University of Waterloo

Department of Applied Mathematics

AMATH 391: FROM FOURIER TO WAVELETS

Fall 2021

Lecture Notes

E.R. Vrscay

Department of Applied Mathematics

Lecture 1

Introduction: An overview of the course

The first sentence in the calendar description of this course is,

> "An introduction to contemporary mathematical concepts in signal analysis."

In other words, this course deals with "signals." (It will also look at images, which may be considered as two- or higher-dimensional signals. More on this later.) Of course, we're all familiar with signals: audio signals, seismic signals, electrocardiograms, electroencephalograms. Mathematically, it seems reasonable to represent signals as functions $f(t)$ of a continuous time variable t, viz.,

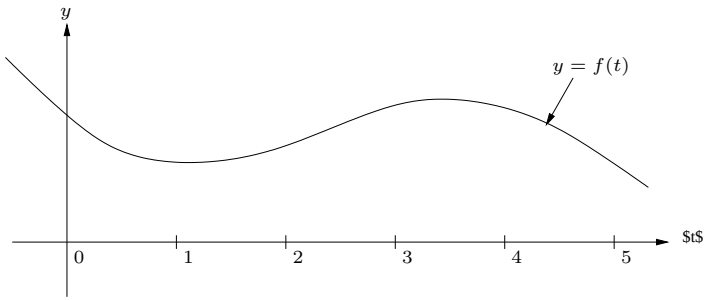

A generic signal $f(t)$

In some cases, we may wish to let the independent variable be denoted by x instead of t – it doesn't matter. And then there is the question of the **domain of definition** of the function $f(t)$. It could be a bounded set, e.g., the interval $[a, b]$, or it could be the entire real line $\mathbb{R}$. These are details that can be left for later.

The above figure represents a generic **analog** signal, where the function $f(t)$ is defined for **continuous** values of t. In today's digital age, signals are generally **discrete**. For example, they have the form of **time series**, which are obtained by measuring a particular property (e.g., temperature, voltage, etc.) at particular times $t_1, t_2, \cdots$. The result is a discrete signal, $f(t_n)$, $n = 1, 2, \cdots$.

Often, for convenience, the measurements are made at equal time steps, i.e., $t_n = n\Delta t$, $n = 1, 2, \cdots$, with $\Delta t > 0$, and we simply let "$f(t_n)$" be represented by the notation $f(n)$, $n = 1, 2, \cdots$. Such a situation is sketched below, with $\Delta t = 1$. This is an example of **sampling**: The particular process that is described by a function $f(t)$ that is sampled, i.e., measured, at discrete time values t_n. Later in this course, we shall study the famous **Shannon Sampling Theorem** which was an important

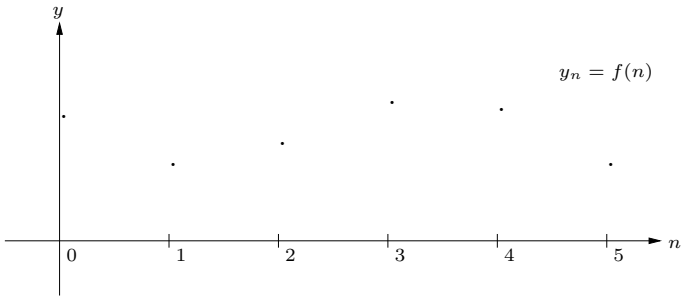

Discrete signal $f(n)$ obtained from measurements of $f(t)$

step towards today's digital age.

It is possible to obtain a discrete series $g(n)$ from the continuous signal $f(t)$ by some other kind of **digitization** process, for example, by letting $g(n)$ be the **mean value** of $f(t)$ over the interval $[t_{n-1}, t_n]$. This is a standard procedure in today's world of digital signals and images and will form an important component of this course.

Furthermore, we may use this discrete series $g(n)$ to define a new function $f_0(t)$ of the continuous variable t which is a piecewise-constant approximation to $f(t)$. We let

$$f_0(t) = g(n), \qquad t_{n-1} \leq t < t_n. \tag{1}$$

This situation is sketched schematically below.

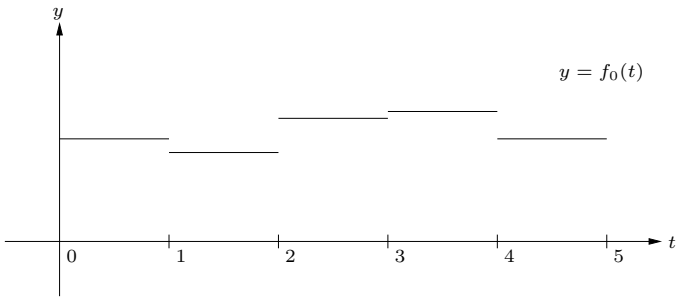

Piecewise-constant approximation $f_0(t)$ to the signal $f(t)$ over unit intervals.

The function $f_0(t)$ may be considered as the approximation to $f(t)$ obtained by viewing it at the resolution of unit time intervals. If we **increased** the resolution by employing the mean values of $f(t)$ over **half-intervals**, the result is a function $f_1(t)$ which has the following form:

Visually, it appears that $f_1(t)$ yields a better approximation to $f(t)$ than $f_0(t)$ does. And this is

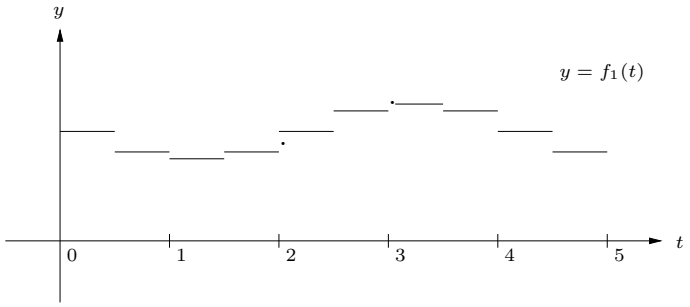

Piecewise-constant approximation $f_1(t)$ to the signal $f(t)$ over half-unit intervals.

certainly the case.

Let us now ask the following question: Given that $f_0(t)$ is a lower-resolution approximation of $f(t)$, what do we have to add to it in order to obtain the higher-resolution approximation $f_1(t)$? (This is the idea behind **progressive transmission** of signals images: Sometimes, when you download an image, you obtain a lower-resolution approximation to it - a very blurred image - which is then updated to provide better, higher-resolution approximations.)

The answer to the above question is as follows: We add the function

$$f_d(t) = f_1(t) - f_0(t) \tag{2}$$

to $f_0(t)$ to obtain $f_1(t)$. As a check:

$$f_0(t) + f_d(t) = f_0(t) + [f_1(t) - f_0(t)] = f_1(t)\,. \tag{3}$$

The function $f_d(t)$ is known as the **detail function** associated with $f_0(t)$. The detail function $f_d(t)$ obtained from the resolution functions $f_0(t)$ and $f_1(t)$ is plotted below.

