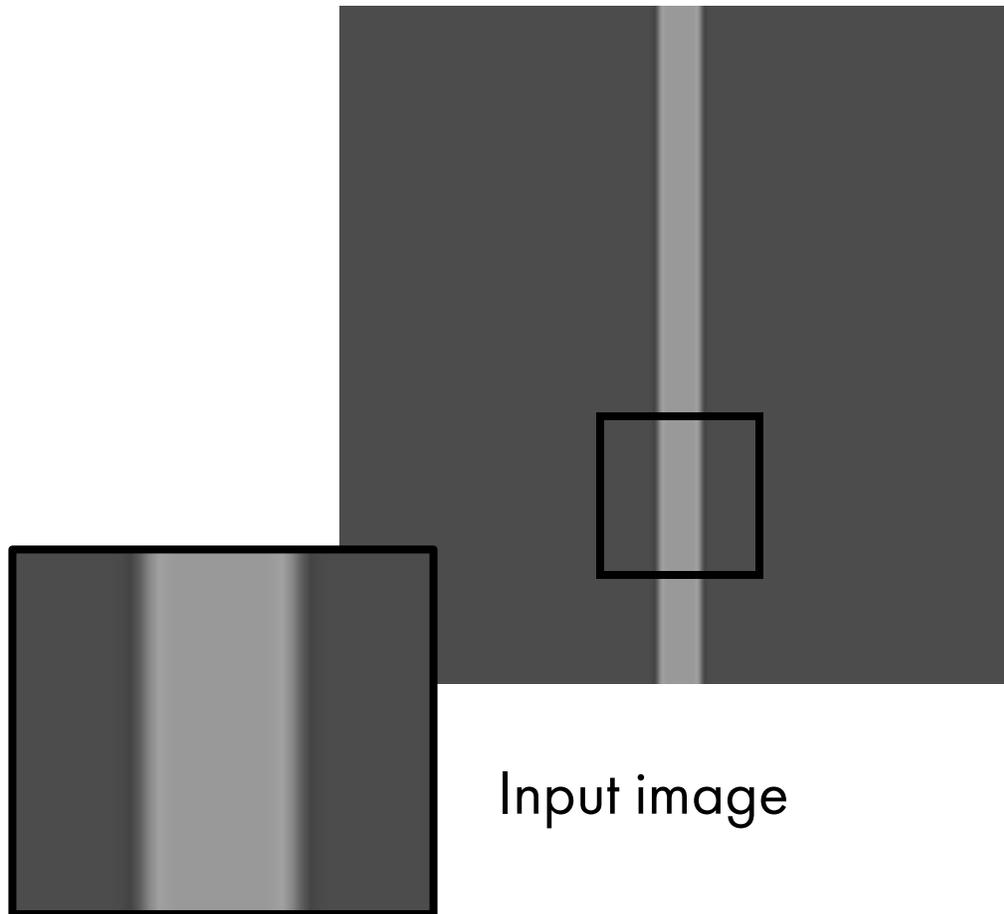


image slice

# Bicubic interpolation

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



# Image Interpolation

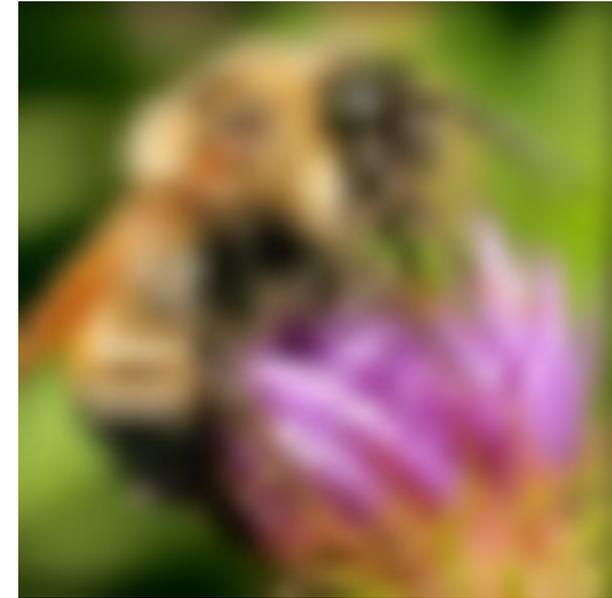
Original image



Nearest-Neighbor  
Interpolation



Bilinear interpolation



Bicubic interpolation

# Summary – Stuff You Should Know

- To down-scale an image: blur it with a small Gaussian (e.g.,  $\sigma = 1.4$ ) 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