

Chapter 1

Introduction

Abstract “Medical imaging refers to several different technologies that are used to view the human body in order to diagnose, monitor, or treat medical conditions”. All imaging modalities have in common that the medical condition becomes visible by some form of contrast, meaning that the feature of interest (such as a tumor) can be recognized in the image and examined by a trained radiologist. The image can be seen as a *model* of the imaged tissue. Images in the context of this book are digital. This implies a finite resolution with the pixel as the smallest element. Furthermore, all imaging modalities lead to some degradation of the image when compared to the original object. Primarily, the degradation consists of blur (loss of detail) and noise (unwanted contrast). Some underlying principles are common to all imaging modalities, such as the interpretation as a *system* and its mathematical treatment. The image itself can be seen as a multidimensional *signal*. In many cases, the steps in image formation can be seen as *linear systems*, which allow simplified mathematical treatment.

“Medical imaging refers to several different technologies that are used to view the human body in order to diagnose, monitor, or treat medical conditions. Each type of technology gives different information about the area of the body being studied or treated, related to possible disease, injury, or the effectiveness of medical treatment”.

This concise definition by the US Food and Drug Administration illuminates the goal of medical imaging: To make a specific condition or disease visible. In this context, *visible* implies that the area of interest is distinguishable in some fashion (for example, by a different shade or color) from the surrounding tissue and, ideally, from healthy, normal tissue. The difference in shade or color can be generalized with the term *contrast*.

The process of gathering data to create a visible model (i.e., the image) is common to all medical imaging technologies and can be explained with the simple example of a visible-light camera. The sample is probed with incident light, and reflected light carries the desired information. For example, a melanoma of the skin would reflect less light than the surrounding healthy skin. The camera lens collects some of the reflected light and—most importantly—focuses the light onto the film or

image sensor in such a way that a spatial relationship exists between the origin of the light ray and its location on the image sensor. The ability to spatially resolve a signal (in this example, light intensity) is fundamental to every imaging method. The ability to spatially resolve a signal can be fairly straightforward (for example, following an X-ray beam along a straight path) or fairly complex (for example in magnetic resonance imaging, where a radiofrequency signal is encoded spatially by its frequency and its phase).

In the next step of the process, the spatially resolved data are accumulated. Once again, the camera analogy is helpful. At the start of the exposure, the sensor array is reset. Over the duration of the exposure, incoming light creates a number of electrical charges that depends on the light intensity. At the end of the exposure, the charges are transferred from the sensor to a storage medium. From here, the image would typically be displayed in such a fashion that higher charge read-outs correspond to higher screen intensity. In the camera example, the relationship between reflected light intensity and displayed intensity is straightforward. In other cases, intensity relates to different physical properties. Examples include X-ray absorption (which gives X-ray images the characteristic negative appearance with bones appearing bright and air dark), concentration of a radioactively labeled compound, or the time it takes for a proton to regain its equilibrium orientation in a magnetic field.

The physical interpretation of image intensity is key to interpreting the image, and the underlying physical process is fundamental to achieving the desired contrast. As a consequence, the information encoded in the image varies fundamentally between image modalities and, in some cases (such as MRI), even within the same modality.

The image is evaluated by an experienced professional, usually a radiologist. Even in today's age of automated image analysis and computerized image understanding, the radiologist combines the information encoded in the image with knowledge of the patient's symptoms and history and with knowledge of anatomy and pathology to finally form a diagnosis. Traditional viewing of film over a light box is still prominent, even with purely digital imaging modalities, although more and more radiologists make use of on-the-fly capabilities of the digital imaging workstation to view and enhance images. Furthermore, computerized image processing can help enhance the image, for example, by noise reduction, emphasizing edges, improving contrast, or taking measurements.

1.1 A Brief Historical Overview

X-rays were discovered in 1895. Within less than a decade, which is an astonishingly short time, X-ray imaging became a main-stream diagnostic procedure and was adopted by most major hospitals in Europe and the USA. At that time, sensitivity was low, and exposure times for a single image were very long. The biological effects of X-rays were poorly explored, and radiation burns were common in the early years of diagnostic—and recreational—X-ray use. As the pernicious effects of ionizing radiation became better understood, efforts were made to shield operators

from radiation and to reduce patient exposure. However, for half a century, X-ray imaging did not change in any fundamental fashion, and X-ray imaging remained the only way to provide images from inside the body.

