Lecture 4: Sampling and Convolution

Yuhao Zhu

http://yuhaozhu.com
yzhu@rochester.edu

CSC 292/572, Fall 2022
Mobile Visual Computing

Logistics

- Written assignment 1 is up and is due Sept. 16 11:30 AM.
- You can work in groups of 2.
- Course schedule: https://www.cs.rochester.edu/courses/572/fall2022/schedule.html. You will find reading assignments and slides.
- Start thinking and talking to me about your final project idea.

The Roadmap

Theoretical Preliminaries

Human Visual Systems

Color in Nature, Arts, Tech
(a.k.a., the birth, life, and death of light)

Digital Camera Imaging

Modeling and Rendering

Applications

Geometric Transformations

Fourier Series & Transforms

Sampling & Reconstruction

Signal Sampling and Reconstruction

- Given just a few sparse samples, can we always reconstruct the underlying continuous signal?

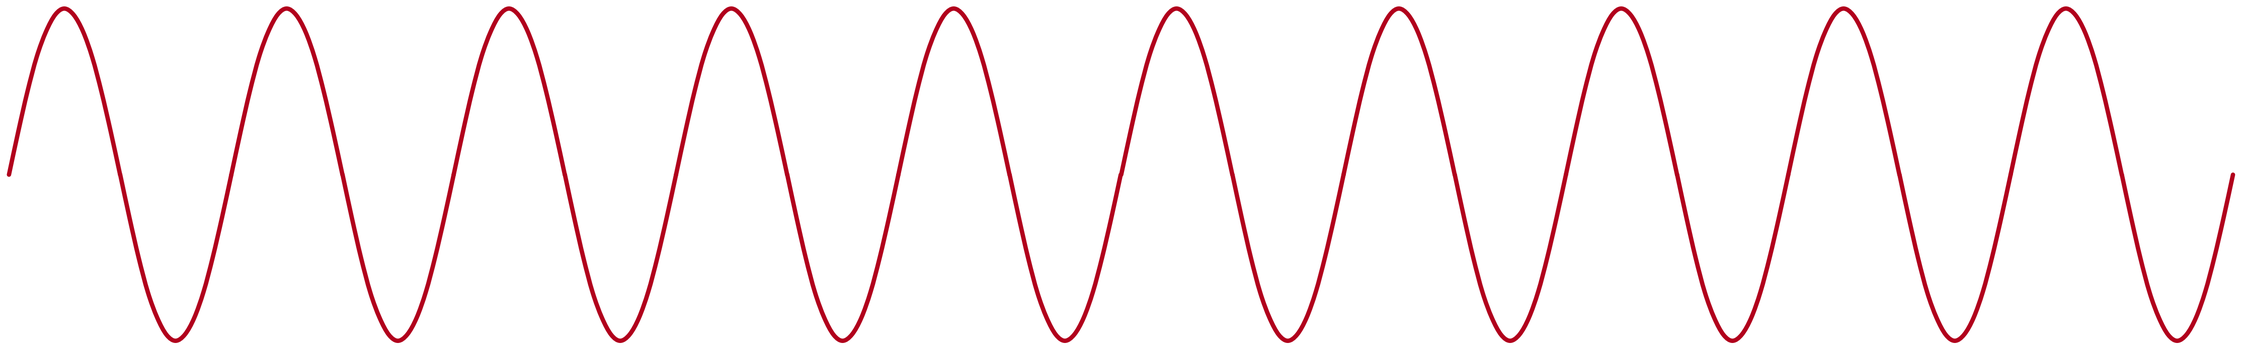

Signal Sampling and Reconstruction

- Given just a few sparse samples, can we always reconstruct the underlying continuous signal?

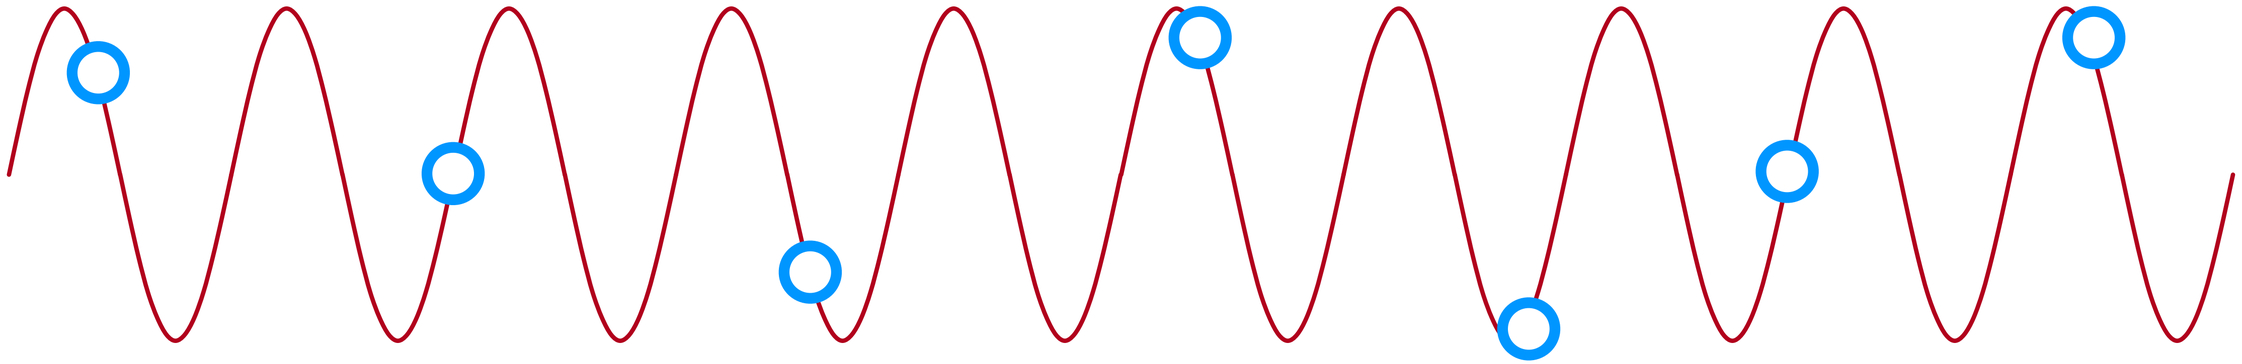

Signal Sampling and Reconstruction

- Given just a few sparse samples, can we always reconstruct the underlying continuous signal?

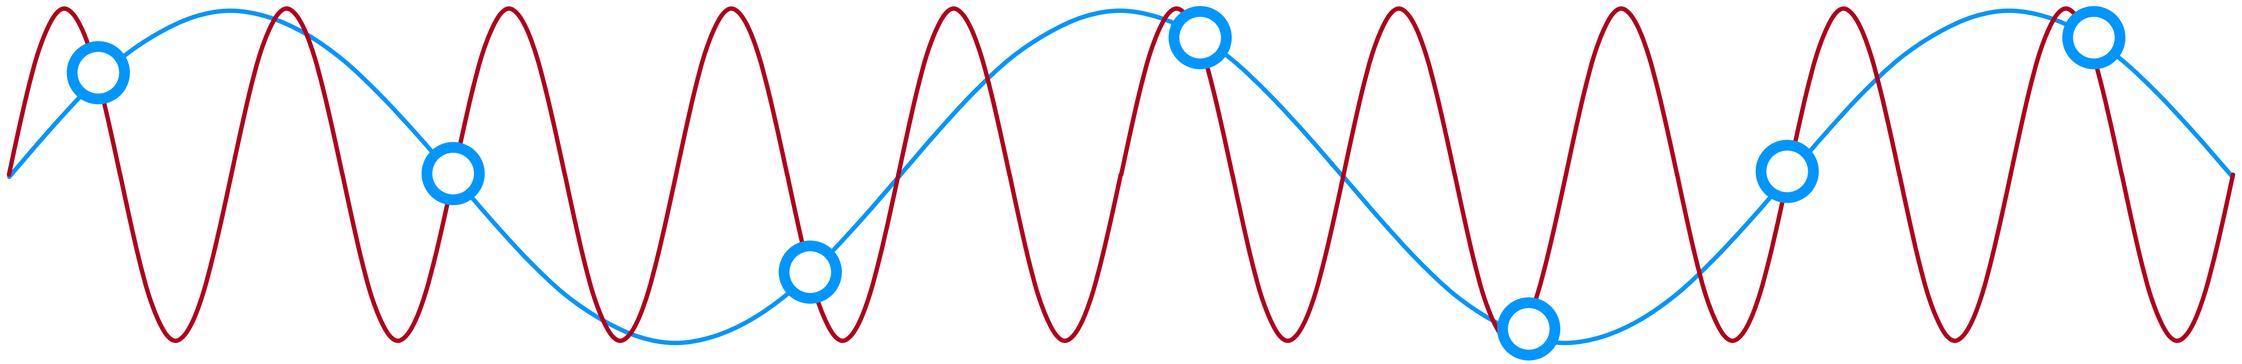

Signal Sampling and Reconstruction

- Given just a few sparse samples, can we always reconstruct the underlying continuous signal?
- If not carefully sampled, reconstructed signals will be aliased: high frequencies signals masquerading as low-frequency signals.

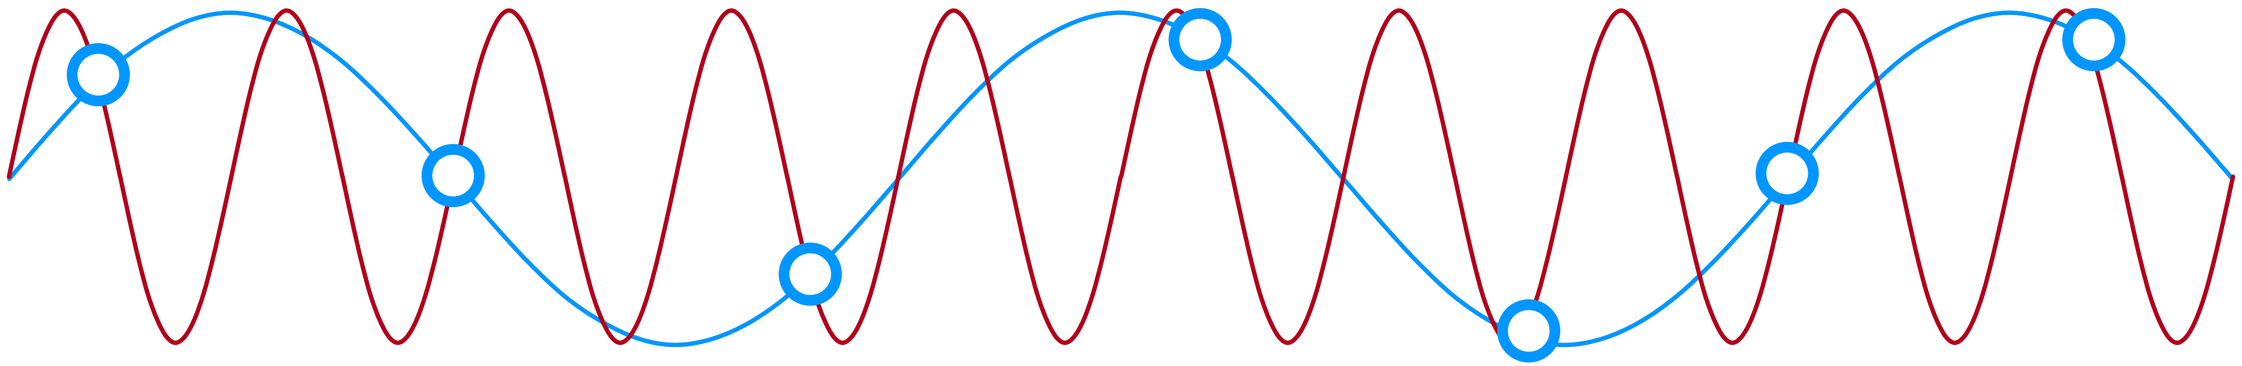

Aliasing

https://svi.nl/AntiAliasing https://distance.ufhealth.org/preparing-for-lecture-capture/ https://www.youtube.com/watch?v=Neh2biiex1A

Fourier Theory of Sampling and Reconstruction

Impulse Train Function

Impuse/Delta/Dirac Delta Function

$$\delta(t) = \begin{cases} 0, & t \neq 0, \\ +\infty, & t = 0. \end{cases} \qquad \int_{-\infty}^{\infty} \delta(t) = 1$$

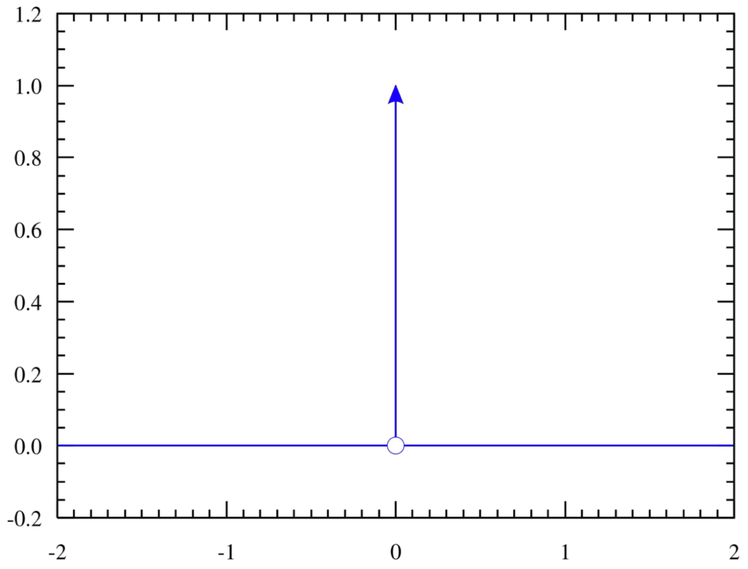

https://en.wikipedia.org/wiki/Dirac_comb

Impulse Train Function

Impuse/Delta/Dirac Delta Function

$$\delta(t) = \begin{cases} 0, & t \neq 0, \\ +\infty, & t = 0. \end{cases} \qquad \int_{-\infty}^{\infty} \delta(t) = 1$$

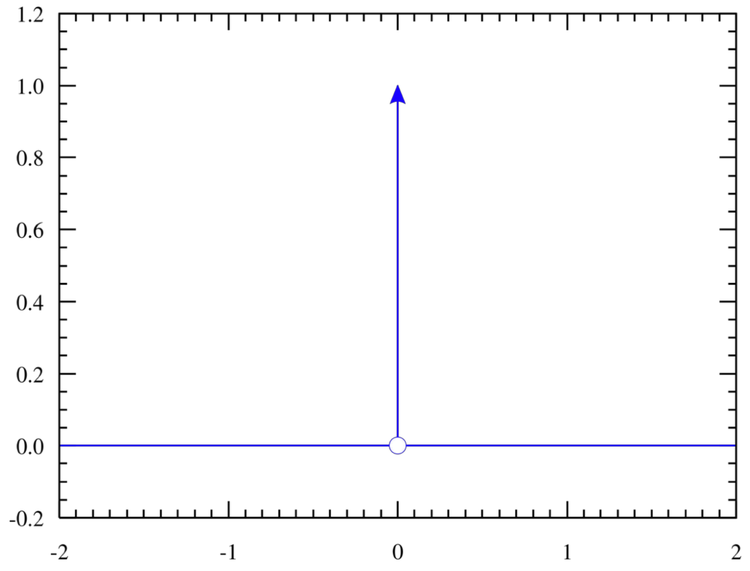

Impulse Train/Dirac Comb/Shah Function (with a period of T and frequency of 1/T)

$$III(t) = \sum_{i=-\infty}^{\infty} \delta(t - iT)$$

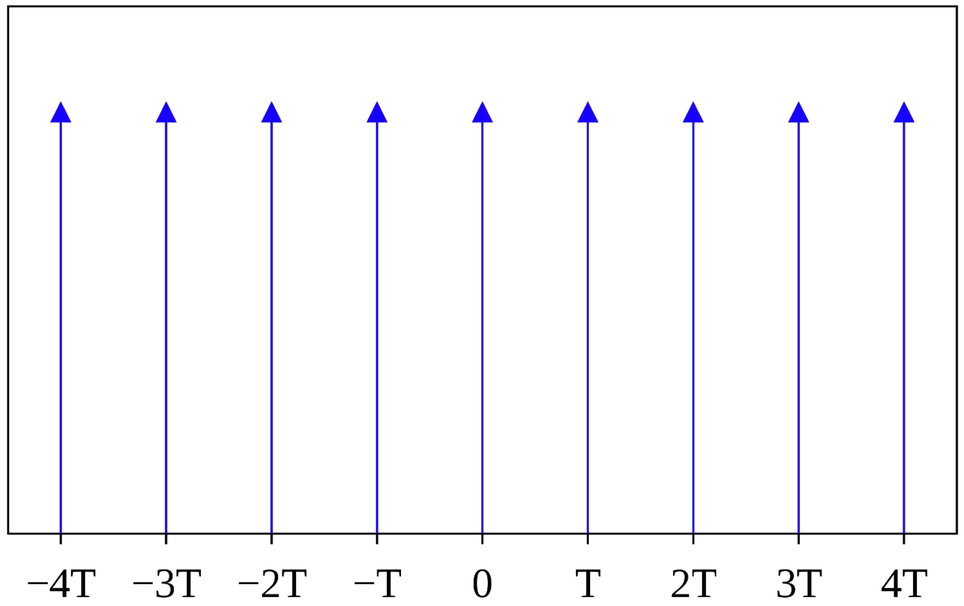

https://en.wikipedia.org/wiki/Dirac_comb

Mathematically Express Sampling

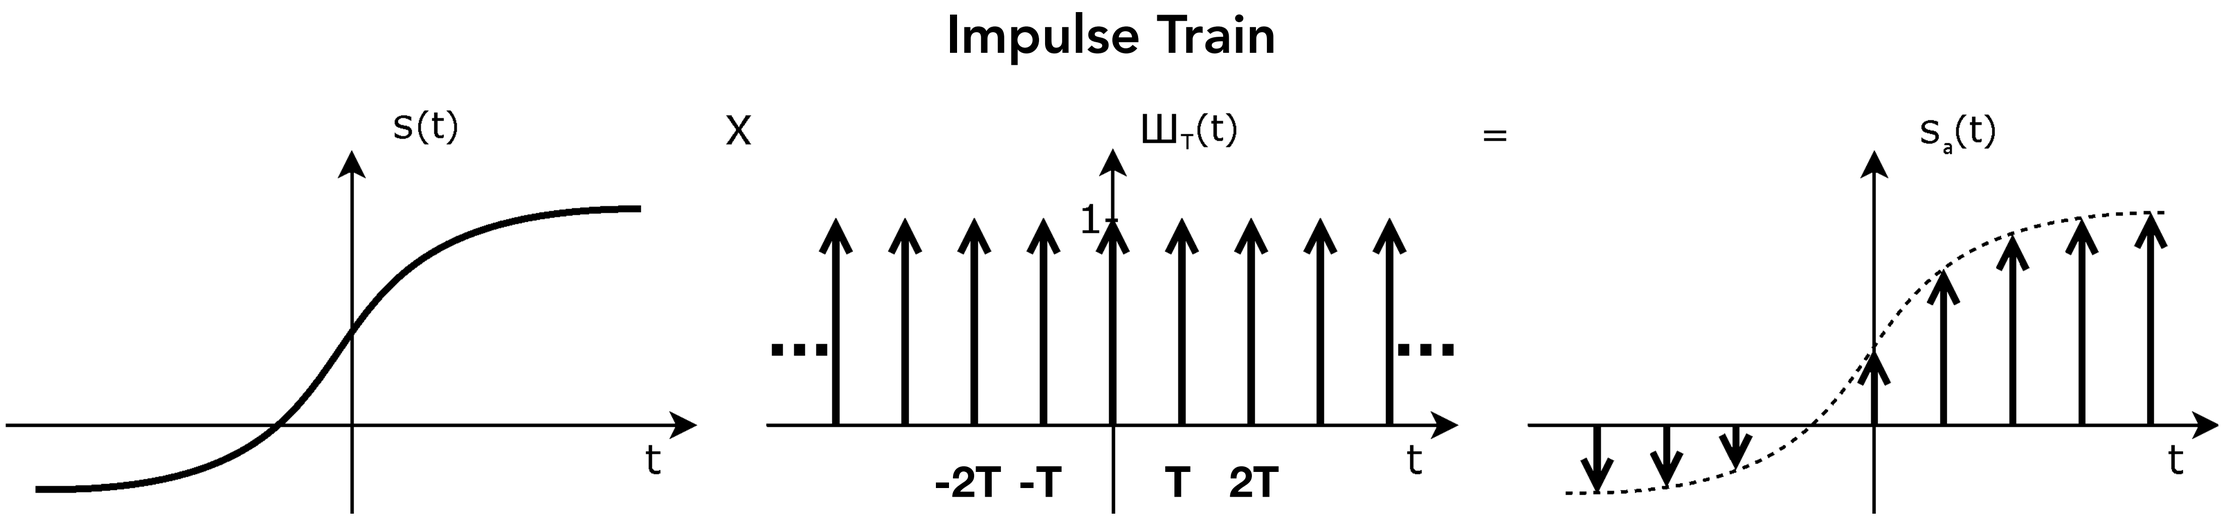

$$s_a(t) = III_T(t)s(t)$$

$S_a(t)$ is defined over a continuous domain; its value is 0 except at integer multiples of **T**.

Fourier Transform a Sampled Signal

$$F(\omega) = \mathcal{F}(s(t))$$

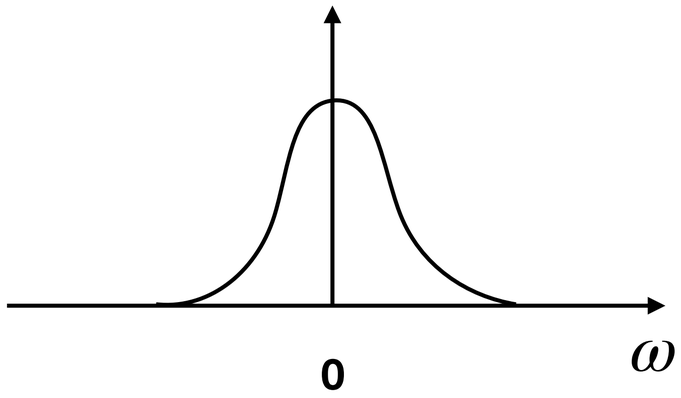

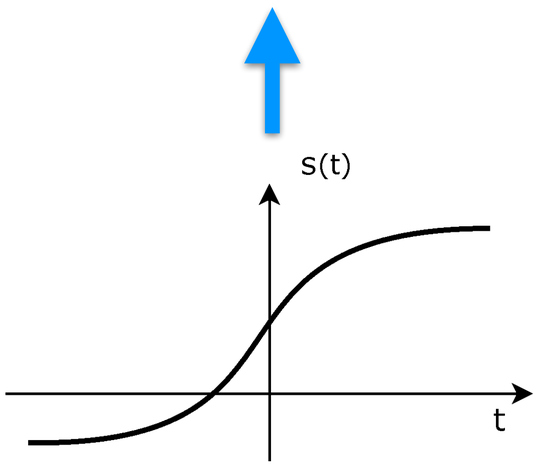

Fourier Transform a Sampled Signal

$$F(\omega) = \mathcal{F}(s(t))$$

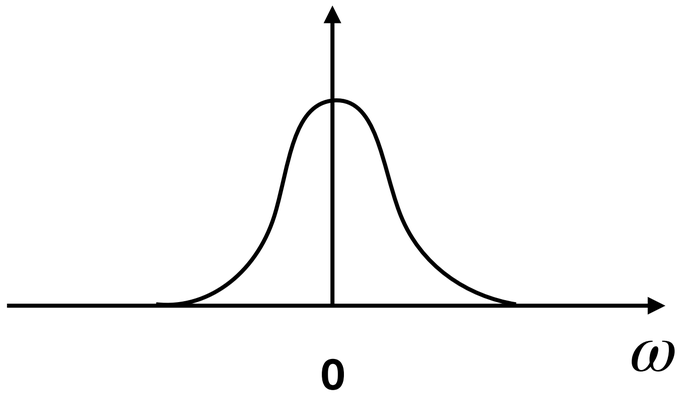

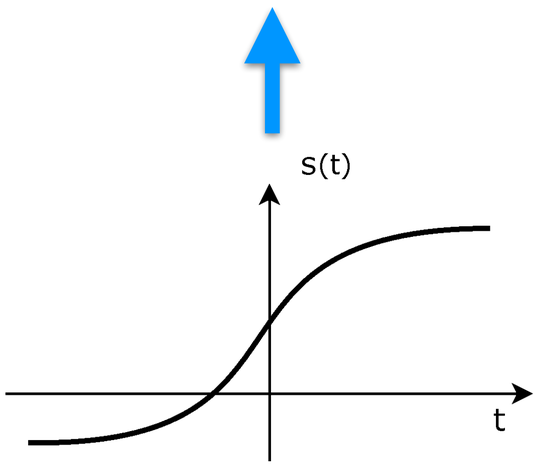

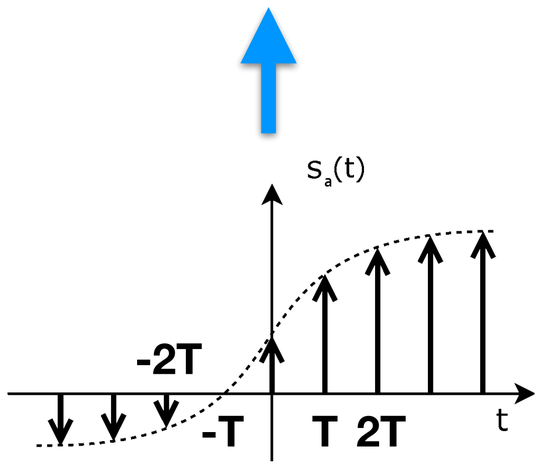

Fourier Transform a Sampled Signal

$$F(\omega) = \mathcal{F}(s(t))$$

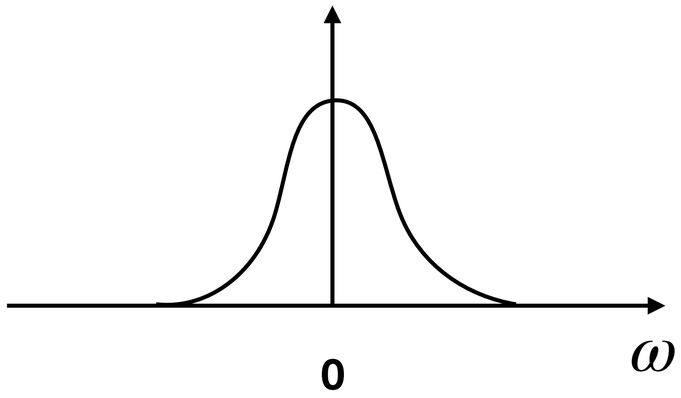

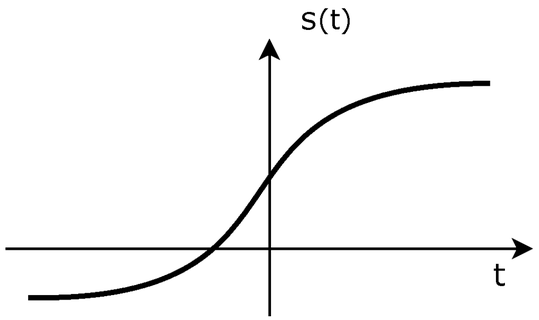

$$\mathcal{F}(III_T(t)s(t))$$

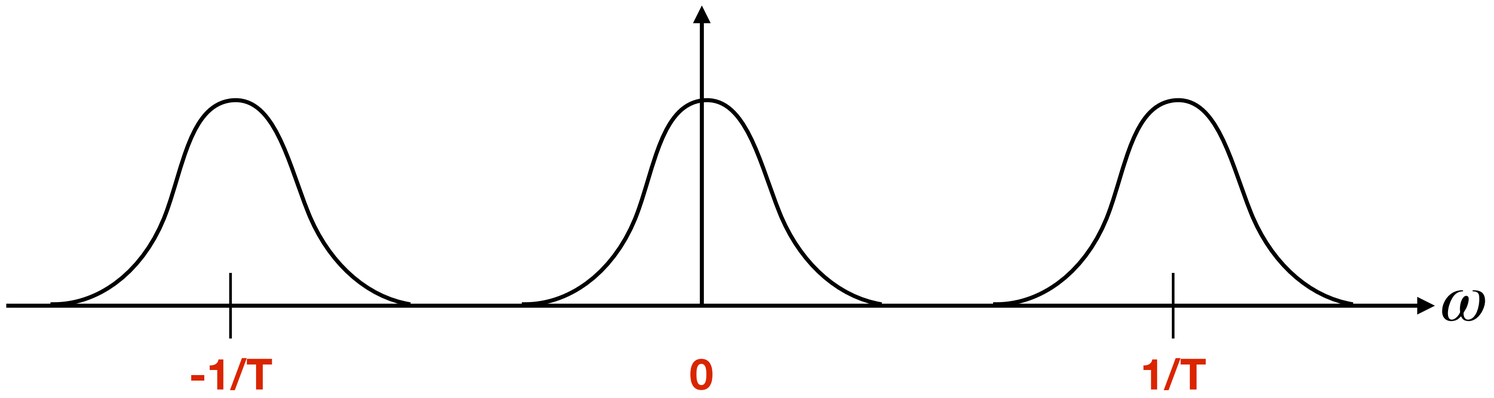

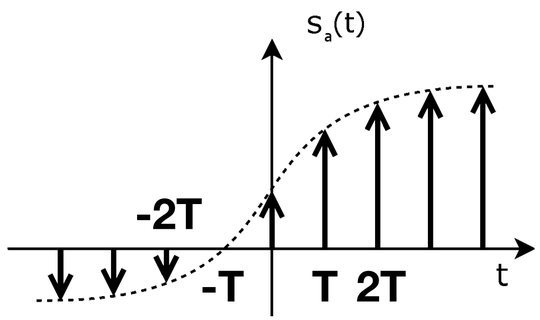

Fourier Transform a Sampled Signal

$$F(\omega) = \mathcal{F}(s(t))$$

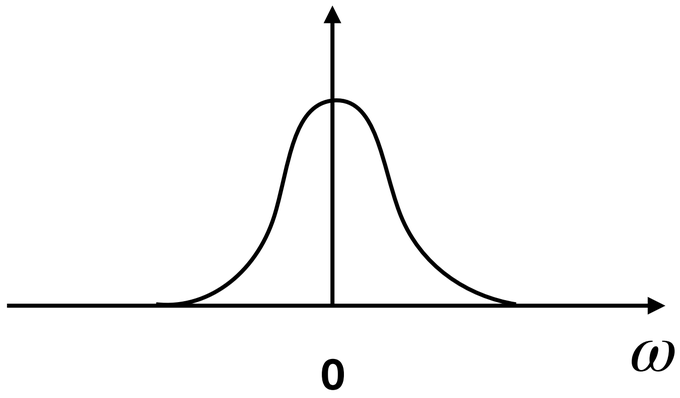

$$\mathcal{F}(III_T(t)s(t))$$

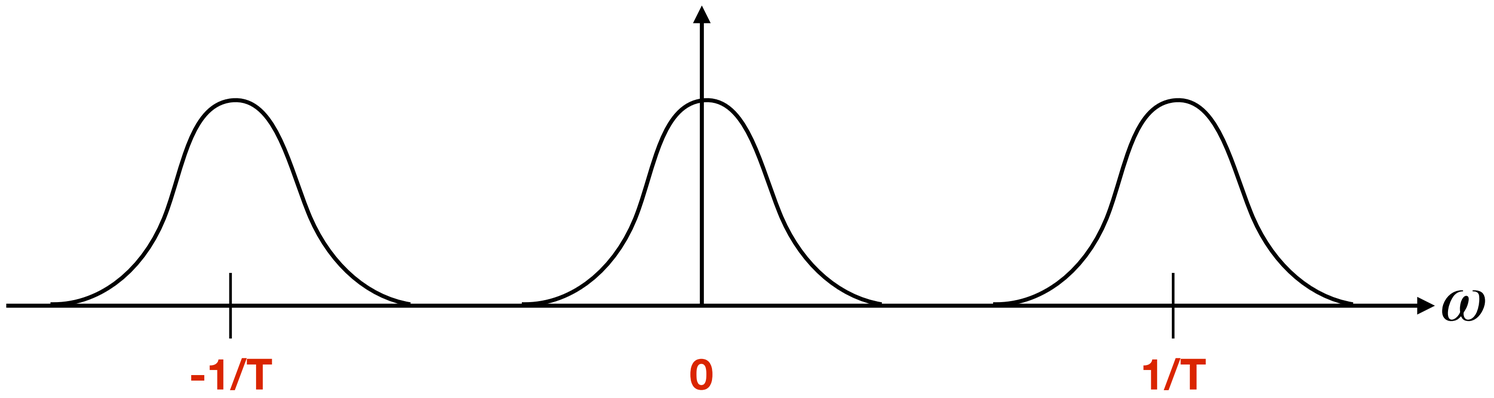

$$\mathcal{F}(III_T(t)s(t)) = \sum_{i=-\infty}^{\infty} F(\omega - i\frac{1}{T})$$