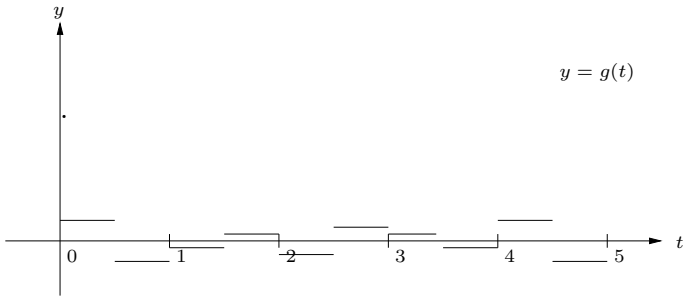

Detail function $f_d(t)$ associated with the functions $f_0(t)$ and $f_1(t)$ shown earlier.

The plot of $f_d(t)$ is certainly interesting. Firstly, it is composed of functions on the unit intervals $(n, n+1)$ which are piecewise constant over half-unit intervals. (This can actually be proved rather easily, but we'll leave it for now.) What is even more interesting is that the two pieces of the graph of $f_d(t)$ over each interval $(n, n+1)$ are symmetrically placed above and below the x-axis. It should not be too difficult to see that each "piece" of the graph of $f_d(t)$ over the interval $(n, n+1)$ can be expressed as an appropriate multiple of an appropriate translation of the following function,

$$\psi(t) = \begin{cases} 1, & 0 \le t < 1/2, \\ -1, & 1/2 \le t < 1. \end{cases} \tag{4}$$

The graph of this function, which is known as the **Haar wavelet function**, is shown below.

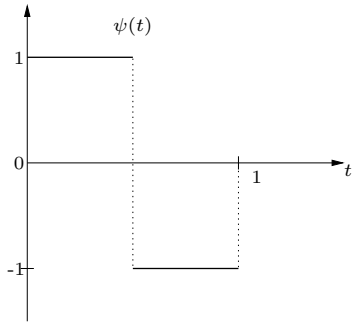

The "Haar wavelet" function $\psi(t)$.

Later, we shall prove that the set of all integer-translations of the function $\psi(t)$, i.e.,

$$\psi_{0k} = \psi(t-k)\,, \quad k \in \mathbb{Z}\,, \tag{5}$$

forms an orthonormal basis for a particular subset of functions on the real line $\mathbb{R}$. And we shall look at higher resolutions of $f(t)$ and show that the above Haar wavelet function can be easily transformed, by means of both **scaling** as well as **translation** to provide basis functions for appropriate higher-resolution spaces.

We conclude this section by showing a more realistic signal, namely, the first (almost) 9 seconds of the famous *Hallelujah Chorus* from Georg Friedrich Handel's *Messiah*. The plot at the left of the figure below is a digital signal composed of $N = 73113$ points. These points were obtained by sampling the continuous audio signal of the Chorus at a frequency of $8192 = 2^{13}$ Hertz (samplings per second). The plot at the right of the figure shows the so-called **power spectrum** of this audio signal – the amplitudes of the complex-valued components of the **discrete Fourier transform** (DFT) of

the signal. The DFT of a discrete signal may be viewed as the discrete analogue of Fourier series of functions of continuous time.

Later in the course, we shall use this audio signal in order to illustrate some aspects of digital signal processing, for example, denoising.

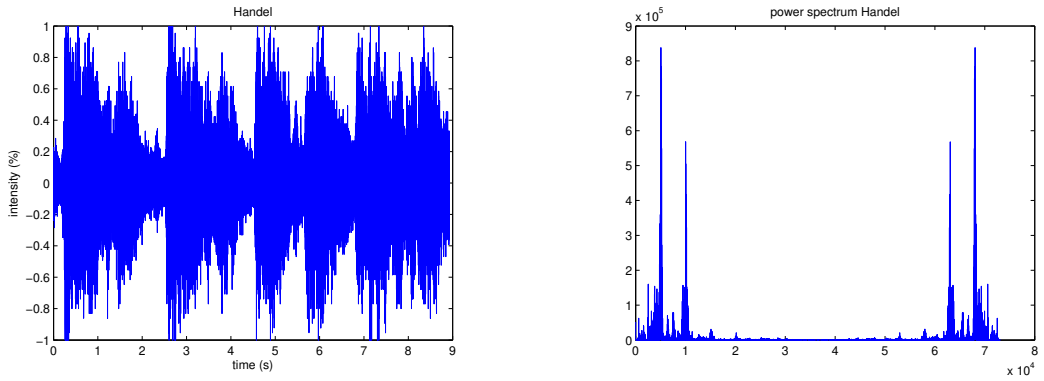

Left: The digital signal representing the first 9 seconds of the *Hallelujah chorus* of Handel's *Messiah*. **Right:** The power spectrum of this digital signal, composed of the amplitudes of the complex-valued components of its discrete Fourier transform (DFT).

Images as functions

Images – at least the ones with which you are familiar – may be considered to be two-dimensional signals. An "ideal image," as approximated by a photograph, may be represented mathematically by a function of two spatial variables, i.e., $f(x, y)$, where x and y are continuous variables over a finite set $D \subset \mathbb{R}^2$, the domain of the function, i.e., $(x, y) \in D$. For a **black-and-white** image, $f(x, y)$ assumes a real and typically non-negative value – the so-called **greyscale value** – that characterizes the "greyness" at a point (x, y) of the image. Mathematically, f is a real-valued function, i.e., $f : \mathbb{R}^2 \to \mathbb{R}$.

In many mathematical treatments, it is assumed, for simplicity that the domain of f is given by $x, y \in [0, 1]$, which we may also write as $(x, y) \in [0, 1]^2$. As well, assume that the range of f is the interval $[0, 1]$, i.e., $f : [0, 1]^2 \to [0, 1]$. Then the value 0 will represent black and the value 1 will represent white. An intermediate value, i.e., $0 < f(x, y) < 1$, will represent some shaded grey value. The graph of the image function, $z = f(x, y)$, then may be viewed as a representation of the image, as shown in the figure below on the right. Note that the colours in this figure represent function values, not actual colour values since the original image is black-and-white.

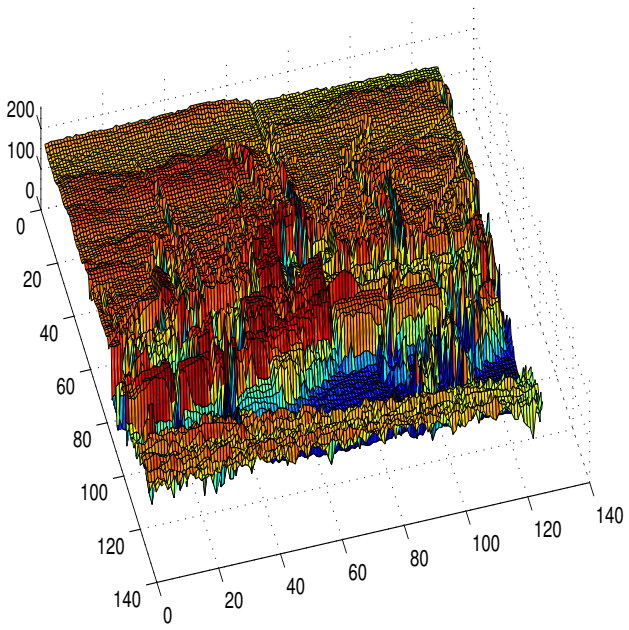

Left: The standard test-image, *Boat*, a 512×512-pixel digital image, 8 bits per pixel. Each pixel assumes one of 256 greyscale values between 0 and 255. **Right:** The *Boat* image, viewed as an image function $z = f(x, y)$. The red-blue spectrum of colours is used to characterize function values: Higher values are more red, lower values are more blue.

