Chapter 8

Poisson Summation, Sampling and Nyquist's Theorem

See: A.6.1, A.5.2.

In Chapters 4 through 7, we developed the mathematical tools needed to describe functions of continuous variables and methods to analyze and reconstruct them. This chapter continues the transition from the world of pure mathematics to its application to problems in image reconstruction. In the first sections of this chapter, we imagine that our "image" is a function, f, of a single real variable, x. In a purely mathematical context, x is a real number that can assume any value along a continuum of numbers. The function also takes values in a continuum, either in $\mathbb{R}$, $\mathbb{C}$, or perhaps $\mathbb{R}^n$. In practical applications we can only evaluate f at a finite set of points $\{x_j\}$. This is called *sampling*. As most of the processing takes place in digital computers, both the points $\{x_j\}$ and the measured values $\{f(x_j)\}$ are forced to lie in the preassigned, finite set of numbers known to the computer. This is called *quantization*. The reader is urged to review Section A.1, where these ideas are discussed in some detail.

Except for a brief discussion of quantization, this chapter is about the consequences of sampling. We examine the fundamental question: How much information about a function is contained in a finite or infinite set of samples? Central to this analysis is the *Poisson summation formula*. This formula is a bridge between the Fourier transform and the Fourier series. While the Fourier transform is well suited to an abstract analysis of image or signal processing, it is the Fourier series that is actually used to do the work. The reason is quite simple: The Fourier transform and its inverse require integrals over the entire real line, whereas the Fourier series is phrased in terms of infinite sums and integrals *over finite intervals*, both of which are eventually approximated by finite sums. Indeed, we also introduce the *finite Fourier transform*, an analogue of the Fourier transform for finite sequences. This chapter covers the next step from the abstract world of the infinite and the infinitesimal to the real world of the finite and discrete.

8.1 Sampling and Nyquist's Theorem*

See: A.5.2, B.1.

Recall that our basic model for a measurement is the evaluation of a function at a point. A set of points $\{x_j\}$ contained in an interval (a, b) is *discrete* if no subsequence converges to a point in (a, b). Evaluating a function on a discrete set of points is called *sampling*. Practically speaking, a function can only be evaluated at points where it is continuous. From the perspective of measurement, the value of a function at points of *dis*continuity is not well defined.

***Definition* 8.1.1.** Suppose that f is a function defined in an interval (a, b) and $\{x_j\}$ is a discrete set of points in (a, b). The points $\{x_j\}$ are called the *sample points*. The values $\{f(x_j)\}$ are called the *samples* of f at the points $\{x_j\}$.

In most applications the discrete set is of the form $\{x_0 + jl \mid j \in \mathbb{Z}\}$, where l is a fixed positive number. These are called *equally spaced samples*; the number l is called the *sample spacing*. The reciprocal l^{-1} of l is called the *sampling rate*. Sampling theory studies the problem of reconstructing functions of a continuous variable from a set of samples and the relationship between these reconstructions and the idealized data.

8.1.1 Bandlimited Functions and Nyquist's Theorem

A model for measurement, more realistic than pointwise evaluation, is the evaluation of a convolution $f * \varphi$. Here φ is an L^1 weight function that models the measuring apparatus. For most reasonable choices of weight function, $f * \varphi$ is a continuous function, so its value is well defined (from the point of view of measurement) at all x. As the Fourier transform of $f * \varphi$ is $\hat{f}\hat{\varphi}$, the Riemann-Lebesgue lemma, Theorem 4.2.2, implies that $\hat{\varphi}$ tends to zero as $|\xi|$ tends to infinity. This means that the measuring apparatus attenuates the high-frequency information in f. In applications we often make the assumption that there is "no high-frequency information" or that is has been filtered out.

***Definition* 8.1.2.** Let f be a function defined on $\mathbb{R}$. If its Fourier transform, $\hat{f}$, is supported in a finite interval, then f is a *bandlimited function*. If $\hat{f}$ is supported in $[-L, L]$, then f is an *L-bandlimited* function.

***Definition* 8.1.3.** Let f be a function defined on $\mathbb{R}$. If its Fourier transform, $\hat{f}$, is supported in a finite interval, then f is said to have *finite bandwidth*. If $\hat{f}$ is supported in an interval of length W, then f is said to have *bandwidth W*.

A bandlimited function is always infinitely differentiable. If f is either L^1 or L^2, then $\hat{f}$ is in L^1 and the Fourier inversion formula states that

$$f(x) = \frac{1}{2\pi} \int_{-L}^{L} \hat{f}(\xi) e^{ix\xi} d\xi. \tag{8.1}$$

As the integrand is supported in a finite interval, the integral in (8.1) can be differentiated as many times as we like. This, in turn, shows that f is infinitely differentiable. If f is in L^2, then so are all of its derivatives.

Nyquist's theorem states that a bandlimited function is determined by a set of uniformly spaced samples, provided that the sample spacing is sufficiently small.

Theorem 8.1.1 (Nyquist's theorem). *If f is a square integrable function and*

$$\hat{f}(\xi) = 0 \textit{ for } |\xi| > L,$$

then f is determined by the samples $\{f(\frac{\pi n}{L}) \,:\, n \in \mathbb{Z}\}$.

***Remark* 8.1.1.** As the proof shows, f is determined by any collection of uniformly spaced samples with sample spacing $\frac{\pi n}{L}$, i.e.$\{f(x_0 + \frac{\pi n}{L})\}$ for any $x_0 \in \mathbb{R}$.

Proof. If we think of $\hat{f}$ as a function *defined* on the interval $[-L, L]$, then it follows from (8.1) that the numbers $\{2\pi f(\frac{\pi n}{L})\}$ are the Fourier *coefficients* of $\hat{f}$. The inversion formula for Fourier series then applies to give

$$\hat{f}(\xi) = \left(\frac{\pi}{L}\right) \underset{N\to\infty}{\mathrm{LIM}} \left[\sum_{n=-N}^{N} f(\frac{n\pi}{L}) e^{-\frac{n\pi i\xi}{L}}\right] \qquad \text{if } |\xi| < L. \tag{8.2}$$

For the remainder of the proof we use the notation

$$\left(\frac{\pi}{L}\right) \sum_{n=-\infty}^{\infty} f(\frac{n\pi}{L}) e^{-\frac{n\pi i\xi}{L}}$$

to denote this LIM. The function defined by this infinite sum is periodic of period $2L$; we can use it to express $\hat{f}(\xi)$ in the form

$$\hat{f}(\xi) = \left(\frac{\pi}{L}\right) \left[\sum_{n=-\infty}^{\infty} f(\frac{n\pi}{L}) e^{-\frac{n\pi i\xi}{L}}\right] \chi_{[-L,L]}(\xi). \tag{8.3}$$

This proves Nyquist's theorem, for a function in $L^2(\mathbb{R})$ is determined by its Fourier transform. $\square$

The exponentials $e^{\pm iLx}$ have period $\frac{2\pi}{L}$ and frequency $\frac{L}{2\pi}$. If a function is L-bandlimited, then $\frac{L}{2\pi}$ is the highest frequency appearing in its Fourier representation. Nyquist's theorem states that we must sample such a function at the rate $\frac{L}{\pi}$; that is, at twice its highest frequency. As we shall see, sampling at a lower rate does not provide enough information to completely determine f.

***Definition* 8.1.4.** The optimal sampling rate for an L-bandlimited function, $\frac{L}{\pi}$, is called the *Nyquist rate*. Sampling at a lower rate is called *undersampling*, and sampling at a higher rate is called *oversampling*.

If this were as far as we could go, Nyquist's theorem would be an interesting result of little practical use. However, the original function f can be explicitly reconstructed using (8.3) in the Fourier inversion formula. To justify our manipulations, we assume that f tends to zero rapidly enough so that

$$\sum_{n-\infty}^{\infty} |f(\frac{n\pi}{L})| < \infty. \tag{8.4}$$

With this understood, we use (8.3) in (8.1) to obtain

$$\begin{aligned}
f(x) &= \frac{1}{2\pi}\int_{-\infty}^{\infty} \chi_{[-L,L]}(\xi)\hat{f}(\xi)e^{i\xi x}d\xi \\
&= \frac{1}{2\pi}\frac{\pi}{L}\int_{-L}^{L}\sum_{n=-\infty}^{\infty} f(\frac{n\pi}{L})e^{ix\xi-\frac{n\pi i\xi}{L}}d\xi \\
&= \frac{1}{2L}\sum_{n=-\infty}^{\infty} f(\frac{n\pi}{L})\int_{-L}^{L} e^{ix\xi-\frac{n\pi i\xi}{L}}d\xi \\
&= \sum_{n=-\infty}^{\infty} f(\frac{n\pi}{L})\operatorname{sinc}(Lx-n\pi).
\end{aligned}$$

This formula expresses the value of f, for every $x \in \mathbb{R}$, in terms of the samples $\{f(\frac{n\pi}{L}) : n \in \mathbb{Z}\}$:

$$f(x) = \sum_{n=-\infty}^{\infty} f(\frac{n\pi}{L})\operatorname{sinc}(Lx-n\pi). \tag{8.5}$$

***Remark* 8.1.2.** If f belongs to $L^2(\mathbb{R})$, then $\hat{f}$ is also square integrable. This, in turn, implies that $\sum_{-\infty}^{\infty}|f(L^{-1}n\pi)|^2$ is finite. Using the Cauchy-Schwarz inequality for l^2, we easily show that the sum on the right-hand side of (8.5) converges locally uniformly. The Fourier transform of a bandlimited function is absolutely integrable. The Riemann-Lebesgue lemma therefore implies that a bandlimited function is bounded and tends to zero as $|x|$ goes to infinity. Indeed, a bounded, integrable function is also square integrable. Note, however, that a square-integrable, bandlimited function need not be absolutely integrable.