The Fourier Transform of the sampled signal is an infinite sequence of the Fourier Transform of the original signal *s* with a period of **1/T.**

Extracting the Original Spectrum

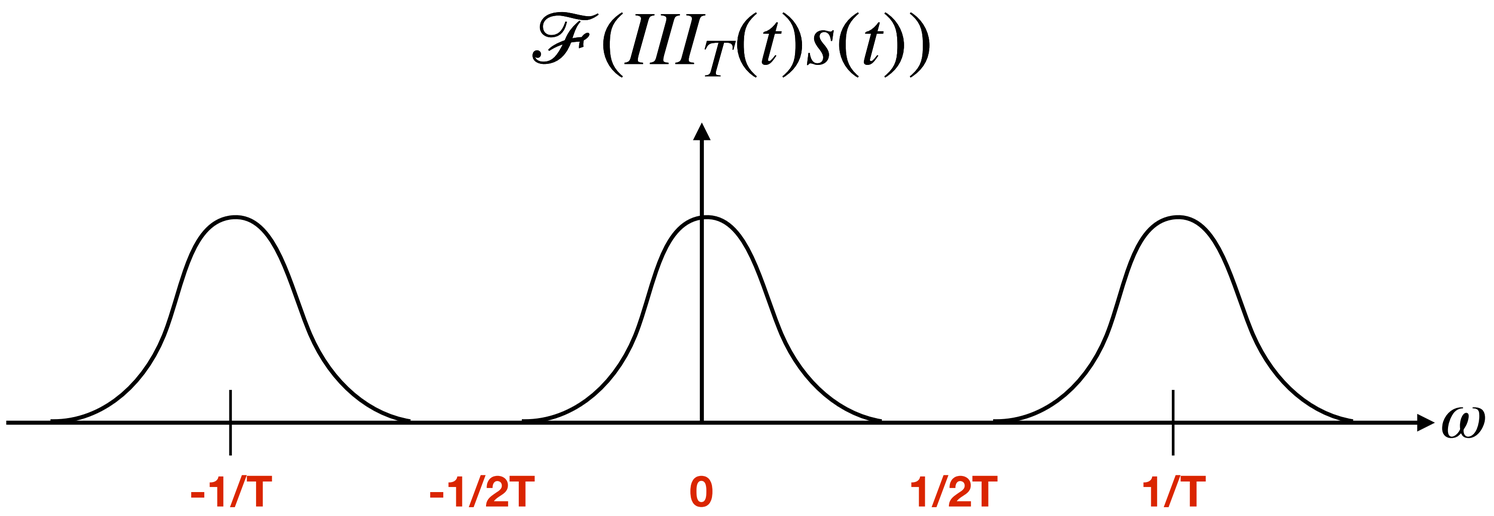

Extracting the Original Spectrum

Box Function

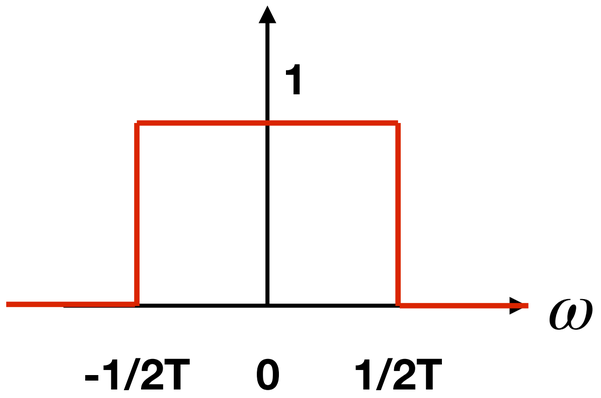

$\mathcal{F}(III_T(t)s(t))$

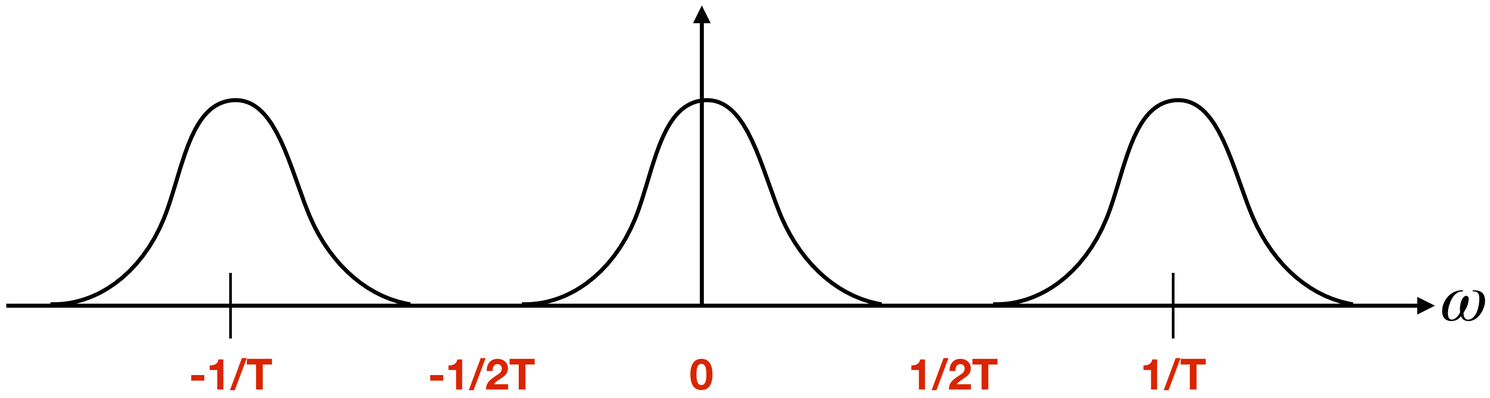

Extracting the Original Spectrum

Box Function

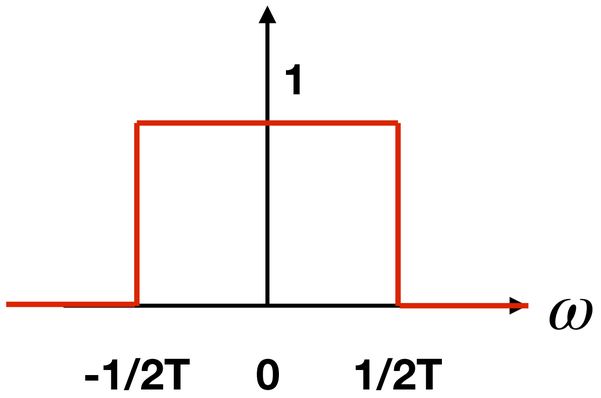

$\mathcal{F}(III_T(t)s(t))$

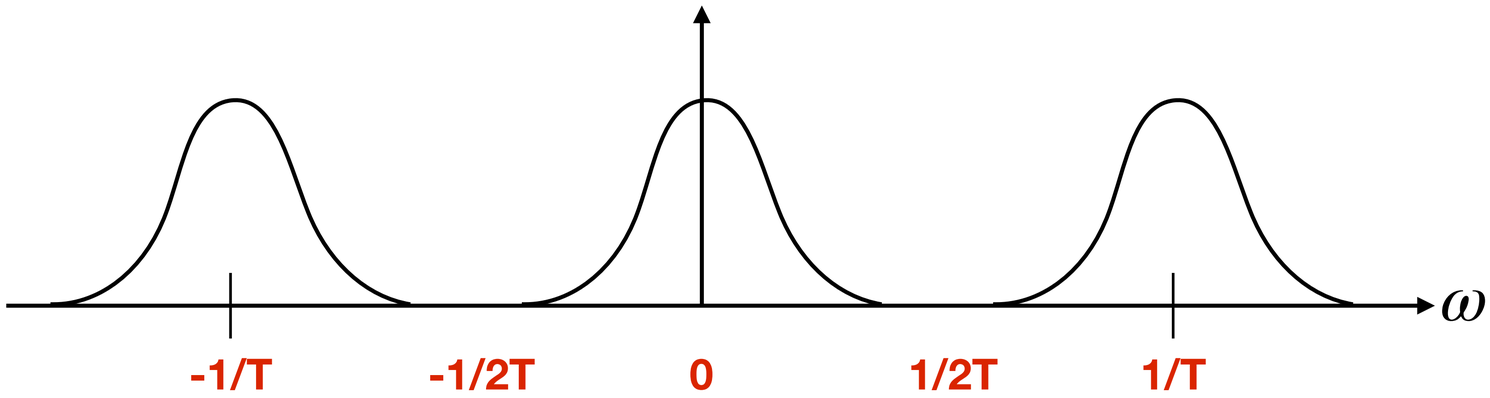

$$B(\omega) = \begin{cases} 1, & |\omega| < \frac{1}{2T}, \\ 0, & \textbf{otherwise}. \end{cases}$$

Extracting the Original Spectrum

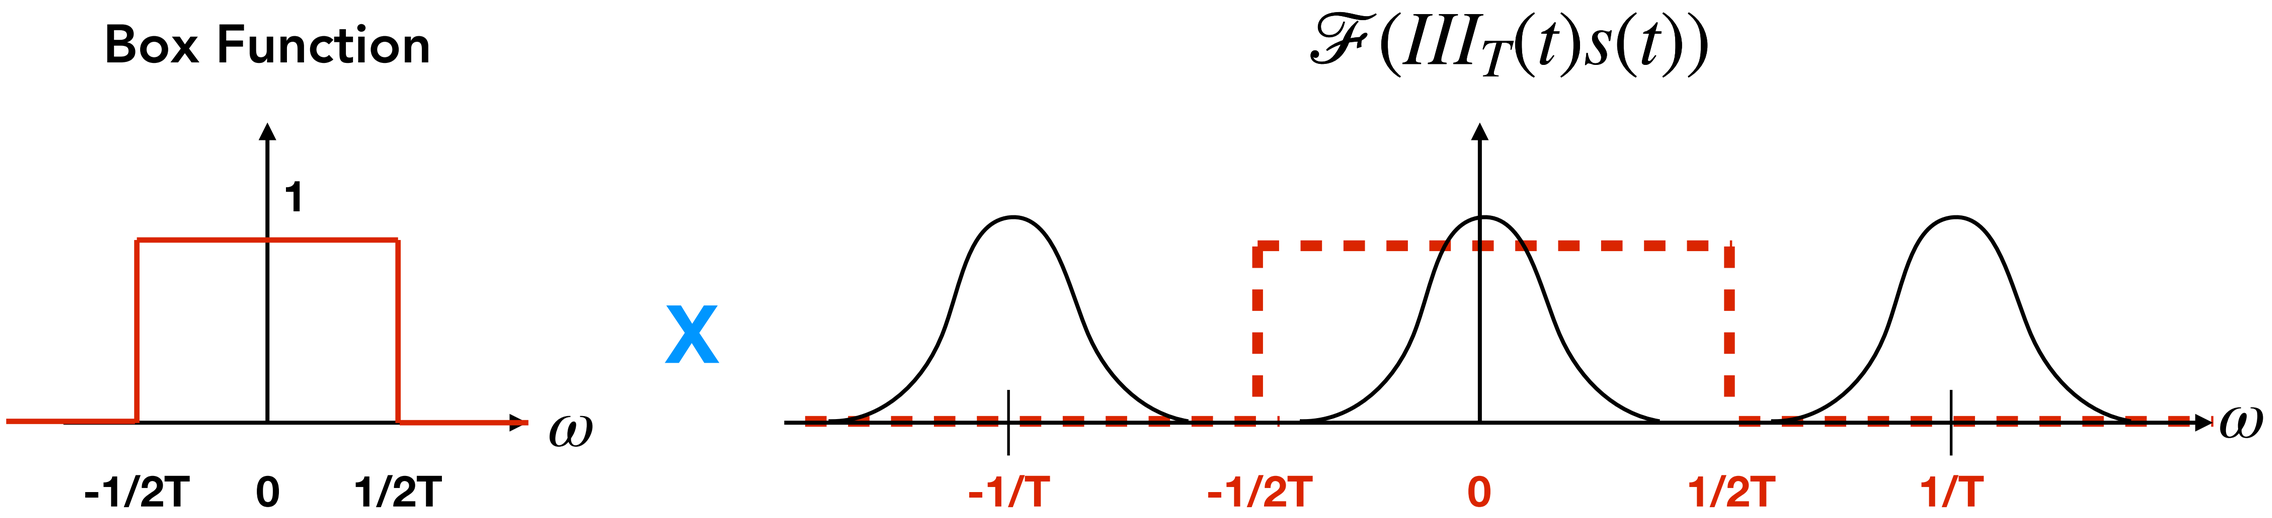

$$B(\omega) = \begin{cases} 1, & |\omega| < \frac{1}{2T}, \\ 0, & \textbf{otherwise}. \end{cases}$$

Extracting the Original Spectrum

Box Function

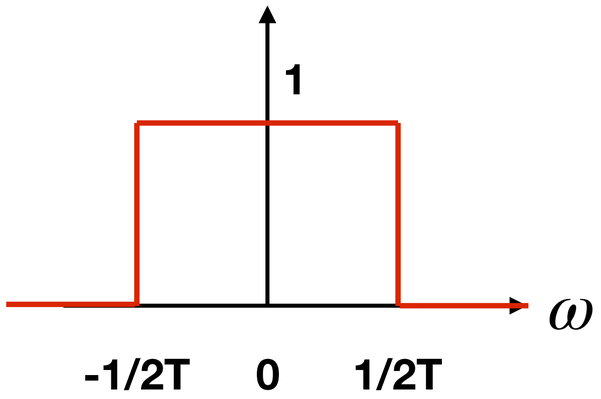

X

$\mathcal{F}(III_T(t)s(t))$

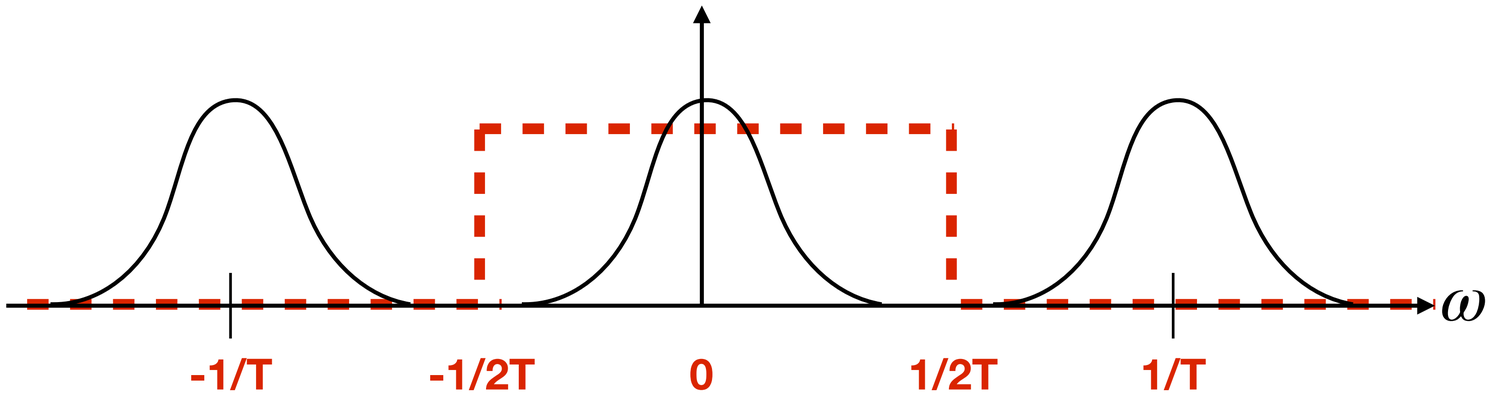

$$B(\omega) = \begin{cases} 1, & |\omega| < \frac{1}{2T}, \\ 0, & \textbf{otherwise}. \end{cases}$$