An ideal **colour image** is represented mathematically by a **vector-valued function**. At each point $(x, y) \in [0, 1]$, are defined three colour values, namely, **red**, **green** and **blue**. The combination of these three **primary colours** produces the colour associated with (x, y). Mathematically, f is a mapping from $\mathbb{R}^2$ to $\mathbb{R}^3$ having the form,

$$f(x, y) = (r(x, y), g(x, y), b(x, y)), \tag{6}$$

where $r(x, y)$, $g(x, y)$ and $b(x, y)$ denote, respectively, the red, green and blue values at position (x, y). We'll examine colour images in the next section.

Digital images

Digital images are two-dimensional arrays that represent samplings of the image function $f(x, y)$.

Black-and-white digital images are represented by $n_1 \times n_2$ matrices, $\mathbf{u} = \{u_{ij}\}$. (As the caption indicates, the *Boat* image of the previous figure is a black-and-white digital image.) The entry u_{ij} of this matrix – usually written as $u[i, j]$ in image processing literature – represents the greyscale values at the (i, j) pixel, $1 \leq i \leq n_1$, $1 \leq j \leq n_2$. The greyscale values of digital images also assume discrete values so that they may easily be stored in digital memory. The standard practice is to allocate n bits of digital memory for each greyscale value so that a total of 2^n values, namely, $\{0, 1, 2, \cdots, 2^n - 1\}$ are possible. In most applications, $n = 8$, i.e., 8-bit images, implying a set of 256 greyscale values ranging from 0 to 255. This is found to be more than sufficient for the human visual system.

Colour digital images will be represented by three matrices, $\mathbf{r} = \{r[i, j]\}$, $\mathbf{g} = \{g[i, j]\}$ and $\mathbf{b} = \{b[i, j]\}$, which represent, respectively, the red, green and blue values at the (i, j) pixel. As you may recall from an earlier Physics or Science courses, any colour can be generated by means of an appropriate combination of these three **primary colours**.

In the top left of the figure below is presented the digital colour image, *Sailboat-lake*, a 512×512-pixel image. At each pixel, 8 bits, i.e., 256 values ranging from 0 to 255, are used to store the intensity of each of the three colours – red, green and blue. This means that a total of 24 bits of memory are used for each pixel. We say that the *Sailboat-lake* image is a 512×512 image, 24 bits per pixel.

The red, green and blue component images are shown in the same figure. Note that they are displayed as black-and-white images, with 0 = black and 255 = white.

A closer look at the component images will show why particular regions of particular primary colours are either low or high in magnitude. One would expect that the blue components for pixels representing the blue sky in the colour image would have a higher intensity than red and green components. This, in turn, would imply that the black-and-white image representing the blue component would be lighter/whiter in the blue sky regions, which is observed to be the case. One can also draw similar conclusions for the red components in reddish regions of the colour image as well as green components in greener regions of the colour image.

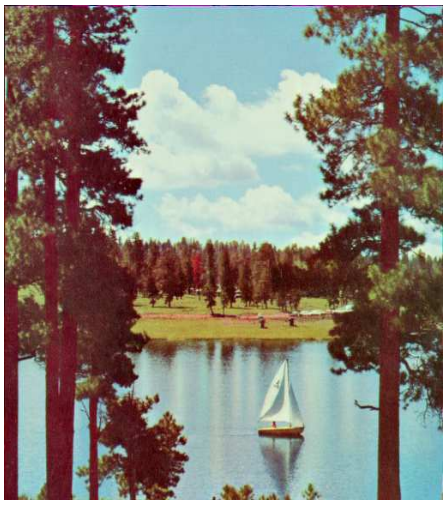

RGB colour image

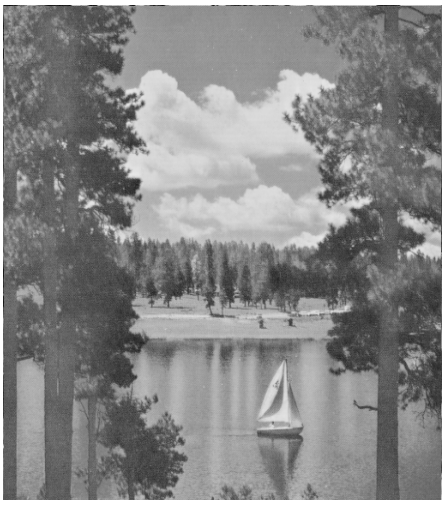

Red component

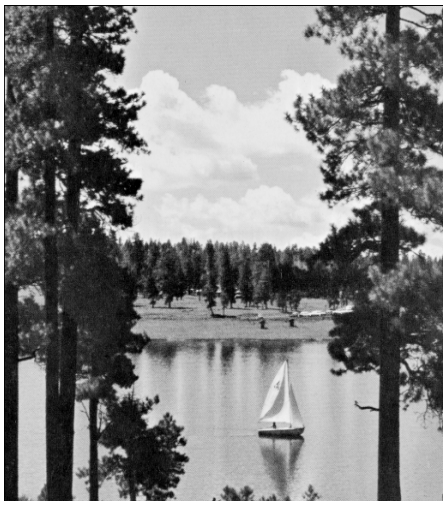

Green component

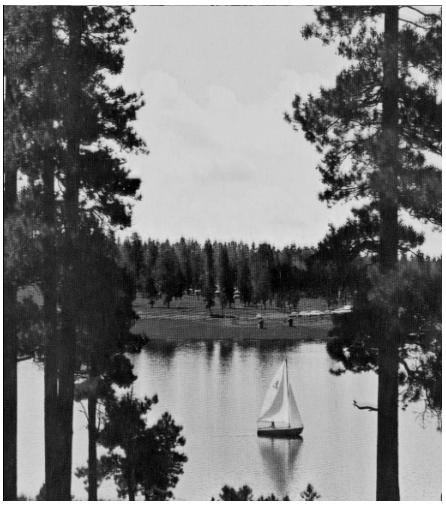

Blue component

The standard digital colour test-image, *Sailboat-lake*, 512 × 512-pixel, 24 bits per pixel (8 bits per colour), along with its red, green and blue component images.

Hyperspectral images

The red, green and blue values of an image at a point/pixel may be viewed as the reflectance values for a particular **sampling** of the visual electromagnetic spectrum at three wavelengths, $\lambda_r > \lambda_g > \lambda_b$. As mentioned earlier, three colours are sufficient since any visible colour may be generated from a combination of these three primary colours.