Exercises

***Exercise* 8.1.1.** Suppose that f is an L-bandlimited function. Show that it is determined by the set of samples $\{f(x_0 + \frac{\pi n}{L}) : n \in \mathbb{Z}\}$, for any $x_0 \in \mathbb{R}$.

***Exercise* 8.1.2.** Explain why, after possible modification on a set of measure zero, a bandlimited L^2-function is given by (8.1).

***Exercise* 8.1.3.** Show that a bounded integrable function defined on $\mathbb{R}$ is also square integrable.

***Exercise* 8.1.4.** Give an example of a bandlimited function in $L^2(\mathbb{R})$ that is not absolutely integrable.

***Exercise* 8.1.5.** Suppose that f is a bandlimited function in $L^2(\mathbb{R})$. Show that the infinite sum in (8.5) converges locally uniformly.

8.1.2 Shannon-Whittaker Interpolation

See: A.5.2.

The explicit interpolation formula, (8.5), for f in terms of its samples at $\{\frac{n\pi}{L} \mid n \in \mathbb{Z}\}$ is sometimes called the Shannon-Whittaker interpolation formula. In Section A.5.2 we consider other methods for interpolating a function from sampled values. These formulæ involve finite sums and only give exact reconstructions for a finite-dimensional family of functions. The Shannon-Whittaker formula gives an exact reconstruction for all L-bandlimited functions. Since it requires an infinite sum, it is mostly of theoretical significance. In practical applications only a finite part of this sum can be used. That is, we would set

$$f(x) \approx \sum_{n=-N}^{N} f(\frac{n\pi}{L}) \operatorname{sinc}(Lx - n\pi).$$

Because

$$\operatorname{sinc}(Lx - n\pi) \simeq \frac{1}{n}$$

and $\sum n^{-1} = \infty$, the partial sums of the series in (8.5) may converge to $f(x)$ very slowly. In order to get a good approximation to $f(x)$, we would therefore need to take N very large. This difficulty can be often be avoided by oversampling.

Formula (8.5) is only one of an infinite family of similar interpolation formulæ . Suppose that f is an $(L-\eta)$-bandlimited function for an $\eta > 0$. Then it is also an L-bandlimited function. This makes it possible to use oversampling to obtain more rapidly convergent interpolation formulæ . To find such formulæ select a function φ such that

1. $\hat{\varphi}(\xi) = 1$ for $|\xi| \leq L - \eta$
2. $\hat{\varphi}(\xi) = 0$ for $|\xi| > L$

A function of this sort is often called a *window function.*

From (8.2) it follows that

$$\hat{f}(\xi) = \left(\frac{\pi}{L}\right) \sum_{n=-\infty}^{\infty} f(\frac{\pi n}{L}) e^{-\frac{n\pi i\xi}{L}} \text{ for } |\xi| < L. \tag{8.6}$$

Since $\hat{f}$ is supported in $[\eta - L, L - \eta]$ and $\hat{\varphi}$ satisfies condition (1), it follows that

$$\hat{f}(\xi) = \hat{f}(\xi)\hat{\varphi}(\xi).$$

Using this observation and (8.2) in the Fourier inversion formula gives

$$\begin{aligned} f(x) &= \frac{1}{2\pi}\int_{-\infty}^{\infty} \hat{f}(\xi)\hat{\varphi}(\xi)e^{i\xi x}d\xi \\ &= \frac{1}{2\pi}\frac{\pi}{L}\sum_{n=-\infty}^{\infty} f(\frac{n\pi}{L})\int_{-L}^{L}\hat{\varphi}(\xi)e^{ix\xi-\frac{n\pi i\xi}{L}}d\xi \\ &= \frac{1}{2L}\sum_{n=-\infty}^{\infty} f(\frac{n\pi}{L})\varphi(x-\frac{n\pi}{L}). \end{aligned} \tag{8.7}$$

This is a different interpolation formula for f; the sinc-function is replaced by $[2L]^{-1}\varphi$. The Shannon-Whittaker formula, (8.5), corresponds to the choice $\hat{\varphi} = \chi_{[-L,L]}$.

Recall that more *smoothness* in the Fourier transform of a function is reflected in faster decay of the function itself. Using a smoother function for $\hat{\varphi}$ therefore leads to a more rapidly convergent interpolation formula for f. There is a small price to pay for using a different choice of $\hat{\varphi}$. The first issue is that φ, given by

$$\varphi(x) = \frac{1}{2\pi}\int_{-\infty}^{\infty}\hat{\varphi}(\xi)e^{ix\xi}d\xi,$$

may be more difficult to accurately compute than the sinc function. The second is that we need to sample f above the Nyquist rate. In this calculation, f is an $(L-\eta)$-bandlimited function, but we need to use a sample spacing

$$\frac{\pi}{L} < \frac{\pi}{L-\eta}.$$

On the other hand, a little oversampling and additional computational overhead often leads to superior results.

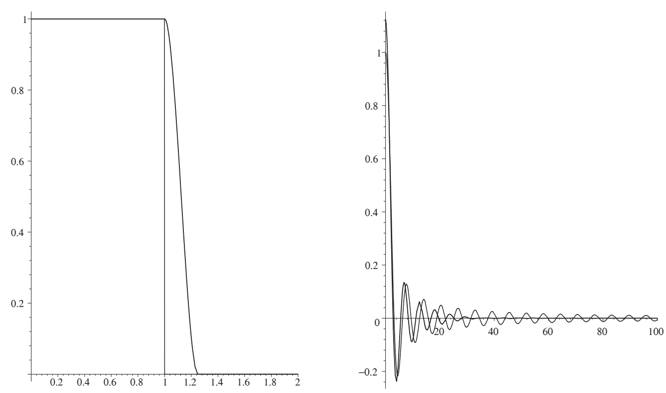

(a) Window functions in Fourier space.

(b) The corresponding interpolation functions.

Figure 8.1. Window functions in Fourier space and their inverse Fourier transforms.

***Example* 8.1.1.** We compare the Shannon-Whittaker interpolant to the interpolant obtained using a second order smoothed window, with 25% oversampling. The second-order window is defined by

$$s_2(\xi) = \begin{cases} 1 & \text{if } |\xi| \leq 1, \\ 128|\xi|^3 - 432|\xi|^2 + 480|\xi| - 175 & \text{if } 1 < |\xi| < \frac{5}{4}, \\ 0 & \text{if } |\xi| \geq \frac{5}{4}. \end{cases} \tag{8.8}$$

It has one continuous derivative and a weak second derivative, its inverse Fourier transform is

$$\mathcal{F}^{-1}(s_2)(x) = 768\left[\frac{\cos(x) - \cos(5x/4)}{\pi x^4}\right] - 96\left[\frac{\sin(5x/4) + \sin(x)}{\pi x^3}\right].$$

Figure 8.1(a) shows the two window functions, and Figure 8.1(b) shows their inverse Fourier transforms. Observe that $\mathcal{F}^{-1}(s_2)$ converges to zero as x tends infinity much faster than the sinc function.

The graphs in Figure 8.2 compare the results of reconstructing *non*-bandlimited functions using the Shannon-Whittaker interpolant and that obtained using (8.7) with $\hat{\varphi}$ a suitably scaled version of s_2. The sample spacing for the Shannon-Whittaker interpolation is .1, for the "oversampled" case we use .08. The original function is shown in the first column; the second column shows both interpolants; and the third column is a detailed view near to ± 1. This is a point where the original function or one of its derivatives is discontinuous. Two things are evident in these graphs: The error in the oversampled interpolant tends to zero much faster as x tends to infinity, and the smoother the function, the easier it is to interpolate.

Exercises

***Exercise* 8.1.6.** Use the Shannon-Whittaker formula to reconstruct the function

$$f(x) = \frac{\sin(Lx)}{\pi x}$$

from the samples $\{f(\frac{n\pi}{L})\}$.

***Exercise* 8.1.7.** How should the Shannon-Whittaker formula be modified if, instead of $\{f(\frac{\pi n}{L}) : n \in \mathbb{Z}\}$, the samples $\{f(x_0 + \frac{\pi n}{L}) : n \in \mathbb{Z}\}$ are collected?

***Exercise* 8.1.8.** Show that for each $n \in \mathbb{N}$, function, $\operatorname{sinc}(Lx - n\pi)$ is L-bandlimited. The Shannon-Whittaker formula therefore expresses an L-bandlimited function as a sum of such functions.

***Exercise* 8.1.9.** The Fourier transform of

$$f(x) = \frac{1 - \cos(x)}{2\pi x}$$

is $\hat{f}(\xi) = i \operatorname{sgn} \xi \chi_{[-1,1]}(\xi)$. Use the Shannon-Whittaker formula to reconstruct f from the samples $\{f(n\pi)\}$.