$$\mathcal{F}(III_T(t)s(t))B(\omega) = F(\omega)$$

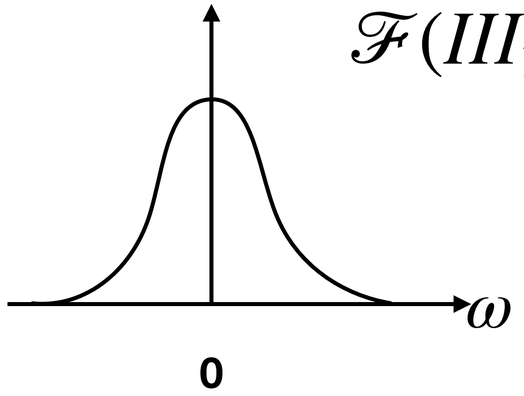

Reconstructing the Original Signal

Box Function

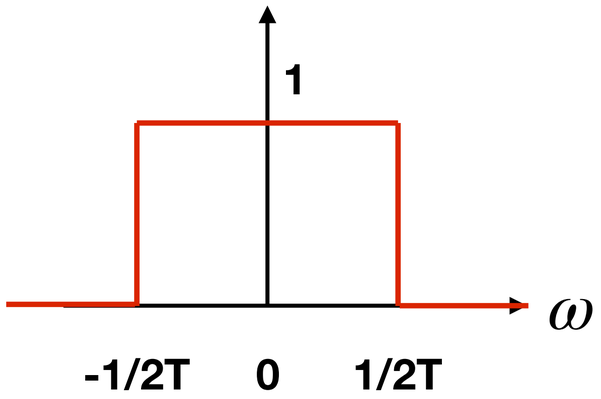

X

$\mathcal{F}(III_T(t)s(t))$

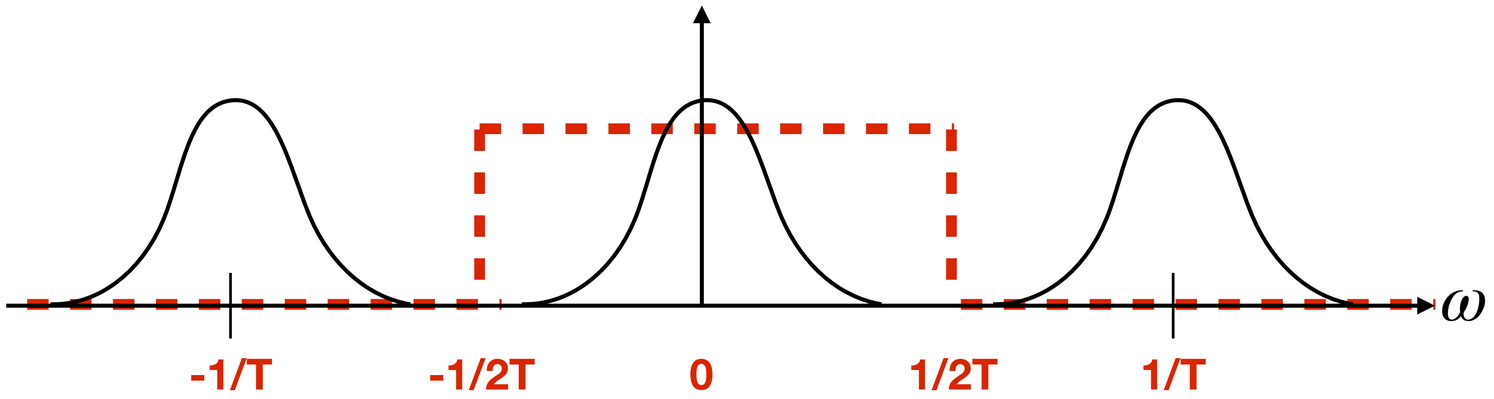

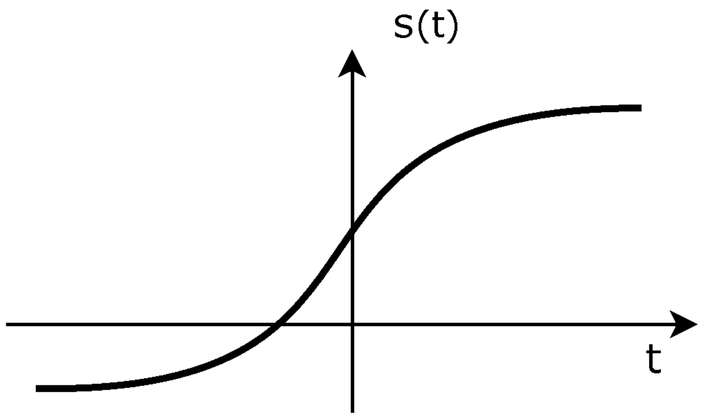

Inverse Fourier Transform

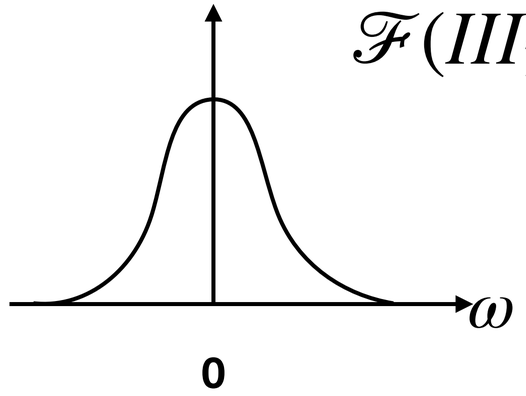

$\mathcal{F}(III_T(t)s(t))B(\omega) = F(\omega)$

Ideal Signal Reconstruction Summary

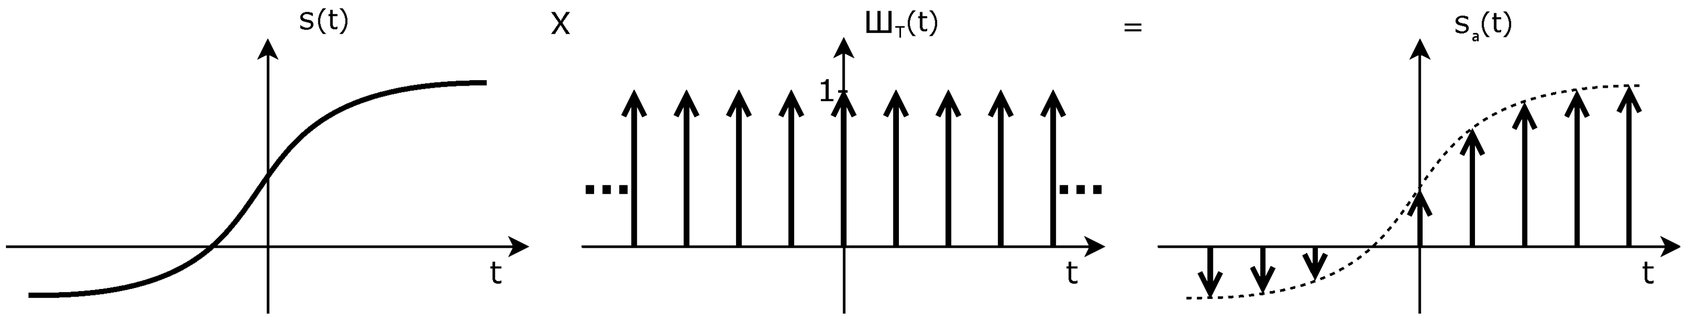

Ideal Signal Reconstruction Summary

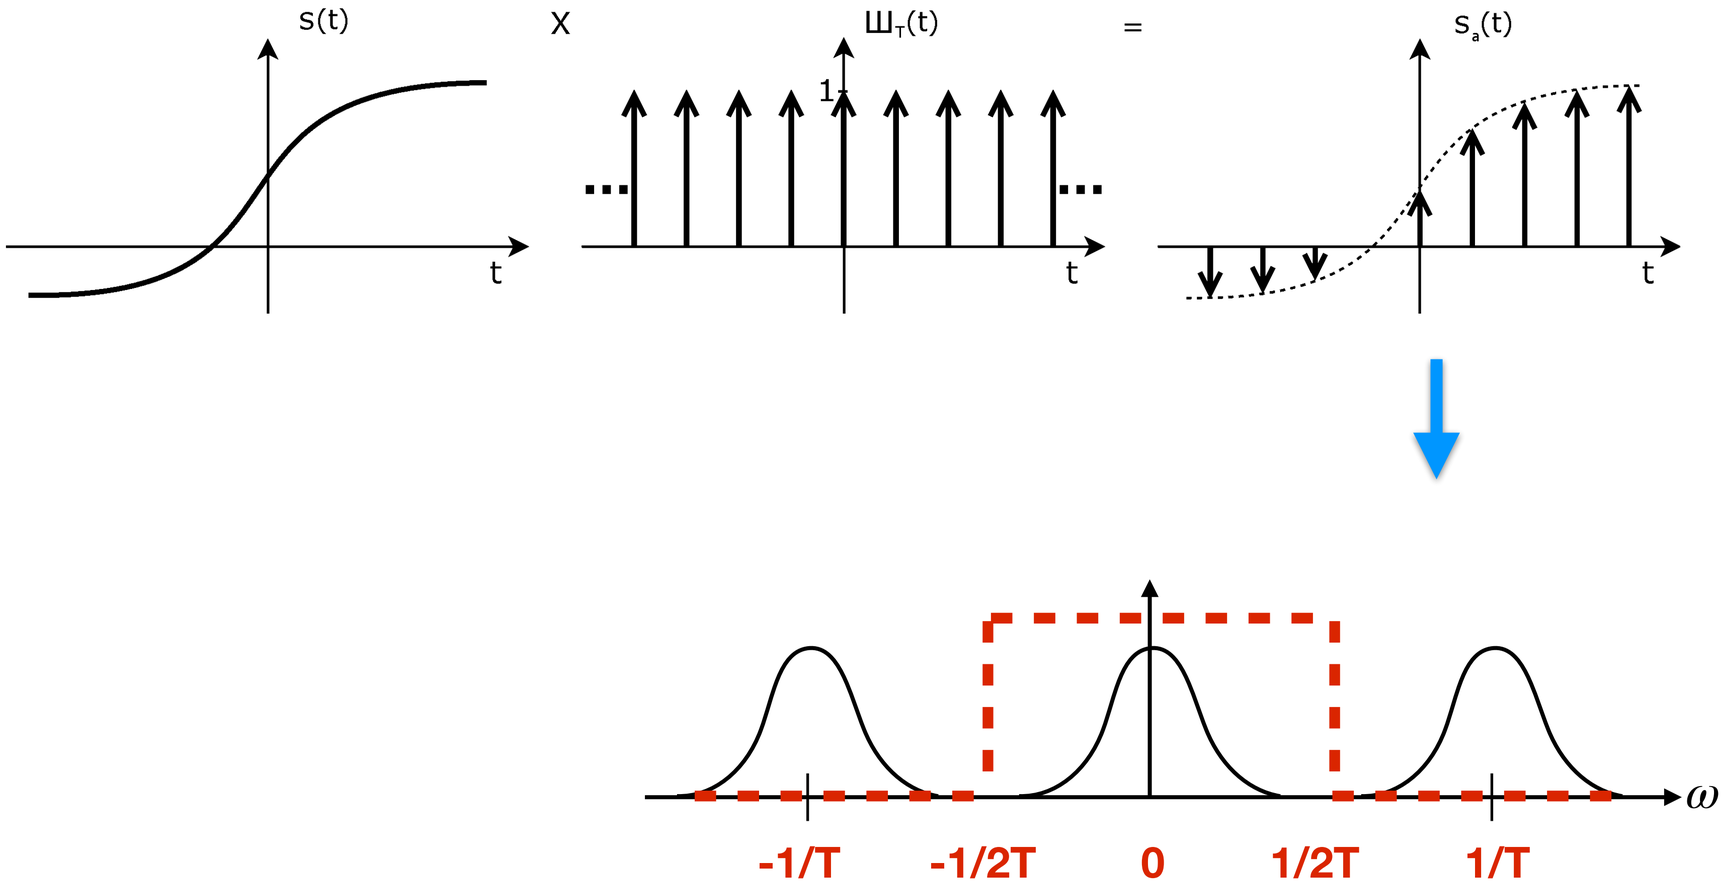

Ideal Signal Reconstruction Summary

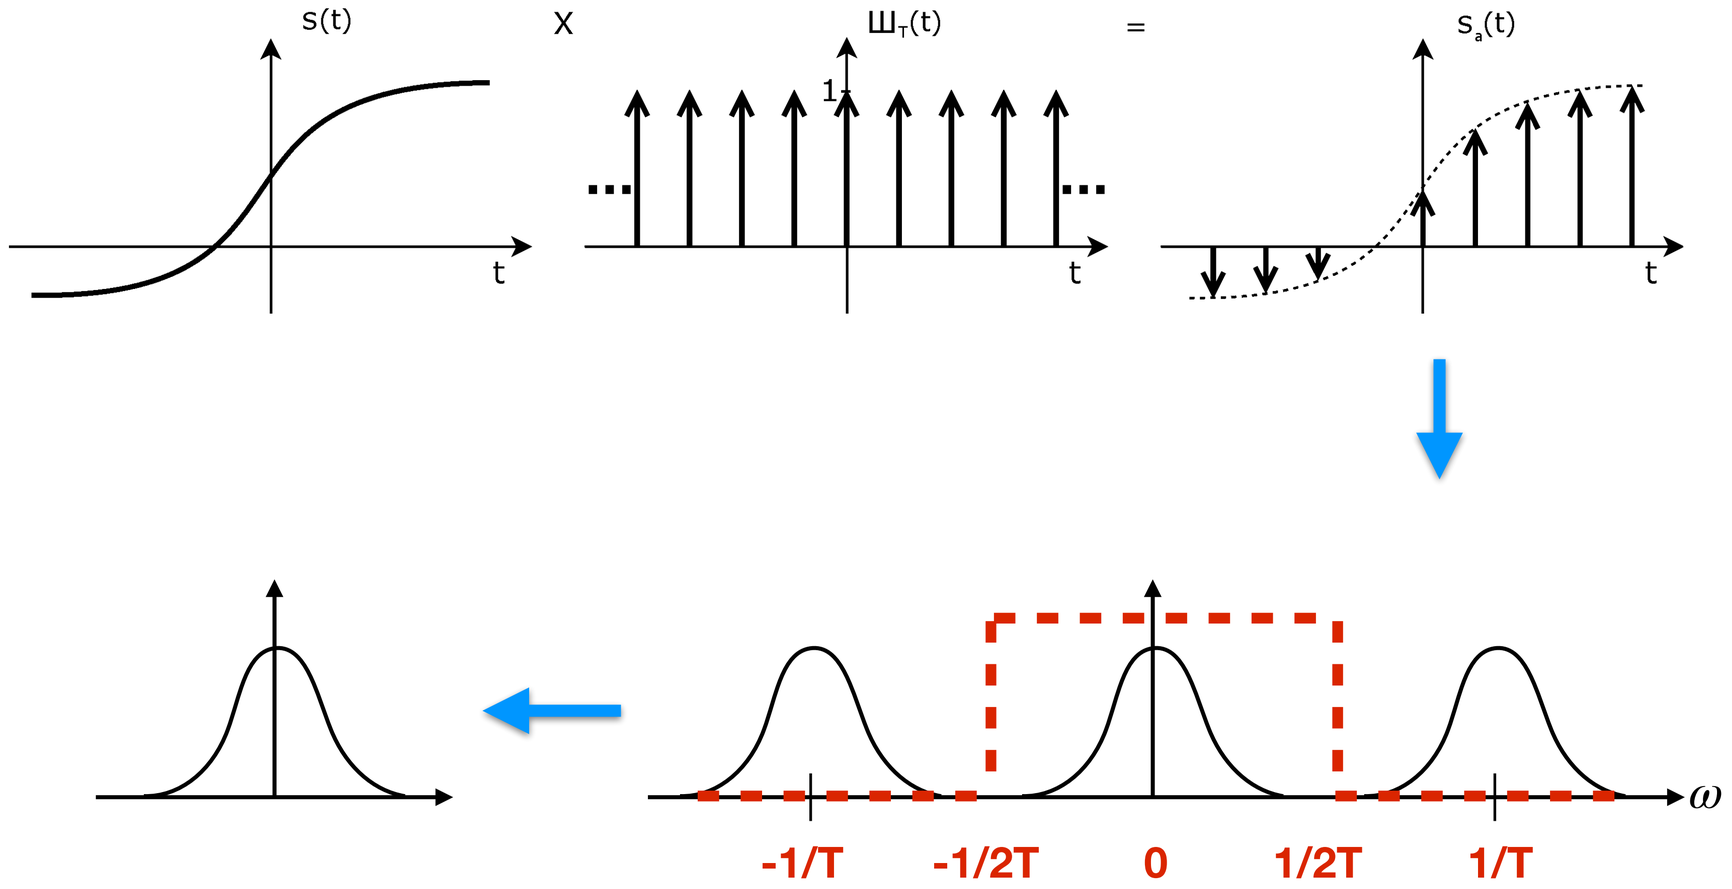

Ideal Signal Reconstruction Summary

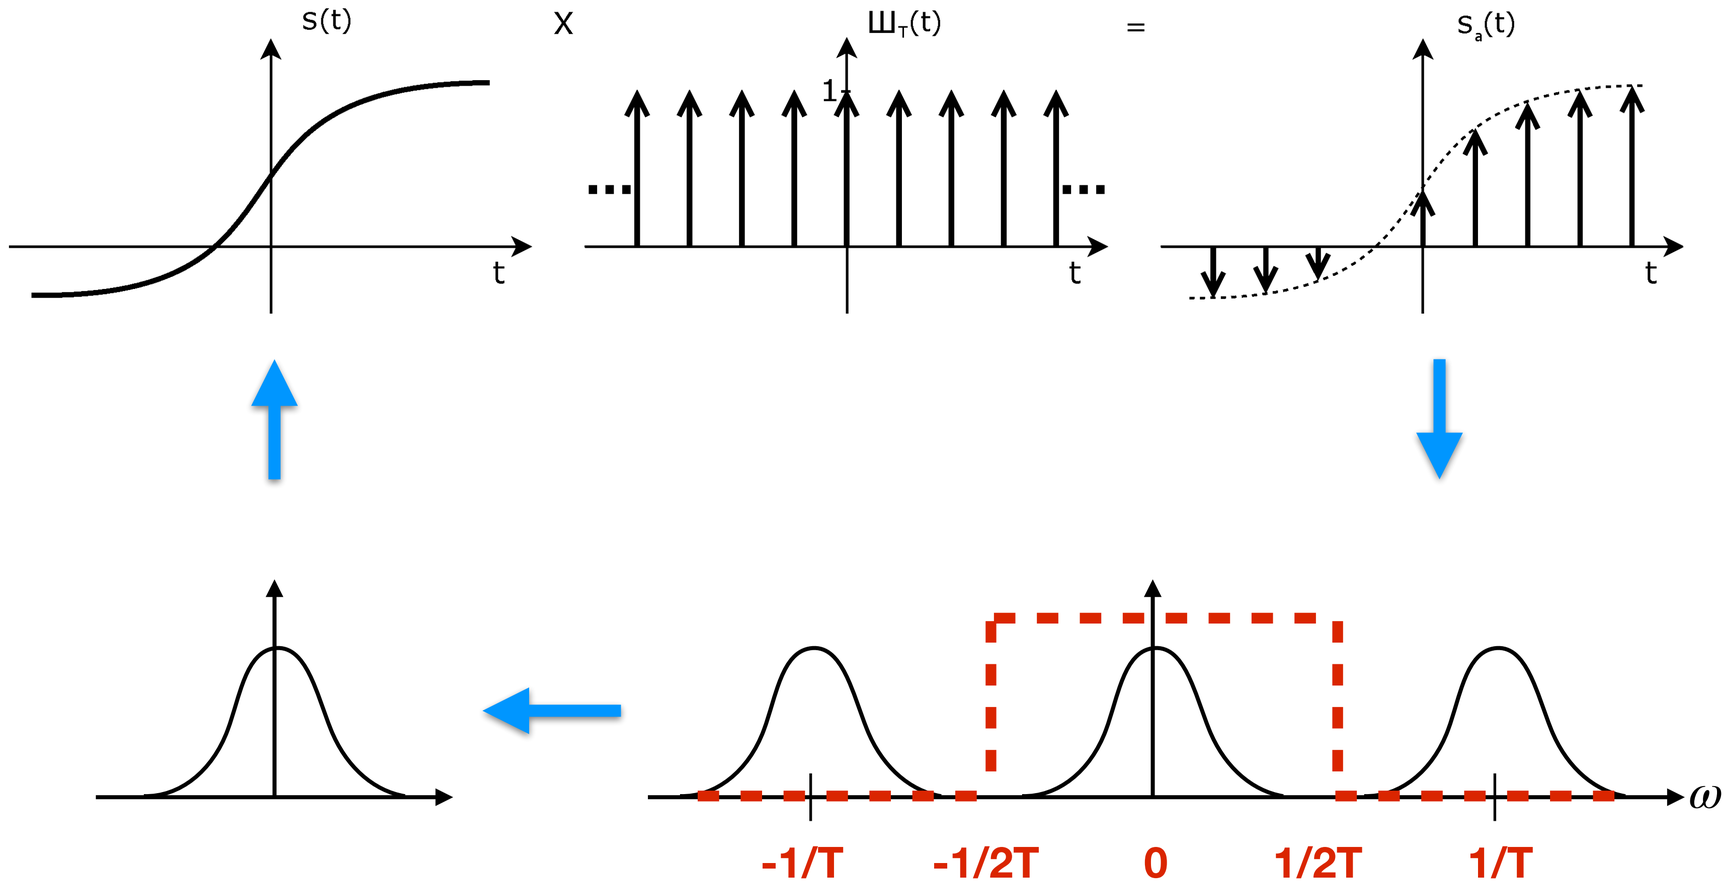

Ideal Signal Reconstruction Summary

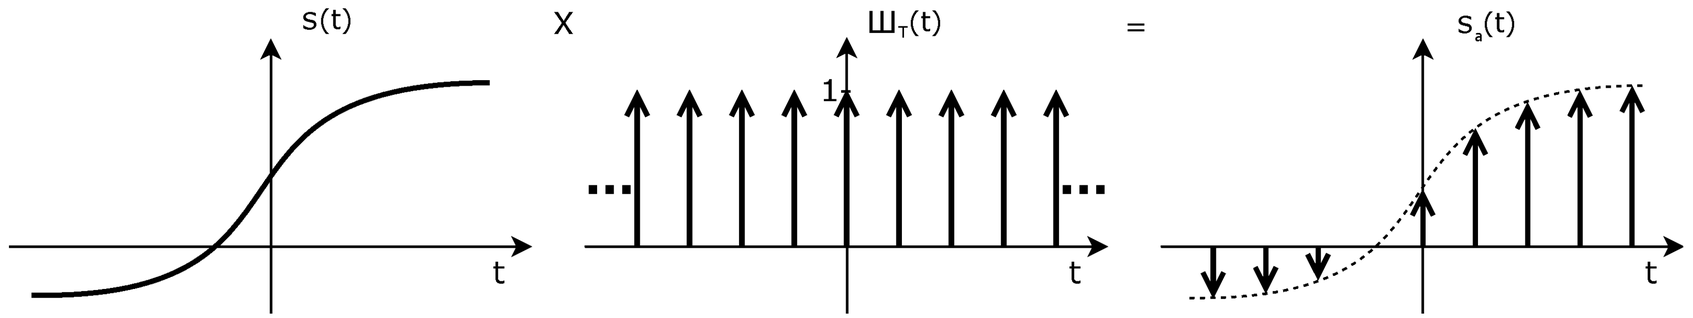

The reconstruction equation

$$s(t) = \mathcal{F}^{-1}[\ \mathcal{F}(S_a(t))B_{1/2T}(\omega)\ ]$$

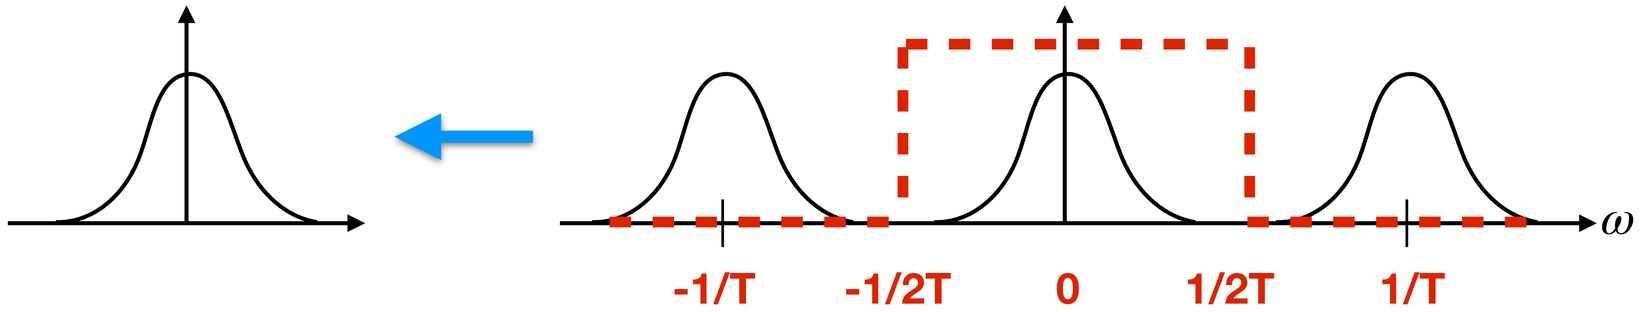

Aliasing

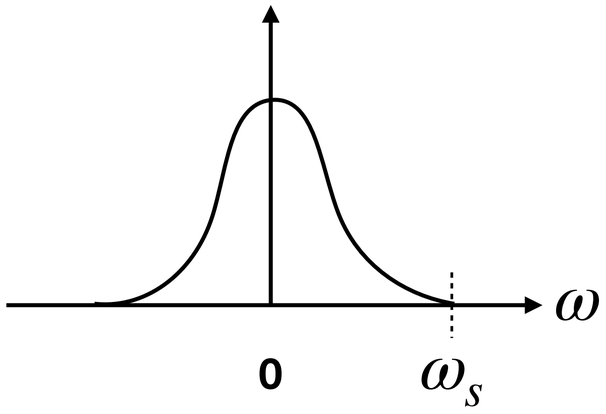

ω_s is the maximum frequency of the spectrum.
$F(\omega) = 0$ for all $|\omega| > \omega_s$.

Aliasing

$$\mathcal{F}(s_a(t)) = \sum_{i=-\infty}^{\infty} F(\omega - i\frac{1}{T})$$

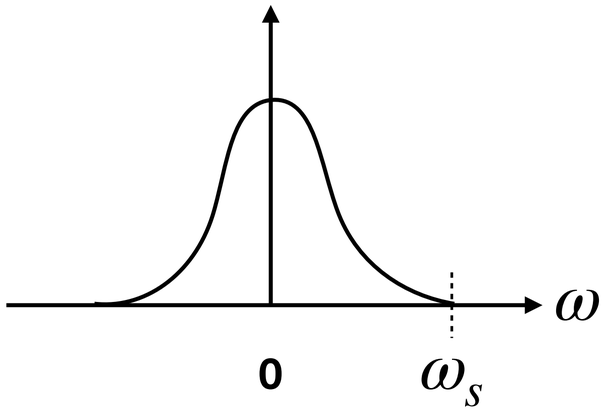

Case 1: $\omega_s < 1/2T$

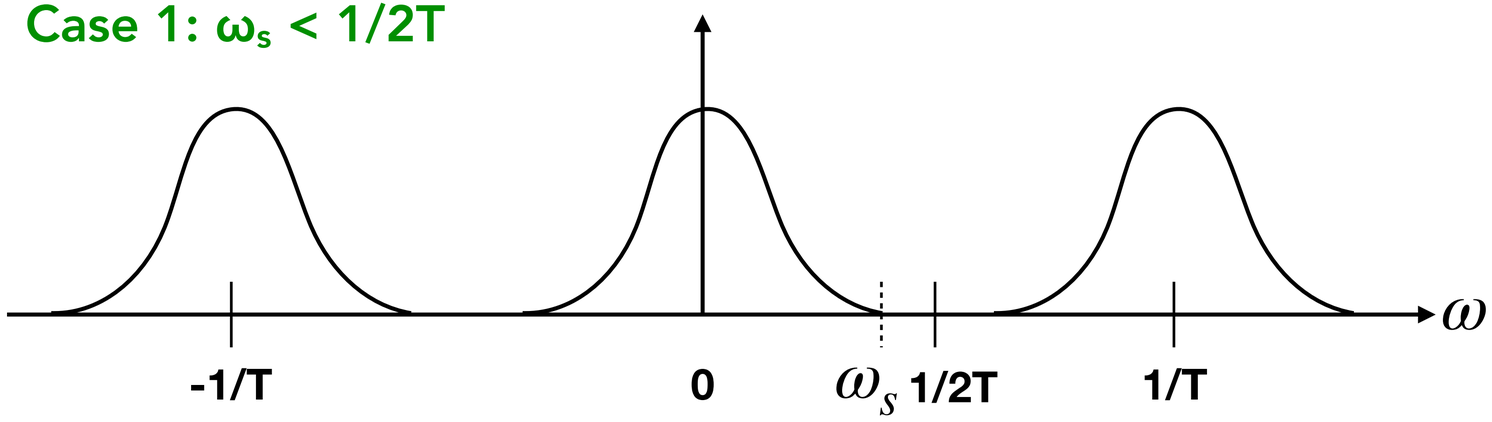

ω_s is the maximum frequency of the spectrum.
$F(\omega) = 0$ for all $|\omega| > \omega_s$.

Aliasing

$$\mathcal{F}(s_a(t)) = \sum_{i=-\infty}^{\infty} F(\omega - i\frac{1}{T})$$

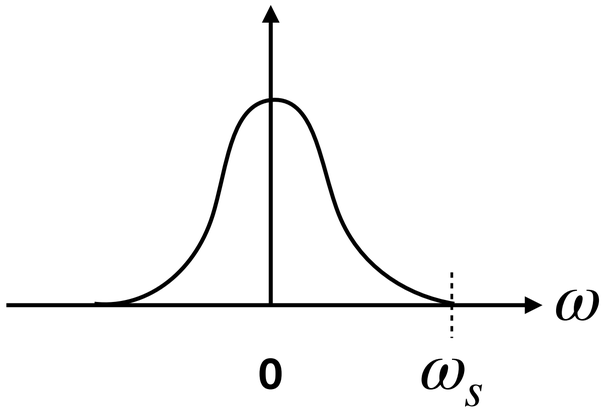

Case 1: ω_s < 1/2T

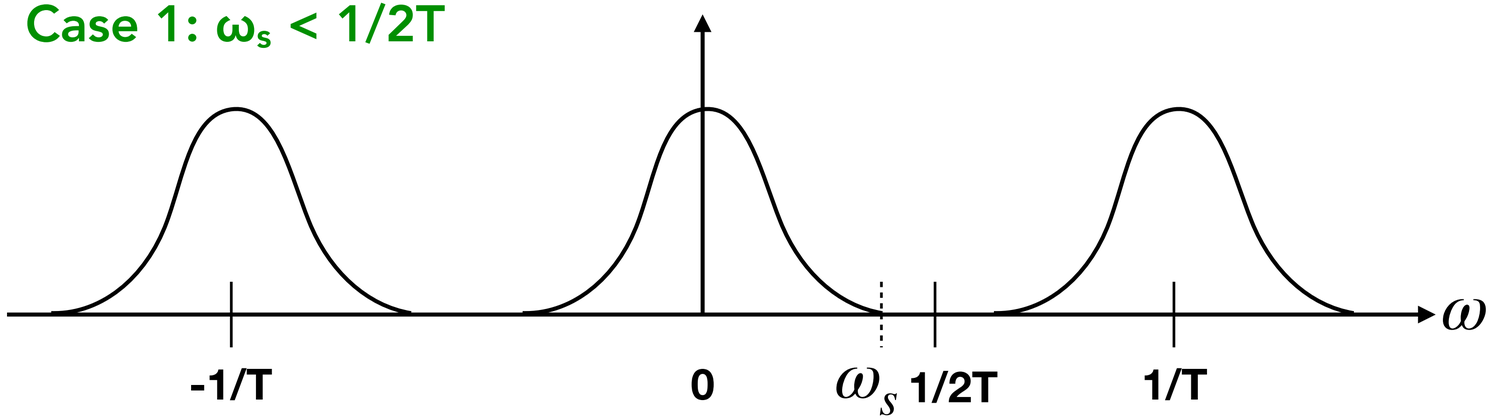

ω_s is the maximum frequency of the spectrum.
F(**ω**) = 0 for all |**ω**| > **ω_s**.

Case 2: ω_s >= 1/2T

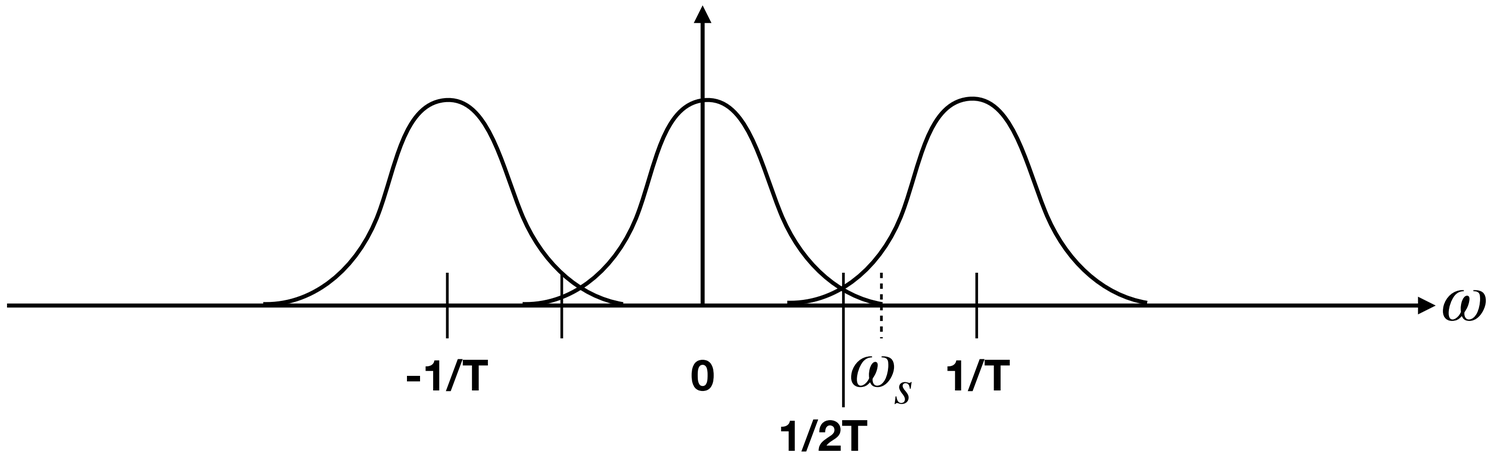

Aliasing

$$\mathcal{F}(s_a(t)) = \sum_{i=-\infty}^{\infty} F(\omega - i\frac{1}{T})$$

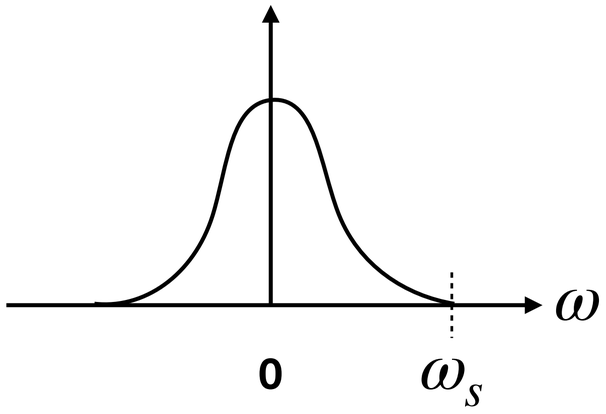

Case 1: $\omega_s < 1/2T$

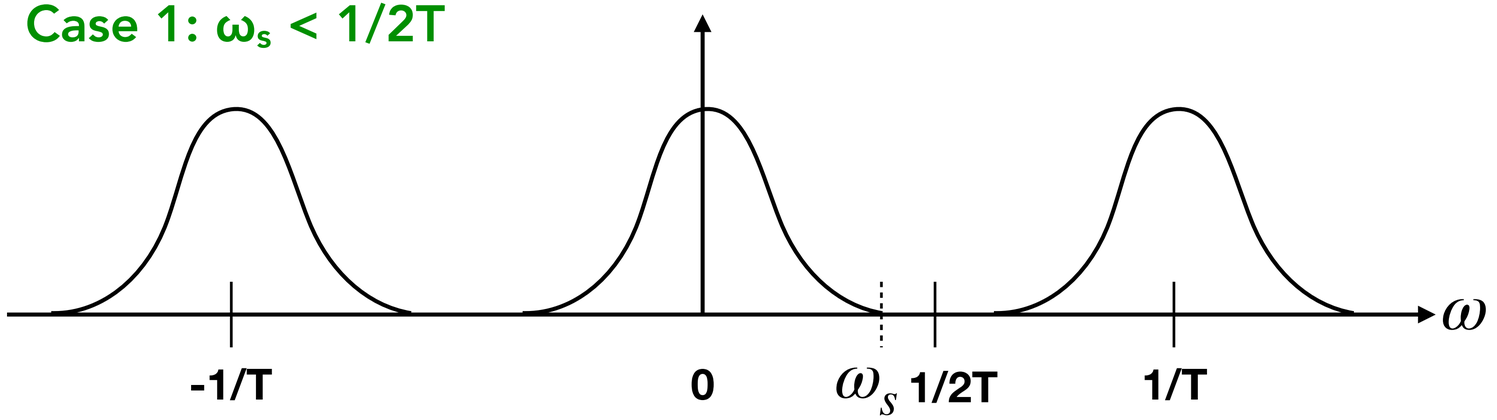

ω_s is the maximum frequency of the spectrum.
F(**ω**) = 0 for all |**ω**| > **ω_s**.

Case 2: $\omega_s >= 1/2T$

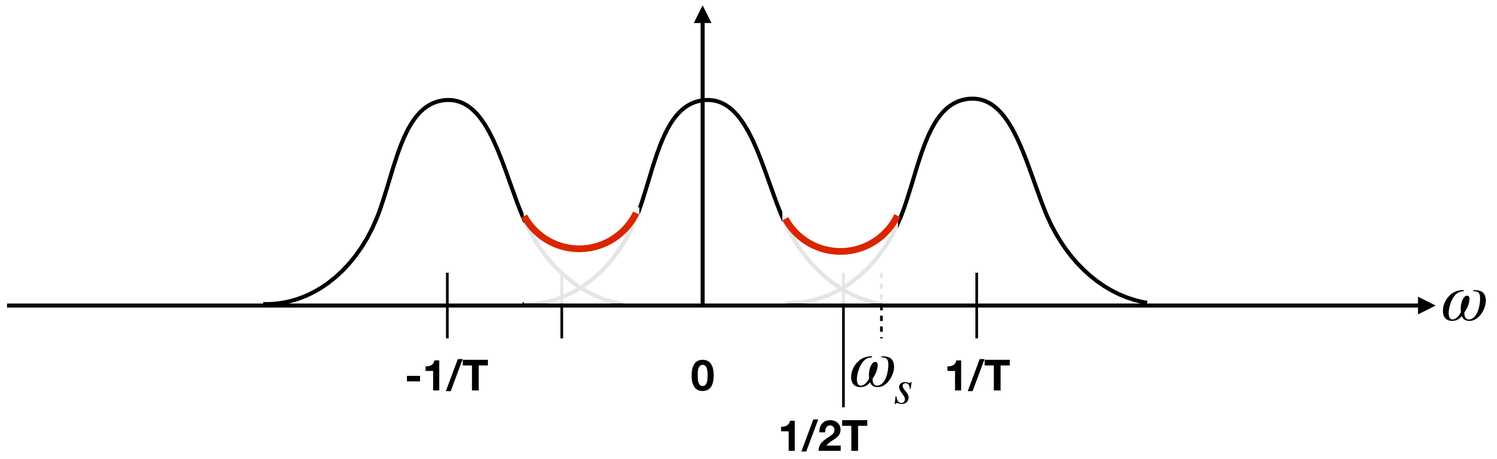

Aliasing

$$\mathcal{F}(s_a(t)) = \sum_{i=-\infty}^{\infty} F(\omega - i\frac{1}{T})$$

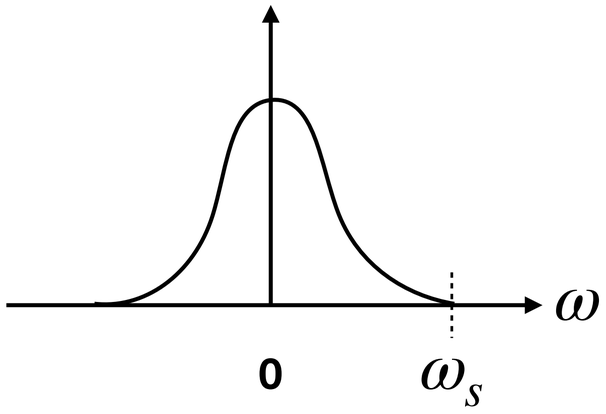

Case 1: $\omega_s < 1/2T$

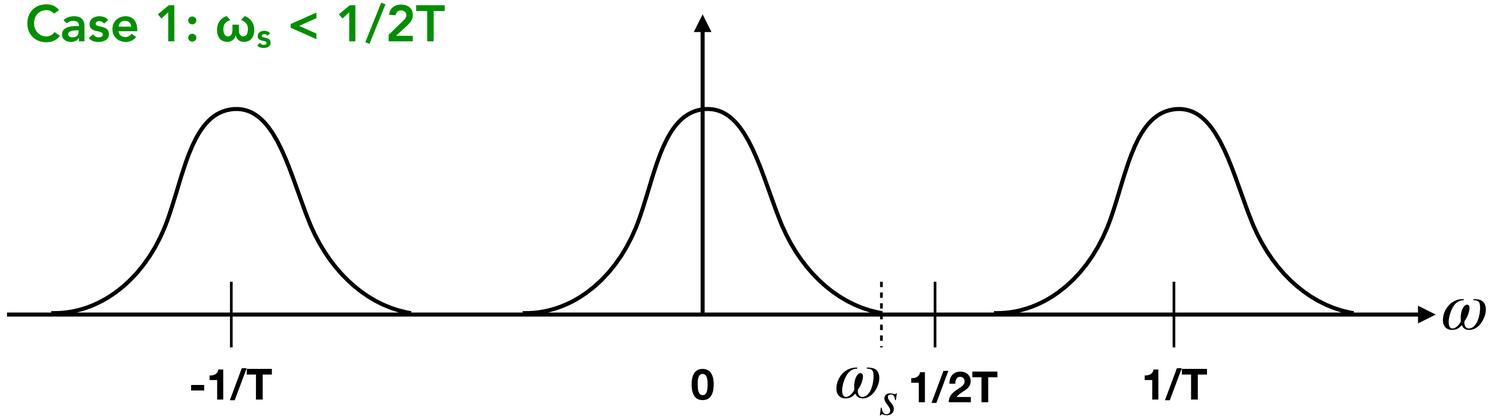

ω$_s$ is the maximum frequency of the spectrum.
F(**ω**) = 0 for all |**ω**| > **ω$_s$**.