That being said, in other applications, e.g., remote sensing, a much greater sampling of the electromagnetic spectrum is performed. For example, the Airborne Visible/Infrared Integrating Spectrometer (AVIRIS), performs a sampling of 224 wavelengths in the visible and infrared regions of the electromagnetic spectrum. Such high sampling is performed in order to determine the composition or nature of regions being photographed. Various soils, for example, based upon their mineral composition exhibit different reflectance characteristics which will be manifested in the "shape" of their 224-vectors. The same may be said for different types of vegetation, etc.. Some years back, when satellite imagery was in its infancy, and a much lower degree of sampling was performed, such images were known as **multispectral images**. Now, with much greater degrees of sampling, these images are known as **hyperspectral images**. Once again, a hyperspectral image may be viewed as a vector-valued function: At each pixel location $[i, j]$, which represents a particular region of the earth's surface, the hyperspectral image function f is an M-vector, i.e.,

$$f[i, j] = (f_1[i, j], f_2[i, j], \cdots, f_M[i, j]) \,. \tag{7}$$

This M-vector defines the **spectral function** of the region represented by the pixel $[i, j]$. As mentioned earlier, spectral functions can contain a great deal of information about the chemical composition of regions.

The component functions $f_k(x, y)$ corresponding to different wavelengths are often referred to as **channels**. Each channel represents an image of the particular region on the earth taken at a particular wavelength. A pictorial representation of the "stacking up" of many channels to form a hyperspectral **"data cube"** is shown in the next figure.

Some basic aspects of **spectral functions** are discussed and presented pictorially in the figure after the next.

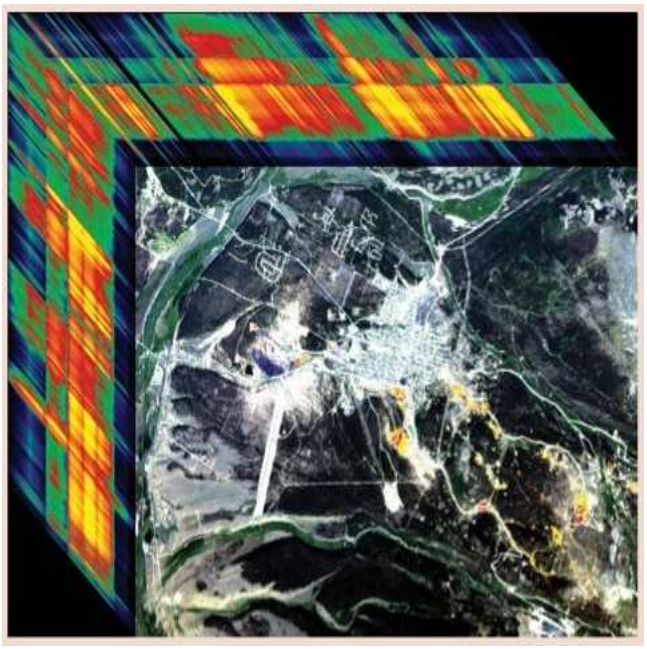

A pictorial representation of the "stacking up" of images – or channels – corresponding to different wavelengths to form a hyperspectral image.

Another type of hyperspectral imaging: Diffusion magnetic resonance imaging (MRI)

You have most probably heard of **magnetic resonance imaging** (MRI), which is based on the so-called **magnetic moment** of the hydrogen atom nucleus, namely, the proton. Very briefly, a proton will interact with an external magnetic field due to its intrinsic **magnetic moment**. As you know, water, "H_2O", which is composed of hydrogen and oxygen atoms, is present virtually everywhere in living tissue. Different regions in a human body, e.g., different organs, tissues, etc., represent different structural and biochemical environments for the water molecules within those regions. As such, protons in water molecules from different regions will respond differently to an external magnetic field. (A little more precisely – protons from different regions will have different rates of **spin relaxation** in the presence of a constant external magnetic field.) In a very clever manner, magnetic resonance imaging uses these differences in response (relaxation) to produce a two- or three-dimensional pictorial representation of the interior of a human body (or whatever is being imaged). In fact, the underlying basis of this procedure is the **Fourier transform** which will be discussed later in this course.

The more recently developed technique of **diffusion magnetic resonance imaging** (DMRI)

The nature of spectral functions

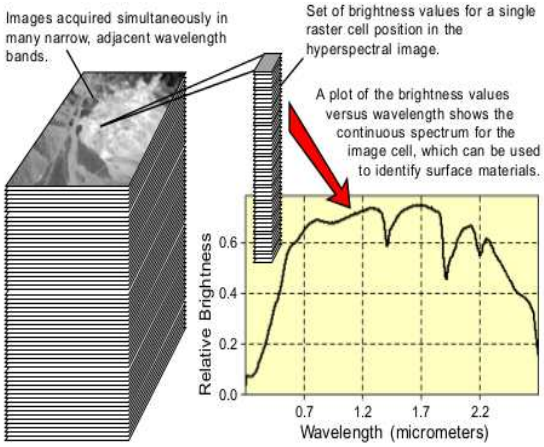

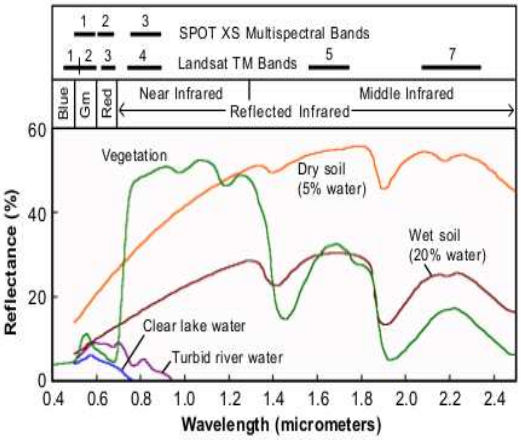

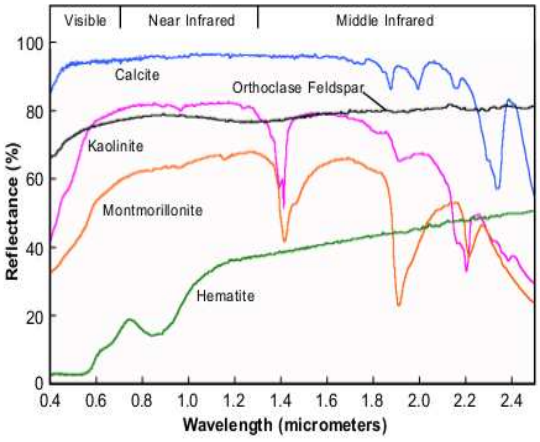

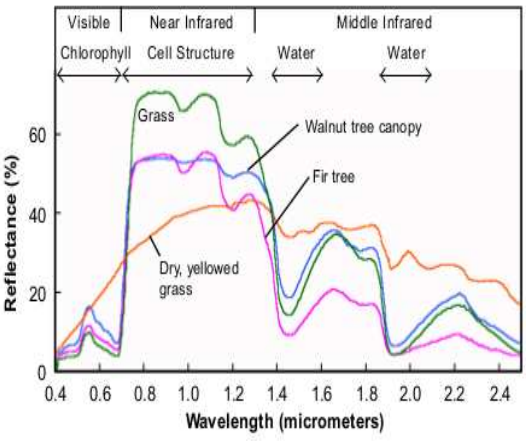

also exploits the magnetic moment of protons. DMRI is able to detect the motion of collections of water molecules in local regions of the body under observation. Realistically, because of limitations in resolution, the characterization of the motion is limited to a finite number of directions. The net result of this procedure is that at a 3D pixel location (i, j, k) in the interior of the body being observed, one can estimate the probability that water molecules at (i, j, k) will move (diffuse) in each of M directions. Once again, the result is a vector-valued image function,