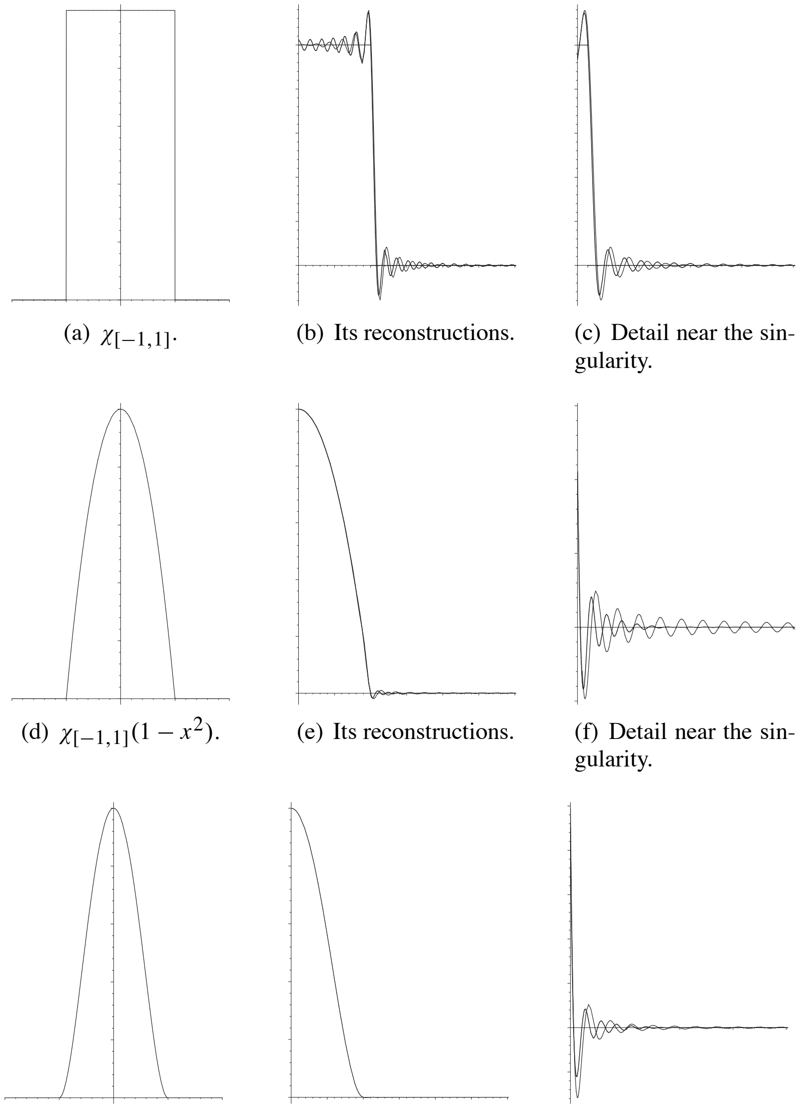

(a) $\chi_{[-1,1]}$. (b) Its reconstructions. (c) Detail near the singularity.

(d) $\chi_{[-1,1]}(1-x^2)$. (e) Its reconstructions. (f) Detail near the singularity.

(g) $\chi_{[-1,1]}(1-x^2)^2$. (h) Its reconstructions. (i) Detail near the singularity.

Figure 8.2. Shannon-Whittaker and generalized Shannon-Whittaker interpolation for several functions

8.2 The Poisson Summation Formula

What happens if we do not have enough samples to satisfy the hypotheses of Nyquist's theorem? For example, what if our signal is not bandlimited? Functions that describe images in medical applications generally have bounded support, so they cannot be bandlimited and therefore we are always undersampling (see Chapter 4, Proposition 4.4.1). To analyze the effects of undersampling, we introduce the Poisson summation formula. It gives a relationship between the Fourier transform and the Fourier series.

8.2.1 The Poisson Summation Formula

Assume that f is a continuous function that decays reasonably fast as $|x| \to \infty$. We construct a periodic function out of f by summing the values of f at its integer translates. Define f_p by

$$f_p(x) = \sum_{n=-\infty}^{\infty} f(x+n). \tag{8.9}$$

This is a periodic function of period 1, $f_p(x+1) = f_p(x)$. If f is absolutely integrable on $\mathbb{R}$, then it follows from Fubini's theorem that f_p is absolutely integrable on $[0, 1]$.

The Fourier *coefficients* of f_p are closely related to the Fourier *transform* of f:

$$\begin{aligned}\hat{f}_p(m) &= \int_0^1 f_p(x)e^{-2\pi imx}dx \\ &= \int_0^1 \sum_{n=-\infty}^{\infty} f(x+n)e^{-2\pi imx}dx = \sum_{n=-\infty}^{\infty}\int_n^{n+1} f(x)e^{-2\pi imx}dx \\ &= \int_{-\infty}^{\infty} f(x)e^{-2\pi imx}dx = \hat{f}(2\pi m).\end{aligned}$$

The interchange of the integral and summation is easily justified if f is absolutely integrable on $\mathbb{R}$.

Proceeding formally, we use the Fourier inversion formula for periodic functions, Theorem 7.1.2, to find Fourier series representation for f_p,

$$f_p(x) = \sum_{n=-\infty}^{\infty} \hat{f}_p(n)e^{2\pi inx} = \sum_{n=-\infty}^{\infty} \hat{f}(2\pi n)e^{2\pi inx}.$$

Note that $\{\hat{f}_p(n)\}$ are the Fourier *coefficients* of the 1-periodic function f_p whereas $\hat{f}$ is the Fourier *transform* of the absolutely integrable function f defined on all of $\mathbb{R}$. To justify these computations, it is necessary to assume that the coefficients $\{\hat{f}(2\pi n)\}$ go to zero sufficiently rapidly. If f is smooth enough, then this will be true. The *Poisson summation formula* is a precise formulation of these observations.

Theorem 8.2.1 (Poisson summation formula). *If f is an absolutely integrable function such that*

$$\sum_{n=-\infty}^{\infty} |\hat{f}(2\pi n)| < \infty$$

then, at points of continuity of f_p, we have

$$\sum_{n=-\infty}^{\infty} f(x+n) = \sum_{n=-\infty}^{\infty} \hat{f}(2\pi n)e^{2\pi inx}. \tag{8.10}$$

***Remark* 8.2.1.** The hypotheses in the theorem are not quite optimal. Some hypotheses are required, as there are examples of absolutely integrable functions f such that both sums,

$$\sum_{n=-\infty}^{\infty} |f(x+n)| \text{ and } \sum_{n=-\infty}^{\infty} |\hat{f}(2\pi n)|$$

converge but (8.10) does *not* hold. A more detailed discussion can be found in [79].

Using the preceding argument and rescaling, we easily find the Poisson summation formula for $2L$-periodic functions:

$$\sum_{n=-\infty}^{\infty} f(x+2nL) = \frac{1}{2L}\sum_{n=-\infty}^{\infty} \hat{f}(\frac{\pi n}{L})e^{\frac{\pi inx}{L}}. \tag{8.11}$$

Suitably scaled, the hypotheses are the same as those in Theorem 8.2.1.

As an application of (8.10) we can prove an x-space version of Nyquist's theorem. Suppose that f equals 0 outside the interval $[-L, L]$ (i.e., f is a *space-limited function*). For each $x \in [-L, L]$, only the $n = 0$ term on the left-hand side of (8.11) is nonzero. The Poisson summation formula states that

$$f(x) = \frac{1}{2L}\sum_{n=-\infty}^{\infty} \hat{f}(\frac{\pi n}{L})e^{\frac{\pi inx}{L}} \qquad \text{for } x \in [-L, L].$$

Therefore, if f is supported in $[-L, L]$, then it can be reconstructed from the samples of its Fourier transform

$$\{\hat{f}(\frac{\pi n}{L}) \mid n \in \mathbb{Z}\}.$$

This situation arises in magnetic resonance imaging (MRI). In this modality, we directly measure samples of the Fourier transform of the image function. That is, we measure $\{\hat{f}(n\Delta\xi)\}$. On the other hand, the function is known, a priori, to be supported in a fixed, bounded set $[-L, L]$. In order to reconstruct f without introducing errors, we need to take

$$\Delta\xi \leq \frac{\pi}{L}. \tag{8.12}$$

Thus, if we measure samples of the Fourier transform of a space-limited function, then Nyquist's theorem places a constraint on the sample spacing in the *Fourier domain*. An extensive discussion of sampling in MRI can be found in [50]. Figure 8.3 shows the result, in MRI, of undersampling the Fourier transform. Note that portions of the original image have been "folded over" at the top and right-hand side of the reconstructed image.

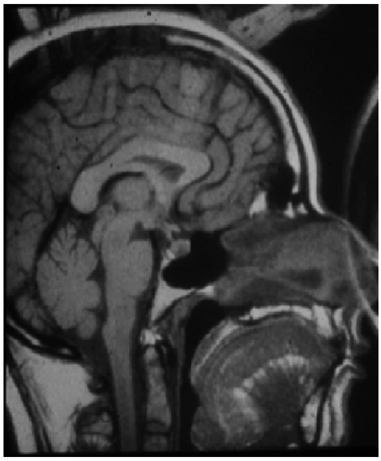

Figure 8.3. Aliasing artifacts produced by undersampling in magnetic resonance imaging. (Image courtesy of Dr. Felix Wehrli.)