Case 2: $\omega_s >= 1/2T$

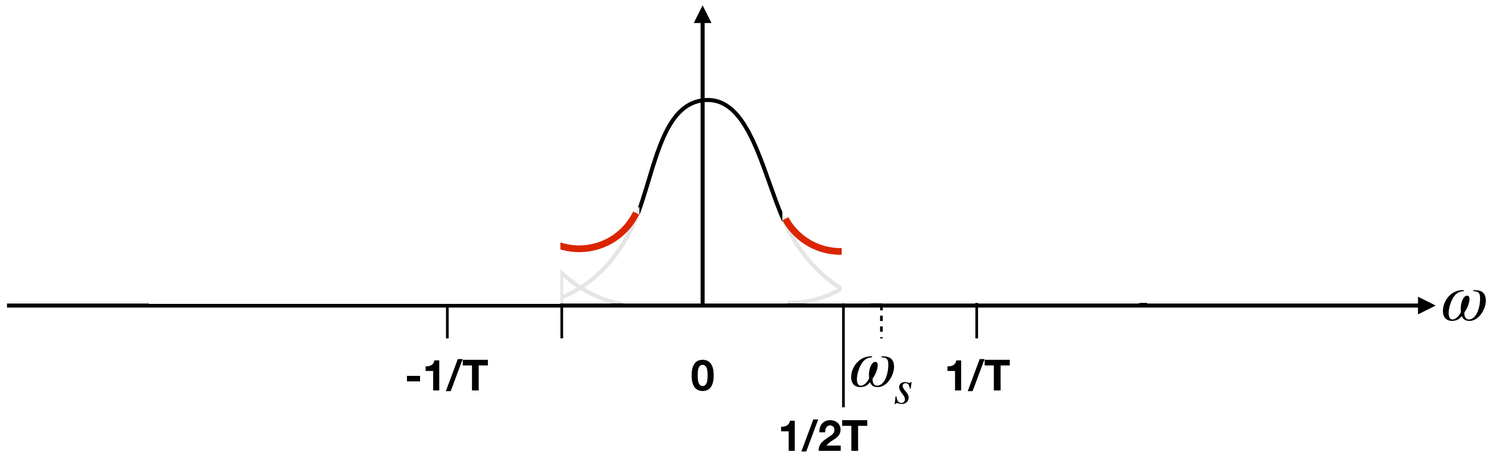

Aliasing

$$\mathcal{F}(s_a(t)) = \sum_{i=-\infty}^{\infty} F(\omega - i\frac{1}{T})$$

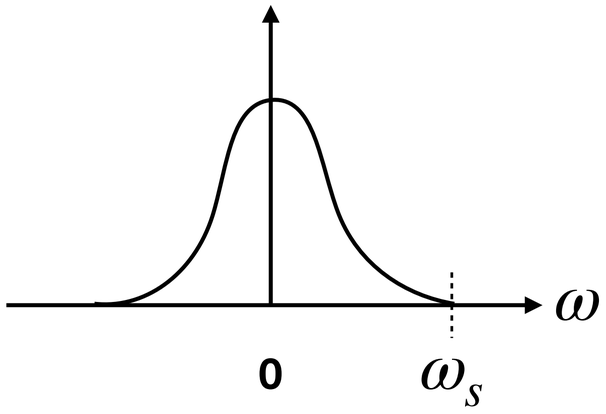

Case 1: ω_s < 1/2T

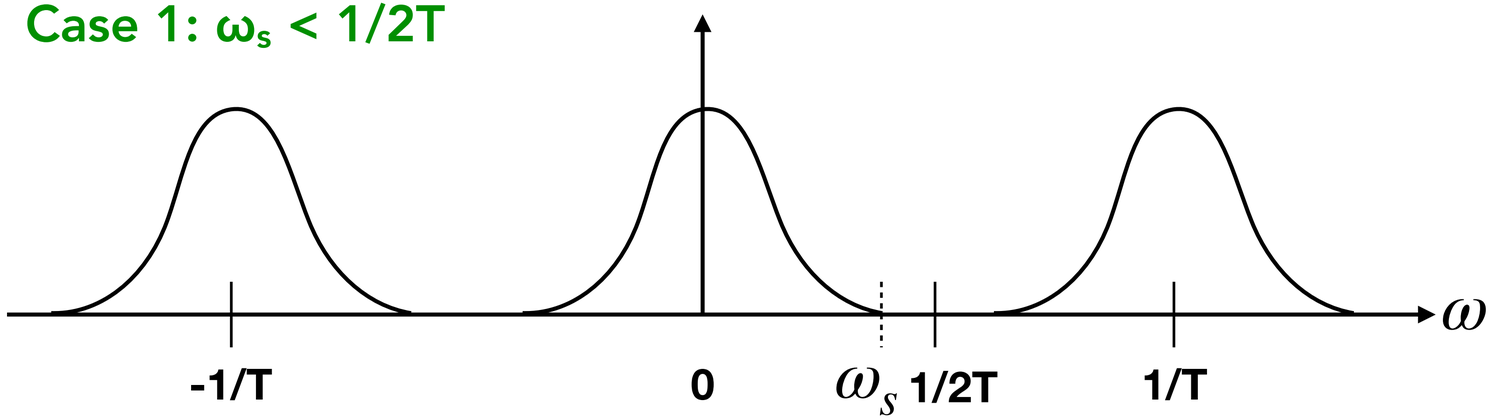

ω_s is the maximum frequency of the spectrum.
F(**ω**) = 0 for all |**ω**| > **ω_s**.

The spectrum doesn't look like many copies of the original spectrum, which can never be extracted now!

Case 2: ω_s >= 1/2T

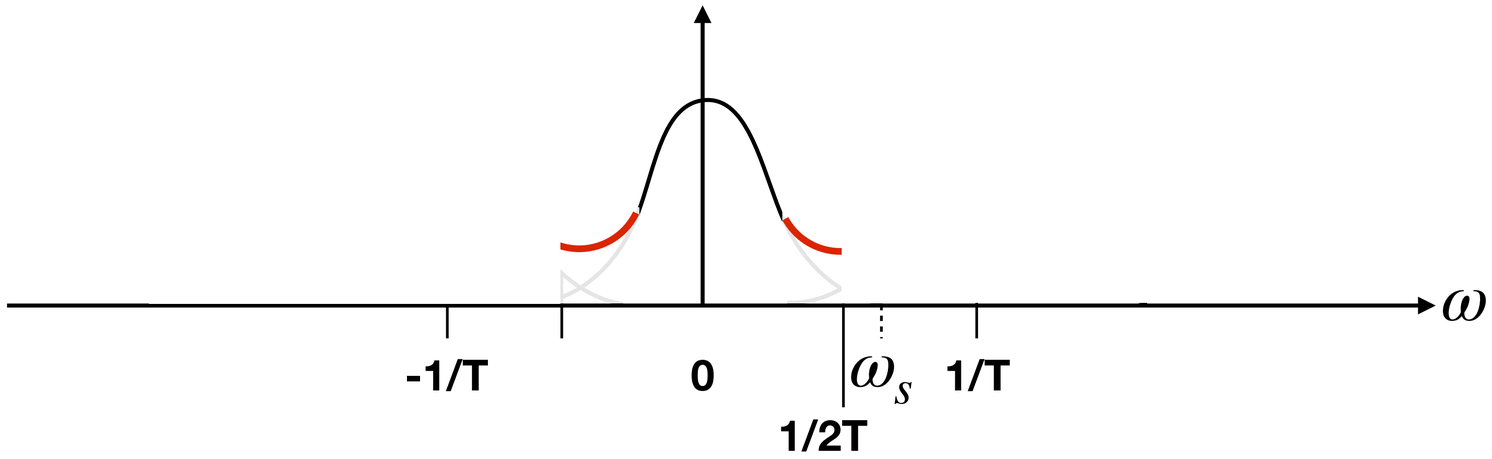

Another Aliasing Example

No Aliasing

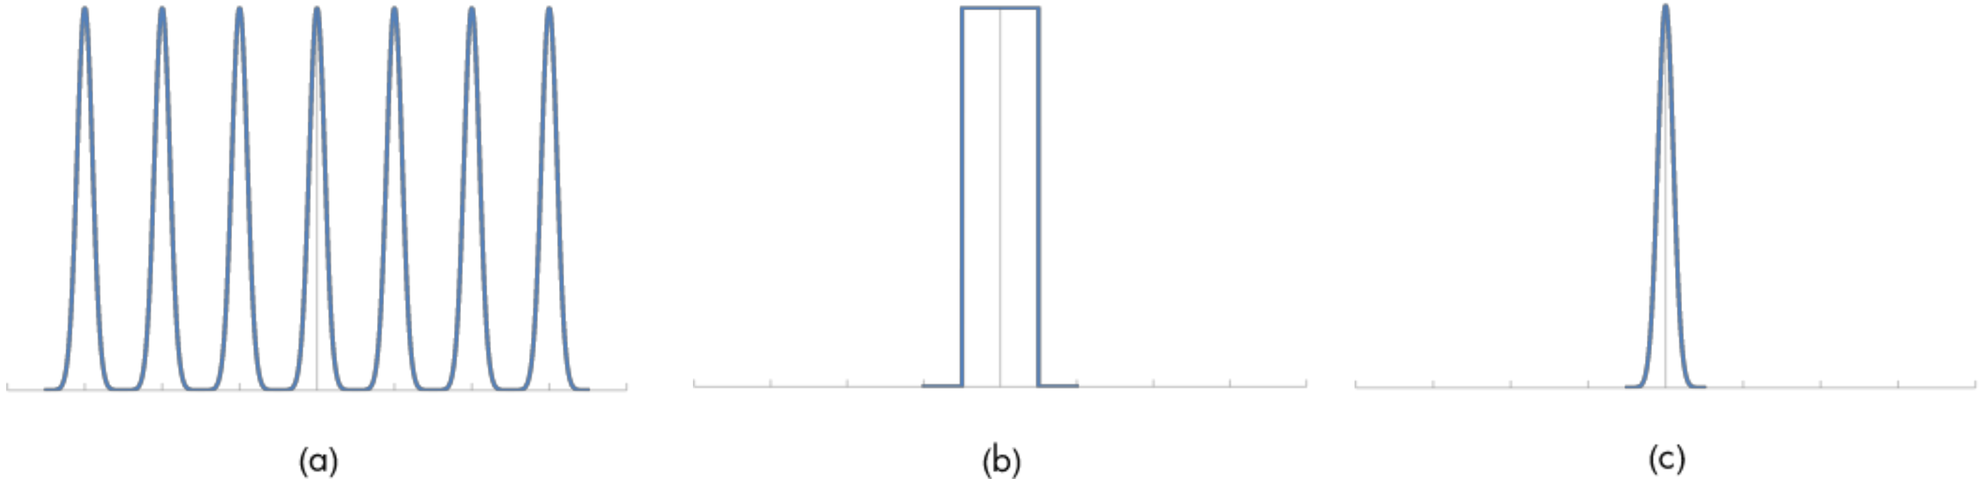

Aliased

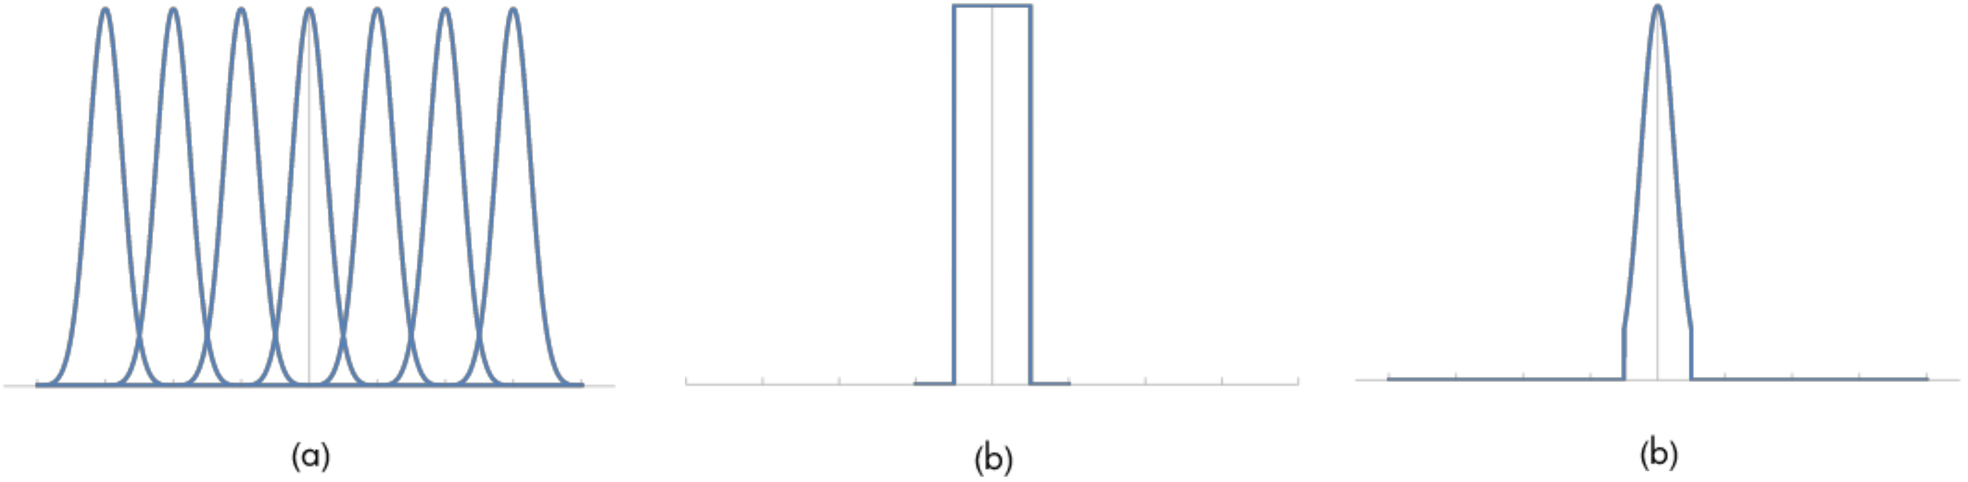

https://pbr-book.org/3ed-2018/Sampling_and_Reconstruction/Sampling_Theory

Reconstruction with Aliasing

Original Signal

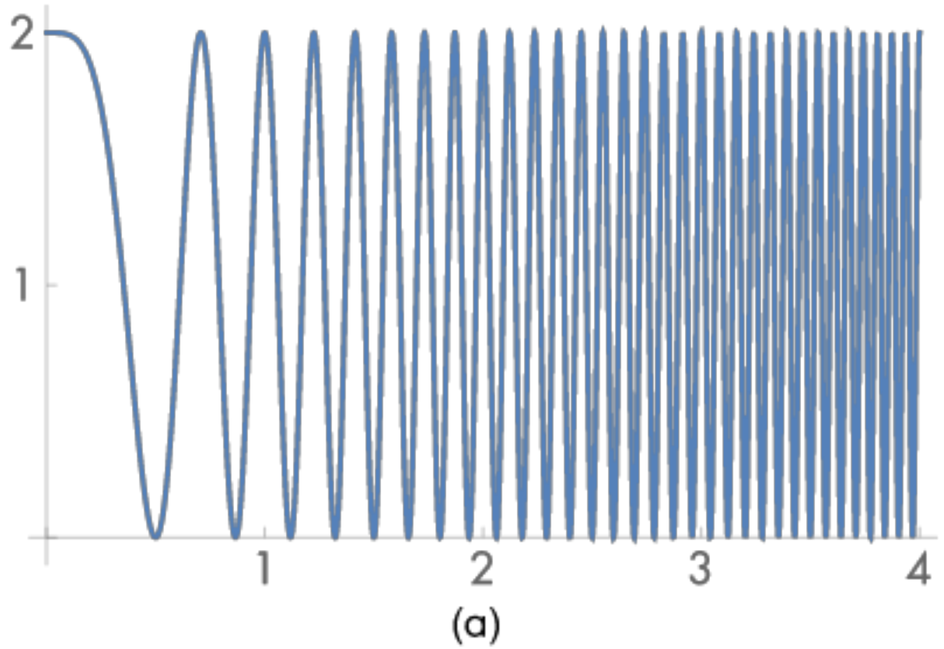

Reconstructed Signal

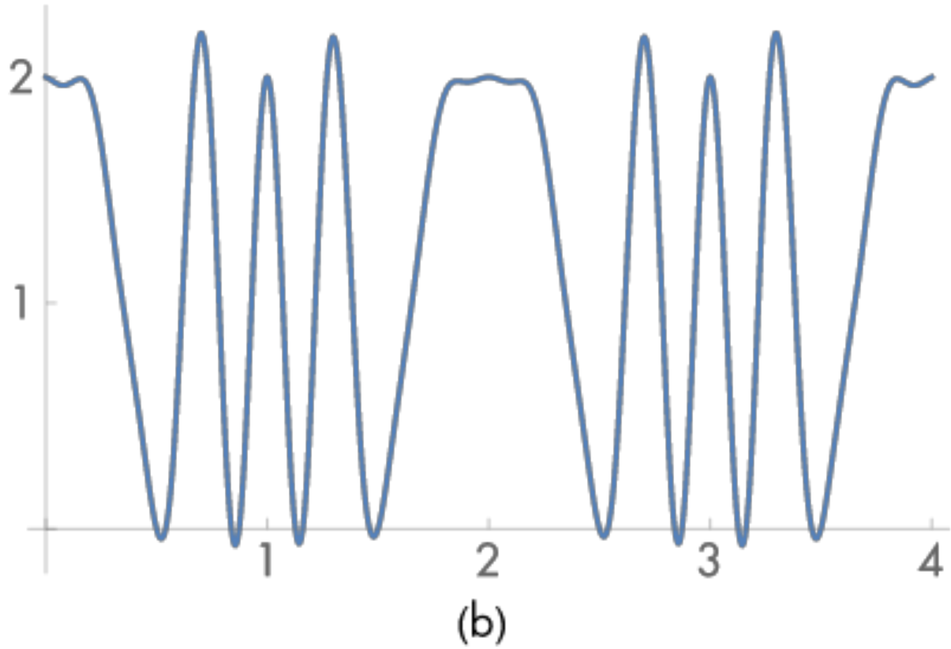

High-frequency information in the original signal is lost and shows up as low-frequency errors.

https://pbr-book.org/3ed-2018/Sampling_and_Reconstruction/Sampling_Theory

How To Guarantee Alias-Free Reconstruction?

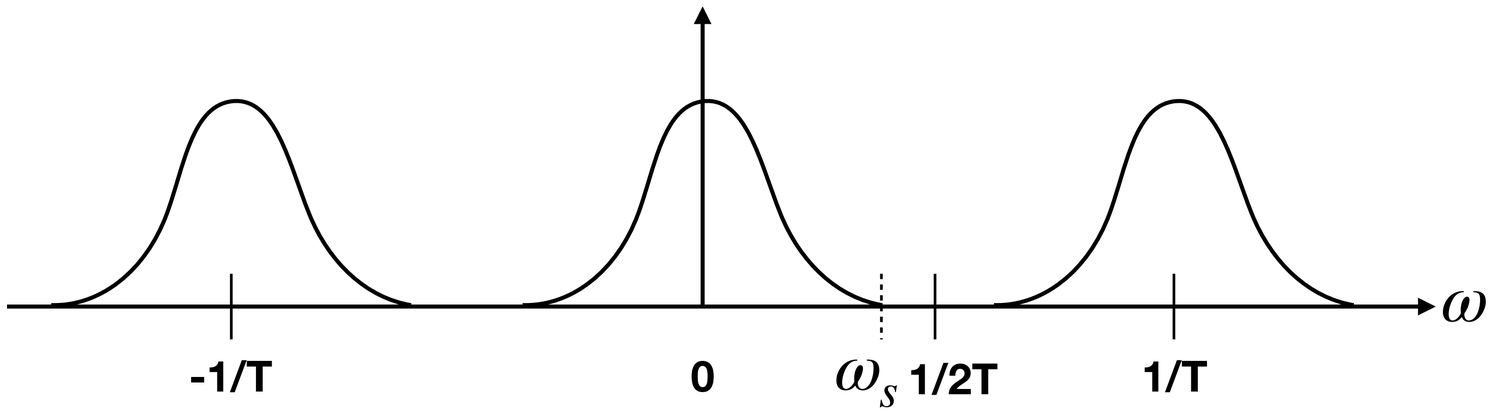

https://en.wikipedia.org/wiki/Nyquist_frequency

How To Guarantee Alias-Free Reconstruction?

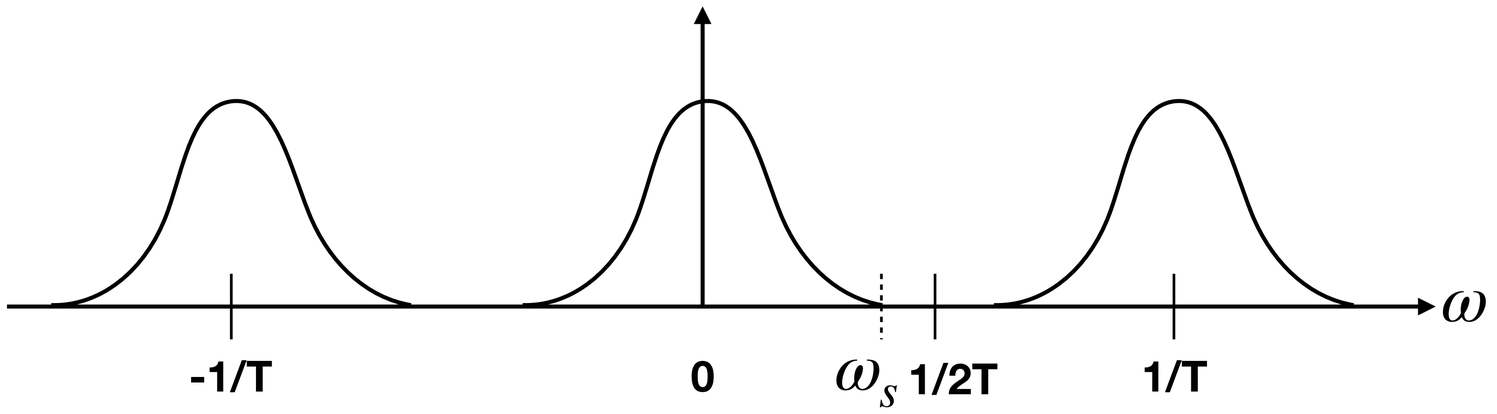

$$\omega_s < \frac{1}{2T} = \frac{f_s}{2}$$

https://en.wikipedia.org/wiki/Nyquist_frequency

How To Guarantee Alias-Free Reconstruction?

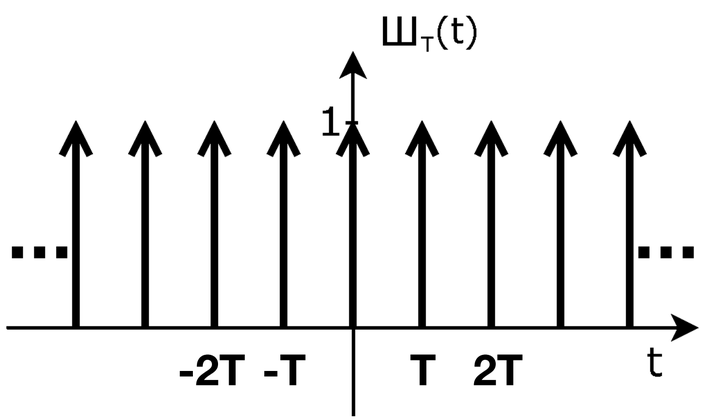

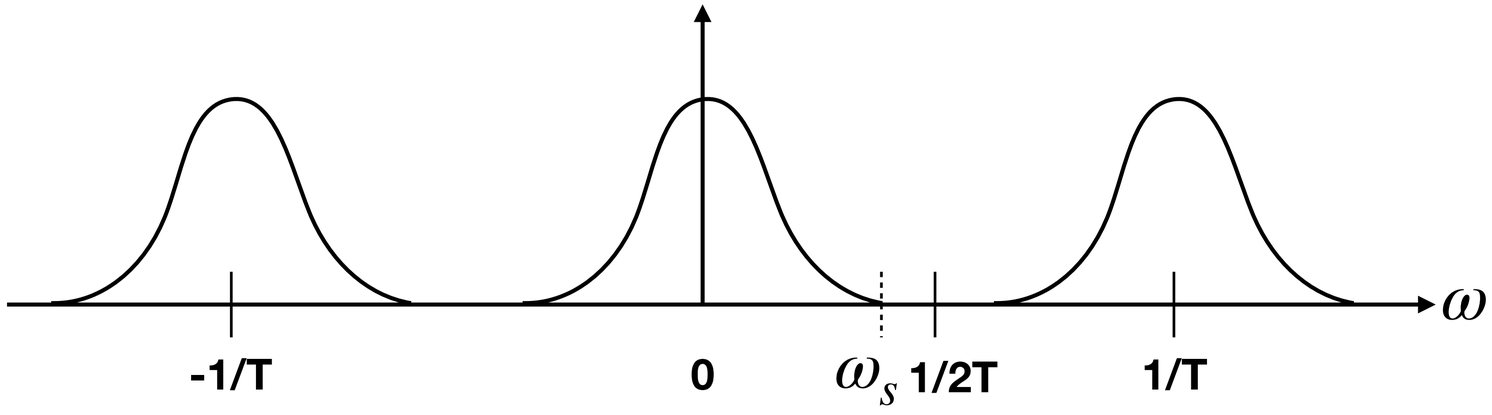

Sampling period: **T**
Sampling frequency f_s: **1/T**

$$\omega_s < \frac{1}{2T} = \frac{f_s}{2}$$

https://en.wikipedia.org/wiki/Nyquist_frequency

How To Guarantee Alias-Free Reconstruction?

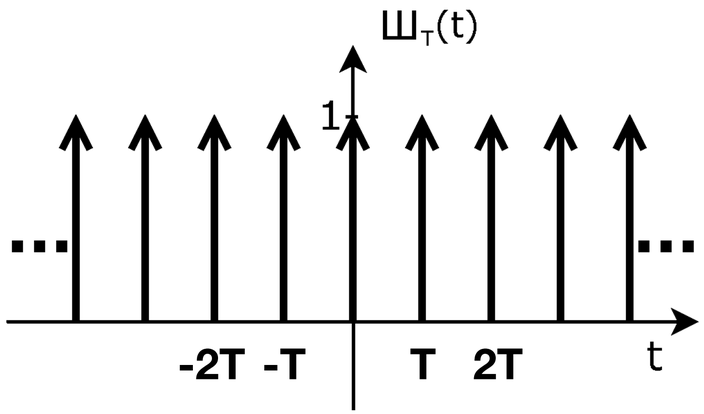

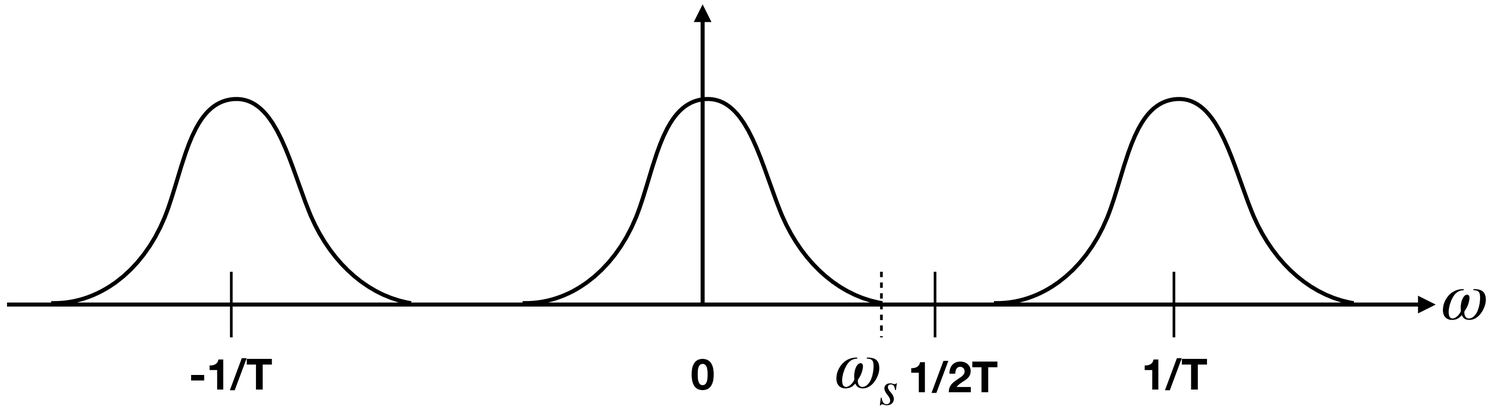

Sampling period: **T**
Sampling frequency f_s: **1/T**

$$\omega_s < \frac{1}{2T} = \frac{f_s}{2}$$

$$f_s > 2\omega_s$$

https://en.wikipedia.org/wiki/Nyquist_frequency

How To Guarantee Alias-Free Reconstruction?