The development of sonar (sound navigation and ranging) eventually led to the next major discovery in biomedical imaging: ultrasound imaging. After World War II, efforts were made, in part with surplus military equipment, to use sound wave transmission and sound echoes to probe organs inside the human body. Ultrasound imaging is unique in that image formation can take place with purely analog circuits. As such, ultrasound imaging was feasible with state-of-the-art electronics in the 1940s and 1950s (meaning: analog signal processing with vacuum tubes). Progress in medical imaging modalities accelerated dramatically with the advent of digital electronics and, most notably, digital computers for data processing. In fact, with the exception of film-based radiography, all modern modalities rely on computers for image formation. Even ultrasound imaging now involves digital filtering and computer-based image enhancement.

In 1972, Geoffrey Hounsfield introduced a revolutionary new device that was capable of providing cross-sectional, rather than planar, images with X-rays. He called the method *tomography*, from the Greek words to cut and to write [7]. The imaging modality is known as computed tomography (CT) or computer-aided tomography (CAT), and it was the first imaging modality that *required* the use of digital computers for image formation. CT technology aided the development of emission tomography, and the first CT scanner was soon followed by the first positron emission tomography scanner.


The next milestone, magnetic resonance imaging (MRI), was introduced in the late 1970s. MRI, too, relies on digital data processing, in part because it uses the Fourier transform to provide the cross-sectional image. Since then, progress became more incremental, with substantial advances in image quality and acquisition speed. The resolution and tissue discrimination of both CT and MRI, for example, that today's devices are capable of, was literally unthinkable at the time these devices were introduced. In parallel, digital image processing and the digital imaging workstation provided the radiologist with new tools to examine images and provide a diagnosis. Three-dimensional image display, multi-modality image matching, and preoperative surgery planning were made possible by computerized image processing and display.

A present trend exists toward the development of imaging modalities based on visible or infrared light. Optical coherence tomography (OCT) became widely known in the 1990s and has evolved into a mainstream method to provide cross-sectional scans of the retina and skin. Other evolving optical modalities, such as diffuse optical tomography, have not reached the maturity level that would allow its use in medical practice.

1.2 Image Resolution and Contrast

Digital images are discretely sampled on a rectangular grid. A digital camera again illustrates the nature of a digital image: the camera sensor is composed of

Fig. 1.1 Sketch of a magnified part of a digital camera image sensor. Each sensor cell consists of a light-sensitive photodiode (gray-shaded area) and associated amplifier and driver circuitry (hatched region). Each sensor cell averages the light across its sensitive surface and provides one single intensity value




millions of light-sensitive cells. A sketch of a few cells, strongly magnified, is shown in Fig. 1.1. Each single sensor cell is composed of a light-sensitive semiconductor element (photodiode) and its associated amplifier and drive circuitry. The cells are spaced Δx apart in the horizontal direction, and Δy in the vertical direction. The actual light-sensitive area is smaller, x_s by y_s . To illustrate these dimensions, let us assume a 12-megapixel sensor with 4,000 cells in the horizontal and 3,000 cells in the vertical direction. When the overall dimensions of the sensor chip are 24 by 18 mm, we know $\Delta x = \Delta y = 6 \mu\text{m}$. Depending on the chip design, the photodiode occupies most of the space, for example, $x_s = y_s = 5 \mu\text{m}$. Irrespective of the amount of detail in the image projected onto the sensor, detail information smaller than the size of a sensor cell is lost, because the photodiode averages the intensity over its surface, and the surrounding driver is not sensitive to light. Each cell (i.e., each pixel), therefore, provides one single intensity value that is representative of the area it occupies.

The spatial resolution of the most important medical imaging modalities spans a wide range. Planar X-ray imaging can achieve a spatial resolution of up to $10 \mu\text{m}$, in part limited by the film grain. Digital X-ray sensors can achieve a similarly high resolution, although $20\text{--}50 \mu\text{m}$ pixel size is more common. With CT, in-plane pixel sizes between 0.1 and 0.5 mm are common in whole-body scanners. MRI scanners have typical in-plane pixels of 0.5–1 mm. Due to the different detector system, radionuclide imaging modalities (SPECT and PET) have pixel sizes in the centimeter range. Ultrasound resolution lies between CT and MRI.