It most applications we sample a function rather than its Fourier transform. The analysis of undersampling in this situation requires the *dual Poisson summation formula*. Let f be a function such that the sum

$$\sum_{-\infty}^{\infty} \hat{f}(\xi + 2nL)$$

converges. Considering

$$\hat{f}_p(\xi) = \sum_{-\infty}^{\infty} \hat{f}(\xi + 2nL)$$

and its Fourier coefficients in the same manner as previously we obtain the following:

Theorem 8.2.2 (The dual Poisson summation formula). *If f is a function such that $\hat{f}$ is absolutely integrable and*

$$\sum_{n=-\infty}^{\infty} |f(\frac{\pi n}{L})| < \infty$$

then, at a point of continuity of $\hat{f}_p$,

$$\sum_{n=-\infty}^{\infty} \hat{f}(\xi + 2nL) = \left(\frac{\pi}{L}\right) \sum_{n=-\infty}^{\infty} f(\frac{\pi n}{L}) e^{-\frac{n\pi i\xi}{L}}. \tag{8.13}$$

Exercises

***Exercise* 8.2.1.** Give the details for the proof of Theorem 8.2.2. You may assume that f is smooth and rapidly decreasing.

***Exercise* 8.2.2.** Explain formula (8.12).

Exercise* 8.2.3. This exercise requires a knowledge of the Fourier transform for generalized functions (see Section 4.4.4). Suppose that f is a periodic function of period 1. The generalized function l_f has a Fourier transform that is a generalized function. Using the dual Poisson summation formula, show that

$$\widehat{l_f} = 2\pi \sum_{n=-\infty}^{\infty} \hat{f}(n)\delta(2\pi n - \xi); \tag{8.14}$$

here $\{\hat{f}(n)\}$ are the Fourier coefficients defined in (7.1).

Exercise* 8.2.4. What is the analogue of formula (8.14) for a $2L$-periodic function?

8.2.2 Undersampling and Aliasing*

Using the Poisson summation formula, we analyze the errors introduced by undersampling. Whether or not f is an L-bandlimited function, the Shannon-Whittaker formula defines an L-bandlimited function:

$$F_L(x) = \sum_{n=-\infty}^{\infty} f(\frac{n\pi}{L}) \operatorname{sinc}(Lx - n\pi).$$

As $\operatorname{sinc}(0) = 1$ and $\operatorname{sinc}(n\pi) = 0$, for nonzero integers it follows that F_L interpolates f at the sample points,

$$F_L(\frac{n\pi}{L}) = f(\frac{n\pi}{L}) \qquad \text{for } n \in \mathbb{Z}.$$

Reversing the steps in the derivation of the Shannon-Whittaker formula, and applying formula (8.13) we see that Fourier transform of F_L is given by

$$\widehat{F_L}(\xi) = \sum_{n=-\infty}^{\infty} \hat{f}(\xi + 2nL)\chi_{[-L,L]}(\xi). \tag{8.15}$$

If f is L-bandlimited then for all ξ, we have

$$\hat{f}(\xi) = \widehat{F_L}(\xi).$$

On the other hand, if f is *not* L-bandlimited, then

$$\hat{f}(\xi) - \widehat{F_L}(\xi) = \begin{cases} \hat{f}(\xi) & \text{if } |\xi| > L, \\ -\sum_{n\neq 0} \hat{f}(\xi + 2nL) & \text{if } |\xi| \leq L. \end{cases} \tag{8.16}$$

The function F_L or its Fourier transform $\hat{F}_L$ encodes all the information present in the sequence of samples. Formula (8.16) shows that there are two distinct sources of error in F_L. The first is *truncation error*; as F_L is L-bandlimited, the high-frequency information in f is no longer available in F_L. The second source of error arises from the fact that the high-frequency information in f *reappears* at low frequencies in the function F_L. This latter type of distortion is called *aliasing*. The high-frequency information in the original signal is not simply "lost" but resurfaces, corrupting the low frequencies. Hence F_L faithfully reproduces neither the high-frequency nor the low-frequency information in f.

Aliasing is familiar in everyday life: If we observe the rotation of the wheels of a fast-moving car in a movie, it appears that the wheels rotate very slowly. A movie image is actually a sequence of samples (24 frames/second). This sampling rate is below the Nyquist rate needed to accurately reproduce the motion of the rotating wheel.

***Example* 8.2.1.** If a car is moving at 60 mph and the tires are 3 ft in diameter, then the angular velocity of the wheels is

$$\omega = 58\frac{1}{3}\frac{\text{rotations}}{\text{second}}.$$

We can model the motion of a point on the wheel as $(r\cos((58\frac{1}{3})2\pi t), r\sin((58\frac{1}{3})2\pi t)$. The Nyquist rate is therefore

$$2 \cdot 58\frac{1}{3}\frac{\text{frames}}{\text{second}} \simeq 117\frac{\text{frames}}{\text{second}}.$$

Sampling only 24 times per second leads to aliasing. As $58\frac{1}{3} = 10\frac{1}{3} + 2 * 24$, the aliased frequencies are $\pm(10\frac{1}{3})$.

The previous example is useful to conceptualize the phenomenon of aliasing but has little direct bearing on imaging. To better understand the role of aliasing in imaging we rewrite F_L in terms of its Fourier transform,

$$F_L(x) = \frac{1}{2\pi}\int_{-L}^{L} \hat{f}(\xi)e^{ix\xi}d\xi + \frac{1}{2\pi}\int_{-L}^{L}\sum_{n\neq 0}\hat{f}(\xi + 2nL)e^{ix\xi}d\xi.$$

The first term is the partial Fourier inverse of f. For a function with jump discontinuities this term produces Gibbs oscillations. The second term is the aliasing error itself. In many examples of interest in imaging, the Gibbs artifacts and the aliasing error are about the same size. What makes either term a problem is slow decay of the Fourier transform of f.

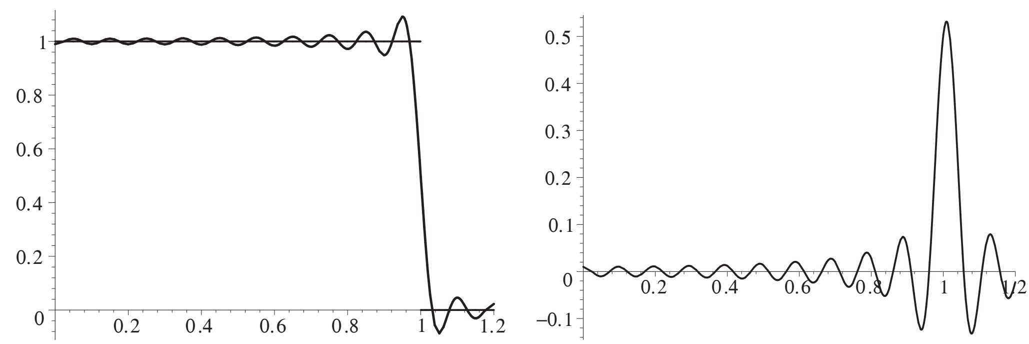

(a) Partial Fourier inverse. (b) Pure aliasing contribution.

Figure 8.4. The two faces of aliasing, $d = .05$.

***Example* 8.2.2.** In Figure 8.4 the two contributions to $f - F_L$ are shown separately, for the rectangle function $f = \chi_{[-1,1]}$. Figure 8.4(a) shows the Gibbs contribution; Figure 8.4(b) shows the "pure aliasing" part. Figure 8.5 shows the original function, its partial Fourier inverse, $\mathcal{F}^{-1}(\hat{f}\chi_{[-L,L]})$, and its Shannon-Whittaker interpolant. The partial Fourier inverse is the medium weight line. The curve slightly to the right of this line is the Shannon-Whittaker interpolant. In this example, the contributions of the Gibbs artifact and the pure aliasing error are of about the same size and have same general character. It is evident that the Shannon-Whittaker interpolant is more distorted than the partial inverse of the Fourier transform, though visually they are quite similar.

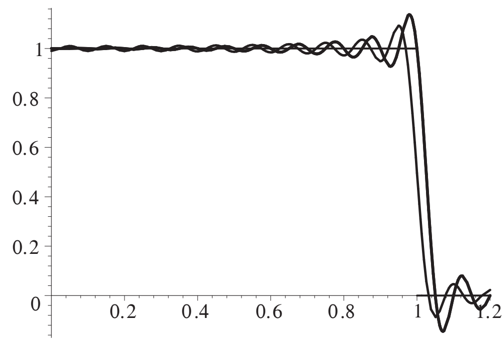

Figure 8.5. Partial Fourier inverse and Shannon-Whittaker interpolant.

***Example* 8.2.3.** For comparison, consider the continuous function $g(x) = \chi_{[-1,1]}(x)(1 - x^2)$ and its reconstruction using the sample spacing $d = .1$. In Figure 8.6(a) it is just barely possible to distinguish the original function from its approximate reconstruction. The worst errors occur near the points ± 1, where g is finitely differentiable. Figure 8.6(b) shows the graph of the difference, $g - G_L$ (note the scale along the y-axis).

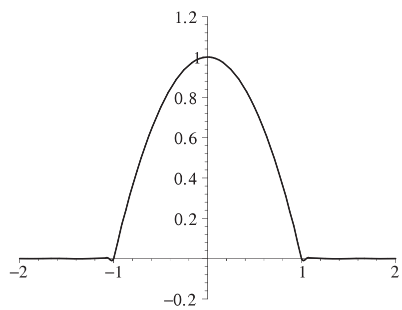

(a) The Shannon-Whittaker interpolation.