$$u[i, j, k] = (u_1[i, j, k], u_2[i, j, k], \cdots, u_M[i, j, k]) \,. \tag{8}$$

This is illustrated in the figure below.

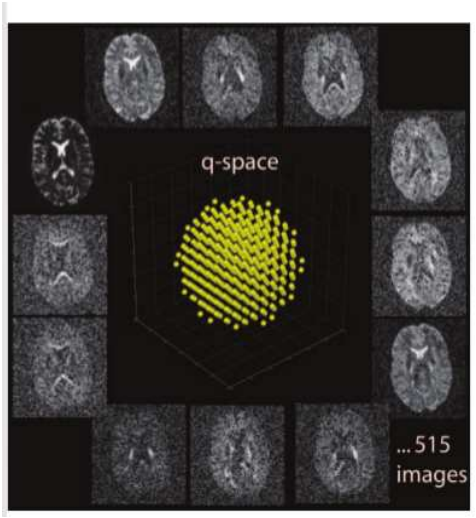

Figure 12. Series of diffusion-weighted MR brain images obtained with variations in the direction and strength of the diffusion gradient in the pulsed gradient SE sequence. Each image shows the signal sampled at one point in q-space (one yellow dot). Every sampling point in q-space corresponds to a specific direction and strength of the diffusion gradient.

From *Understanding Diffusion MR Imaging Techniques: From Scalar Diffusion-weighted imaging to Diffusion Tensor Imaging and Beyond*, by P. Hagmann *et al.*, Radiographics 2006, 26:S205-S223. Published online `10.1148/rg.26si065510`.

One fascinating application of DMRI is in neurobiology – the ability to produce maps of neural connections in the brain. It seems natural that there will be a greater probability for water molecules inside a neuron to travel in the tubular direction of the neuron as opposed to through its boundaries. Using this information, and a kind of "connect the arrows", connectivity maps such as the one in the figure below can be obtained. Such connectivity maps are called **connectomes**.

Signal and image processing

Signal processing and **image processing** are the terms used to describe procedures that are generally designed to achieve specific tasks, for example: (i) to "improve" signals and images (deblurring,

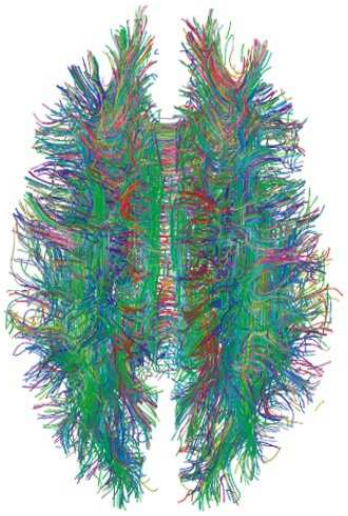

A typical "connectome," a pictorial representation of the connectivity of neurons in the human brain.

denoising or both) and (ii) to "compress" them, i.e., to reduce the amount of computer memory required to store them. Associated with most, if not all, signal and image processing procedures are underlying mathematical principles that account for their efficacy. One of the goals of this course is to examine some of these mathematical principles.

A fundamental procedure in signal/image processing is the **digitization** of signals and images. A signal $f(t)$ in continuous time may be stored on an audio tape. However, if it is digitized into a discrete series $f(n)$, it can be stored on a computer hard drive any other digital device (e.g., CD, DVD, USB stick) at a fraction of the storage requirement. In addition, such digital data can be more easily accessed. We'll discuss this later in the course.

The idea of "distances" between signals/images

One of the most fundamental mathematical concepts that underlie signal and image analysis is that of "distance." Given two signal functions, $f(t)$ and $g(t)$, what is the "distance" between them – or, put another way, how "close" are they?

For example, suppose that we have a "pure audio signal" $f(t)$ that we wish to store in the form of a digital signal, i.e., a discrete series $g(n)$. It is possible to digitize the signal with (almost) zero loss of fidelity, but the storage requirements are huge. As such, one tries to **compress** the digital signal by removing redundant information. There is a trade-off, however – the greater the degree of compression, the greater the loss of information, implying that the error in approximating the original

signal $f(t)$ with the discrete series $g(n)$ – in other words, the "distance" between f and g – is greater.

As another case, suppose that we once again have a "pure audio signal" $f(t)$ recorded in a sound studio. We transmit this signal over some "channel", e.g., a cable, only to find that the signal recorded at some other location, to be denoted as $\tilde{f}(t)$, is **degraded**, for example, with the presence of noise. The distance between f and $\tilde{f}$ can be used to characterize the degree of degradation of the signal by the transmission process.

Now suppose that the observer who records the degraded signal $\tilde{f}$ attempts to restore the original signal from it, for example, by applying a standard denoising algorithm "D". The result of applying algorithm D to the noisy signal $\tilde{f}$ is a new signal g. In the case of a perfect denoising algorithm, which is never achieved, $g = f$, i.e., the distance between the original signal f and the denoised signal g is zero. In practice, f is never retrieved, which means that the distance between g and f is nonzero. Of course, the smaller the distance between f and g, the better the denoiser D.

In both applications described above, i.e., compression and denoising, we stated that the distance between two functions can be used to characterize the effectiveness of the procedure. The smaller the distance, the better the procedure. But there remains the question, "What distance do we use?" This is an extremely important question, and one which is still being researched. Different procedures may require different distance functions. That being said, the most commonly used distance function is the so-called L^2 or **Euclidean** distance – analogous to the distance function employed in your course on Fourier series and which will be reviewed very shortly. We shall discuss this matter later in the course.