The sensor is not the only limiting factor for the spatial resolution. An ideally focused light source is spread out by the camera lens, primarily as a consequence of lens shape approximations and light diffraction. The image of a point source is called the *point-spread function*. The importance of the point-spread function is demonstrated in Fig. 1.2. The image shows photos of tightly focused laser beams taken with a digital SLR camera from 2 m distance. It can be seen that the image of a single beam shows a Gaussian profile (Fig. 1.2a). An ideal imaging apparatus would provide a delta function (i.e., a cylinder of one pixel width). The point-spread function can be quantified by its full width at half-maximum (FWHM), that is, the

Fig. 1.2 Point-spread function of a digital camera, shown in grayscale representation and as an elevation map where intensity translates into height. The pixel size Δx and Δy is indicated. **a** The image of a highly focused laser beam has a Gaussian shape in the image. The full width at half-maximum (FWHM) spread is 6 pixels. **b** Two closely spaced sources can be distinguished if their distance is larger than the FWHM. **c** Two sources that are more closely spaced than the FWHM become indistinguishable from a single source



width of the point image where it drops to one half of its peak value (Fig. 1.2a). In this example, we observe a FWHM of 6 pixels. As long as two closely spaced point sources are further apart than the FWHM, they can be distinguished as two separate peaks (Fig. 1.2b), which is no longer possible when the point sources are closer than the FWHM (Fig. 1.2c).

Clearly, the point-spread function poses a limit on the spatial resolution, often more so than the detector size. In X-ray imaging, for example, one factor that determines the point-spread function is the active area of the X-ray tube. In ultrasound imaging, factors are the length of the initial ultrasound pulse and the diameter of the ultrasound beam. Furthermore, the wavelength of the sound wave itself is a limiting factor.

The image values are stored digitally. A certain number of bits is set aside for each cell (each pixel). Since each bit can hold two values (one and zero), a n -bit pixel can hold 2^n discrete intensity levels. Color photos are commonly stored with 24 bits per pixel, with 8 bits each for the three fundamental colors, red, green, and blue. For each color, 256 intensity levels are possible. Most magnetic resonance and ultrasound images are also stored with 8 bits depth, whereas computed tomography normally provides 12 bits.

The pixel size determines the absolute limit for the spatial resolution, and the bit depth determines the contrast limit. Consider an 8-bit image: the intensity increase from one discrete image value to the next is 0.39% of the maximum value. Any smaller intensity variations cannot be represented. The error that is associated with rounding of a continuous signal to the next possible image value is referred to as *digitization noise*.

Noise is introduced in several steps of the acquisition and processing chain. Both the sensors and the amplifiers introduce noise components, particularly when weak signals need to be amplified by a large gain factor. Examples are the RF echo signal in MRI and the ultrasound echo in ultrasound imaging. To some extent noise can be suppressed with suitable filters, but the side-effect is a broadening of the point-spread function and the associated loss of detail. Conversely, any filter that tries to counteract the point-spread function increases the noise component. The noise component is critical for the overall image quality, because noise can “bury” detail information from small objects or objects with low contrast.

The ability to provide a specific, desired contrast depends strongly on the modality. X-ray imaging, for example, provides very strong contrast between bone and soft tissue, and between soft tissue and air (e.g., in images of the lung or the chest). Magnetic resonance imaging shows high contrast between different types of soft tissue (e.g., gray and white matter of the brain), but bone and air are dark due to the absence of water. Ultrasound generally provides good tissue contrast, but suffers from a high noise component, visible as characteristic ultrasound speckles.