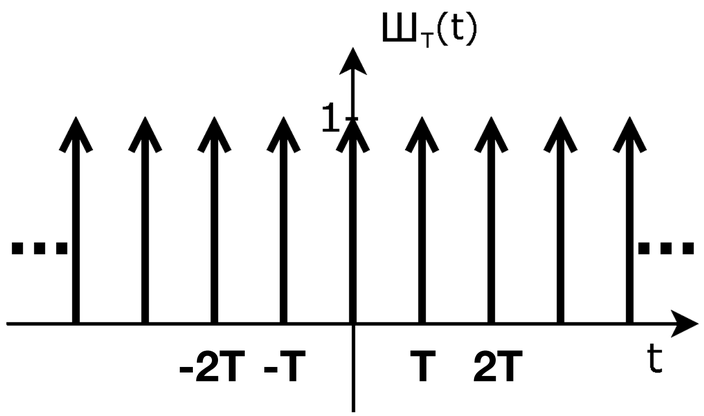

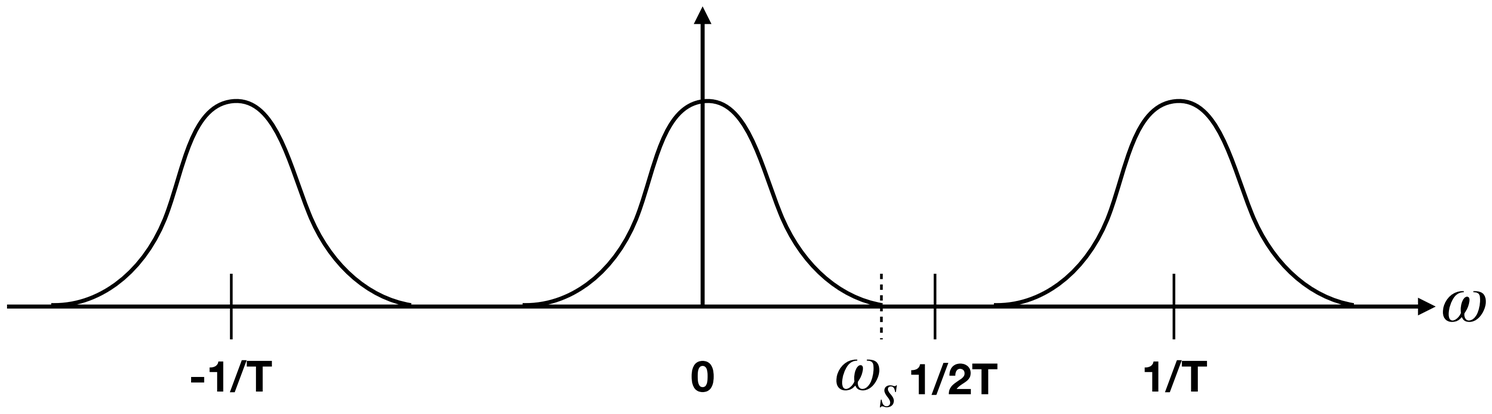

Sampling period: **T**
Sampling frequency f_s: **1/T**

$$\omega_s < \frac{1}{2T} = \frac{f_s}{2}$$

$$f_s > 2\omega_s$$

Nyquist–Shannon sampling theorem: if a signal contains no frequency higher than ω_s, it can be perfectly reconstructed from sampling it at a rate higher than $2\omega_s$.

https://en.wikipedia.org/wiki/Nyquist_frequency

How To Guarantee Alias-Free Reconstruction?

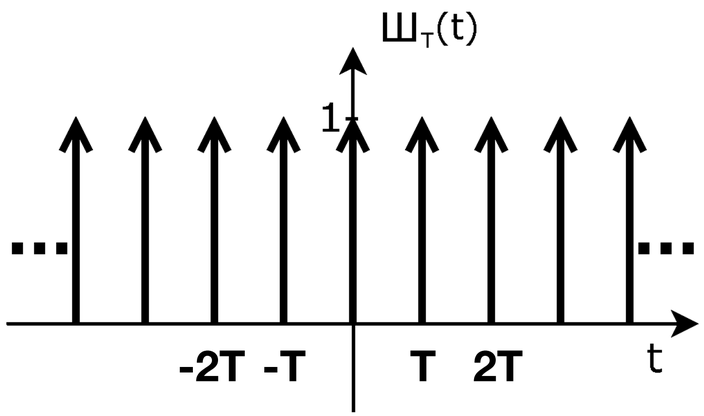

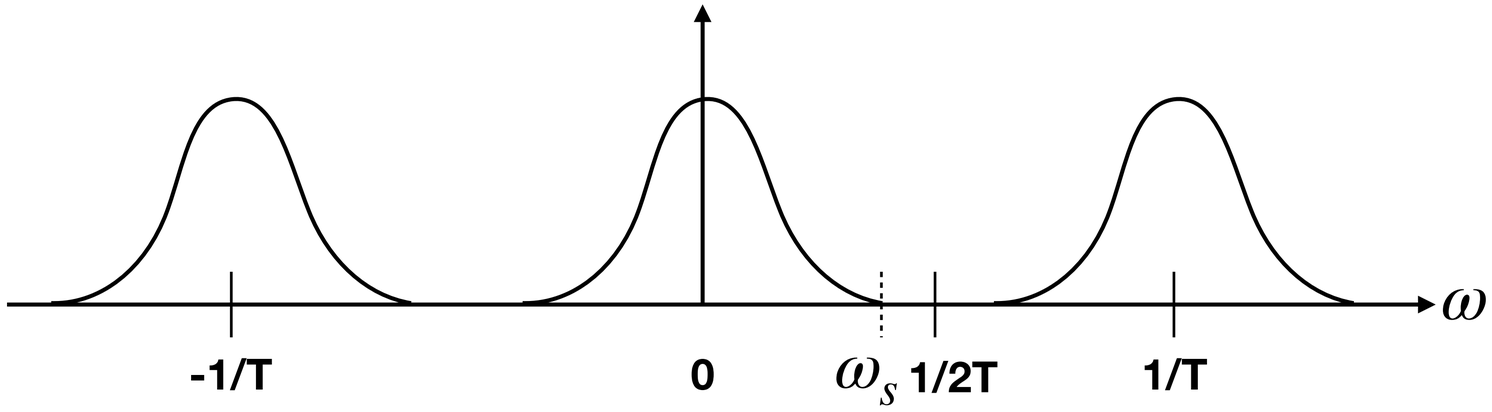

Sampling period: **T**
Sampling frequency f_s: **1/T**

$$\omega_s < \frac{1}{2T} = \boxed{\frac{f_s}{2}}$$

Nyquist frequency

$$f_s > 2\omega_s$$

Nyquist–Shannon sampling theorem: if a signal contains no frequency higher than ω_s, it can be perfectly reconstructed from sampling it at a rate higher than $2\omega_s$.

https://en.wikipedia.org/wiki/Nyquist_frequency

How To Guarantee Alias-Free Reconstruction?

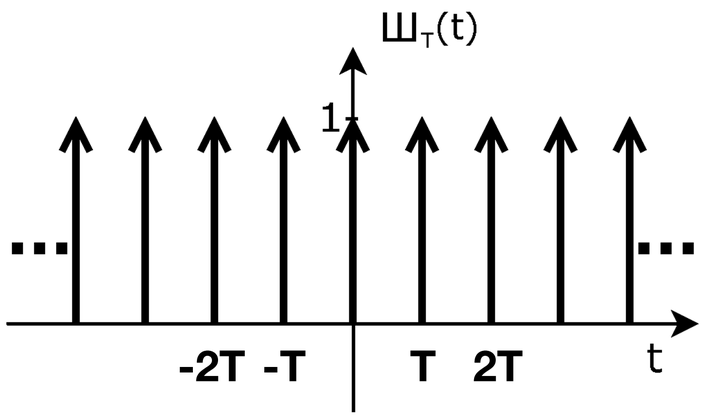

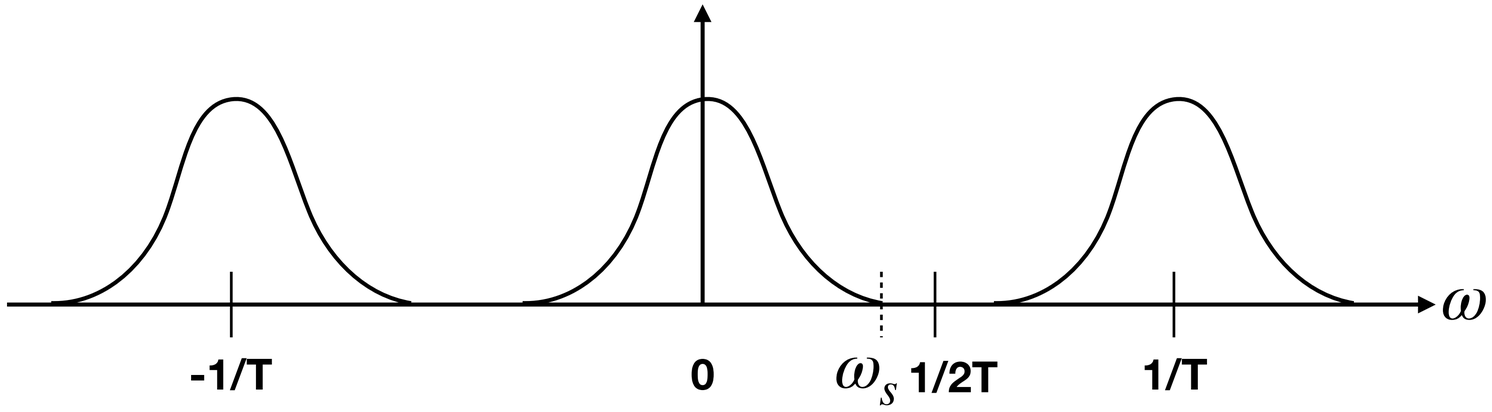

Sampling period: **T**
Sampling frequency f_s: **1/T**

$$\omega_s < \frac{1}{2T} = \boxed{\frac{f_s}{2}}$$

Nyquist frequency

$$f_s > \boxed{2\omega_s}$$

Nyquist rate

Nyquist–Shannon sampling theorem: if a signal contains no frequency higher than ω_s, it can be perfectly reconstructed from sampling it at a rate higher than $2\omega_s$.

https://en.wikipedia.org/wiki/Nyquist_frequency

Band-Limited Signal

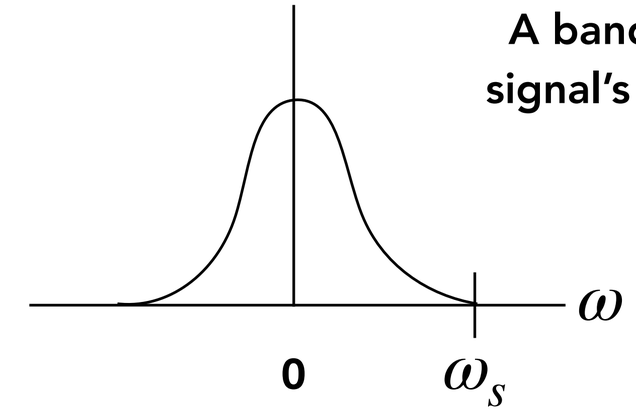

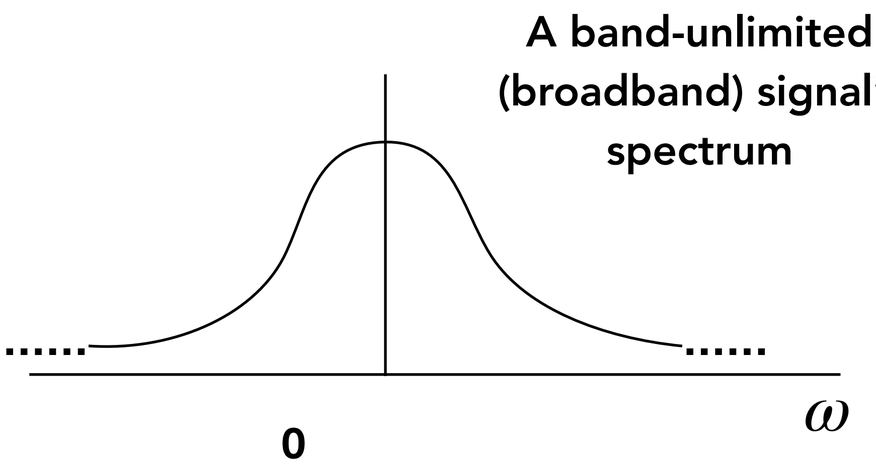

- Nyquist-Shannon sampling theorem is only useful if the original signal **does have** a maximum frequency ω_s. Otherwise the sampling rate would have to be infinite.
- Band-limited signal: a signal where there exists a frequency ωs such that $F(\omega) = 0$ for all $|\omega| > \omega_s$.
- Most real-world signals are "broad-band" signals; they are **not** band-limited.

2D Sampling

“Brush” Function

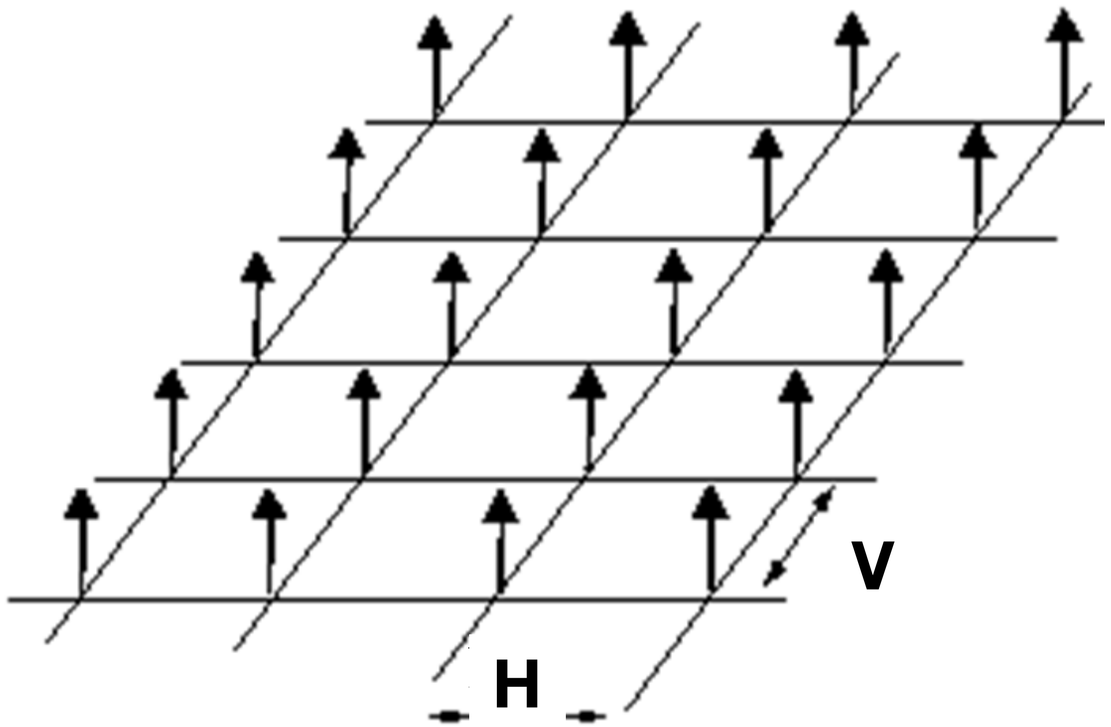

2D sampling theorem: if a signal contains no horizontal frequency higher than ω_u and no vertical frequency higher than ω_v, it can be completely reconstructed from sampling it at a horizontally rate higher than $2\omega_u$ and a vertical rate higher than $2\omega_v$.

The horizontal sampling distance $H < 1/(2\ \omega_u)$
The vertical sampling distance $V < 1/(2\ \omega_v)$

https://www.di.univr.it/documenti/Occorrenzalns/matdid/matdid346761.pdf

Everything is Signal
Sampling & Reconstruction

Sampling and Reconstruction in Camera and Display

Light in the physical scene (3D continuous signal)

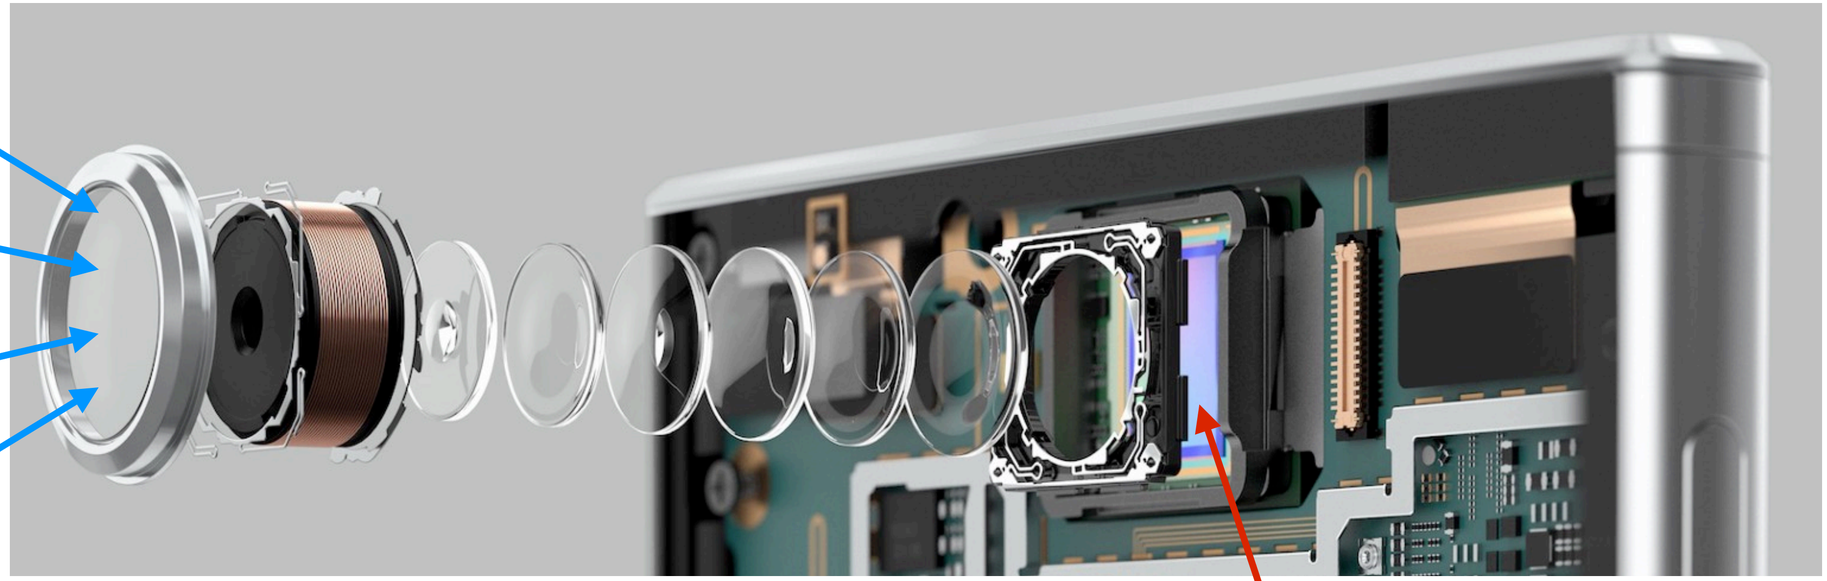

Light on the sensor plane (2D continuous signal)