Lecture 2

A brief review of Fourier series and some important concepts from analysis

In AMATH 231 (Vector Calculus and Fourier Series) and possibly other courses, e.g., AMATH 353 (Partial Differential Equations I), you saw the following result: If $f(x)$ is a real-valued Riemann integrable function defined on $[-\pi, \pi]$, then we can express it as follows,

$$f(x) = a_0 + \sum_{n=1}^{\infty} (a_n \cos nx + b_n \sin nx). \tag{9}$$

The expression on the right-hand side is known as the "Fourier series expansion" of f on $[-\pi, \pi]$. The coefficients of the Fourier expansion are given by

$$\begin{aligned} a_0 &= \frac{1}{2\pi} \int_{-\pi}^{\pi} f(x)\, dx, \\ a_n &= \frac{1}{\pi} \int_{-\pi}^{\pi} f(x) \cos nx\, dx, \qquad n = 1, 2, \cdots, \\ b_n &= \frac{1}{\pi} \int_{-\pi}^{\pi} f(x) \sin nx\, dx, \qquad n = 1, 2, \cdots. \end{aligned} \tag{10}$$

These coefficients are oftern referred to as the **Fourier coefficients** of the function $f(x)$.

> **Note:** In the AMATH 231 notes, there is a factor of $1/2$ multiplying the coefficient a_0 in Eq. (9) which, in turn, implies that the factor 2π must be replaced by π. This is a standard definition that is employed in the literature, motivated by the fact that the expression for a_0 does not sit apart from the expressions for the other a_n. In this course, however, we shall adopt the notation used above, since it is used by many books, both in mathematics and signal and image processing, including the textbook by Boggess and Narcowich used for this course. In some sense, it is the more "natural" definition for reasons that we shall see below.

These formulas were obtained – as they are in many standard textbooks – by exploiting several integrals relations involving sine and cosine functions. The first and simplest relations are the following,

$$\begin{aligned} \int_{-\pi}^{\pi} \sin mx\, dx &= 0\,, \qquad m = 1, 2, \cdots. \\ \int_{-\pi}^{\pi} \cos mx\, dx &= \begin{cases} 0\,, & m = 1, 2, \cdots, \\ 2\pi\,, & m = 0\,. \end{cases} \end{aligned} \tag{11}$$

From the above, and the following identities,

$$\begin{aligned}
\sin A\,\cos B &= \frac{1}{2}\left[\sin(A-B)+\sin(A+B)\right] \\
\sin A\,\sin B &= \frac{1}{2}\left[\cos(A-B)-\cos(A+B)\right] \\
\cos A\,\cos B &= \frac{1}{2}\left[\cos(A-B)+\cos(A+B)\right]\,, \qquad (12)
\end{aligned}$$

the following additional relations can be derived: For m and n positive integers,

$$\int_{-\pi}^{\pi} \sin mx \cos nx\, dx = 0\,, \qquad (13)$$

$$\int_{-\pi}^{\pi} \sin mx \sin nx\, dx = \begin{cases} 0 & \text{if } m \neq n\,, \\ \pi & \text{if } m = n\,, \end{cases} \qquad (14)$$

and

$$\int_{-\pi}^{\pi} \cos mx \cos nx\, dx = \begin{cases} 0 & \text{if } m \neq n\,, \\ \pi & \text{if } m = n\,. \end{cases} \qquad (15)$$

Note: It is highly recommended that you derive these results at least once during your academic career.

Now perform the following operations on Eq. (9):

1. Considering each side of the equation as a function of $x \in [-\pi, \pi]$, integrate both sides w.r.t. x from $-\pi$ to π:
$$\int_{-\pi}^{\pi} f(x)\, dx = \int_{-\pi}^{\pi} \left[a_0 + \sum_{n=1}^{\infty} (a_n \cos nx + b_n \sin nx) \right] dx\,. \qquad (16)$$

2. Assume, for the moment, that the limiting operations of integration and infinite summation can be interchanged and that the two terms inside the summation can be separated, i.e.,
$$\begin{aligned}
\int_{-\pi}^{\pi} f(x)\, dx &= \int_{-\pi}^{\pi} a_0\, dx + \sum_{n=1}^{\infty} \int_{-\pi}^{\pi} a_n \cos nx\, dx + \sum_{n=1}^{\infty} \int_{-\pi}^{\pi} b_n \sin nx\, dx \\
&= a_0 \int_{-\pi}^{\pi} dx + \sum_{n=1}^{\infty} a_n \int_{-\pi}^{\pi} \cos nx\, dx + \sum_{n=1}^{\infty} b_n \int_{-\pi}^{\pi} \sin nx\, dx\,. \qquad (17)
\end{aligned}$$

(Note: To do this properly or "rigorously," as you may have seen in a course in Analysis, one must consider partial sums of the infinite summations in the above equation. In this case, the interchanging of integration and summation is permitted - the integral of a finite sum is equal to the finite sum of integrals. The result is a partial sum S_N which is a real number. It then

remains to show that as $N \to \infty$, the partial sum S_N converges to the real number on the left side of Eq. (17).)

3. The first integral on the RHS of (17) is 2π. From Eq. (11) all other integrals on the RHS of (17) vanish. As a result, we obtain the first expression in (10).

In order to obtain the second expression in (10), select an integer $p \geq 1$ and perform the following operations on Eq. (9) for this fixed value of p:

1. Multiply both sides of Eq. (9) by $\cos px$ and integrate the result from $-\pi$ to $-\pi$:

$$\int_{-\pi}^{\pi} f(x) \cos px \, dx = \int_{-\pi}^{\pi} \left[a_0 \cos px + \sum_{n=1}^{\infty} (a_n \cos nx \cos px + b_n \sin nx \cos px) \right] dx \, . \tag{18}$$

2. Once again assume that the limiting operations of integration and infinite summation can be interchanged and that the two terms inside the summation can be separated, i.e.,

$$\begin{aligned} & \int_{-\pi}^{\pi} f(x) \cos px \, dx \\ = & \int_{-\pi}^{\pi} a_0 \cos px \, dx + \sum_{n=1}^{\infty} \int_{-\pi}^{\pi} a_n \cos nx \cos px \, dx + \sum_{n=1}^{\infty} \int_{-\pi}^{\pi} b_n \sin nx \cos px \, dx \\ = & \; a_0 \int_{-\pi}^{\pi} \cos px \, dx + \sum_{n=1}^{\infty} a_n \int_{-\pi}^{\pi} \cos nx \cos px \, dx + \sum_{n=1}^{\infty} b_n \int_{-\pi}^{\pi} \sin nx \cos px \, dx \, . \end{aligned} \tag{19}$$

3. From Eq. (11), the first integral on the RHS vanishes. From Eq. (13), all of the integrals in the final summation vanish. From Eq. (15), only the integral for which $n = p$ does not vanish – its value is π. We therefore obtain the result,

$$a_p = \frac{1}{\pi} \int_{-\pi}^{\pi} f(x) \cos px \, dx \, , \tag{20}$$

which, up to the label of the index, the second result in Eq. (10).

The third result in Eq. (10) can be obtained by replacing $\cos px$ with $\sin px$, $p \geq 1$, in the above series of steps.

Appendix: Some linear algebra revisited

The above series of operations should bring back some (pleasant?) memories from linear algebra. Let us suppose that $\mathbf{u}_1, \mathbf{u}_2, \cdots, \mathbf{u}_n$ are n-vectors and that they form an **orthogonal set** in the vector space $\mathbb{R}^n$. This means that

$$\langle \mathbf{u}_i, \mathbf{u}_j \rangle = 0 \quad \text{if} \quad i \neq j \,, \tag{21}$$

where $\langle \cdot, \cdot \rangle$ denotes the usual inner product of n-vectors in $\mathbb{R}^n$.