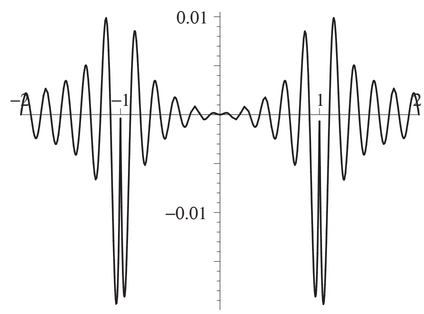

(b) The difference.

Figure 8.6. What aliasing looks like for a continuous, piecewise differentiable function, $d = 0.1$.

***Example* 8.2.4.** As a final example we consider the effect of sampling on a "fuzzy function." Here we use a function of the sort introduced in Example4.2.5. These are continuous functions with "sparse," but slowly decaying, Fourier transforms. Figure 8.7(a) is the graph of such a function, and Figure 8.7(b) shows the Shannon-Whittaker interpolants with $d = .1, .05$, and $.025$. For a function of this sort, Shannon-Whittaker interpolation appears to produce *smoothing*.

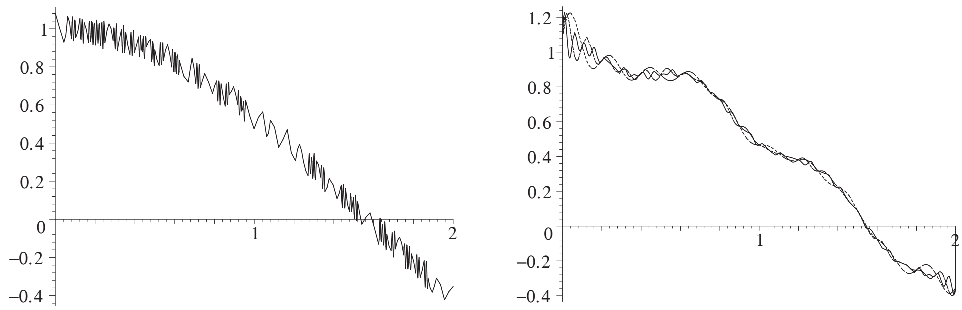

(a) A fuzzy function. (b) Shannon-Whittaker interpolants.

Figure 8.7. What aliasing looks like for a fuzzy function, $d = .1, .05, .025$.

***Remark* 8.2.2.*** The functions encountered in imaging applications are usually spatially limited and therefore cannot be bandlimited. However, if a function f is smooth enough, then its Fourier transform decays rapidly and therefore, by choosing L sufficiently large, the difference, $\hat{f} - \hat{F}_L$ can be made "small." If this is so, then the *effective support* of $\hat{f}$ is contained in $[-L, L]$, and f is an *effectively bandlimited function*. Though effective support is an important concept in imaging, it does not have a precise definition. In most applications it does not suffice to have $\hat{f}$ itself small outside of $[-L, L]$. Usually the sampling rate must be large enough so that the aliasing error,

$$\sum_{n \neq 0} \hat{f}(\xi + 2nL),$$

is also small. This is a somewhat heuristic principle because the meaning of *small* is dictated by the application.

Examples 8.2.2 and 8.2.3 illustrate what is meant by effective bandlimiting. Neither function is actually bandlimited. No matter how large L is taken, a Shannon-Whittaker interpolant for f, in Example 8.2.2, displays large oscillatory artifacts. In most applications this function would not be considered effectively L-bandlimited, for any L. However, it should be noted that, away from the jumps, the Shannon-Whittaker interpolant does a good job reconstructing f. For most purposes, the function g would be considered effectively bandlimited, though the precise effective bandwidth would depend on the application.

To diminish the effects of aliasing, an analogue signal may be passed through a "lowpass filter" *before* it is sampled. In general terms, a lowpass filter is an operation that attenuates the high-frequency content of a signal without introducing too much distortion into the low-frequency content. In this way the sampled data accurately represent the low-frequency information present in the original signal without corruption from the high frequencies. An *ideal* lowpass filter removes all the high-frequency content in a signal outside of given band, leaving the data within the *passband* unchanged. An ideal lowpass filter replaces f with the signal f_L, defined by the following properties:

$$\begin{aligned} \hat{f}_L(\xi) &= \hat{f}(\xi) \text{ if } |\xi| \leq L, \\ \hat{f}_L(\xi) &= 0 \text{ if } |\xi| \geq L. \end{aligned} \tag{8.17}$$

The samples $\{f_L(\frac{n\pi}{L})\}$ contain all the low frequency-information in f *without* the aliasing errors. Using the Shannon-Whittaker formula to reconstruct a function, with these samples, gives f_L for all x. This function is just the partial Fourier inverse of f,

$$f_L(x) = \frac{1}{2\pi} \int_{-L}^{L} \hat{f}(\xi) e^{ix\xi} d\xi,$$

and is still subject to unpleasant artifacts like the Gibbs phenomenon.

A realistic measurement consists of samples of a convolution $\varphi * f$, with φ a function with total integral one. If it has support in $[-\eta, \eta]$, then the measurement at $\frac{n\pi}{L}$ is the average,

$$\varphi * f(\frac{n\pi}{L}) = \int_{-\eta}^{\eta} f(\frac{n\pi}{L} - x)\varphi(x)\,dx.$$

Hence "measuring" the function f at $x = \frac{n\pi}{L}$ is the same thing as sampling the convolution $\varphi * f$ at $\frac{n\pi}{L}$. The Fourier transform of φ goes to zero as the frequency goes to infinity; the smoother φ is, the faster this occurs. As the Fourier transform of $\varphi * f$ is

$$\widehat{\varphi * f}(\xi) = \hat{\varphi}(\xi)\hat{f}(\xi),$$

the measurement process itself attenuates the high-frequency content of f. On the other hand,

$$\hat{\varphi}(0) = \int_{-\infty}^{\infty} \varphi(x)\,dx = 1$$

and therefore $\widehat{\varphi * f}$ resembles $\hat{f}$ for sufficiently low frequencies. Most measurement processes provide a form of lowpass filtering.

The more sharply peaked φ is, the larger the interval over which the "measurement error,"

$$\widehat{\varphi * f}(\xi) - \hat{f}(\xi) = (1 - \hat{\varphi}(\xi))\hat{f}(\xi)$$

can be controlled. The aliasing error in the measured samples is

$$\sum_{n \neq 0} \hat{\varphi}(\xi + 2nL)\hat{f}(\xi + 2nL).$$

By choosing φ to be smooth, this can be made as small as we like. If φ is selected so that $\hat{\varphi}(nL) = 0$ for $n \in \mathbb{Z} \setminus \{0\}$, then the Gibbs-like artifacts that result from truncating the Fourier transform to the interval $[-L, L]$ can also be eliminated.

***Remark* 8.2.3.** A detailed introduction to *wavelets* that includes interesting generalizations of the Poisson formula and the Shannon-Whittaker formula can be found in [58].

Exercises

***Exercise* 8.2.5.** Derive formula (8.15) for $\widehat{F}_L$.

***Exercise* 8.2.6.** Compute the Fourier transform of $g(x) = \chi_{[-1,1]}(x)(1 - x^2)$.

***Exercise* 8.2.7.** What forward velocity of the car in Example8.2.1 corresponds to the apparent rotational velocity of the wheels? What if the car is going 40 mph?

***Exercise* 8.2.8.** Sometimes in a motion picture or television image the wheels of a car appear to be going clockwise, even though the car is moving forward. Explain this by giving an example.

***Exercise* 8.2.9.** Explain why the artifact produced by aliasing looks like the Gibbs phenomenon. For the function $\chi_{[-1,1]}$, explain why the size of the pointwise error, in the Shannon-Whittaker interpolant, does not go to zero as the sample spacing goes to zero.

***Exercise* 8.2.10.** Experiment with the family of functions

$$f_\alpha(x) = \chi_{[-1,1]}(x)(1 - x^2)^\alpha$$

to understand effective bandlimiting. For a collection of $\alpha \in [0, 2]$, see whether there is a Gibbs-like artifact in the Shannon-Whittaker interpolants and, if not, at what sample spacing is the Shannon-Whittaker interpolant visually indistinguishable from the original function (over $[-2, 2]$).

***Exercise* 8.2.11.** The ideal lowpass filtered function, f_L, can be expressed as a convolution

$$f_L(x) = f * k_L(x).$$

Find the function k_L. If the variable x is "time," explain the difficulty in implementing an ideal lowpass filter.

***Exercise* 8.2.12.** Suppose that φ is a non-negative, even, real valued function such that $\hat{\varphi}(0) = 1$; explain why the interval over which $|\hat{f}(\xi) - \widehat{\varphi * f}(\xi)|$ is small is controlled by

$$\int_{-\infty}^{\infty} x^2 \varphi(x)\, dx.$$

***Exercise* 8.2.13.** Show that if

$$\psi(x) = \varphi * \chi_{[-\frac{\pi}{L}, \frac{\pi}{L}]}(x),$$

then $\hat{\psi}(nL) = 0$ for all $n \in \mathbb{Z} \setminus \{0\}$.

