Master of Information and Data Science
DATASCI 281: Computer Vision
Spring 2022

Professor Hany Farid
University of California, Berkeley

Sanpling

A discrete-time signal, $f[x]$, is formed from a continuous-time signal, $f(x)$, by the following relationship:

$$f[x] \;=\; f(xT) \qquad -\infty < x < \infty, \tag{1}$$

for integer values x. In this expression, the quantity T is the *sampling period*. I will denote continuous-time signals with rounded parenthesis (e.g., $f(\cdot)$), and discrete-time signals with square parenthesis (e.g., $f[\cdot]$). As shown on the right, this sampling operation may be considered as a multiplication of the continuous time signal with an *impulse train*. The impulse train is defined as:

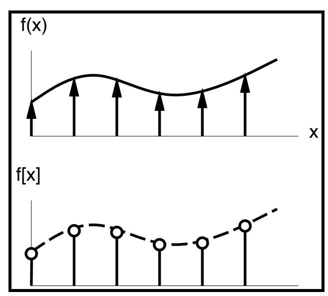

$$s(x) \;=\; \sum_{k=-\infty}^{\infty} \delta(x - kT), \tag{2}$$

where $\delta(\cdot)$ is the unit-impulse, and T is the sampling period. Note that the impulse train is a continuous-time signal. Multiplying the impulse train with a continuous-time signal gives a sampled signal:

$$f_s(x) \;=\; f(x)s(x), \tag{3}$$

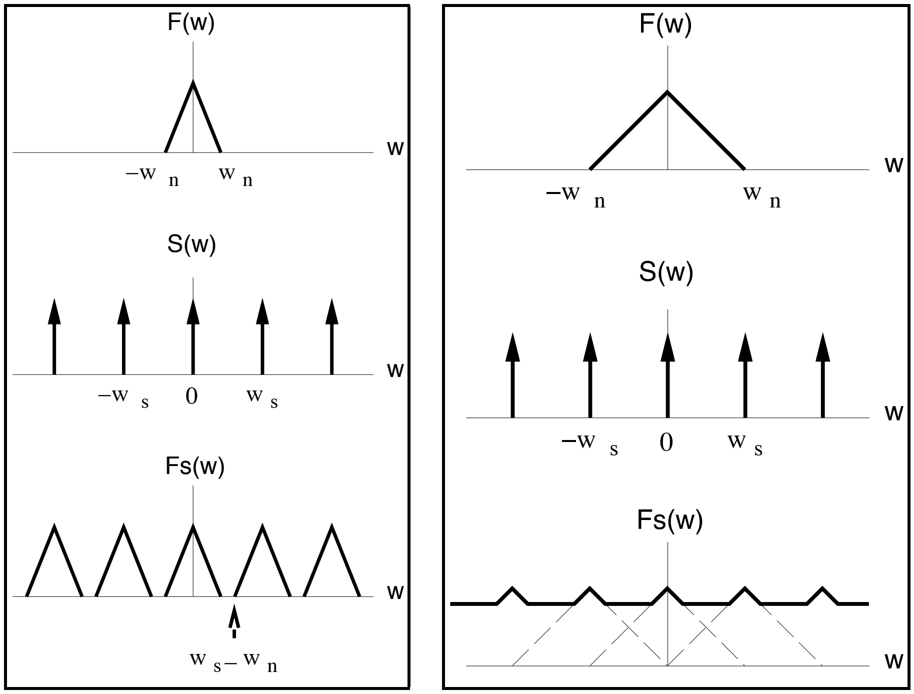

(a) (b)

Figure 1: Sampling in the Fourier domain without (a) and with (b) aliasing.

Note that the sampled signal, $f_s(x)$, is indexed on the *continuous* variable x, while the final discrete-time signal, $f[x]$ is indexed on the *integer* variable x. It will prove to be mathematically convenient to work with this intermediate sampled signal, $f_s(x)$.

In the space domain, sampling was described as a product between the impulse train and the continuous-time signal, Equation (3). In the frequency domain, this operation amounts to a convolution between the Fourier transform of these two signals:

$$F_s(\omega) \quad = \quad F(\omega) \star S(\omega) \tag{4}$$

For example, shown in Figure 1(a) (from top to bottom) are the Fourier transforms of the continuous-time function, $F(\omega)$, the impulse train, $S(\omega)$, itself an impulse train, and the results of convolving these two signals, $F_s(\omega)$. Notice that the Fourier transform of the *sampled* signal contains multiple (yet exact) copies of the Fourier transform of the original *continuous* signal. Note however the conditions under which an exact replica is preserved depends on the maximum frequency response ω_n of the original continuous-time signal, and the sampling interval of the impulse train, ω_s which, not surprisingly, is related to the sampling period T as $\omega_s = 2\pi/T$. More precisely, the copies of the frequency response will not overlap if:

$$\begin{aligned} \omega_n &< \omega_s - \omega_n \qquad \text{or} \\ \omega_s &> 2\omega_n, \end{aligned} \tag{5}$$

The frequency ω_n is called the *Nyquist frequency* and $2\omega_n$ is called the *Nyquist rate*. Shown in Figure 1(b) is another example of this sampling process in the frequency domain, but this time, the Nyquist rate is not met, and the copies of the frequency response overlap. In such a case, the signal is said to be *aliased*.