Wait! Let's review this by writing the vectors $\mathbf{u}_i$ and $\mathbf{u}_j$ in terms of their components, i.e.,

$$\mathbf{u}_i = (u_{i1}, u_{i2}, \cdots, u_{in}) \qquad \mathbf{u}_j = (u_{j1}, u_{j2}, \cdots, u_{jn}) \,. \tag{22}$$

Then the inner product in (21) is defined as

$$\langle \mathbf{u}_i, \mathbf{u}_j \rangle = \sum_{k=1}^{n} u_{ik} u_{il} \,. \tag{23}$$

Eq. (21) implies that the set of vectors $\{\mathbf{u}_k\}_{k=1}^n$ forms a **basis** in $\mathbb{R}^n$ – in fact, an **orthogonal basis** in $\mathbb{R}^n$: Any $\mathbf{v} \in \mathbb{R}^n$ may be expressed as a unique linear combination of the $\mathbf{u}_k$, i.e.,

$$\begin{aligned} \mathbf{v} &= c_1 \mathbf{u}_1 + c_2 \mathbf{u}_2 + \cdots c_n \mathbf{u}_n \\ &= \sum_{k=1}^{n} c_k \mathbf{u}_k \,, \end{aligned} \tag{24}$$

for a unique set of coefficients $\{c_k\}_{k=1}^n$.

Given a $\mathbf{v} \in \mathbb{R}^n$, recall how we can find the expansion coefficients c_k. For each $i = 1, 2, \cdots, N$, we take the inner product of both sides of (24) with the vector $\mathbf{u}_i$:

$$\begin{aligned} \langle \mathbf{v}, \mathbf{u}_i \rangle &= \left\langle \sum_{k=1}^{n} c_k \mathbf{u}_k, \mathbf{u}_i \right\rangle \\ &= \sum_{k=1}^{n} \langle c_k \mathbf{u}_k, \mathbf{u}_i \rangle \\ &= \sum_{k=1}^{n} c_k \langle \mathbf{u}_k, \mathbf{u}_i \rangle \\ &= c_i \langle \mathbf{u}_i, \mathbf{u}_i \rangle \quad \text{(by orthogonality of the } \mathbf{u}_i\text{)} \\ &= c_i \|\mathbf{u}_i\|^2 \,, \end{aligned} \tag{25}$$

which we arrange to arrive at the well known result,

$$c_i = \frac{1}{\|\mathbf{u}_i\|^2}\langle \mathbf{v}, \mathbf{u}_i \rangle \,. \tag{26}$$

Note: If the basis is not only orthogonal but also **orthonormal**, i.e.,

$$\langle \mathbf{u}_i, \mathbf{u}_i \rangle = 1 \,, \tag{27}$$

then

$$c_i = \langle \mathbf{v}, \mathbf{u}_i \rangle \,. \tag{28}$$

Recall that from any orthogonal basis $\{\mathbf{u}_k\}_{k=1}^n$, we can always construct an orthonormal basis $\{\hat{\mathbf{u}}_k\}_{k=1}^n$ as follows. Define

$$\hat{\mathbf{u}}_k = \frac{1}{\|\mathbf{u}_k\|}\mathbf{u}_k \,. \tag{29}$$

It is easy to see that these vectors are orthogonal to each other since they are simply constant multiples of the $\mathbf{u}_k$. Furthermore,

$$\begin{aligned} \langle \hat{\mathbf{u}}_k, \hat{\mathbf{u}}_k \rangle &= \left\langle \frac{1}{\|\mathbf{u}_k\|}\mathbf{u}_k, \frac{1}{\|\mathbf{u}_k\|}\mathbf{u}_k \right\rangle \\ &= \frac{1}{\|\mathbf{u}_k\|^2}\langle \mathbf{u}_k, \mathbf{u}_k \rangle \\ &= \frac{\|\mathbf{u}_k\|^2}{\|\mathbf{u}_k\|^2} \\ &= 1 \,. \end{aligned} \tag{30}$$

End of Appendix

The procedure that we used to obtain the expressions for the Fourier coefficients in Eq. (10) appears to be quite analogous to the procedure we used to obtain Eq. (26) in the vector space case. In fact, we claim that the following infinite set of functions,

$$\mathcal{S} = \{1\} \bigcup \left\{ \sin nx \right\}_{n=1}^{\infty} \bigcup \left\{ \cos nx \right\}_{n=1}^{\infty} , \tag{31}$$

forms an **orthogonal basis set** in a particular **inner product space of functions** defined over the interval $[-\pi, \pi]$ which we'll denote as $\mathcal{F}[-\pi, \pi]$. The inner product between two functions f and g in this space is defined as follows,

$$\langle f, g \rangle = \int_{-\pi}^{\pi} f(x)g(x)\, dx \,. \tag{32}$$

Two functions f and g in this space are said to be **orthogonal** to each other if

$$\langle f, g\rangle = \int_{-\pi}^{\pi} f(x)g(x)\,dx = 0\,. \tag{33}$$

(Of course, we can generalize the above definition to functions defined over an interval $[a, b]$, and we shall do this later.)

A side note: This might require a little stretching of the brain cells, but the inner product between functions in (32) may be viewed as a continuous version of the inner product between n-vectors in (21). In the case of functions, the integration variable $x \in [a, b]$ is a continuous version of the discrete index $k \in \{1, 2, \cdots, n\}$ in (21). A connection between the two can be made if we consider a Riemann sum approximation to the inner product in (32) using the partition $a = x_0 < x_1 < \cdots < x_{n-1} < x_n = b$ for any $n > 0$. Now take sample points $\bar{x}_k \in I_k = [x_{k-1}, x_k]$ for $1 \le k \le n$ and construct the Riemann sum approximation to the integral in (32). This approximation may be viewed as an inner product between two vectors, the components of which are composed of the values of the functions f and g at the sample points $\bar{x}_k$.

Well that's fine. But the next step is to go from these finite approximations with n-vectors to the limit $n \to \infty$ to arrive at the integral. This is left as an exercise.