1.3 Systems and Signals: A Short Introduction

System is a broad term that encompasses any assembly of interconnected and interacting components that have measurable behavior and a defined response to a defined manipulation of its parts. Any device that provides a medical image is a system in this definition, and it consists in turn of several components that can be seen as systems themselves. Systems have inputs and outputs. One example for a system is an X-ray detector. The number of X-ray photons hitting the conversion layer can be interpreted as the system input. The detector provides a voltage that is proportional to the incident photon flux, and this voltage is the output. Similarly, a computer algorithm for image reconstruction is a system. In a computed tomography scanner, for example, the input to the image reconstruction algorithm is the X-ray intensity as a function of the scan angle and position, and the output is a two-dimensional map (i.e., cross-section) of X-ray absorption coefficients.

The input and output to a system can be interpreted as *signals*. Often, a signal is understood as a function of time, but in imaging devices, signals are functions of a spatial coordinate. In the most general form, imaging devices process signals of the form $f(x, y, z, t)$, that is, a quantity that depends on a location in (three-dimensional) space and on time. Often, simplifications can be made when a signal is approximately constant over the image acquisition time, or when a signal is obtained only within one plane. An example is shown in Fig. 1.3, where components of the scanning and image reconstruction process are shown as blocks with signals represented by arrows.

In any system, the output signal can be described mathematically for a given input signal. The X-ray detector, for example, converts X-ray photon flux $I(t, \theta)$ into a proportional voltage $U(t, \theta)$:




Fig. 1.3 Systems interpretation of a computed tomography scanner. Components of the system (in itself systems) are represented by blocks, and signals represented by arrows. The original object has some property A , for example, X-ray absorption that varies within the x, y -plane. The detector collects X-ray intensity $I(t, \theta)$ as a function of scan direction t and scan angle θ in the (x, y) -plane, and provides a proportional voltage $U(t, \theta)$. In the image formation stage, these data are transformed into a cross-sectional map of apparent X-ray opaqueness, $\mu(x, y)$. Finally, the display outputs a light intensity $I'(x, y)$ that is proportional to $\mu(x, y)$ and approximates $A(x, y)$

$$U(t, \theta) = \alpha \cdot I(t, \theta) \quad (1.1)$$

where α is the gain of the X-ray detector. Similarly, the image reconstruction stage approximates the inverse Radon transform \mathcal{R}^{-1} (see Sect. 3.1.1):

$$\mu(x, y) = \mathcal{R}^{-1} \{U(t, \theta)\}. \quad (1.2)$$

A special group of systems are *linear, time-invariant systems*. These systems are characterized by three properties,

- **Linearity:** If y is the output for a given input x , then a change of the magnitude of the input signal by a constant factor a (i.e., we input ax) leads to a proportional output signal ay .
- **Superposition:** If a system responds to an input signal x_1 with the output signal y_1 and to a different input signal x_2 with y_2 , then the sum of the input signals $x_1 + x_2$ will elicit the response $y_1 + y_2$.
- **Time-invariance:** If $y(t)$ is the time-dependent output signal for a given input signal $x(t)$, then the application of the delayed signal $x(t - \tau)$ causes an identical, but equally delayed response $y(t - \tau)$. In images, time-invariance translates into shift-invariance. This means that an operator that produces an image $I(x, y)$ from an input image produces the same image, but shifted by $\Delta x, \Delta y$, when the input image is shifted by the same distance.

Figure 1.3 provides a different view of the point-spread function: we can see that an object (the original tissue property $A(x, y)$) is probed by some physical means. The image formation process leads to the display of an image $I'(x, y)$, which differs from $A(x, y)$. Referring back to Fig. 1.2, we can see that the image functions of the laser dots are superimposed (i.e., added together). With the superposition principle, we can examine each individual pixel separately and subject it to the point-spread function, then add the results. Very often, the point-spread function has Gaussian character, and we can model the peak seen in Fig. 1.2a as

$$g(r) = \frac{1}{\sigma \sqrt{2\pi}} \exp\left(-\frac{r^2}{2\sigma^2}\right) \quad (1.3)$$

where r is the Euclidean distance from the center pixel of the point source (x_0, y_0) . If we know the signal of the idealized point source S , we can now predict the measured (i.e., blurred with the PSF) intensity for each pixel (x, y) :

$$I(x, y) = \frac{S}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x-x_0)^2 + (y-y_0)^2}{2\sigma^2}\right) = S \cdot g(x-x_0, y-y_0) \quad (1.4)$$