Not surprisingly, the Nyquist rate depends on both the characteristics of the continuous-time signal, and the sampling rate. More precisely, as the maximum frequency, ω_n, of the continuous-time signal increases, the sampling period, T must be made smaller (i.e., denser sampling), which in turn increases ω_s, preventing overlap of the frequency responses. In other words, a signal that changes slowly and smoothly can be sampled fairly coarsely, while a signal that changes quickly requires more dense sampling.

If the Nyquist rate is met, then a discrete-time signal fully characterizes the continuous-time signal from which it was sampled. On the other hand, if the Nyquist rate is not met, then the sampling leads to aliasing, and the discrete-time signal does not accurately represent its continuous-time counterpart. In the former case, it is possible to reconstruct the original continuous-time signal, from the discrete-time signal. In particular since the frequency response of the discrete-time signal contains exact copies of the original continuous-time signals frequency response, we need only extract one of these copies, and inverse transform the result. The result will be identical to the original signal.

In order to extract a single copy, the Fourier transform of the sampled signal is multiplied by an *ideal reconstruction filter* as shown on the right. This filter has unit value between the frequencies $-\pi/T$ to π/T and is zero elsewhere. This frequency band is

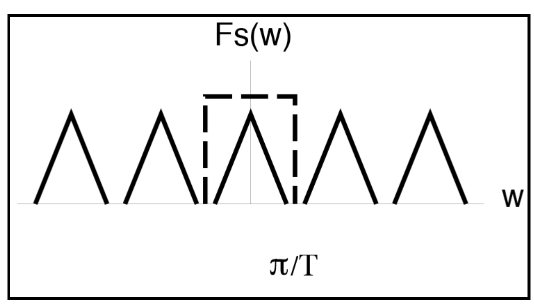

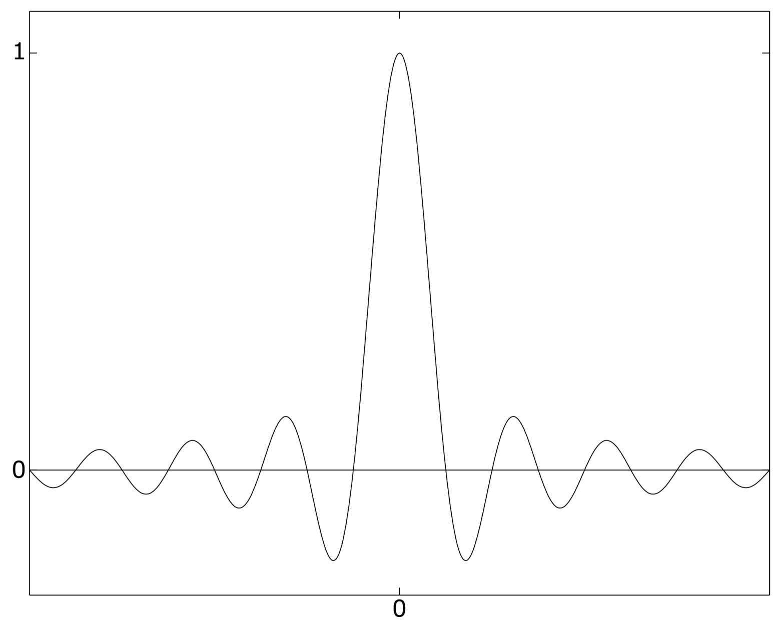

Figure 2: The ideal sync function.

guaranteed to be greater than the Nyquist frequency, ω_n (i.e., $\omega_s = 2\pi/T > 2\omega_n$, so that $\pi/T > \omega_n$). In the space domain, this ideal reconstruction filter has the form:

$$h(x) = \frac{\sin(\pi x/T)}{\pi x/T}, \tag{6}$$

and is often referred to as the *ideal sync* function, as shown in Figure 2. Since reconstruction in the frequency domain is accomplished by multiplication with the ideal reconstruction filter, we could equivalently reconstruct the signal by convolving with the ideal sync in the space domain.