Sampling and Reconstruction in Camera and Display

Light in the physical continu

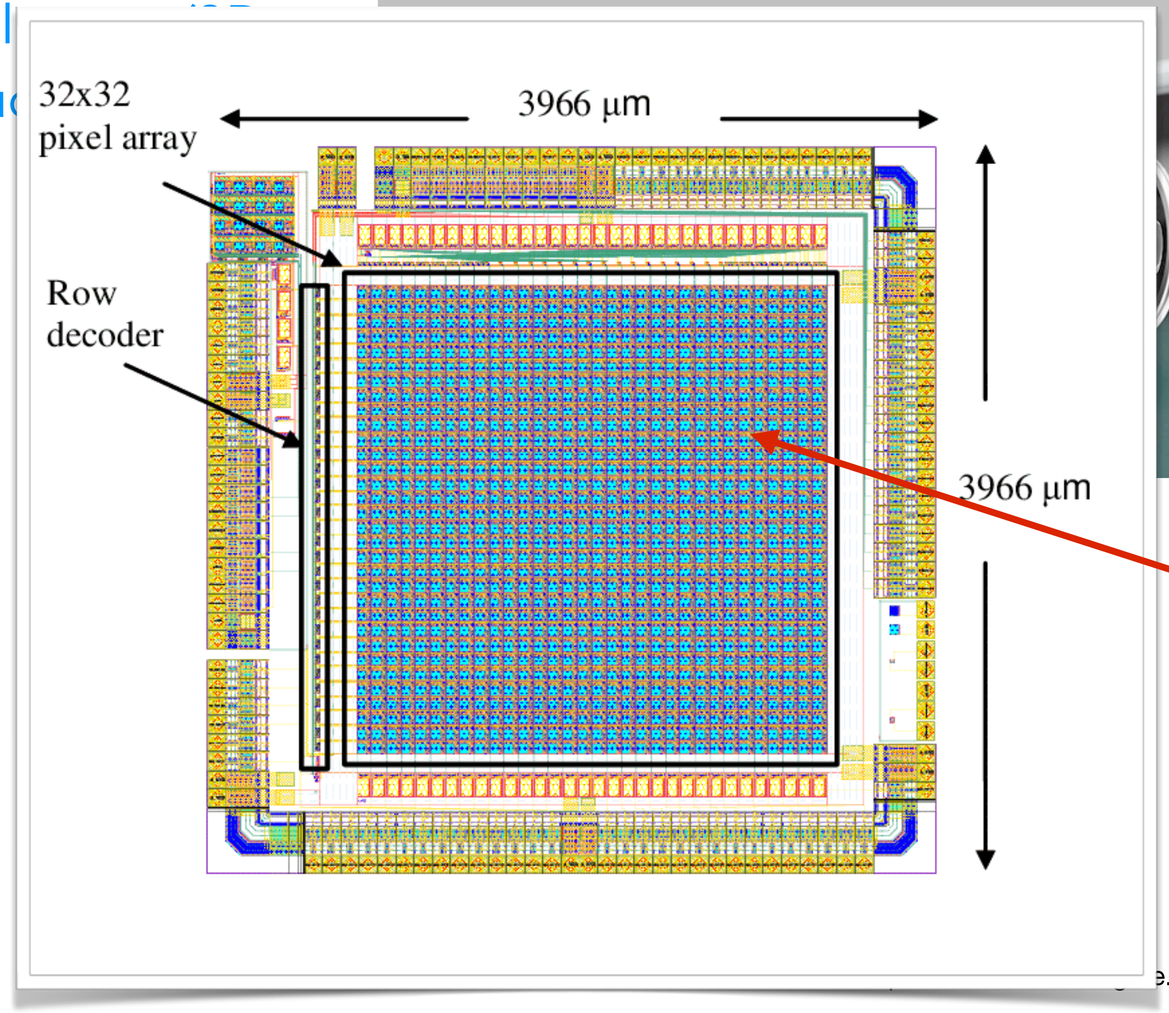

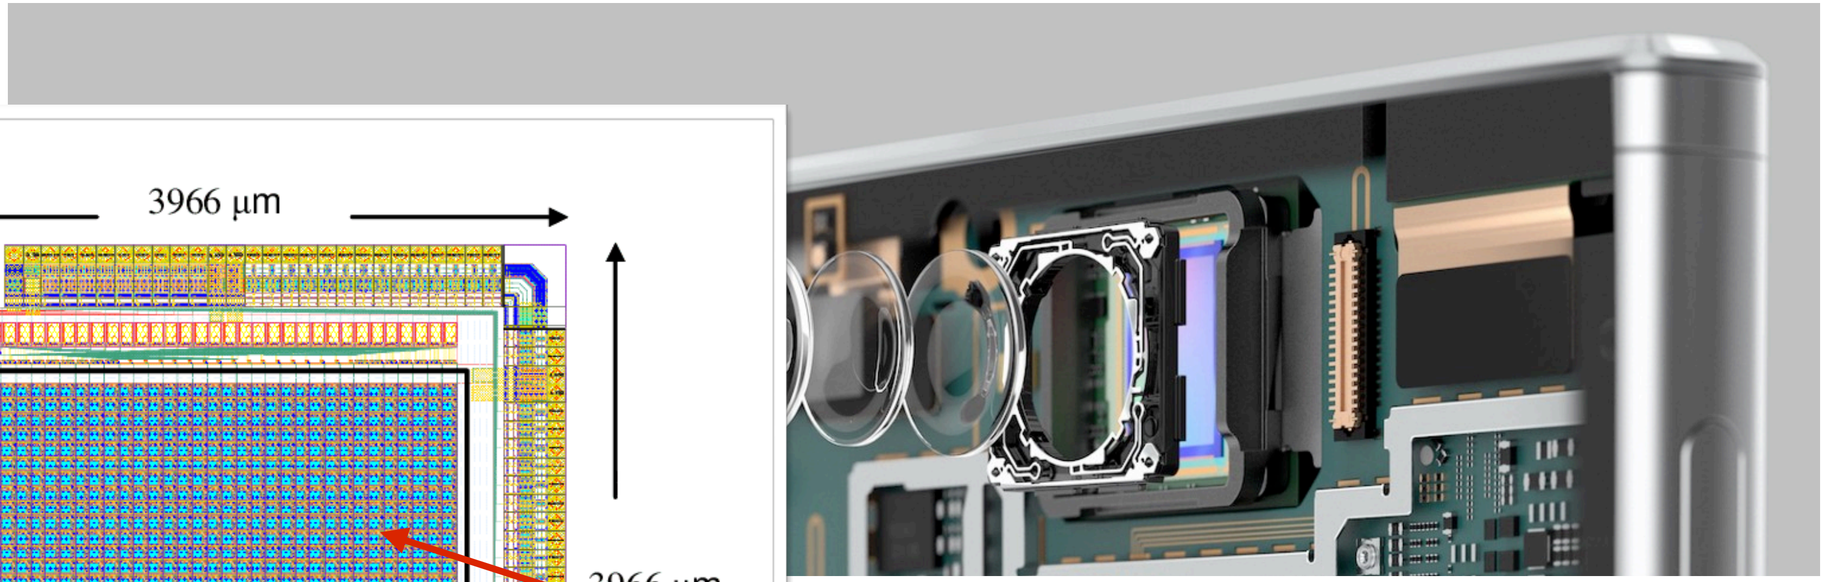

Pixels on the image sensor **sample** the continuous 2D function

Sampling and Reconstruction in Camera and Display

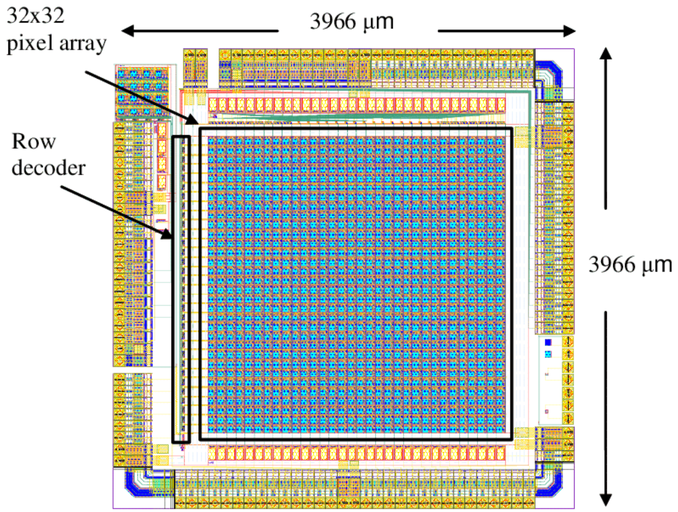

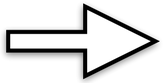

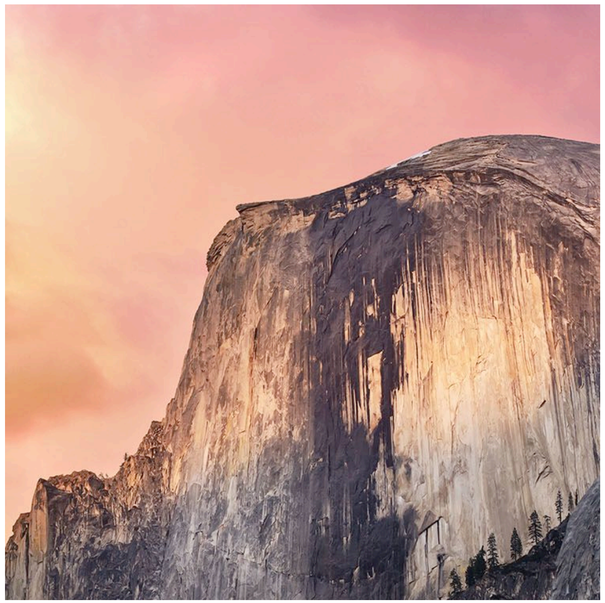

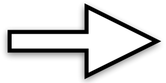

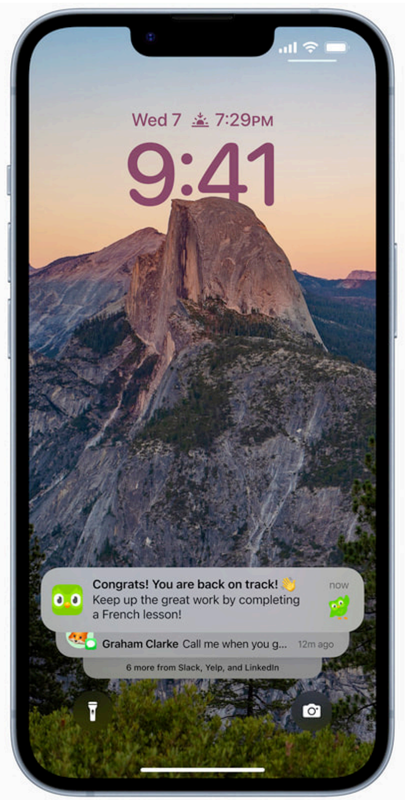

https://www.pinterest.com/pin/866943615545355525/

Sampling and Reconstruction in Camera and Display

Displays reconstruct continuous 2D functions from 2D samples

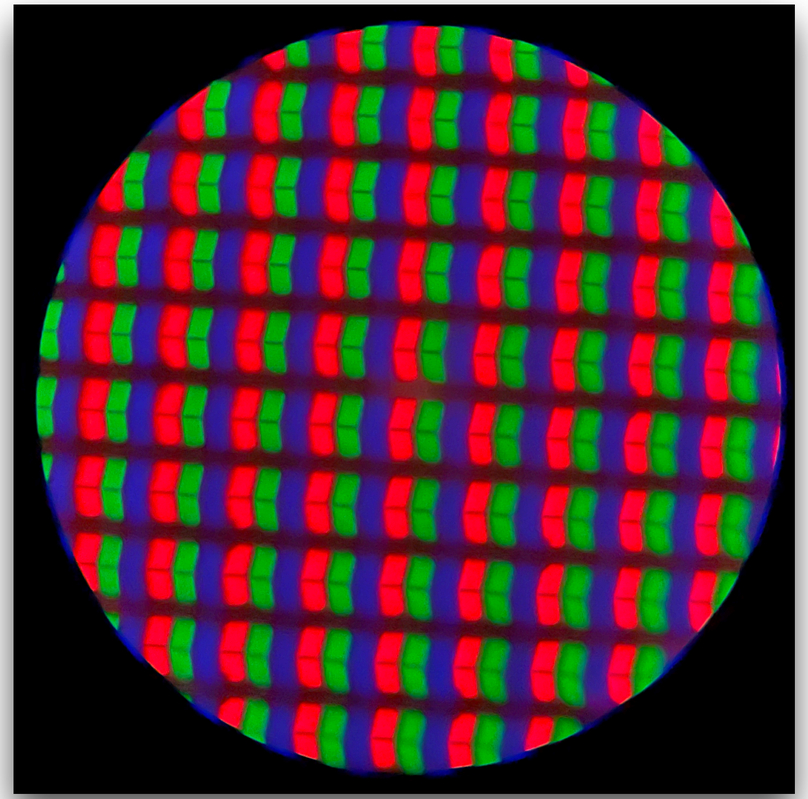

iPhone 6

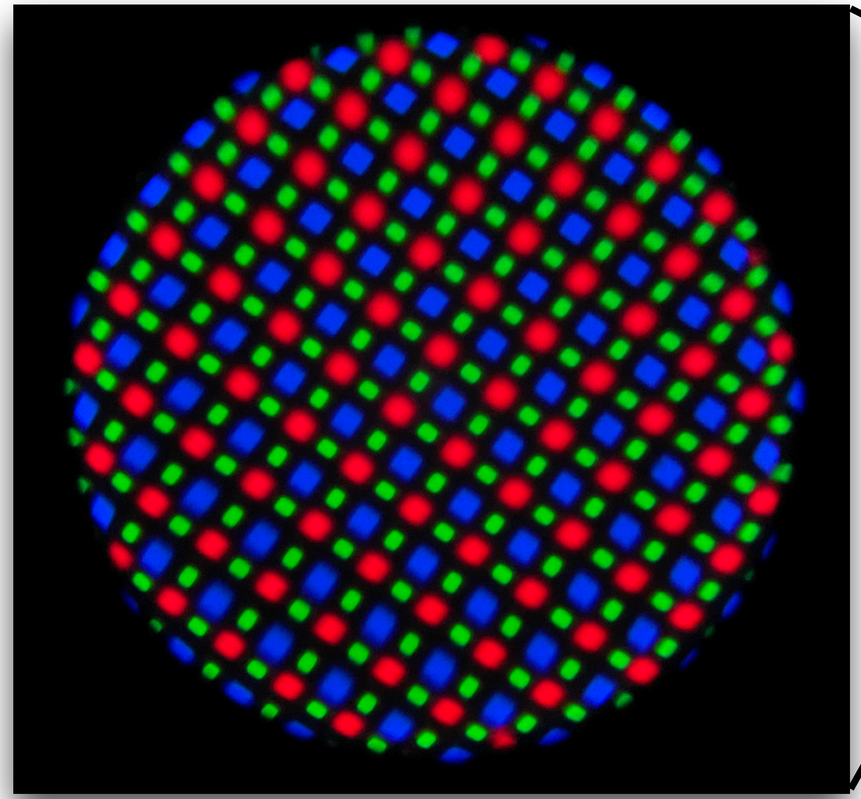

iPhone 11

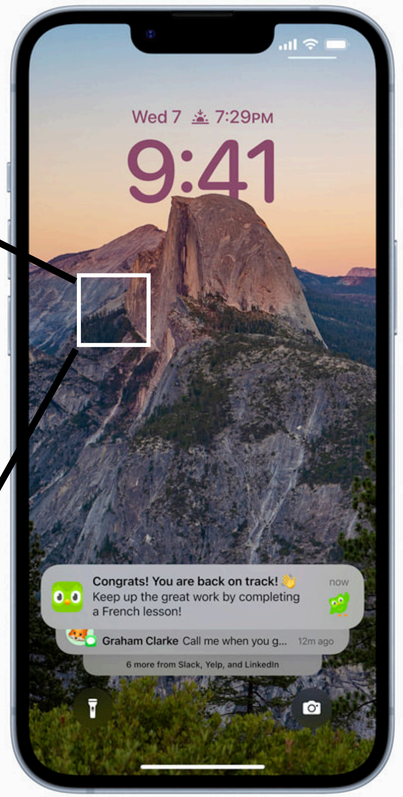

https://www.pinterest.com/pin/866943615545355525/

Sampling and Reconstruction in Camera and Display

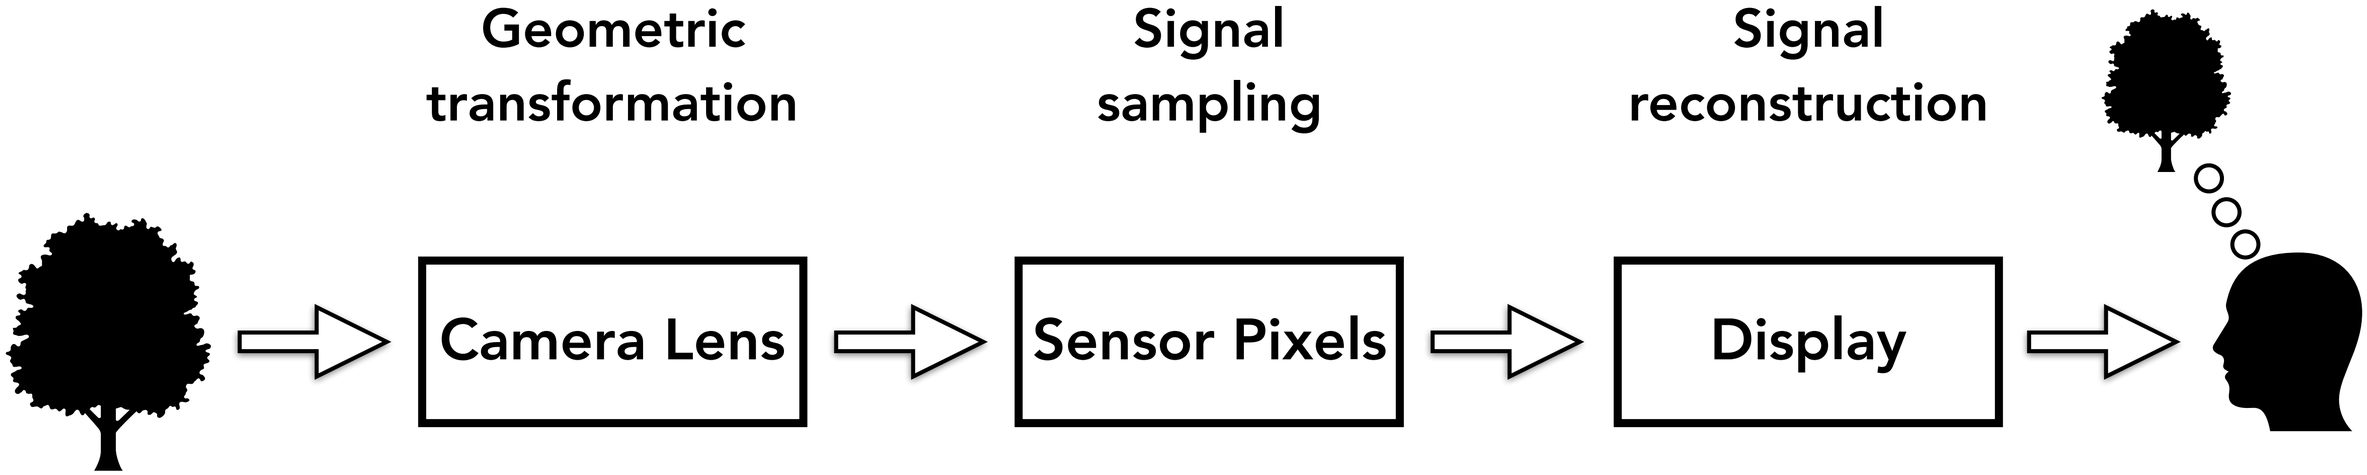

3D continuous signal

2D continuous signal (on the sensor plane)

2D sampled signal (i.e., image)

Reconstructed 3D continuous signal

Everything is Signal Sampling and Reconstruction

- **Physical scene**: continuous (at least) 3D function
- **Camera lens ⇒ Image sensor ⇒ Display**
- **Camera lens ⇒ Image sensor ⇒ Print/Paint**
- **Eye lens ⇒ Retina**

Anti-Aliasing Techniques

Insufficient Sampling and Reconstruction

https://svi.nl/AntiAliasing https://distance.ufhealth.org/preparing-for-lecture-capture/

“Just-Enough” Sampling

- Anti-aliasing by sampling at a high rate. But most real-world signals are not band-limited. No hope?

“Just-Enough” Sampling

- Anti-aliasing by sampling at a high rate. But most real-world signals are not band-limited. No hope?
- In many cases we just need to sample at a rate “just high enough”, because humans can’t sense signals above certain frequency anyways.
  - Acoustic; Vision

“Just-Enough” Sampling

Sampling rate	Use
8,000 Hz	Telephone and encrypted walkie-talkie, wireless intercom and wireless microphone transmission; adequate for human speech but without sibilance (*ess* sounds like *eff* (/s/, /f/)).
11,025 Hz	One quarter the sampling rate of audio CDs; used for lower-quality PCM, MPEG audio and for audio analysis of subwoofer bandpasses.[*citation needed*]
16,000 Hz	Wideband frequency extension over standard telephone narrowband 8,000 Hz. Used in most modern VoIP and VVoIP communication products.[14]
22,050 Hz	One half the sampling rate of audio CDs; used for lower-quality PCM and MPEG audio and for audio analysis of low frequency energy. Suitable for digitizing early 20th century audio formats such as 78s and AM Radio.[15]
32,000 Hz	miniDV digital video camcorder, video tapes with extra channels of audio (e.g. DVCAM with four channels of audio), DAT (LP mode), Germany's Digitales Satellitenradio, NICAM digital audio, used alongside analogue television sound in some countries. High-quality digital wireless microphones.[16] Suitable for digitizing FM radio.[*citation needed*]
37,800 Hz	CD-XA audio
44,056 Hz	Used by digital audio locked to NTSC *color* video signals (3 samples per line, 245 lines per field, 59.94 fields per second = 29.97 frames per second).
44,100 Hz	Audio CD, also most commonly used with MPEG-1 audio (VCD, SVCD, MP3). Originally chosen by Sony because it could be recorded on modified video equipment running at either 25 frames per second (PAL) or 30 frame/s (using an NTSC *monochrome* video recorder) and cover the 20 kHz bandwidth thought necessary to match professional analog recording equipment of the time. A PCM adaptor would fit digital audio samples into the analog video channel of, for example, PAL video tapes using 3 samples per line, 588 lines per frame, 25 frames per second.
47,250 Hz	world's first commercial PCM sound recorder by Nippon Columbia (Denon)
48,000 Hz	The standard audio sampling rate used by professional digital video equipment such as tape recorders, video servers, vision mixers and so on. This rate was chosen because it could reconstruct frequencies up to 22 kHz and work with 29.97 frames per second NTSC video – as well as 25 frame/s, 30 frame/s and 24 frame/s systems. With 29.97 frame/s systems it is necessary to handle 1601.6 audio samples per frame delivering an integer number of audio samples only every fifth video frame.[9] Also used for sound with consumer video formats like DV, digital TV, DVD, and films. The professional Serial Digital Interface (SDI) and High-definition Serial Digital Interface (HD-SDI) used to connect broadcast television equipment together uses this audio sampling frequency. Most professional audio gear uses 48 kHz sampling, including mixing consoles, and digital recording devices.

https://en.wikipedia.org/wiki/Sampling_(signal_processing)

Anti-Aliasing By Pre-Filtering

- If we can only sample at a rate of f_{sample}, pre-filter the signal to remove the frequency higher than $f_{sample}/2$.

Anti-Aliasing By Pre-Filtering

- If we can only sample at a rate of f_{sample}, pre-filter the signal to remove the frequency higher than $f_{sample}/2$.
- Then sample; won't see aliasing, but the reconstructed signal is blurred.

Anti-Aliasing By Pre-Filtering

- If we can only sample at a rate of f_{sample}, pre-filter the signal to remove the frequency higher than $f_{sample}/2$.
- Then sample; won't see aliasing, but the reconstructed signal is blurred.
- Blur is more acceptable visually than aliasing.

Pre-Filtering In the Frequency Domain

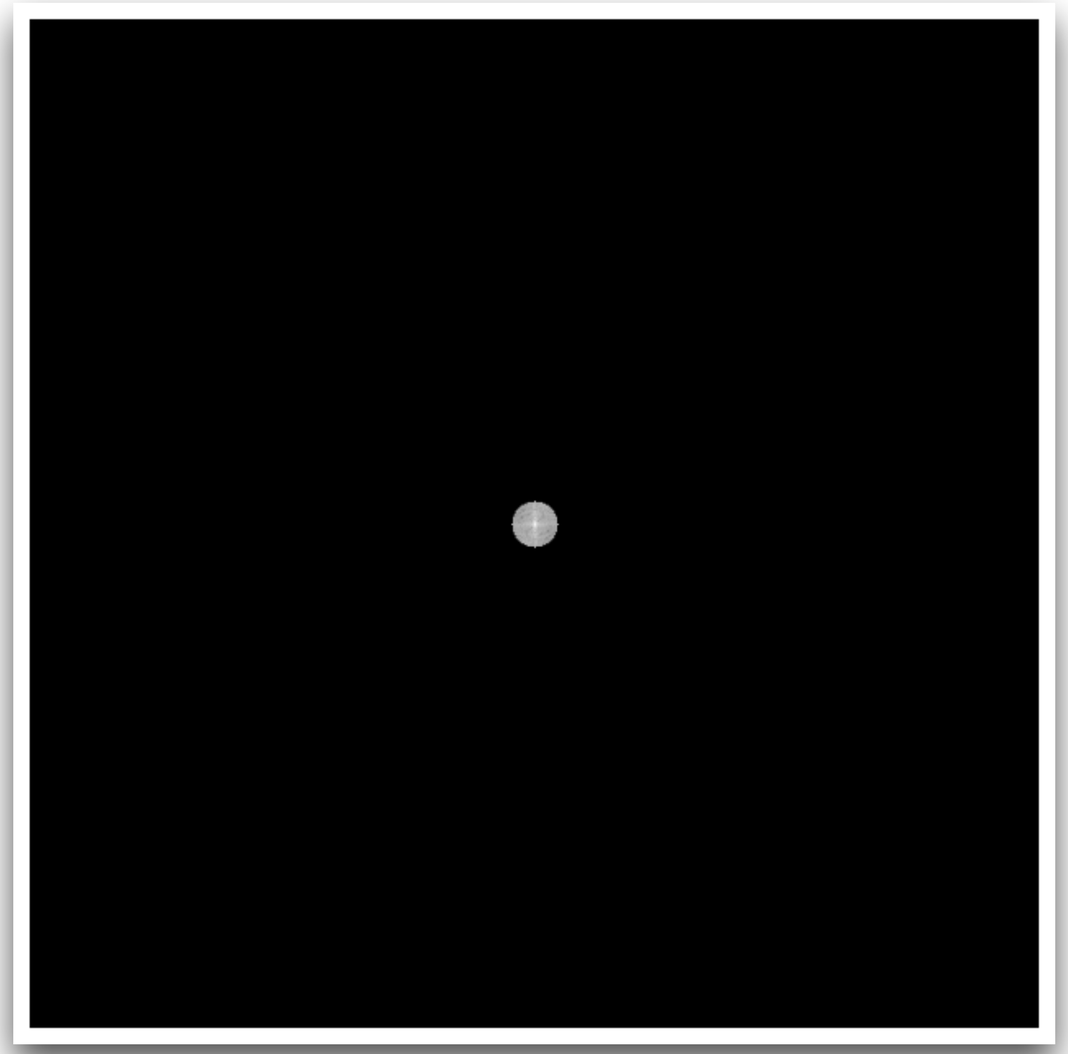

\

1D Discrete Convolution

1D Discrete Signal

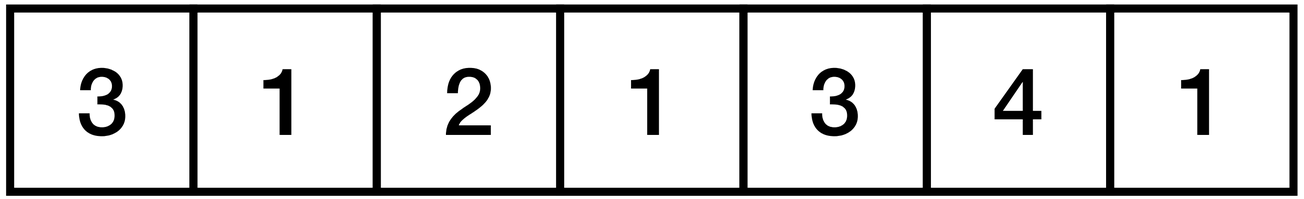

Filter/Kernel

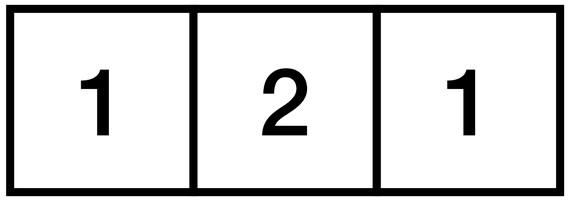

Convolution 3x1 + 1x2 + 2x1

Filtered Signal

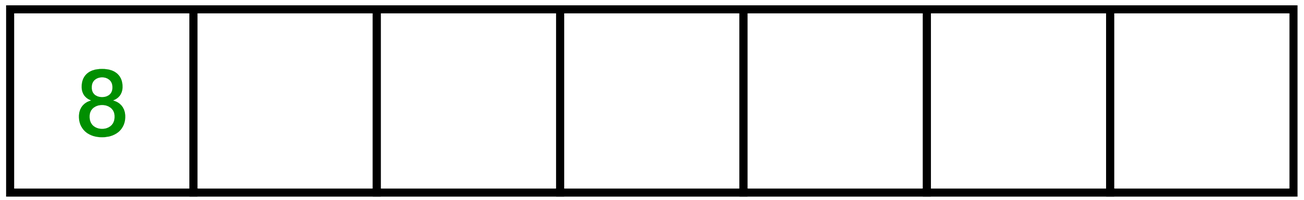

1D Discrete Convolution

1D Discrete Signal

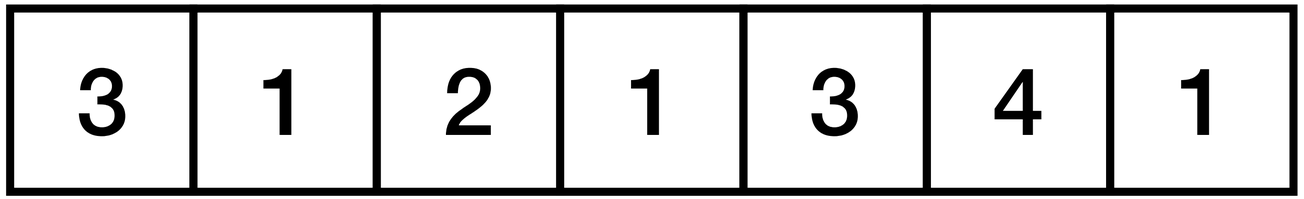

Filter/Kernel

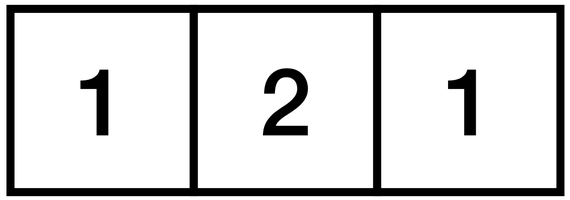

Convolution

1x1 + 2x2 + 1x1

Filtered Signal

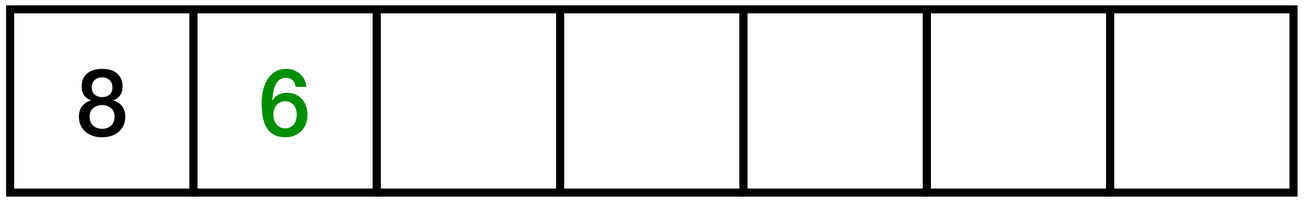

1D Discrete Convolution

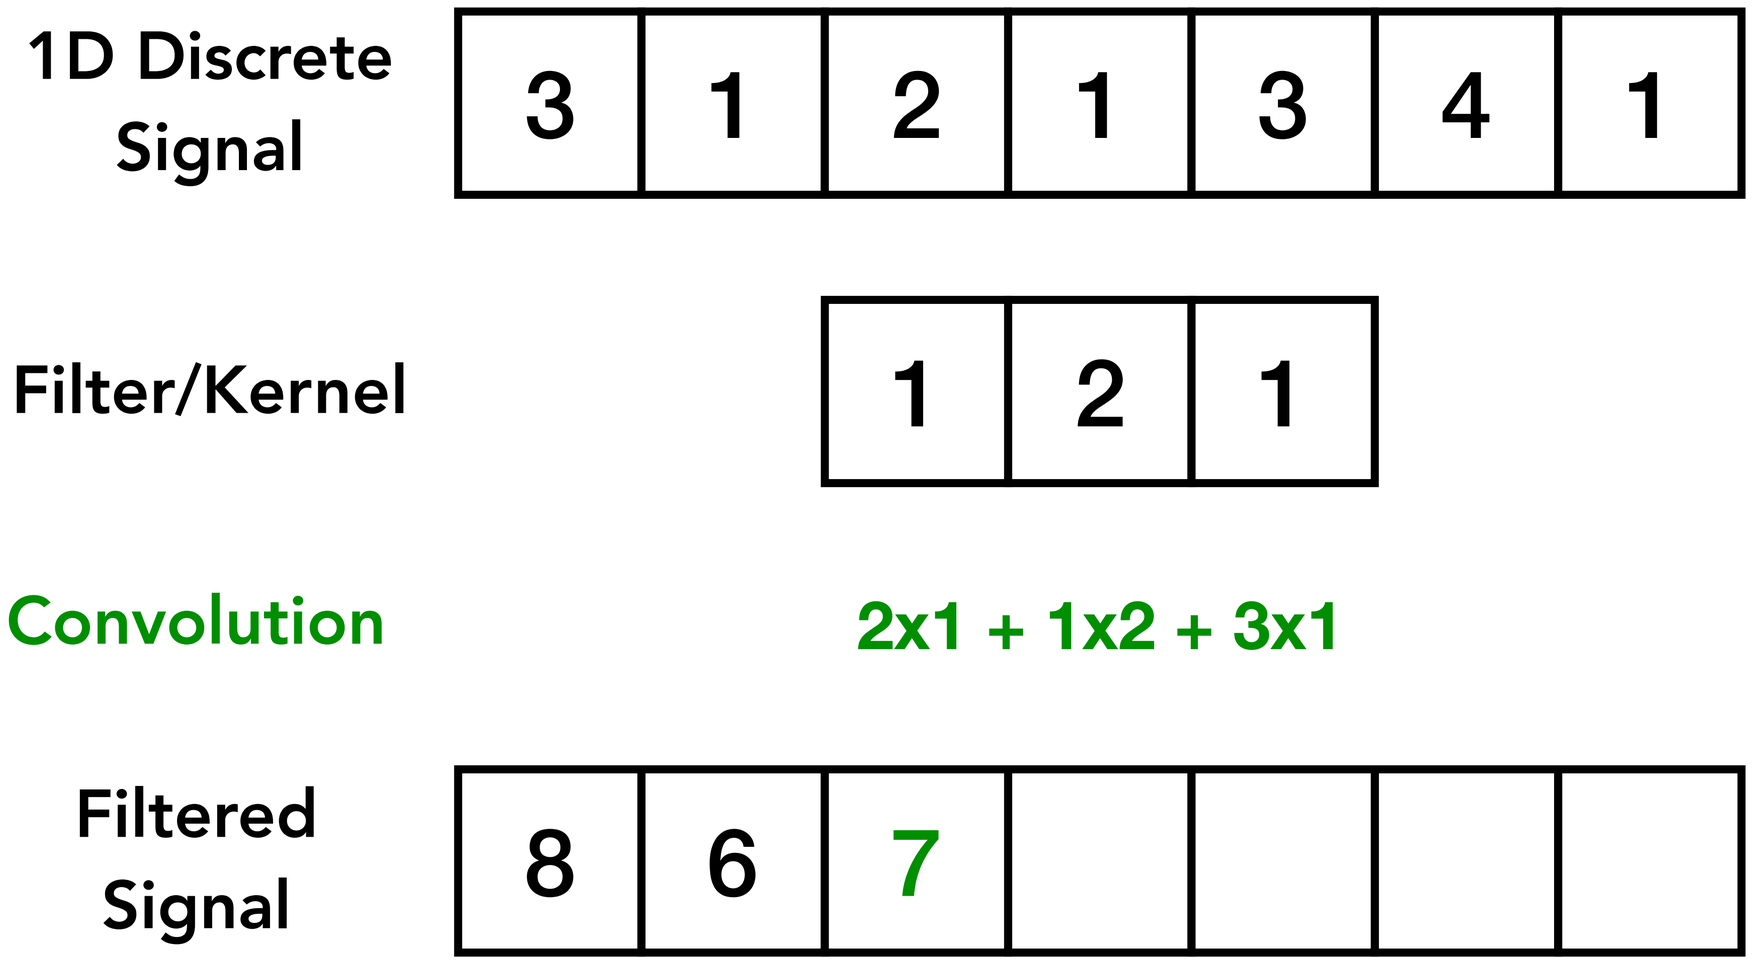

2D Discrete Convolution

2D Discrete Signal

3	1	2	1	3	4	1
2	4	0	1	10	2	0
0	2	4	21	9	1	14
34	5	4	7	8	90	34
54	6	8	9	13	36	4
6	8	14	2	4	8	52
32	14	54	3	6	8	0

2D Filter

1/9	1/9	1/9
1/9	1/9	1/9
1/9	1/9	1/9

Filtered Signal

2D Discrete Convolution

2D Discrete Signal

3	1	2	1	3	4	1
2	4	0	1	10	2	0
0	2	4	21	9	1	14
34	5	4	7	8	90	34
54	6	8	9	13	36	4
6	8	14	2	4	8	52
32	14	54	3	6	8	0

2D Filter

1/9	1/9	1/9
1/9	1/9	1/9
1/9	1/9	1/9

Filtered Signal

2						

2D Discrete Convolution

2D Discrete Signal

3	1	2	1	3	4	1
2	4	0	1	10	2	0
0	2	4	21	9	1	14
34	5	4	7	8	90	34
54	6	8	9	13	36	4
6	8	14	2	4	8	52
32	14	54	3	6	8	0

2D Filter

1/9	1/9	1/9
1/9	1/9	1/9
1/9	1/9	1/9

Filtered Signal

2	4					

Box Filter

2D Discrete Signal

3	1	2	1	3	4	1
2	4	0	1	10	2	0
0	2	4	21	9	1	14
34	5	4	7	8	90	34
54	6	8	9	13	36	4
6	8	14	2	4	8	52
32	14	54	3	6	8	0

Box Filter

1/9	1/9	1/9
1/9	1/9	1/9
1/9	1/9	1/9

Box Filter

2D Discrete Signal

3	1	2	1	3	4	1
2	4	0	1	10	2	0
0	2	4	21	9	1	14
34	5	4	7	8	90	34
54	6	8	9	13	36	4
6	8	14	2	4	8	52
32	14	54	3	6	8	0

Box Filter

1/9	1/9	1/9
1/9	1/9	1/9
1/9	1/9	1/9

A box filter makes pixels more similar to its neighbors, i.e., blur.