8.2.3 Subsampling

Subsampling is a way to take advantage of aliasing to "demodulate" a bandlimited signal whose Fourier transform is supported in a set of the form $[-\omega-B, -\omega+B]$ or $[\omega-B, \omega+B]$. In this context ω is called the *carrier frequency* and $2B$ the bandwidth of the signal. This situation arises in FM radio as well as in MR imaging. For simplicity, suppose that there is a positive integer N so that

$$\omega = NB.$$

Let f be a function whose Fourier transform is supported in $[\omega - B, \omega + B]$. If we sample this function at the points $\{\frac{n\pi}{B}, : \ n \in \mathbb{Z}\}$ and use formula (8.5) to obtain F_L, then (8.13) implies that

$$\widehat{F}_L(\xi) = \hat{f}(\omega + \xi). \tag{8.18}$$

The function F_L is called the demodulated version of f; the two signals are very simply related:

$$f(x) = e^{-i\omega x} F_L(x).$$

From the formula relating f and F_L, it is clear that if F_L is real valued, then, in general, the measured signal, f, is not. A similar analysis applies to a signal with Fourier transform supported in $[-\omega - B, -\omega + B]$.

Exercises

***Exercise* 8.2.14.** Suppose that ω is not an integer multiple of B and that f is a signal whose Fourier transform is supported in $[\omega-B, \omega+B]$. If F_L is constructed as previously from the samples $\{f(\frac{n\pi}{B})\}$, then determine $\widehat{F}_L$. What should be done to get a faithfully demodulated signal by sampling? Keep in mind that normally $\omega >> B$.

***Exercise* 8.2.15.** Suppose that f is a real-valued function whose Fourier transform is supported in $[-\omega - B, -\omega + B] \cup [\omega - B, \omega + B]$. Assuming that $\omega = NB$ and f is sampled at $\{\frac{n\pi}{B}\}$, how is $\widehat{F}_L$ related to $\hat{f}$?

8.3 Periodic Functions and the Finite Fourier Transform*

See: 2.3.2.

We now adapt the discussion of sampling to the context of periodic functions. In this case a natural *finite* analogue of the Fourier series plays a role analogous to that played, in the study of sampling for functions defined on $\mathbb{R}$, by the Fourier series itself. We begin with the definition and basic properties of the *finite Fourier transform.*

***Definition* 8.3.1.** Suppose that $< x_0, \dots, x_{m-1} >$ is a sequence of complex numbers. The *finite Fourier transform* of this sequence is the sequence $< \hat{x}_0, \dots, \hat{x}_{m-1} >$ of complex numbers defined by

$$\hat{x}_k = \frac{1}{m} \sum_{j=0}^{m-1} x_j e^{-\frac{2\pi ijk}{m}}. \tag{8.19}$$

Sometimes it is denoted by

$$\mathcal{F}_m(< x_0, \dots, x_{m-1} >) = (< \hat{x}_0, \dots, \hat{x}_{m-1} >).$$

Using the formula for the sum of a geometric series and the periodicity of the exponential function, we easily obtain the formulæ

$$\sum_{j=0}^{m-1} \exp\left(\frac{2\pi ij}{m}(k-l)\right) = \begin{cases} m & k = l, \\ 0 & k \neq l. \end{cases} \tag{8.20}$$

These formulæ have a nice geometric interpretation: The set of vectors

$$\{(1, e^{\frac{2\pi ik}{m}}, e^{\frac{4\pi ik}{m}}, \dots, e^{\frac{2(m-1)\pi ik}{m}}) \, : \, k = 0, \dots, m-1\}$$

is an orthogonal basis for $\mathbb{C}^m$. These vectors are obtained by sampling the functions $\{e^{2\pi ikx} : \quad k = 0, \dots, m-1\}$ at the points $\{\frac{j}{m} \, : \, j = 0, \dots, m-1\}$.

The computations in (8.20) show that the inverse of the finite Fourier transform is given by

$$x_j = \sum_{k=0}^{m-1} \hat{x}_k e^{\frac{2\pi ijk}{m}}. \tag{8.21}$$

The formulæ (8.19) and (8.21) defining the sequences $< \hat{x}_k >$ and $< x_j >$, respectively, make sense with k or j any integer. In much the same way, as it is often useful to think of a function defined on $[0, L]$ as an L-periodic functions, it is useful to think of $< x_j >$ and $< \hat{x}_k >$ as bi-infinite sequences satisfying

$$x_{j+m} = x_j \text{ and } \hat{x}_{k+m} = \hat{x}_k. \tag{8.22}$$

Such a sequence is called an *m-periodic sequence*.

The summation in (8.21) is quite similar to that in (8.19); the exponential multipliers have been replaced by their complex conjugates. This means that a fast algorithm for computing $\mathcal{F}_m$ automatically provides a fast algorithm for computing $\mathcal{F}_m^{-1}$. Indeed, if m is a power of 2, then there is a fast algorithm for computing both $\mathcal{F}_m$ and $\mathcal{F}_m^{-1}$. Either transformation requires about $3m \log_2 m$ computations, which should be compared to the $O(m^2)$ computation generally required to multiply an m-vector by an $m \times m$-matrix. This algorithm, called the *fast Fourier transform* or *FFT*, is outlined in Section 10.5.

We now return to our discussion of sampling for periodic functions. Let f be an L-periodic function with Fourier coefficients $< \hat{f}(n) >$.

***Definition* 8.3.2.** A periodic function f is called N-bandlimited if $\hat{f}(n) = 0$ for all n with $|n| \geq N$.

In this case, the inversion formula for Fourier series implies that

$$f(x) = \frac{1}{L} \sum_{n=1-N}^{N-1} \hat{f}(n) e^{\frac{2\pi i n x}{L}}.$$

This is a little simpler than the continuum case since f already lies in the finite-dimensional space of functions spanned by

$$\{e^{\frac{2\pi i n x}{L}} \ : \ 1 - N \leq n \leq N - 1\}.$$

Suppose that f is sampled at $\{\frac{jL}{2N-1} \ : \ j = 0, \ldots, 2N-2\}$. Substituting the Fourier representation of f into the sum defining the finite Fourier transform gives

$$\begin{aligned}
\sum_{j=0}^{2N-2} f(\frac{jL}{2N-1}) e^{-\frac{2\pi i k j}{2N-1}} &= \sum_{j=0}^{2N-2} \frac{1}{L} \sum_{n=1-N}^{N-1} \hat{f}(n) \exp\left(\frac{2\pi i n j L}{(2N-1)L} - \frac{2\pi i k j}{2N-1}\right) \\
&= \frac{1}{L} \sum_{n=1-N}^{N-1} \hat{f}(n) \sum_{j=0}^{2N-2} \exp\left(\frac{2\pi i j}{2N-1}(n-k)\right) \\
&= \frac{2N-1}{L} \hat{f}(k), \quad k \in \{1-N, 2-N, \ldots, N-2, N-1\}.
\end{aligned} \tag{8.23}$$

The relations in (8.20), with m replaced by $2N-1$, are used to go from the second to the third line. This computation shows that if f is an N-bandlimited function, then, but for an overall multiplicative factor, the *finite* Fourier transform of the sequence of samples $< f(0), \ldots, f(\frac{(2N-2)L}{2N-1}) >$ computes the nonzero Fourier coefficients of f itself. From the periodicity of $\mathscr{F}_{2N-1}(< f(0), \ldots, f(\frac{(2N-2)L}{2N-1}) >)$ it follows that

$$\begin{aligned}
&\mathscr{F}_{2N-1}(< f(0), \ldots, f(\frac{(2N-2)L}{2N-1}) >) = \\
&\qquad < \hat{f}(0), \hat{f}(1), \ldots, \hat{f}(N-1), \hat{f}(1-N), \ldots, \hat{f}(-2), \hat{f}(-1) > .
\end{aligned} \tag{8.24}$$

The inversion formula for the finite Fourier transform implies the periodic analogue of Nyquist's theorem.

Theorem 8.3.1 (Nyquist's theorem for periodic functions). *If f is an L-periodic function and $\hat{f}(n) = 0$ for $|n| \geq N$, then f can be reconstructed from the equally spaced samples $\{f(\frac{jL}{2N-1}) \ : \ j = 0, 1, \ldots, (2N-2)\}$.*

From equation (8.23) and the Fourier inversion formula, we derive an interpolation formula analogous to (8.5):

$$
\begin{aligned}
f(x) &= \frac{1}{L}\sum_{n=1-N}^{N-1} \hat{f}(n)e^{\frac{2\pi inx}{L}} \\
&= \frac{1}{L}\sum_{n=1-N}^{N-1} \frac{L}{2N-1}\sum_{j=0}^{2N-2} f\Big(\frac{jL}{2N-1}\Big)e^{-\frac{2\pi inj}{2N-1}}e^{\frac{2\pi inx}{L}} \\
&= \frac{1}{2N-1}\sum_{j=0}^{2N-2} f\Big(\frac{jL}{2N-1}\Big)\sum_{n=1-N}^{N-1} e^{-\frac{2\pi inj}{2N-1}}e^{\frac{2\pi inx}{L}} \\
&= \frac{1}{2N-1}\sum_{j=0}^{2N-2} f\Big(\frac{jL}{2N-1}\Big)\frac{\sin\pi(2N-1)(\frac{x}{L}-\frac{j}{2N-1})}{\sin\pi(\frac{x}{L}-\frac{j}{2N-1})}
\end{aligned}
\tag{8.25}
$$

Even if f is not bandlimited, the last line in (8.25) defines an N-bandlimited function,

$$F_N(x) = \frac{1}{2N-1}\sum_{j=0}^{2N-2} f\Big(\frac{jL}{2N-1}\Big)\frac{\sin\pi(2N-1)(\frac{x}{L}-\frac{j}{2N-1})}{\sin\pi(\frac{x}{L}-\frac{j}{2N-1})}.$$

As before, this function interpolates f at the sample points

$$F_N\Big(\frac{jL}{2N-1}\Big) = f\Big(\frac{jL}{2N-1}\Big), \quad j = 0, 1, \ldots, (2N-2).$$

The Fourier coefficients of F_N are related to those of f by

$$\hat{F}_N(k) = \sum_{n=-\infty}^{\infty} \hat{f}(k+n(2N-1)) = \hat{f}(k) + \sum_{n\neq 0}\hat{f}(k+n(2N-1)) \quad 1-N \leq k \leq N-1.$$

If f is not N-bandlimited, then F_N has aliasing distortion: High-frequency data in f distort the low frequencies in F_N. Of course, if f is discontinuous, then F_N also displays Gibbs oscillations.