where $g(x-x_0, y-y_0)$ should be seen as a generalized point-spread function whose center is shifted to the center of the point source. Consequently, we can express the image formed by two point sources of strength S_0 and S_1 and centered on (x_0, y_0) and (x_1, y_1) , respectively, as the superposition of the image functions

$$I(x, y) = S_0 \cdot g(x-x_0, y-y_0) + S_1 \cdot g(x-x_1, y-y_1). \quad (1.5)$$

This concept can be further generalized. Assume that we have an idealized (but inaccessible) source image $S(x, y)$ and we measure the image $I(x, y)$ with an imaging device that makes it subject to the point-spread function g . In this case, we can subject *each individual pixel* of $S(x, y)$ to the point-spread function and recombine them by addition:

$$I(x', y') = \sum_y \sum_x S(x, y) \cdot g(x'-x, y'-y). \quad (1.6)$$

The sum in Eq. 1.6 needs to be evaluated for all pixels (x', y') of the target image I . Equation 1.6 describes the two-dimensional discrete *convolution* of the source image S with a convolution function (often called *convolution kernel*) g . Since any bright pixel spreads out and influences its neighbors (thus, point *spread* function), sharp transitions are softened, and detail is lost. The effect is demonstrated in Fig. 1.4, where the idealized image $S(x, y)$ (Fig. 1.4a) has been subjected to a simulated point-spread function, in this case, a Gaussian function with $\sigma \approx 3.5$




Fig. 1.4 Illustration of the effects of a Gaussian point-spread function. **a** Idealized image (note the two added white dots in the top left and top right corners indicated by arrows). **b** Image obtained after a process with a Gaussian point-spread function. **c** Intensity profiles along the horizontal dashed line in **a**. It can be seen that sharp transitions are softened, because higher image values also influence their neighbors. Point sources assume a Gaussian-shaped profile

pixels to reveal the actual image $I(x, y)$ (Fig. 1.4b). The line profiles (Fig. 1.4c) help illustrate how sharp transitions are blurred and how isolated points assume a Gaussian shape.

1.4 The Fourier Transform

The Fourier transform is one of the most important linear operations in image processing, and it is fundamental to most imaging modalities. Intuitively, a *transform* reveals a different aspect of the data. In the case of the Fourier transform, it shows the distribution of harmonic content—how the signal is composed of periodic oscillations of different frequency and amplitude. For example, a time-dependent oscillation $f_\omega(t) = A \sin(\omega t)$ could be described by its amplitude A and its frequency ω . In a diagram $f_\omega(t)$ over t , we obtain an oscillation. In a diagram of amplitude over frequency, the same signal is defined by a single point at (A, ω) . Superimposed sine waves would be represented by multiple points in the diagram of amplitude over frequency. In this simplified explanation, a phase shift $f(t) = A \sin(\omega t + \varphi)$ cannot be considered, because a third dimension becomes necessary to include A , ω , and φ . The Fourier transform uses sine and cosine functions to include the phase shift, and each harmonic oscillation becomes $f(t) = a \cos(\omega t) + b \sin(\omega t)$.

Fourier's theorem states that *any periodic signal $s(t)$ can be represented as an infinite sum of harmonic oscillations*, and the Fourier synthesis of the signal $s(t)$ can be written as

$$s(t) = \frac{a_0}{2} + \sum_{k=1}^{\infty} a_k \cdot \cos(kt) + b_k \cdot \sin(kt) \quad (1.7)$$

where a_k and b_k are the Fourier coefficients that determine the contribution of the k th harmonic to the signal $s(t)$. For any given signal $s(t)$, the Fourier coefficients can be obtained by *Fourier analysis*,

$$\begin{aligned} a_k &= \frac{1}{\pi} \int_{-\pi}^{\pi} s(t) \cos(kt) dt \\ b_k &= \frac{1}{\pi} \int_{-\pi}^{\pi} s(t) \sin(kt) dt. \end{aligned} \quad (1.8)$$