Box (Mean/Moving Average) Filter

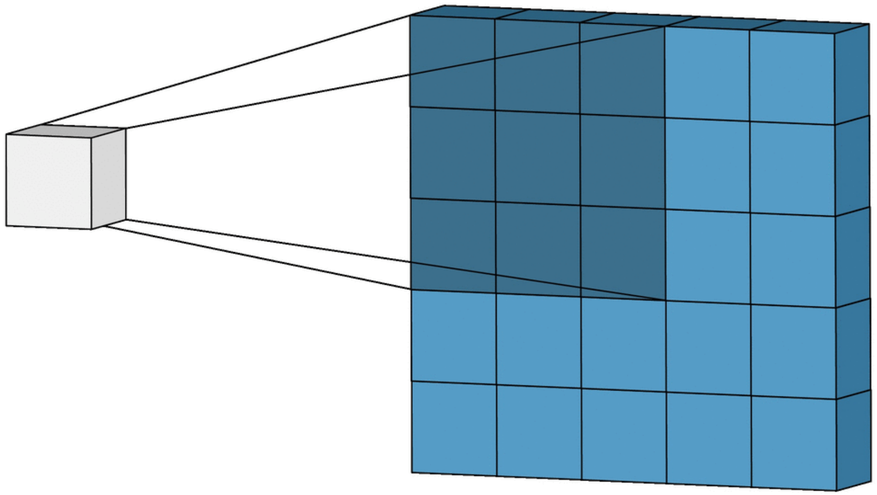

2D Box Filter

1/9	1/9	1/9
1/9	1/9	1/9
1/9	1/9	1/9

* Output size is smaller than input size without padding.

https://towardsdatascience.com/intuitively-understanding-convolutions-for-deep-learning-1f6f42faee1

Box (Mean/Moving Average) Filter

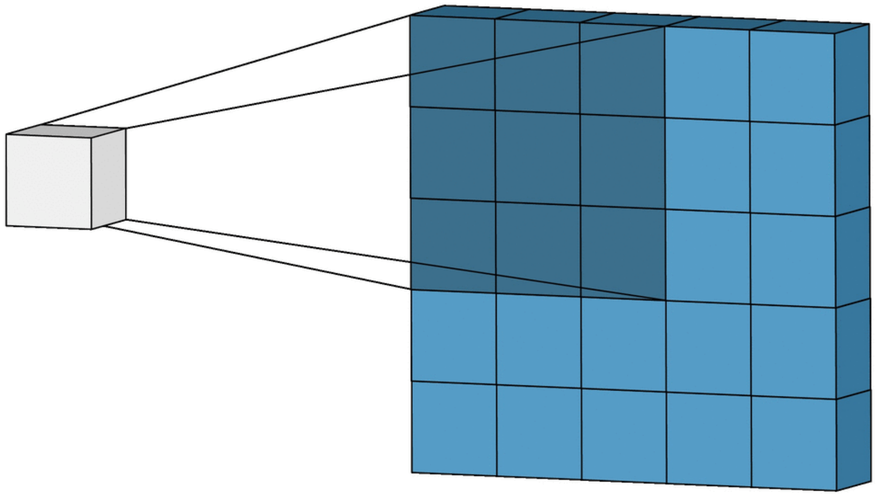

2D Box Filter

1/9	1/9	1/9
1/9	1/9	1/9
1/9	1/9	1/9

* Output size is smaller than input size without padding.

https://towardsdatascience.com/intuitively-understanding-convolutions-for-deep-learning-1f6f42faee1

Gaussian Filter

$$G(x,y) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}}$$

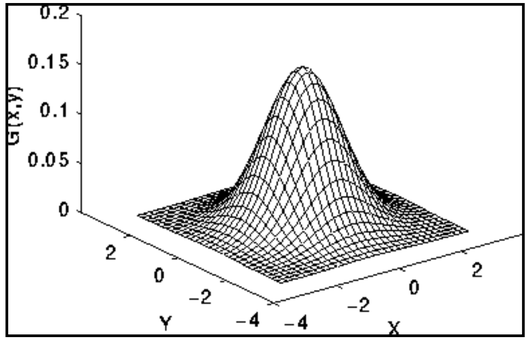

2D Gaussian distribution

$$\frac{1}{273}$$

1	4	7	4	1
4	16	26	16	4
7	26	41	26	7
4	16	26	16	4
1	4	7	4	1

A sample 2D Gaussian kernel with mean [0, 0] and σ=1

A Gaussian filter also averages neighboring pixels, but gives more weight to closer neighbors. It's still a low-pass filter.

https://homepages.inf.ed.ac.uk/rbf/HIPR2/gsmooth.htm

Convolution

1D Discrete Convolution

$$f[x] \star g[x] = \sum_{k=-\infty}^{k=\infty} f[k]g[x-k]$$

1D Continuous Convolution

$$f(x) \star g(x) = \int_{-\infty}^{\infty} f(\tau)g(x-\tau)d\tau$$

2D Discrete Convolution

$$f[x,y] \star g[x,y] = \sum_{i=-\infty}^{i=\infty} \sum_{j=-\infty}^{j=\infty} f[i,j]g[x-i,y-j]$$

2D Continuous Convolution

$$f(x) \star g(x) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(\tau,\eta)g(x-\tau,y-\eta)d\tau d\eta$$

Convolution Theorem

$$\mathcal{F}(f \star g) = \mathcal{F}(f)\mathcal{F}(g) \qquad f \star g = \mathcal{F}^{-1}(\mathcal{F}(f)\mathcal{F}(g))$$

- Spatial domain convolution is equivalent to frequency domain multiplication, and vice versa.
- Why useful? In many cases the spatial domain signal is not given analytically, so we can't have its spectrum analytically.
  - The only thing we can do is to manipulate the spatial domain data through convolution.

Convolution Theorem

Spatial Domain

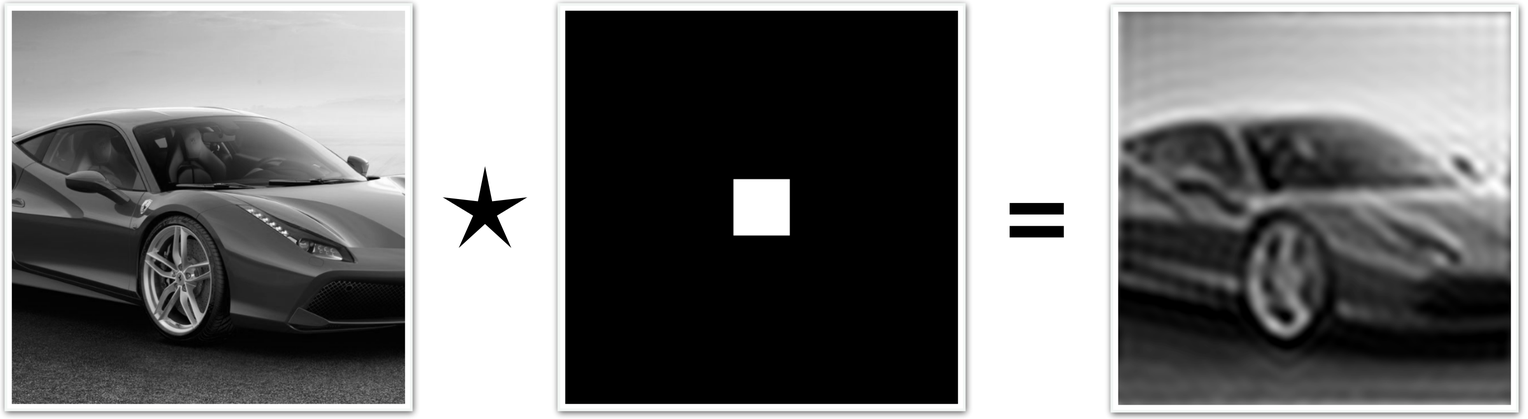

Fourier Transform

Inv. Fourier Transform

Frequency Domain

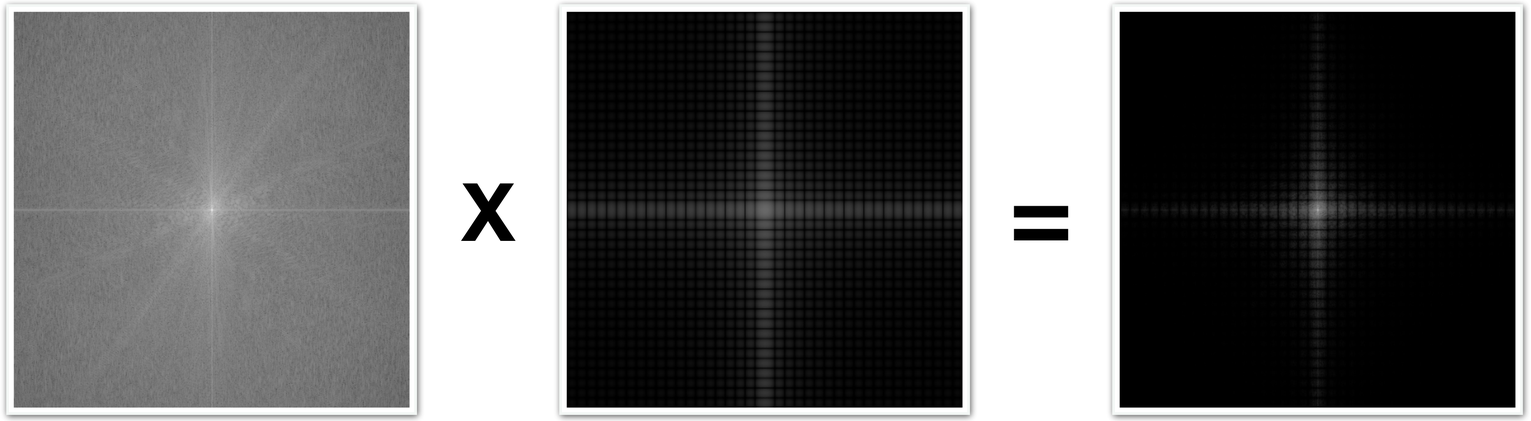

Box Filter in Spatial and Frequency Domains

Multiplying with this spectrum
attenuates high-frequency components.

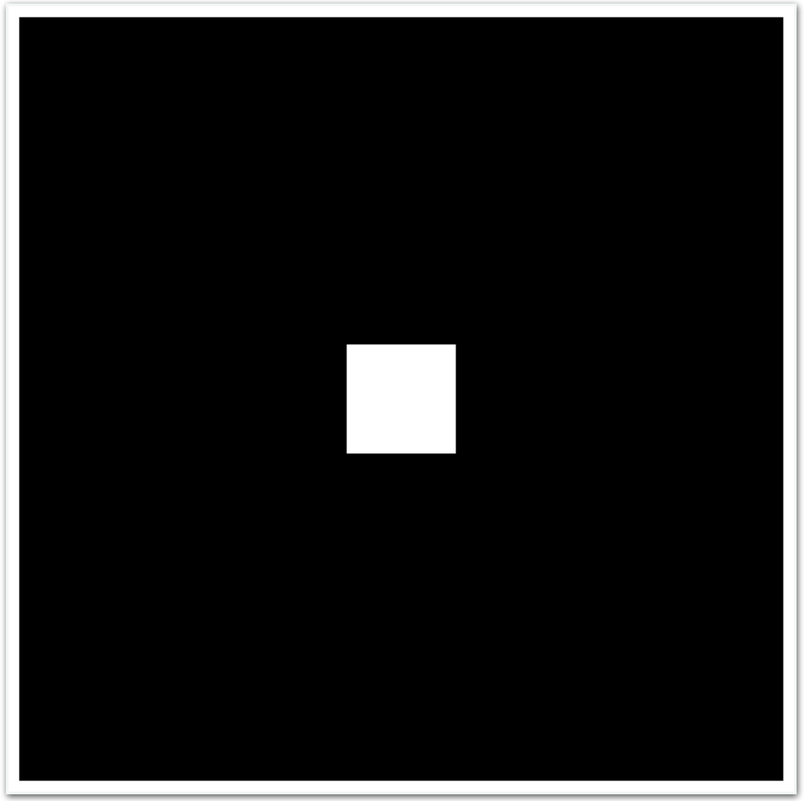

Spatial Domain

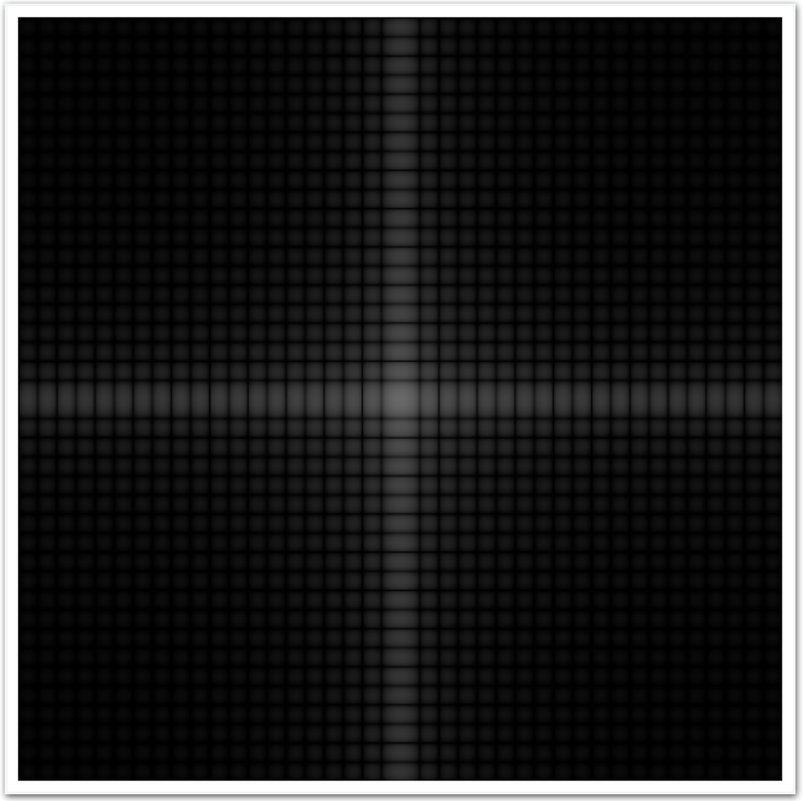

Frequency Domain

Box Filter in Spatial and Frequency Domains

Wider box attenuates high frequencies even more (averaging over a larger window)

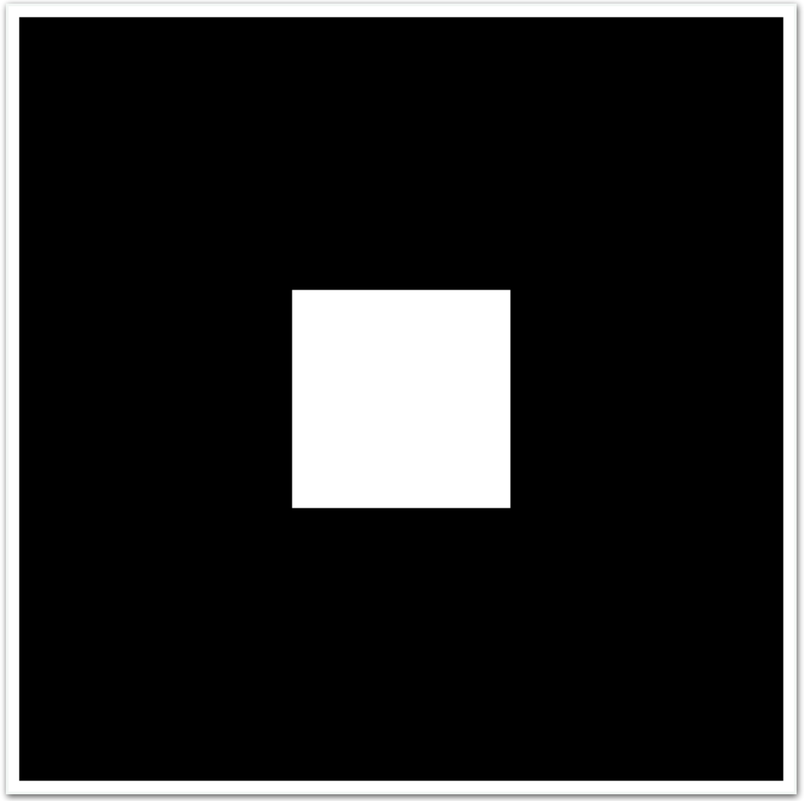

Spatial Domain

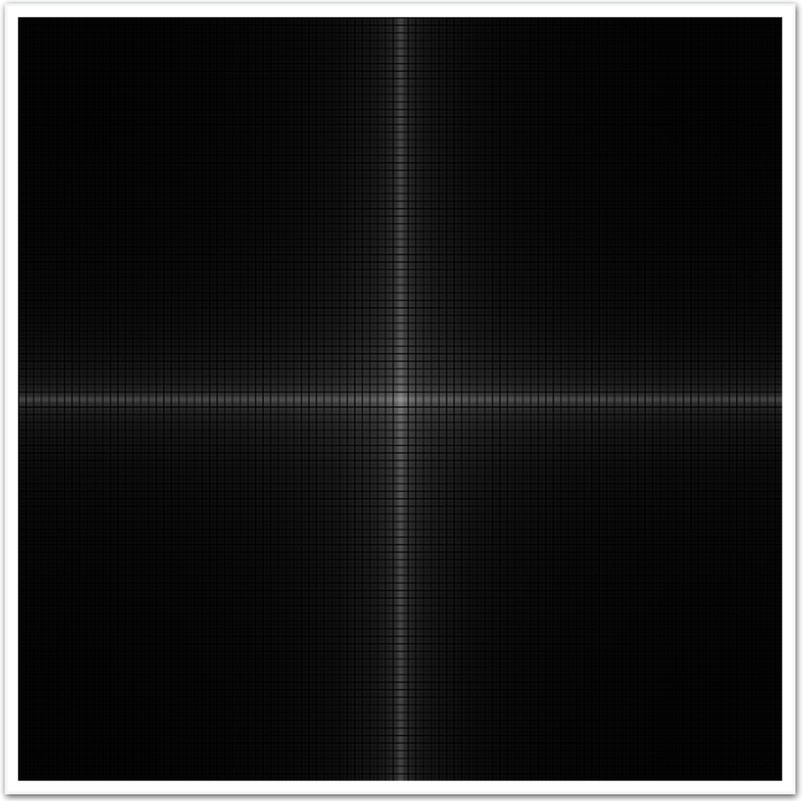

Frequency Domain

Revisiting Ideal Signal Reconstruction

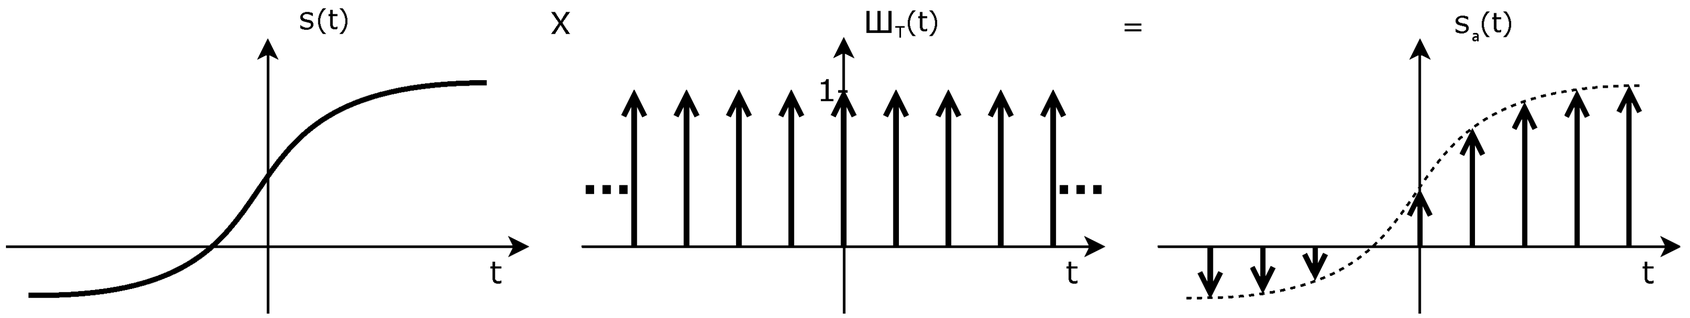

The reconstruction equation

$$s(t) = \mathcal{F}^{-1}[\ \mathcal{F}(S_a(t))B_{1/2T}(\omega)\ ]$$

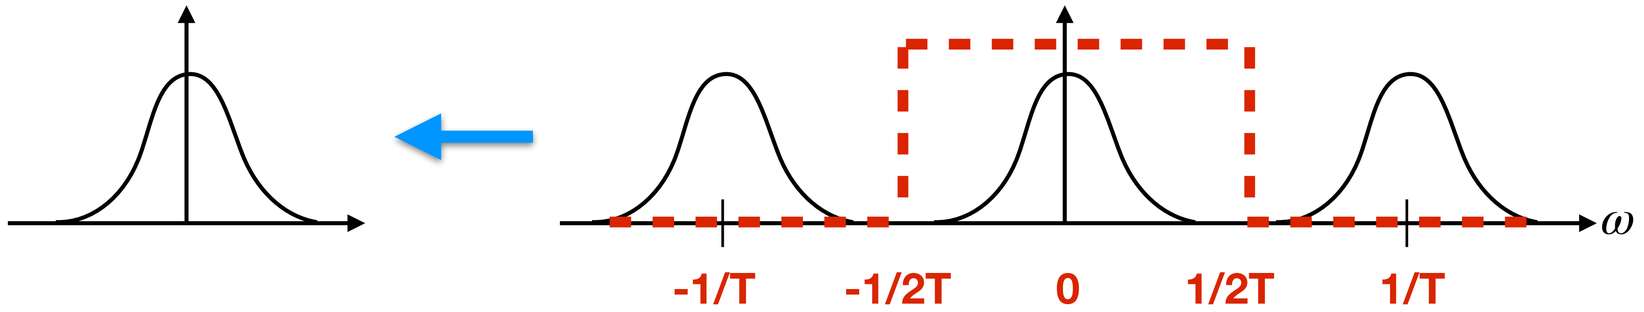

Revisiting Ideal Signal Reconstruction

Reconstructed signal

Sampled signal

$$s(t) = \mathcal{F}^{-1}[\ \mathcal{F}(S_a(t))B(\omega)\ ]$$

$$f \star g = \mathcal{F}^{-1}(\mathcal{F}(f)\mathcal{F}(g))$$

Box $\xrightarrow[\text{IFFT}]{}$ sinc

https://web.cs.ucdavis.edu/~okreylos/PhDStudies/Winter2000/SamplingTheory.html

Revisiting Ideal Signal Reconstruction

Reconstructed signal

Sampled signal

$$s(t) = \mathcal{F}^{-1}[\ \mathcal{F}(S_a(t))B(\omega)\ ]$$

$$f \star g = \mathcal{F}^{-1}(\mathcal{F}(f)\mathcal{F}(g))$$

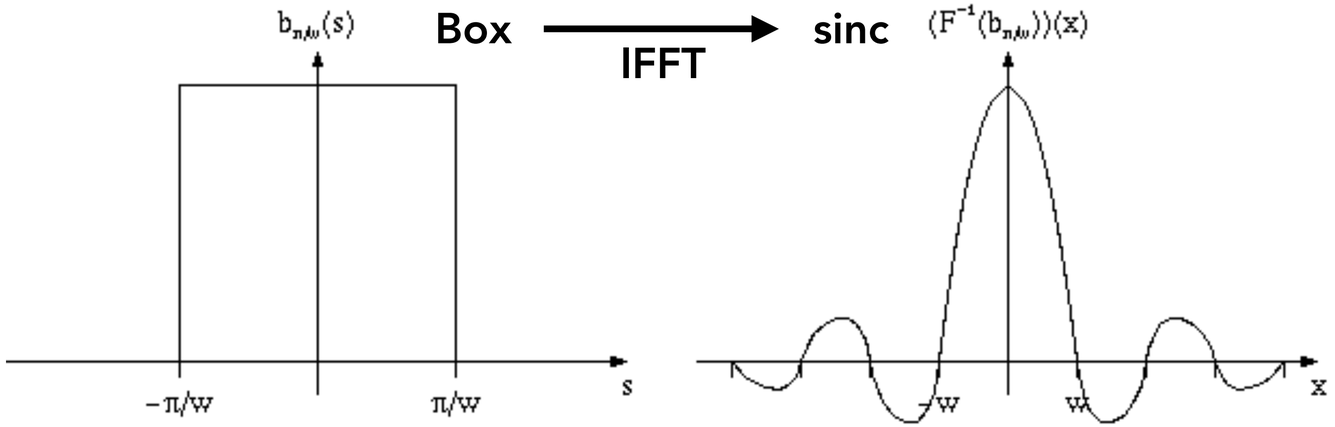

https://web.cs.ucdavis.edu/~okreylos/PhDStudies/Winter2000/SamplingTheory.html

Revisiting Ideal Signal Reconstruction

Reconstructed signal

Sampled signal

$$f \star g = \mathcal{F}^{-1}(\mathcal{F}(f)\mathcal{F}(g))$$

$$s(t) = \mathcal{F}^{-1}[\ \mathcal{F}(S_a(t))B(\omega)\ ]$$

$$s(t) = S_a(t) \star sinc(t)$$

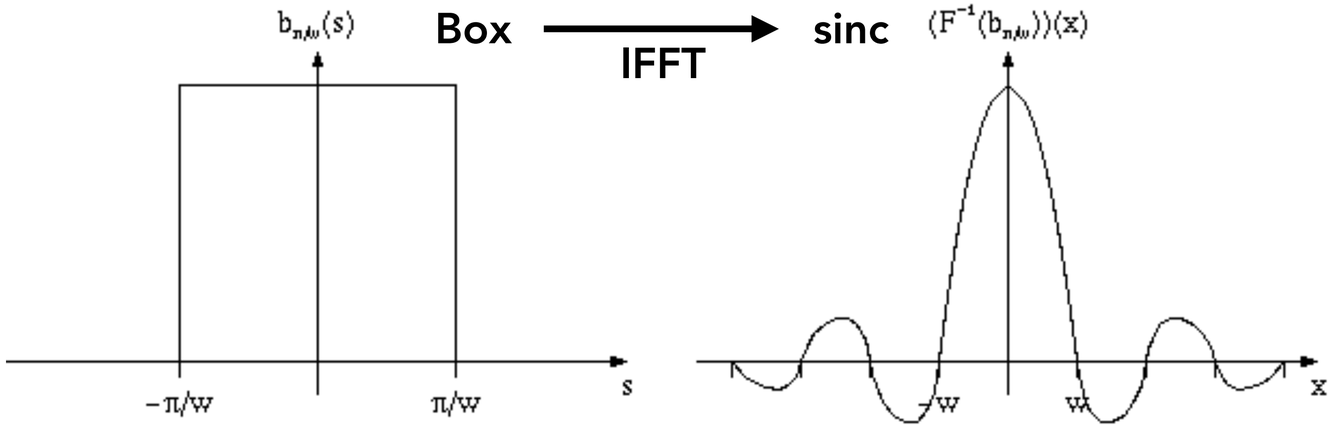

Anti-Aliased Signal Reconstruction

Anti-aliasing filter
(e.g., a box filter)

Reconstruction filter
(e.g., a sinc filter)

$$s(t) = S_a(t) \star F_1(t) \star F_2(t)$$

- Convolving with multiple filters is equivalent to convolving with one composite filter
- People come up many empirical filters as a composite filter

Beating Nyquist-Shannon Sampling Theorem

- Nyquist-Shannon sampling theorem applies to generic signals

https://onlinelibrary.wiley.com/doi/10.1002/jmri.25547

Beating Nyquist-Shannon Sampling Theorem

- Nyquist-Shannon sampling theorem applies to generic signals
- If signals have strong patterns (mostly sparse in the Fourier domain), we can sample at a much lower rate than the Nyquist rate but still obtain a good reconstruction

https://onlinelibrary.wiley.com/doi/10.1002/jmri.25547

Beating Nyquist-Shannon Sampling Theorem

- Nyquist-Shannon sampling theorem applies to generic signals
- If signals have strong patterns (mostly sparse in the Fourier domain), we can sample at a much lower rate than the Nyquist rate but still obtain a good reconstruction
  - Through compressive sensing

https://onlinelibrary.wiley.com/doi/10.1002/jmri.25547

Beating Nyquist-Shannon Sampling Theorem

- Nyquist-Shannon sampling theorem applies to generic signals
- If signals have strong patterns (mostly sparse in the Fourier domain), we can sample at a much lower rate than the Nyquist rate but still obtain a good reconstruction
  - Through compressive sensing

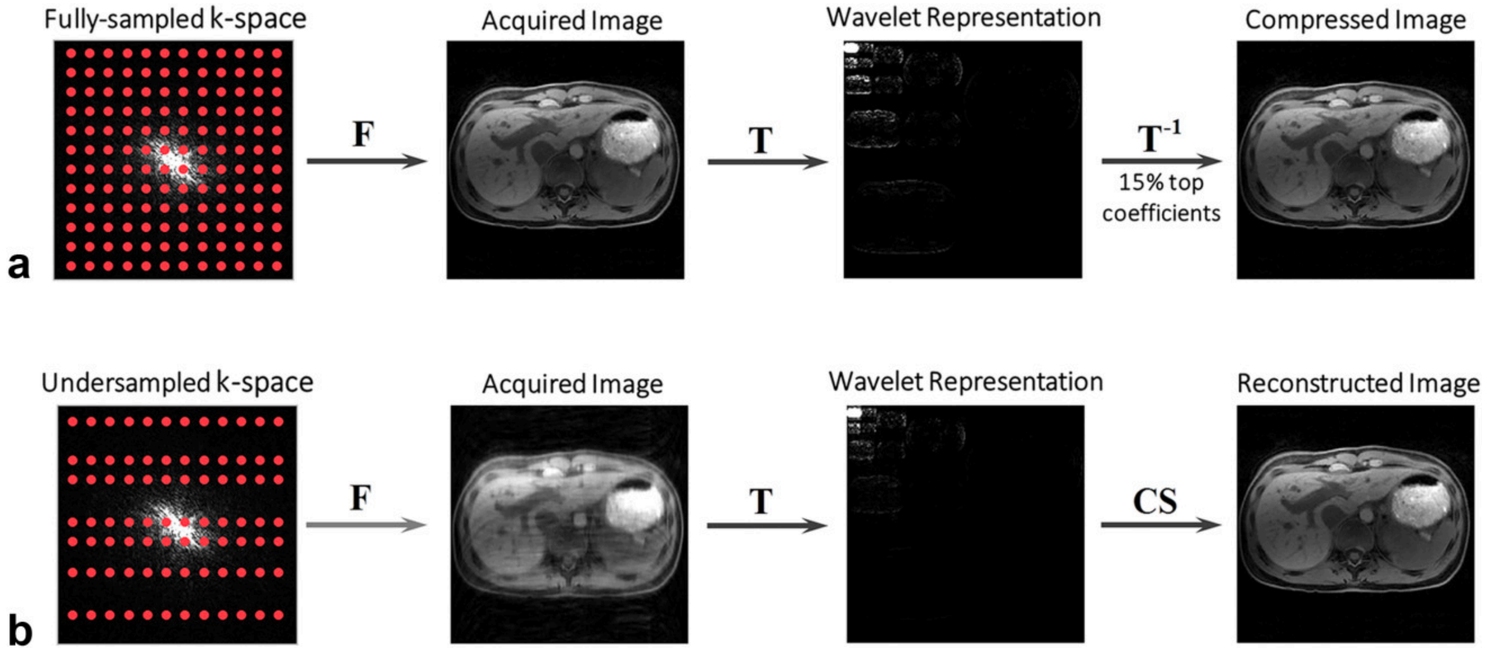

https://onlinelibrary.wiley.com/doi/10.1002/jmri.25547

Beating Nyquist-Shannon Sampling Theorem

- Nyquist-Shannon sampling theorem applies to generic signals
- If signals have strong patterns (mostly sparse in the Fourier domain), we can sample at a much lower rate than the Nyquist rate but still obtain a good reconstruction
  - Through compressive sensing
  - "Single-pixel" camera