Exercises

***Exercise* 8.3.1.** Prove (8.20). Remember to use the Hermitian inner product!

***Exercise* 8.3.2.** Explain formula (8.24). What happens if f is not N-bandlimited?

***Exercise* 8.3.3.**⋆ As an m-periodic sequence, $< x_0, \ldots, x_{m-1} >$ is even if

$$x_j = x_{m-j} \text{ for } j = 1, \ldots, m$$

and odd if

$$x_j = -x_{m-j} \text{ for } j = 1, \ldots, m.$$

Show that the finite Fourier transform of a real-valued, even sequence is real valued and the finite Fourier transform of a real-valued, odd sequence is imaginary valued.

***Exercise* 8.3.4.** Suppose that f is an N-bandlimited, L-periodic function. For a subset $\{x_1, \ldots, x_{2N-1}\}$ of $[0, L)$ such that

$$x_j \neq x_k \text{ if } j \neq k,$$

show that f can be reconstructed from the samples

$$\{f(x_j) \,:\, j = 1, \ldots, 2N-1\}.$$

From the point of view of computation, explain why equally spaced samples are preferable.

***Exercise* 8.3.5.** Prove that F_N interpolates f at the sample points:

$$F_N(\frac{jL}{2N-1}) = f(\frac{jL}{2N-1}),\ j = 0, 1, \ldots, 2N-2.$$

***Exercise* 8.3.6.** Find analogues of the generalized Shannon-Whittaker formula in the periodic case.

8.4 Quantization Errors

See: A.1.

In the foregoing sections it is implicitly assumed that we have a continuum of numbers at our disposal to make measurements and do computations. As digital computers are used to implement the various filters, this is not the case. In Section A.1.2 we briefly discuss how numbers are actually stored and represented in a computer. For simplicity we consider a base 2, fixed-point representation of numbers. Suppose that we have $(n+1)$ bits and let the binary sequence $(b_0, b_1, \ldots, b_n)$ correspond to the number

$$(b_0, b_1, \ldots, b_n) \leftrightarrow (-1)^{b_0} \frac{\sum_{j=0}^{n-1} b_{j+1} 2^j}{2^n}. \tag{8.26}$$

This allows us to represent numbers between -1 and $+1$ with a maximum error of 2^{-n}. There are several ways to map the continuum onto $(n+1)$-bit binary sequences. Such a correspondence is called a *quantization* map. In essentially any approach, numbers greater than or equal to 1 are mapped to $(0, 1, \ldots, 1)$, and those less than or equal to -1 are mapped to $(1, 1, \ldots, 1)$. This is called clipping and is very undesirable in applications. To avoid clipping, the data are usually scaled before they are quantized.

The two principal quantization schemes are called *rounding* and *truncation*. For a number x between -1 and $+1$, its rounding is defined to be the number of the form in (8.26) closest to x. If we denote this by $Q_r(x)$, then clearly

$$|Q_r(x) - x| \leq \frac{1}{2^{n+1}}.$$

There exist finitely many numbers that are equally close to two such numbers; for these values a choice simply has to be made. If

$$x = (-1)^{b_0} \frac{\sum_{j=-\infty}^{n-1} b_{j+1} 2^j}{2^n}, \quad \text{where } b_j \in \{0, 1\},$$

then its $(n+1)$-bit truncation corresponds to the binary sequence $(b_0, b_1, \dots, b_n)$. If we denote this quantization map by $Q_t(x)$, then

$$0 \le x - Q_t(x) \le \frac{1}{2^n}.$$

We use the notation $Q(x)$ to refer to either quantization scheme.

Not only are measurements quantized, but arithmetic operations are as well. The usual arithmetic operations must be followed by a quantization step in order for the result of an addition or multiplication to fit into the same number of bits. The machine uses

$$Q(Q(x) + Q(y)) \text{ and } Q(Q(x) \cdot Q(y))$$

for addition and multiplication, respectively. The details of these operations depend on both the quantization scheme and the representation of numbers (i.e., fixed point or floating point). We consider only the fixed point representation. If $Q(x)$, $Q(y)$ and $Q(x) + Q(y)$ all lie between -1 and $+1$, then no further truncation is needed to compute the sum. If $Q(x) + Q(y)$ is greater than $+1$, we have an *overflow* and if the sum is less than -1 an *underflow*. In either case the value of the sum is clipped. On the other hand, if

$$Q(x) = (-1)^{b_0} \frac{\sum_{j=0}^{n-1} b_{j+1} 2^j}{2^n} \text{ and } Q(y) = (-1)^{c_0} \frac{\sum_{j=0}^{n-1} c_{j+1} 2^j}{2^n},$$

then

$$Q(x)Q(y) = (-1)^{b_0+c_0} \frac{\sum_{j,k=1}^{n-1} b_{j+1} c_{k+1} 2^{j+k}}{2^{2n}}.$$

This is essentially a $(2n+1)$-bit binary representation and therefore must be re-quantized to obtain an $(n+1)$-bit representation. Because all numbers lie between $+1$ and -1, overflows and underflows cannot occur in fixed-point multiplication.

It is not difficult to find numbers x and y between -1 and 1 so that $x+y$ is also between -1 and 1 but

$$Q(x + y) \neq Q(x) + Q(y).$$

This means that quantization is not a linear map!

***Example* 8.4.1.** Using truncation as the quantization method and three binary bits, we see that $Q(\frac{3}{16}) = 0$ but $Q(\frac{3}{16} + \frac{3}{16}) = \frac{1}{4}$.

Because it is *nonlinear,* quantization is difficult to analyze. An *exact* analysis requires entirely new techniques. Another approach is to regard the error $e(x) = x - Q(x)$ as *quantization noise*. If $\{x_j\}$ is a sequence of samples, then $\{e_j = x_j - Q(x_j)\}$ is the quantization noise sequence. For this approach to be useful, we need to assume that the sequence $\{e_j\}$ has good statistical properties (e.g., it is of mean zero and the successive values are not highly correlated). If the original signal is sufficiently complex, then this is a good approximation. However, if the original signal is too slowly varying, then these assumptions may not hold. This approach is useful because it allows an analysis of the effect on the signal-to-noise ratio of the number of bits used in the quantization scheme. It is beyond the scope of this text to consider these problems in detail; a thorough treatment and references to the literature can be found in Chapter 9 of [100].

8.5 Higher-Dimensional Sampling

In imaging applications, we usually work with functions of two or three variables. Let f be a function defined on $\mathbb{R}^n$ and let $\{\boldsymbol{x}_k\}$ be a discrete set of points in $\mathbb{R}^n$. As before, the values $\{f(\boldsymbol{x}_k)\}$ are the *samples* of f at the *sample points* $\{\boldsymbol{x}_k\}$. Parts of the theory of sampling in higher dimensions exactly parallel the one-dimensional theory, though the problems of sampling and reconstruction are considerably more complicated.

As in the one-dimensional case, samples are usually collected on a uniform grid. In this case it is more convenient to label the sample points using vectors with integer coordinates. As usual, boldface letters are used to denote such vectors, that is,

$$\boldsymbol{j} = (j_1, \dots, j_n), \text{ where } j_i \in \mathbb{Z}, \quad i = 1, \dots, n.$$

***Definition* 8.5.1.** The *sample spacing* for a set of uniformly spaced samples in $\mathbb{R}^n$ is a vector $\boldsymbol{h} = (h_1, \dots, h_n)$ with positive entries. The index $\boldsymbol{j}$ corresponds to the sample point

$$\boldsymbol{x}_{\boldsymbol{j}} = (j_1 h_1, \dots, j_n h_n).$$

A set of values, $\{f(\boldsymbol{x}_{\boldsymbol{j}})\}$, at these points is a uniform sample set.

A somewhat more general definition of uniform sampling is sometimes useful: Fix n orthogonal vectors $\{\boldsymbol{v}_1, \dots, \boldsymbol{v}_n\}$. For each $\boldsymbol{j} = (j_1, \dots, j_n) \in \mathbb{Z}^n$, define the point

$$\boldsymbol{x}_{\boldsymbol{j}} = j_1 \boldsymbol{v}_1 + \cdots j_n \boldsymbol{v}_n. \tag{8.27}$$

The set of points $\{\boldsymbol{x}_{\boldsymbol{j}} : \boldsymbol{j} \in \mathbb{Z}^n\}$ defines a uniform sample set. This sample set is the result of applying a rotation to a uniform sample set with sample spacing $(\|\boldsymbol{v}_1\|, \dots, \|\boldsymbol{v}_n\|)$. As in the one-dimensional case, the definitions of sample spacing and uniform sampling depend on the choice of coordinate system. A complication in several variables is that there are many different coordinate systems that naturally arise.