Equation 1.7 describes the synthesis of a signal from harmonics with integer multiples of its fundamental frequency. The spectrum (i.e., the sequence of a_k and b_k) is discrete. The continuous Fourier transform is better derived from a different form of Fourier synthesis that uses a continuous spectrum $a(\omega)$ and $b(\omega)$:

$$s(t) = \int_{\omega} a(\omega) \cdot \cos(\omega t) + b(\omega) \cdot \sin(\omega t) d\omega. \quad (1.9)$$

The integration takes place over all possible frequencies ω . Since the basis functions (sin and cos) in Eq. 1.9 are orthogonal, we can express the Fourier synthesis in terms of a complex harmonic oscillation $e^{j\varphi} = \cos(\varphi) + j \sin(\varphi)$. Fourier synthesis restores a signal from its spectrum and corresponds to the inverse Fourier transform \mathcal{F}^{-1} , whereas the Fourier analysis, which provides the spectrum of a signal, is referred to as the actual Fourier transform \mathcal{F} :

$$S(\omega) = \mathcal{F}\{s(t)\} = \int_{-\infty}^{\infty} s(t) \exp(-j\omega t) dt$$

$$s(t) = \mathcal{F}^{-1}\{S(\omega)\} = \frac{1}{2\pi} \int_{-\infty}^{\infty} S(\omega) \exp(j\omega t) d\omega. \quad (1.10)$$

Equation 1.10 defines the Fourier transform in terms of the angular frequency $\omega = 2\pi f$. In some cases, it is more convenient to express the spectrum $S(f)$ as a function of the linear frequency f , for which the Fourier transform becomes

$$S(f) = \mathcal{F}\{s(t)\} = \int_{-\infty}^{\infty} s(t) \exp(-2\pi j f t) dt$$

$$s(t) = \mathcal{F}^{-1}\{S(f)\} = \int_{-\infty}^{\infty} S(f) \exp(2\pi j f t) d\omega. \quad (1.11)$$

To explain the significance of the Fourier transform, let us consider two examples. First, in magnetic resonance imaging, we deal with signals that are caused by protons spinning at different speeds (cf. Sect. 5.4.3). The angular frequency of the protons increases along one spatial axis (let us call it the y -axis), and the protons emit a signal whose strength is determined, among other factors, by the number of protons at any point along the y -axis. The signal can be collected by an antenna, but the antenna only provides the additive mix of all signals. We can, however, obtain the local proton density by using the relationship $\omega = \omega_0 + m \cdot y$, where m is the rate of change of the frequency along the y -axis. The antenna provides a signal $s(t)$, which we subject to the Fourier transform. The resulting harmonic content $S(\omega) = S(\omega_0 + m \cdot y)$ is directly related to the signal strength at any point along the y -axis and therefore to the proton density.

Second, it is sometimes desirable to have a signal that contains all frequencies in a limited range (i.e., a broadband signal). We can ask the question, how would a broadband signal $b(t)$ look like for which the spectral component is unity for all frequencies between $-f_0$ and $+f_0$?¹ To answer this question, we use the inverse

¹ Examining negative frequencies is not unreasonable. Equation 1.9 holds for $\omega < 0$, and the Fourier transform shows some symmetry. The Fourier transform has a number of very interesting properties, but they go beyond the scope of this book.

Fourier transform (Eq. 1.11) with the description of the broadband signal,

$$B(f) = \begin{cases} 1 & \text{for } -f_0 < f < f_0 \\ 0 & \text{otherwise} \end{cases} \quad (1.12)$$

which leads to the following integral where the limits of the bandwidth determine the integration bounds:

$$b(t) = \int_{-f_0}^{f_0} e^{2\pi jft} df = \frac{1}{2\pi jt} \left[e^{2\pi jf_0t} - e^{-2\pi jf_0t} \right]. \quad (1.13)$$

Fortunately, Euler's relationship allows us to simplify the expression in square brackets to $2j \sin(2\pi f_0t)$, and the imaginary unit j cancels out. We therefore obtain our broadband signal as

$$b(t) = \frac{\sin(2\pi f_0t)}{\pi t}. \quad (1.14)$$

For $f_0 = 1/2$, Eq. 1.14 describes the well-known *sinc* function, and it can be shown that the boxcar function (Eq. 1.12) and the sinc-function are a Fourier transform pair, meaning, a square pulse in the time domain has a sinc-like spectrum, and a sinc-like function has a boxcar-type spectrum.