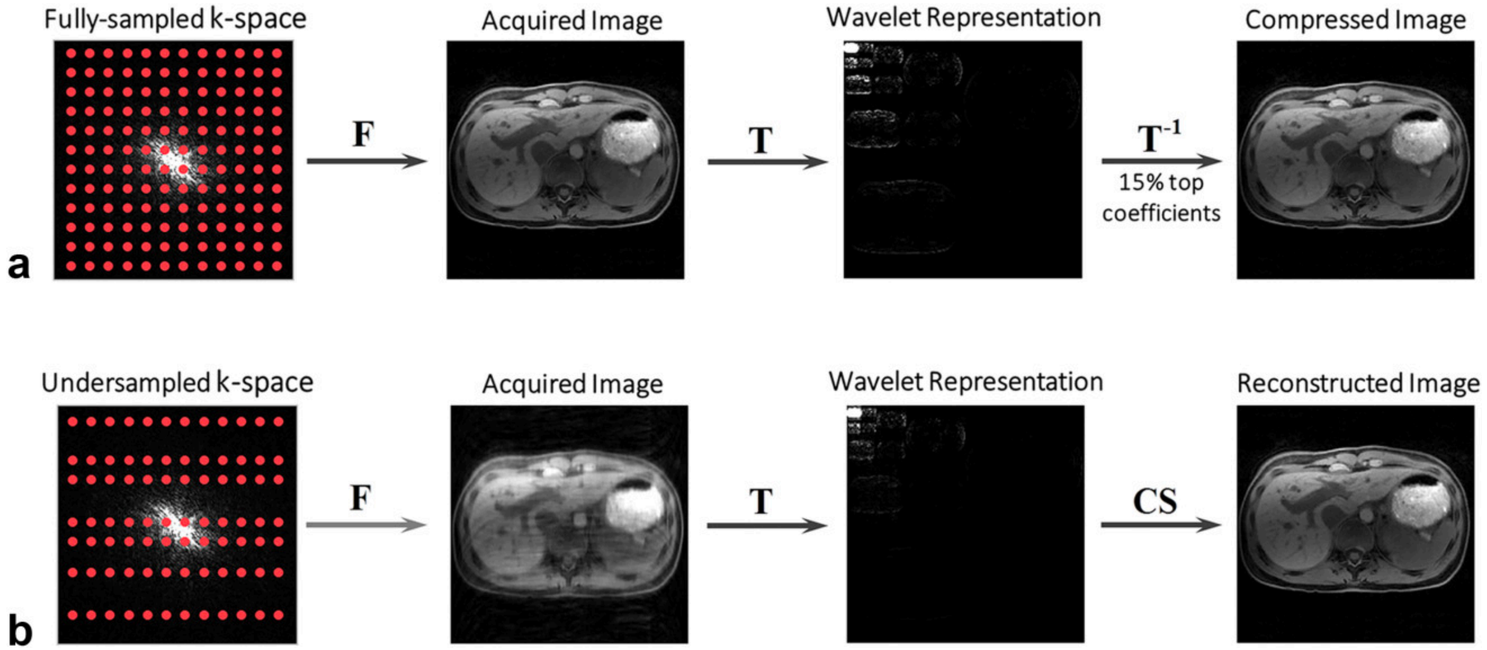

https://onlinelibrary.wiley.com/doi/10.1002/jmri.25547

Beating Nyquist-Shannon Sampling Theorem

- Nyquist-Shannon sampling theorem applies to generic signals
- If signals have strong patterns (mostly sparse in the Fourier domain), we can sample at a much lower rate than the Nyquist rate but still obtain a good reconstruction
  - Through compressive sensing
  - "Single-pixel" camera
  - Must do random sampling

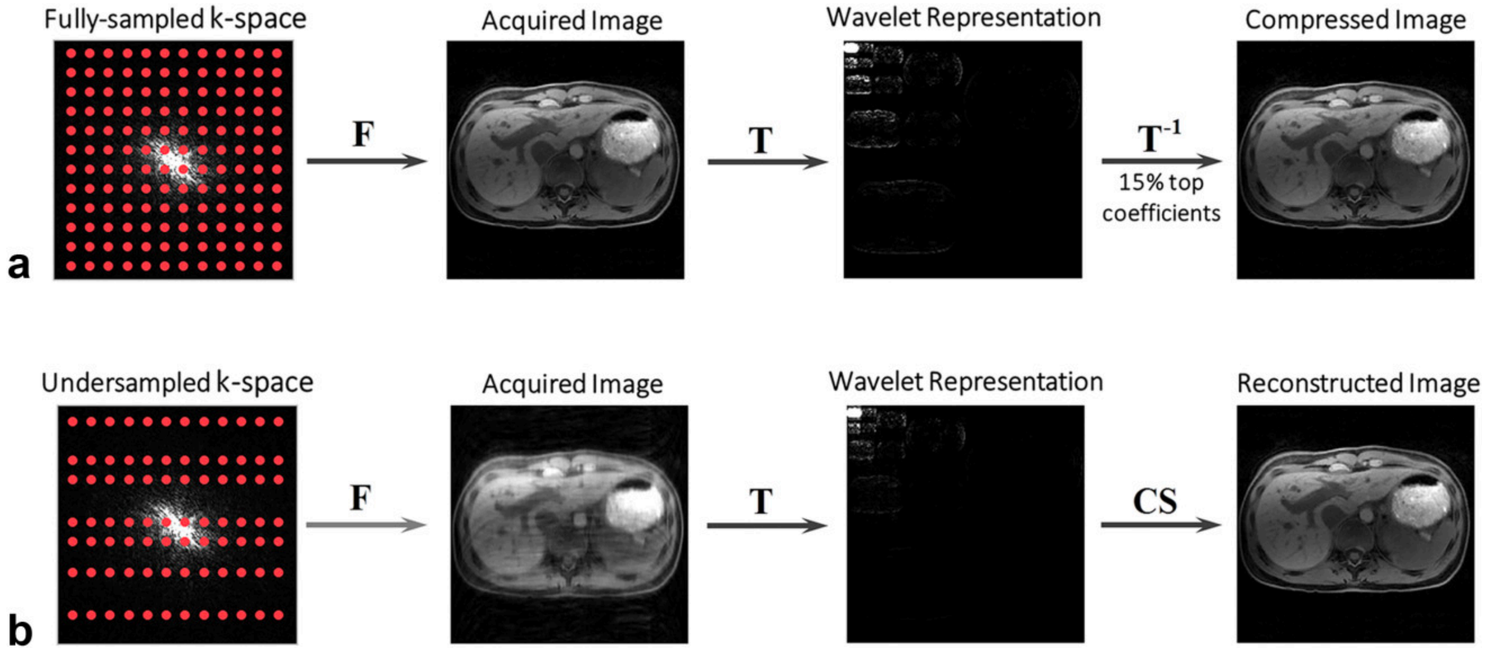

https://onlinelibrary.wiley.com/doi/10.1002/jmri.25547

Beating Nyquist-Shannon Sampling Theorem

- Nyquist-Shannon sampling theorem applies to reconstruction from samples of a single signal without help from “others”

Beating Nyquist-Shannon Sampling Theorem

- Nyquist-Shannon sampling theorem applies to reconstruction from samples of a single signal without help from “others”
- We could also learn from prior data to reconstruct signals
  - Through machine/deep learning
  - Image inpainting, super-resolution

Beating Nyquist-Shannon Sampling Theorem

(a) Face Editing

(b) Object Removal

https://www.slashgear.com/google-pixel-4-camera-review-brilliant-and-frustrating-21596272/

https://pythonawesome.com/aot-gan-for-high-resolution-image-inpainting/

Upsampling and Downsampling

Naive Upsampling

Linear interpolation

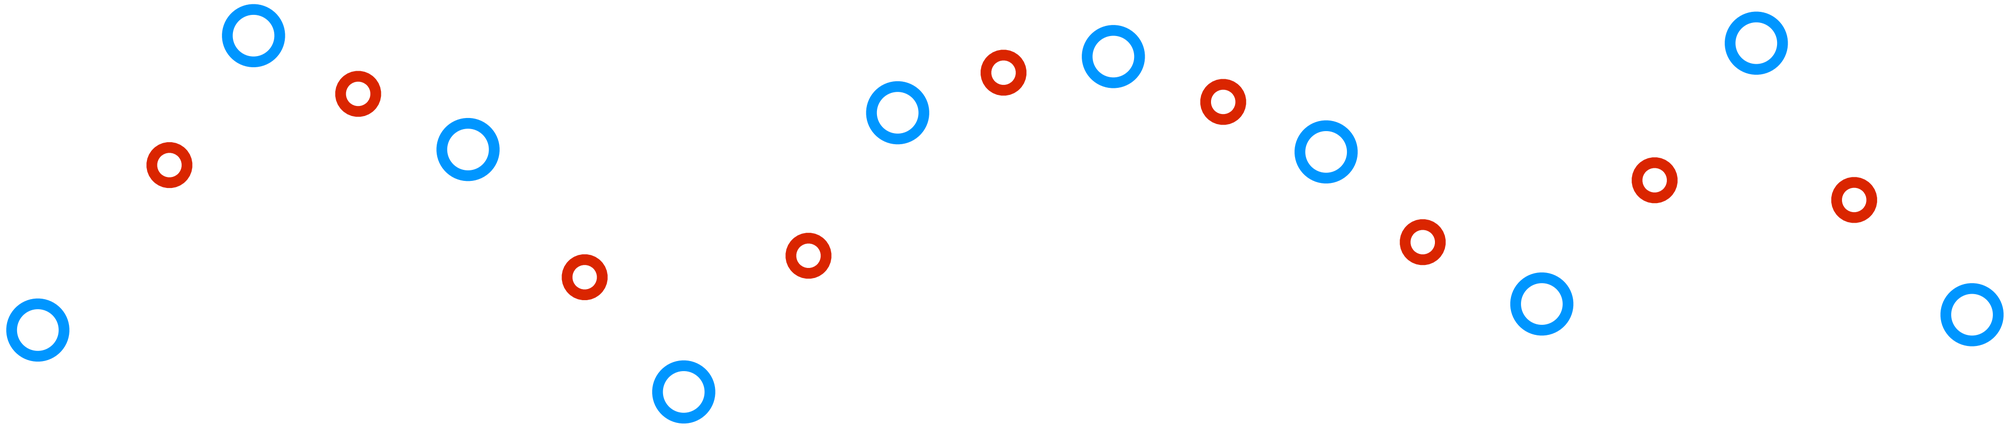

Naive Upsampling

"Sample and hold"

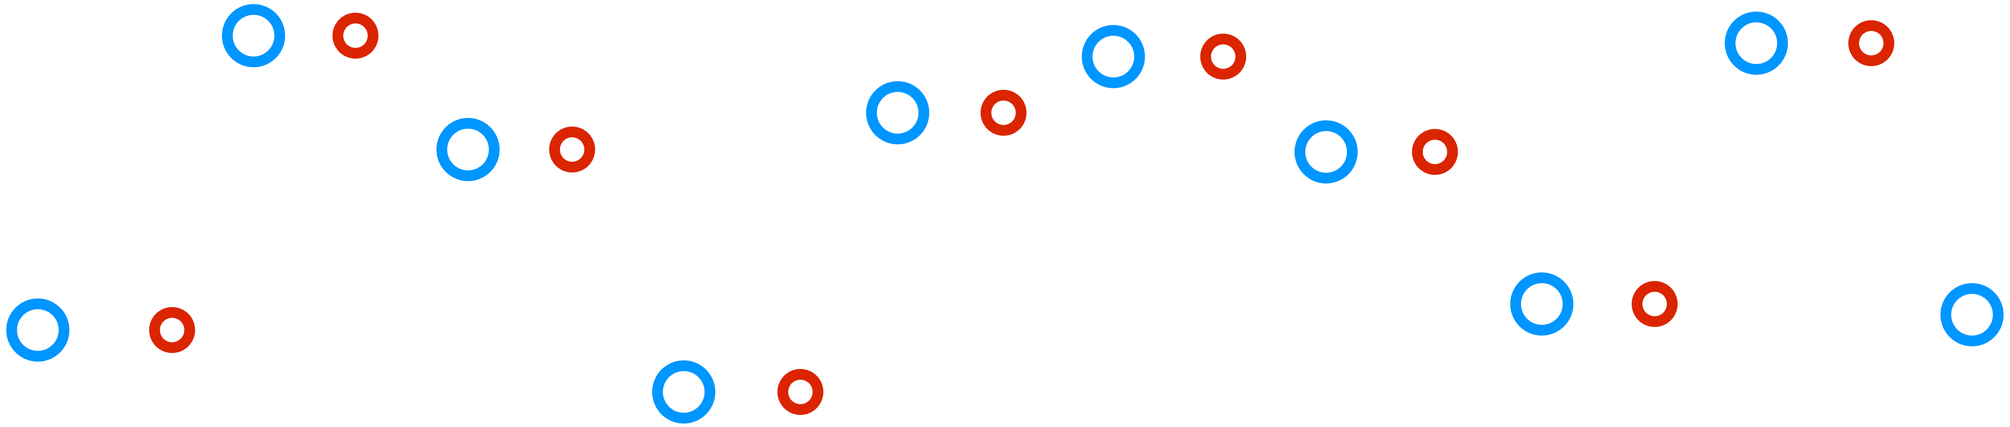

Theoretically Optimal Upsampling

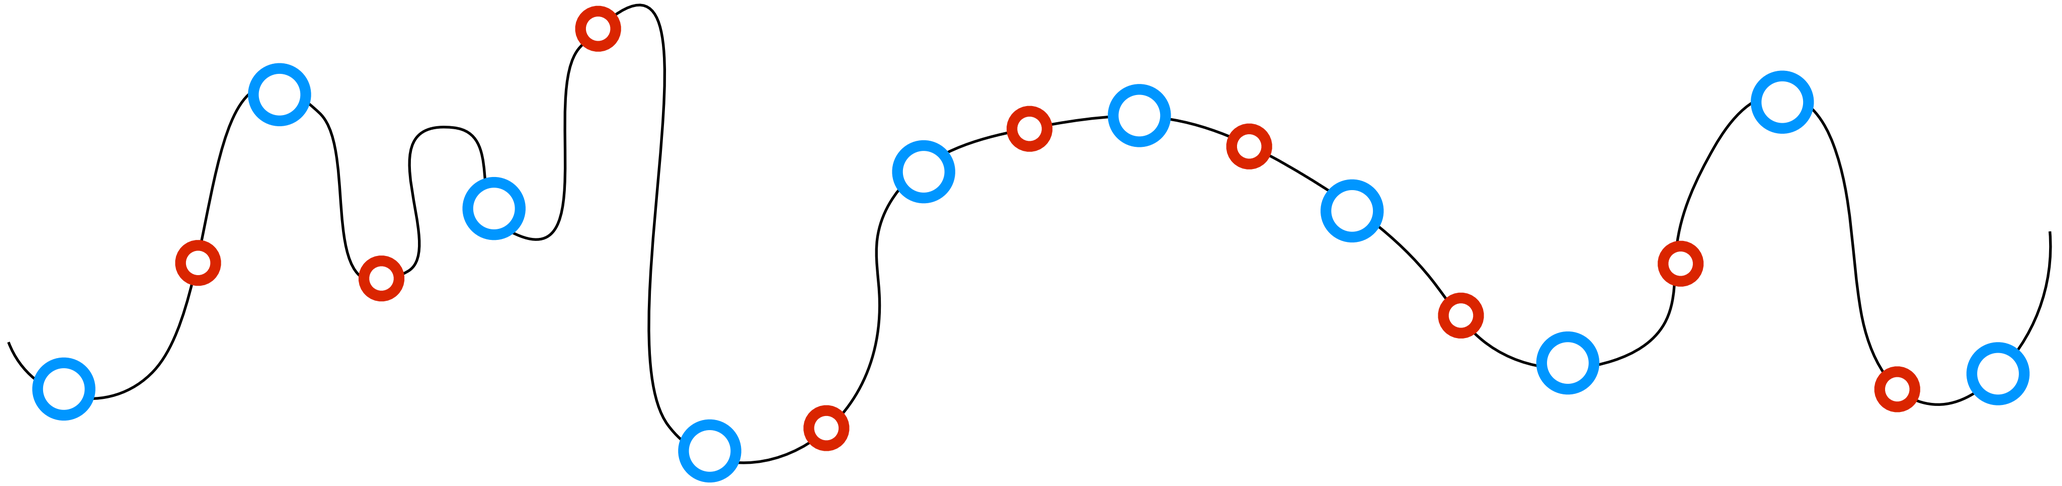

- First reconstruct the underlying continuous signal
  - With potential anti-aliasing
- And then resample at a desired, higher rate

Theoretically Optimal Upsampling

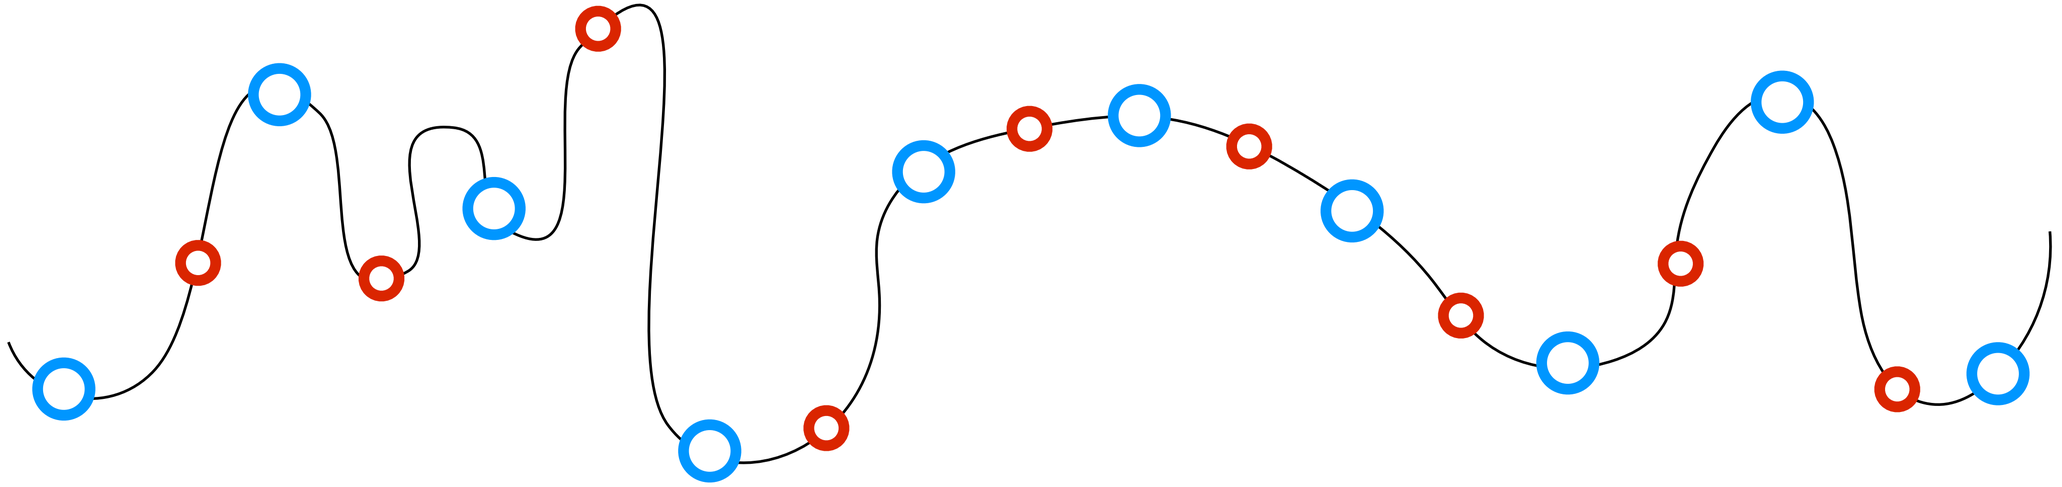

- That is, first convolve the sampled signal with a filter and then sample
- Mathematically that's equivalent to computing the convolution only at the desired re-sampled points!

Upsizing by 16:1

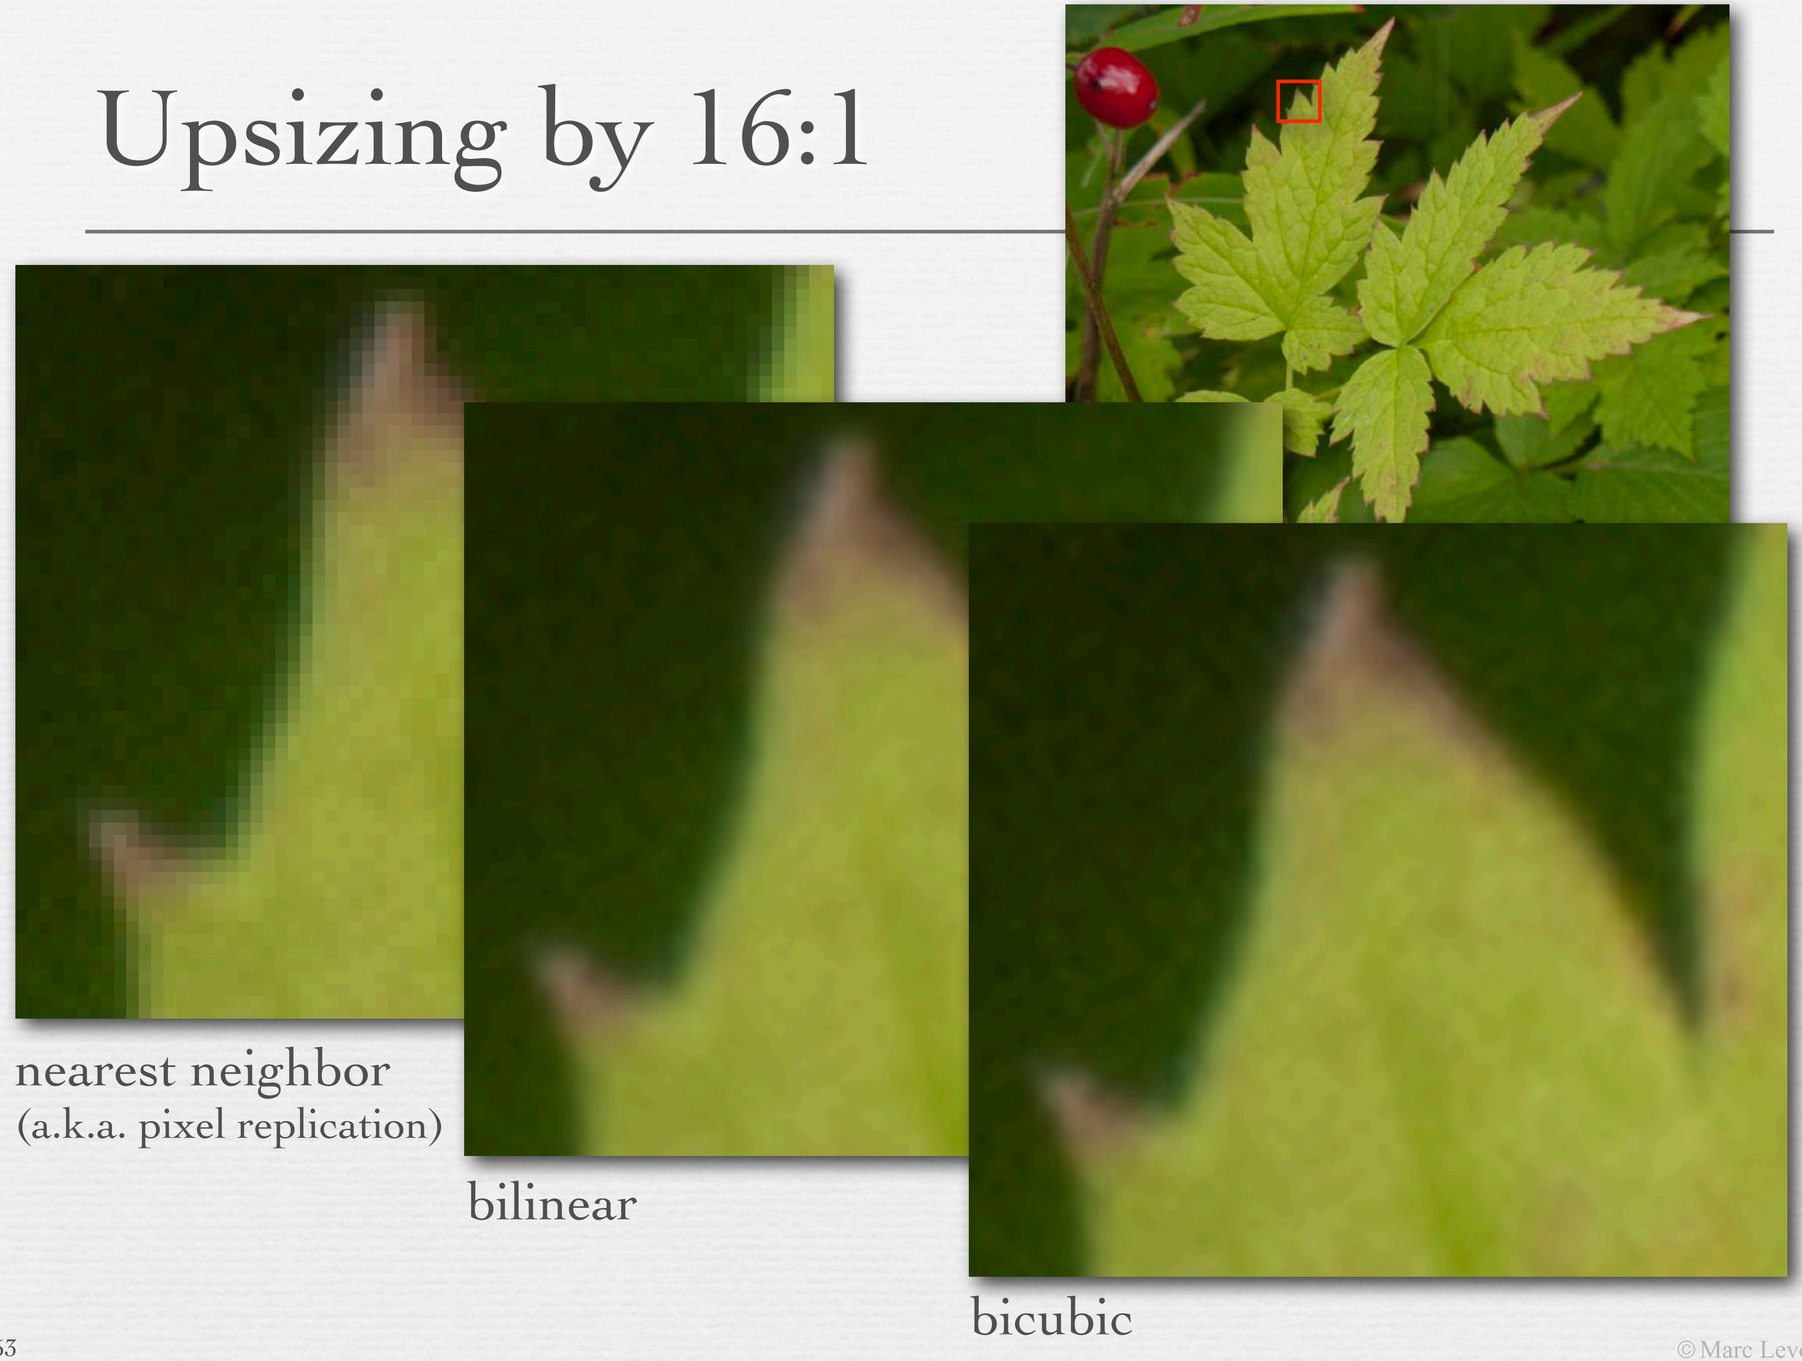

nearest neighbor
(a.k.a. pixel replication)

bilinear

bicubic

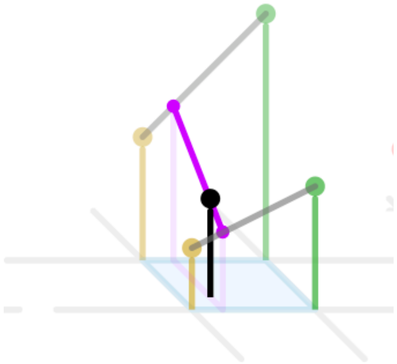

Bilinear

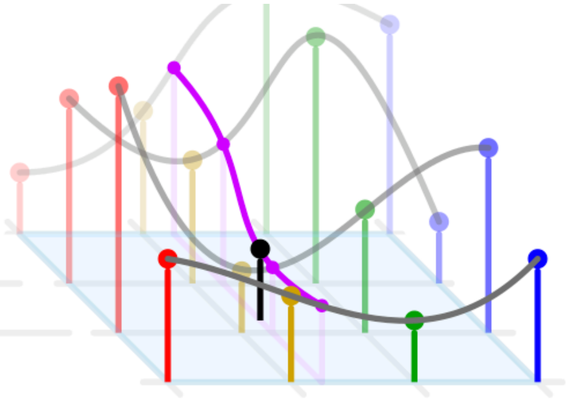

Bicubic

Slide credit: Marc Levoy

https://en.wikipedia.org/wiki/Bicubic_interpolation

Downsampling

- Simply dropping samples/pixels? Why is that bad?
- Dropping pixels is equivalent to sampling the original continuous signal using a lower rate. It would make later reconstruction harder!
- Take a photo of a high-frequency scene, and then try to scale it down on your computer. If your image viewer isn’t smart, the scaled-down version will be aliased. Why?

Summary

- Think of functions as vectors
  - Fourier Series: functions as infinite-dimensional vectors with finite span
  - Fourier Transform: functions as infinite-dimensional vectors with infinite span
  - Discrete Fourier Transform (Series): functions as finite-dimensional vectors with finite span
- Numerous applications in science and engineering
  - Function approximation (mind the Gibbs phenomena)
  - Denoising
  - Compression
  - Sampling and reconstruction
  - …

Summary

- Nyquite-Shannon Sampling Theorem
  - Sample at $2f$ if f is the maximum frequency component.
  - Otherwise aliasing occurs.
- Anti-Aliasing
  - Sample at a higher rate (or just high enough)
  - Pre-filter high-frequency components
  - Pre-filtering could be done either in the frequency domain or in the spatial domain (convolution); they are equivalent by the Convolution Theorem.