***Example* 8.5.1.** Let $(h_1, \dots, h_n)$ be a vector with positive coordinates. The set of points

$$\{(j_1 h_1, \dots, j_n h_n) : (j_1, \dots, j_n \in \mathbb{Z}^n\}$$

is a uniform sample set.

***Example* 8.5.2.** Let (r, θ) denote polar coordinates for $\mathbb{R}^2$; they are related to rectangular coordinates by

$$x = r\cos\theta, \quad y = r\sin\theta.$$

In CT imaging we often encounter functions that are uniformly sampled on a *polar grid.* Let $f(r, \theta)$ be a function on $\mathbb{R}^2$ in terms of polar coordinates and let $\rho > 0$ and $M \in \mathbb{N}$ be fixed. The set of values

$$\{f(j\rho, \frac{2k\pi}{M}) \;:\; j \in \mathbb{Z},\, k = 1, \ldots M\}$$

consists of uniform samples of f, in polar coordinates; however, the points

$$\{(j\rho\cos\left(\frac{2k\pi}{M}\right), \quad j\rho\sin\left(\frac{2k\pi}{M}\right))\}$$

are *not* a uniform sample set as defined previously.

In more than one dimension there are many different, reasonable notions of finite bandwidth, or bandlimited data. If D is any convex subset in $\mathbb{R}^n$ containing 0, then a function is D-bandlimited if $\hat{f}$ is supported in D. The simplest such regions are boxes and balls.

***Definition* 8.5.2.** Let $\boldsymbol{B} = (B_1, \ldots, B_n)$ be an n-tuple of positive numbers. A function f defined in $\mathbb{R}^n$ is $\boldsymbol{B}$-bandlimited if

$$\hat{f}(\xi_1, \ldots, \xi_n) = 0 \text{ if } |\xi_j| > B_j \qquad \text{for } j = 1, \ldots, n. \tag{8.28}$$

***Definition* 8.5.3.** A function f defined in $\mathbb{R}^n$ is R-bandlimited if

$$\hat{f}(\xi_1, \ldots, \xi_n) = 0 \qquad \text{if } \|\xi\| > R \tag{8.29}$$

The Nyquist theorem and Shannon-Whittaker interpolation formula carry over easily to $\boldsymbol{B}$-bandlimited functions. However, these generalizations are often inadequate to handle problems that arise in practice. The generalization of Nyquist's theorem is as follows:

Theorem 8.5.1 (Higher-dimensional Nyquist theorem). *Let $\boldsymbol{B} = (B_1, \ldots, B_n)$ be an n-tuple of positive numbers. If f is a square-integrable function that is $\boldsymbol{B}$-bandlimited, then f can be reconstructed from the samples*

$$\{f(\frac{j_1\pi}{B_1}, \ldots, \frac{j_n\pi}{B_n}) \;:\; (j_1, \ldots, j_n) \in \mathbb{Z}^n\}.$$

This result is "optimal."

In order to apply this result to an R-bandlimited function, we would need to collect the samples:

$$\{f(\frac{j_1\pi}{R}, \ldots, \frac{j_n\pi}{R}) \;:\; (j_1, \ldots, j_m) \in \mathbb{Z}^n\}.$$

As $\hat{f}$ is known to vanish in a large part of $[-R, R]^n$, this would appear to be some sort of oversampling.

Neither Theorem 8.1.1 nor Theorem 8.5.1 say anything about nonuniform sampling. It is less of an issue in one dimension. If the vectors, $\{\boldsymbol{v}_1, \ldots, \boldsymbol{v}_n\}$ are linearly independent but not orthogonal, then formula (8.27) defines a set of sample points $\{\boldsymbol{x}_{\boldsymbol{j}}\}$. Unfortunately, Nyquist's theorem is not directly applicable to decide whether or not the set of samples $\{f(\boldsymbol{x}_{\boldsymbol{j}}) : \boldsymbol{j} \in \mathbb{Z}^n\}$ suffices to determine f. There are many results in the mathematics literature that state that a function whose Fourier transform has certain support properties is determined by its samples on appropriate subsets, though few results give an explicit interpolation formula like (8.5). The interested reader is referred to [66] and [102].

The Poisson summation formula also has higher-dimensional generalizations. If f is a rapidly decreasing function, then

$$f_p(\boldsymbol{x}) = \sum_{\boldsymbol{j} \in \mathbb{Z}^n} f(\boldsymbol{x} + \boldsymbol{j})$$

is a periodic function. The Fourier coefficients of f_p are related to the Fourier transform of f in much the same way as in one dimension:

$$\begin{aligned}\widehat{f_p}(\boldsymbol{k}) &= \int\limits_{[0,1]^n} f_p(\boldsymbol{x}) e^{-2\pi i \langle \boldsymbol{x}, \boldsymbol{k} \rangle} d\boldsymbol{x} \\ &= \int\limits_{\mathbb{R}^n} f(\boldsymbol{x}) e^{-2\pi i \langle \boldsymbol{x}, \boldsymbol{k} \rangle} d\boldsymbol{x} \\ &= \hat{f}(2\pi i \boldsymbol{k}).\end{aligned} \tag{8.30}$$

Applying the Fourier series inversion formula with a function that is smooth enough and decays rapidly enough shows that

$$\sum_{\boldsymbol{j} \in \mathbb{Z}^n} f(\boldsymbol{x} + \boldsymbol{j}) = \sum_{\boldsymbol{k} \in \mathbb{Z}^n} \hat{f}(2\pi i \boldsymbol{k}) e^{2\pi i \langle \boldsymbol{x}, \boldsymbol{k} \rangle}. \tag{8.31}$$

This is the n-dimensional Poisson summation formula.

The set of sample points is sometimes determined by the physical apparatus used to make the measurements. As such, we often have samples $\{f(\boldsymbol{y}_{\boldsymbol{k}})\}$, of a function, on a nonuniform grid $\{\boldsymbol{y}_{\boldsymbol{k}}\}$. To use computationally efficient methods, it is often important to have samples on a uniform grid $\{\boldsymbol{x}_{\boldsymbol{j}}\}$. To that end, *approximate* values for f, at these points, are obtained by interpolation. Most interpolation schemes involve averaging the known values at nearby points. For example, suppose that $\{\boldsymbol{y}_{\boldsymbol{k}_1}, \ldots, \boldsymbol{y}_{\boldsymbol{k}_l}\}$ are the points in the nonuniform grid closest to $\boldsymbol{x}_{\boldsymbol{j}}$ and there are numbers $\{\lambda_i\}$, all between 0 and 1, so that

$$\boldsymbol{x}_{\boldsymbol{j}} = \sum_{i=1}^{l} \lambda_i \boldsymbol{y}_{\boldsymbol{k}_i}.$$

A reasonable way to assign a value to f at $\boldsymbol{x}_{\boldsymbol{j}}$ is to set

$$f(\boldsymbol{x}_{\boldsymbol{j}}) \overset{d}{=} \sum_{i=1}^{l} \lambda_i f(\boldsymbol{y}_{\boldsymbol{k}_i}).$$

This sort of averaging is not the result of convolution with an L^1-function and does not produce smoothing. The success of such methods depends critically on the smoothness of f. A somewhat more robust and efficient method for multi-variable interpolation is discussed in Section 11.8. Another approach to nonuniform sampling is to find a computational scheme adapted to the nonuniform grid. An example of this is presented in Section 11.5.

Exercises

***Exercise* 8.5.1.** Prove Theorem 8.5.1.

***Exercise* 8.5.2.** Find an n-dimensional generalization of the Shannon-Whittaker interpolation formula (8.5).

***Exercise* 8.5.3.** Give a definition of oversampling and a generalization of formula (8.7) for the n-dimensional case.

***Exercise* 8.5.4.** For a set of linearly independent vectors $\{\boldsymbol{v}_1, \dots, \boldsymbol{v}_n\}$, find a notion of $\boldsymbol{V}$-bandlimited so that a $\boldsymbol{V}$-bandlimited function is determined by the samples $\{f(\boldsymbol{x}_{\boldsymbol{j}}) \ : \ \boldsymbol{j} \in \mathbb{Z}^n\}$ where $\boldsymbol{x}_{\boldsymbol{j}} = j_1\boldsymbol{v}_1 + \cdots + j_n\boldsymbol{v}_n$. Show that your result is optimal.

***Exercise* 8.5.5.** Using the results proved earlier about Fourier series, give hypotheses on the smoothness and decay of f that are sufficient for (8.31) to be true.

8.6 Conclusion

Data acquired in medical imaging are usually described as samples of a function of continuous variables. Nyquist's theorem and the Poisson summation formula provide a precise and quantitative description of the errors, known as aliasing errors, introduced by sampling. The Shannon-Whittaker formula and its variants give practical methods for approximating (or reconstructing) a function from a discrete set of samples. The finite Fourier transform, introduced in Section 8.3, is the form in which Fourier analysis is finally employed in applications. In the next chapter we reinterpret (and rename) many of the results from earlier chapters in the context and language of filtering theory. In Chapter 10, we analyze how the finite Fourier transform provides an approximation to the both the Fourier transform and Fourier series and use this analysis to approximately implement (continuum) shift invariant filters on finitely sampled data. This constitutes the final step in the transition from abstract continuum models to finite algorithms.