Since digital signals and digital images are discretely sampled, we need to take a look at the *discrete Fourier transform*. In the one-dimensional case, the signal $s(t)$ exists as a set of N discretely sampled values s_k , obtained at $t = k\Delta t$. Here, Δt is the sampling period. In the discrete world, the integral corresponds to a summation, and the discrete Fourier transform becomes

$$S_u = \mathcal{F}\{s_k\} = \frac{1}{N} \sum_{k=0}^{N-1} s_k \exp\left(-2\pi j \frac{u \cdot k}{N}\right) \quad (1.15)$$

where u is the discrete frequency variable, and the sum needs to be evaluated for $0 \leq u \leq N/2$. Equation 1.15 does not consider the sampling rate, and Δt needs to be known to relate u to any real-world units. Any spectral component S_u has the corresponding frequency f_u ,

$$f_u = \frac{u}{N \cdot \Delta t}. \quad (1.16)$$

Note that Eq. 1.16 is not limited to sampling in time. When Δt is a time interval, f_u has units of frequency (i.e., inverse seconds). However, a signal can be sampled with discrete detectors along a spatial axis (see, for example, Fig. 1.1). In this case, the sampling interval has units of distance, and f_u has units of inverse distance. This is referred to as *spatial frequency*. An example to illustrate spatial frequency is a diffraction grating, which causes interference patterns with a certain spatial distance.

For example, if an interference maximum occurs every 0.2 mm, the corresponding spatial frequency is 5 mm^{-1} (or 5 maxima per mm).

We can see from Eq. 1.15 that choosing $u = N$ yields the same result as $u = 0$. For increasing u , therefore, the spectrum repeats itself. Even more, the symmetry of the complex exponential in Eq. 1.15 provides us with $S_{-u}^* = S_u$, where S^* indicates the conjugate-complex of S . For this reason, we gain no new information from computing the discrete Fourier transform for $u > N/2$. By looking at Eq. 1.16, we can see that the frequency at $u = N/2$ is exactly one half of the sampling frequency. This is the maximum frequency that can be unambiguously reconstructed in a discretely-sampled signal (known as the Shannon sampling theorem). The frequency $f_N = 1/2\Delta t$ is known as the *Nyquist frequency*. In the context of Fig. 1.1, we briefly touched on the loss of detail smaller than the sensor area. Here, we have approached the same phenomenon from the mathematical perspective. The situation becomes worse if the signal actually contains frequency components higher than the Nyquist frequency, because those spectral components are *reflected* into the frequency band below f_N , a phenomenon known as *aliasing*. Aliasing is not limited to signals sampled in time. Spatial discretization of components with a higher spatial frequency than $1/2\Delta x$ leads to Moiré patterns.

For images (i.e., discretely-sampled functions in two dimensions), the Fourier transform can be extended into two dimensions as well. Because of the linearity of the Fourier transform, we can perform the row-by-row Fourier transform in one dimension and subject the result to a column-by-column Fourier transform in the orthogonal dimension:

$$S(u, v) = \mathcal{F}\{s(x, y)\} = \frac{1}{MN} \sum_{x=0}^{N-1} \sum_{y=0}^{M-1} s(x, y) \exp\left(-2\pi j \left[\frac{ux}{N} + \frac{vy}{M}\right]\right). \quad (1.17)$$

The Fourier transform now has two orthogonal frequency axes, u and v . The inverse Fourier transform is

$$s(x, y) = \mathcal{F}^{-1}\{S(u, v)\} = \sum_{u=0}^{N-1} \sum_{v=0}^{M-1} S(u, v) \exp\left(2\pi j \left[\frac{ux}{N} + \frac{vy}{M}\right]\right). \quad (1.18)$$

The two-dimensional inverse Fourier transform finds its application in the reconstruction process in computed tomography and magnetic resonance. In both cases, one-dimensional Fourier-encoded information is gathered and used to fill a 2D Fourier-domain spaceholder. Once the spaceholder is completely filled, the inverse Fourier transform yields the cross-sectional reconstructed image.