Example: Consider the following continuous-time signal:

$$f(x) = \cos(\omega_0 x),$$

a sinusoid with frequency ω_0. We will eventually be interested in sampling this function and seeing how the effects of aliasing are manifested. But first, let's compute the Fourier transform of this signal:

$$\begin{aligned} F(\omega) &= \sum_{k=-\infty}^{\infty} f(k) e^{-i\omega k} \\ &= \sum_{k=-\infty}^{\infty} \cos(\omega_0 k)(\cos(\omega k) - i \sin(\omega k)) \\ &= \sum_{k=-\infty}^{\infty} \cos(\omega_0 k)\cos(\omega k) - i\cos(\omega_0 k)\sin(\omega k) \end{aligned}$$

First let's consider the product of two cosines. It is easy to show from basic trigonometric identities that $\cos(A)\cos(B) = 0$ when $A \neq B$, and is equal to π when $|A| = |B|$. Similarly, one can show that $\cos(A)\sin(B) = 0$ for all A and B. So, the Fourier transform of $\cos(\omega_0 x) = \pi$ for $|\omega| = \omega_0$, and is 0 otherwise (see below). If the sampling rate is greater than $2\omega_0$, then there will be no aliasing, but if the sampling rate is less than $2\omega_0$, then the reconstructed signal will be of the form $\cos((\omega_s - \omega_0)x)$, that is, the reconstructed signal will be appear as a lower frequency sinusoid - it will be aliased.

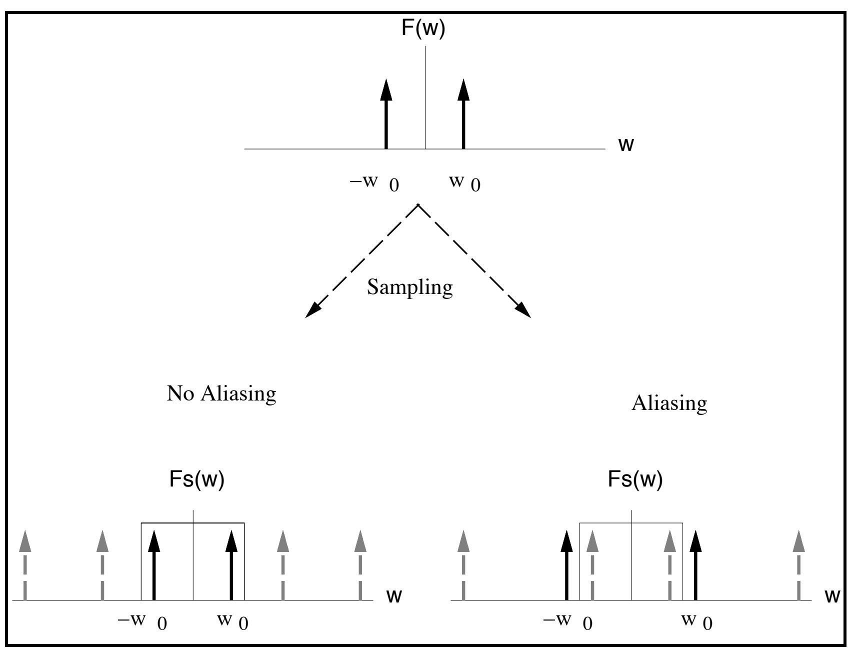

We will close by drawing on the linear algebraic framework for additional intuition on the sampling and reconstruction process. First we will need to restrict ourselves to the sampling of an already sampled signal. Consider a m-dimensional signal sub-sampled to a n-dimensional signal. We may express this operation in matrix form as follows:

$$
\begin{aligned}
\begin{pmatrix} g_1 \\ \vdots \\ g_n \end{pmatrix} &= \begin{pmatrix} 1 & 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & \dots & 0 & 0 & 0 & 0 \\ \vdots & & & & \ddots & & & & \vdots \\ 0 & 0 & 0 & 0 & \dots & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & \dots & 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} f_1 \\ f_2 \\ \vdots \\ f_{m-1} \\ f_m \end{pmatrix} \\
\vec{g}_n &= S_{n \times m} \vec{f}_m, \qquad (7)
\end{aligned}
$$

where the subscripts denote the vector and matrix dimensions, and in this example $n = m/2$. Our goal now is to determine when it is possible to reconstruct the signal $\vec{f}$, from the sub-sampled signal $\vec{g}$. The Nyquist sampling theory tells us that if a signal is band-limited (i.e., can be written as a sum of a finite number of sinusoids), then we can sample it without loss of information. We can express this constraint in matrix notation:

$$
\vec{f}_m = B_{m \times n} \vec{w}_n, \qquad (8)
$$

where the columns of the matrix B contains the basis set of sinusoids - in this case the first n sinusoids. Substituting into the above sampling equation gives:

$$
\begin{aligned}
\vec{g}_n &= S_{n \times m} B_{m \times n} \vec{w}_n \\
&= M_{n \times n} \vec{w}_n. \qquad (9)
\end{aligned}
$$

If the matrix M is invertible, then the original weights (i.e., the representation of the original signal) can be determined by simply left-multiplying the sub-sampled signal $\vec{g}$ by M^{-1}. In other words, Nyquist sampling theory can be thought of as simply a matrix inversion problem. This should not be at all surprising, the trick to sampling and perfect reconstruction is to simply limit the dimensionality of the signal to at most twice the number of samples.