We now return to the the integral relations in Eqs. (11), (13), (14) and (15) involving sine and cosine functions that were used to obtain the expressions in Eq. (10) for the Fourier coefficients. From the above discussion on inner products for functions, we can view these integral relations as orthogonality relations. Let's rewrite the first set as follows,

$$\begin{aligned}
\langle \sin mx, 1\rangle = \int_{-\pi}^{\pi} \sin mx\,dx &= 0\,, \qquad m = 1, 2, \cdots\,. \\
\langle \cos mx, 1\rangle = \int_{-\pi}^{\pi} \cos mx\,dx &= \begin{cases} 0\,, & m = 1, 2, \cdots\,, \\ 2\pi\,, & m = 0\,. \end{cases}
\end{aligned} \tag{34}$$

The next set of three integral relations may be rewritten as follows,

$$\langle \sin mx, \cos nx\rangle = \int_{-\pi}^{\pi} \sin mx \cos nx\,dx = 0\,, \tag{35}$$

$$\langle \sin mx, \sin nx\rangle = \int_{-\pi}^{\pi} \sin mx \sin nx\,dx = \begin{cases} 0 & \text{if } m \neq n\,, \\ \pi & \text{if } m = n\,, \end{cases} \tag{36}$$

and

$$\langle \cos mx, \cos nx \rangle = \int_{-\pi}^{\pi} \cos mx \cos nx \, dx = \begin{cases} 0 & \text{if } m \neq n \,, \\ \pi & \text{if } m = n \,. \end{cases} \tag{37}$$

The procedures employed earlier to obtain expressions for the Fourier coefficients a_n and b_n may now be viewed in terms of inner products. Let's start with the first method to obtain a_0. Recall that we simply integrated Eq. (9) from $-\pi$ to π. This is equivalent to taking the inner product of the function $g(x) = 1$ with both sides of the equation, i.e.,

$$\langle f, 1 \rangle = \left\langle \left[a_0 + \sum_{n=1}^{\infty} (a_n \cos nx + b_n \sin nx) \right], 1 \right\rangle \,. \tag{38}$$

We now assume, once again, that the operations of integration (involved in the inner product) and infinite summation can be interchanged, the infinite summation separated into two terms, i.e.,

$$\langle f, 1 \rangle = \langle a_0, 1 \rangle + \sum_{n=1}^{\infty} \langle a_n \cos nx, 1 \rangle + \sum_{n=1}^{\infty} \langle b_n \sin nx, 1 \rangle \,. \tag{39}$$

Now move the constants outside the inner products,

$$\langle f, 1 \rangle = a_0 \langle 1, 1 \rangle + \sum_{n=1}^{\infty} a_n \langle \cos nx, 1 \rangle + \sum_{n=1}^{\infty} b_n \langle \sin nx, 1 \rangle \,. \tag{40}$$

As before, only the first inner product on the RHS of the equation is nonzero – all other inner products vanish. Thus we have

$$\langle f, 1 \rangle = a_0 \langle 1, 1 \rangle \,. \tag{41}$$

Let's now evaluate these inner products:

$$\langle f, 1 \rangle = \int_{-\pi}^{\pi} f(x) \cdot 1 \, dx = \int_{-\pi}^{\pi} f(x) \, dx \tag{42}$$

and

$$\langle 1, 1 \rangle = \int_{-\pi}^{\pi} 1 \cdot 1 \, dx = \int_{-\pi}^{\pi} 1 \, dx = 2\pi \,. \tag{43}$$

Substitution of these results into (41) yields

$$\int_{-\pi}^{\pi} f(x) \, dx = 2\pi a_0 \,, \tag{44}$$

which produces the expression for a_0 in (10).

The expressions for a_n and b_n in Eq. (10), where $n \geq 1$, can be obtained in the same way. The expression for a_n is obtained by taking the inner product of both sides of Eq. (9) with $\cos px$ for any

integer $p \geq 1$. (We use p instead of n because n is used as a summation index in Eq. (9)). The result is (Exercise),

$$\begin{aligned} \langle f, \cos px \rangle &= a_p \langle \cos px, \cos px \rangle \\ &= a_p \pi \,, \end{aligned} \tag{45}$$

which yields the equation for the a_n in Eq. (10). Likewise, the expression for b_n is obtained by taking the inner product of Eq. (9) with $\sin px$ for an integer $p \geq 1$ and then changing p to n.

Once again, we mention that the formulas for the Fourier coefficients are obtained from the fact that the functions 1, $\cos nx$ and $\sin nx$ for $n \geq 1$ form an orthnormal set in a particular space of functions defined over the interval $[-\pi, \pi]$. There still is one technical point, namely, the validity of the **infinite series** expansion. (Recall that there was no question about the validity in finite dimensions.) We'll have to return to this point later in the course and discuss the fact that the orthogoonal set of functions, $\left\{1, \sin nx, \cos nx\right\}_{n=1}^{\infty}$, forms a **complete set** – in other words, a **basis** - in the infinite-dimensional space of functions.

Let us now return to the Fourier series in Eq. (9). Clearly, it is an **infinite series**. Just as was done for power series in first-year Calculus, we have to make sense of this expression in terms of the convergence of **partial sums** of the series. In this case, however, the partial sums of the Fourier series are **functions**.

Suppose that we have a set of functions $\{u_n(x)\}_{n=1}^{\infty}$ that form a basis over an interval $[a, b]$ on the real line. (We'll eventually have to address the term "form a basis.") And suppose that, given some function $f(x)$ defined on $[a, b]$ we can write the expression,

$$\text{“}f(x) = \sum_{n=1}^{\infty} a_n u_n(x).\text{”} \tag{46}$$

Mathematically, this means that the sequence of partial sums $\{S_N\}$ of the infinite series, i.e.,

$$S_N(x) = \sum_{n=1}^{N} a_n u_n(x), \tag{47}$$

which are **functions** on $[a, b]$, is converging to the function f.

Technically speaking, Eq. (46) gives the impression that the equality is true for all $x \in [a, b]$. As you may have seen in a course in Fourier series, e.g., AMATH 231, this equality may not hold at each point $x \in [a, b]$. As such, Eq. (46) should really be written as follows,

$$f = \sum_{n=1}^{\infty} a_n u_n \,. \tag{48}$$

For any finite $N \geq 1$, the equality in Eq. (47) may hold true for all $x \in [a, b]$ since the sum is finite. But the convergence may not hold pointwise – once again, as you may have seen in AMATH 231 or equivalent. As such, we must write,

$$\lim_{N \to \infty} S_N = f \,, \tag{49}$$

where the limit is understood in terms of an appropriate distance function. More on this later.

In the figure below are shown a couple of examples of such convergence in the particular case of Fourier series. Two functions are considered, and the partial sum approximations, $S_3(x)$ and $S_{25}(x)$ are presented for each case.

For each value of N, it appears that the partial sum for the continuous function $f(x) = \frac{1}{2}(\pi - |x|)$ is providing a "better approximation" to $f(x)$ than its counterpart for the piecewise continuous "square wave" function. This will probably bring back memories from AMATH 231 (or from whatever course you saw Fourier series). We'll return to explore Fourier series in more detail very shortly.

At this point, we could pursue this problem in the way that we did for power series in first-year Calculus, asking the question: For which $x \in [a, b]$ does the sequence of values $S_N(x)$ converge to $f(x)$? (Remember that this led to the idea of the **interval of convergence** of the power series.) This leads to the idea of **pointwise convergence** of the Fourier series expansion in Eq. (9). This is an important concept which was covered, to some extent in AMATH 231, and we shall return to it. Here, however, we wish to look at this problem from the following viewpoint: The partial sums $S_N(x)$ are **functions** that will serve as **approximations** to the function $f(x)$ over the interval $[a, b]$ – approximations that "converge" to $f(x)$ over the interval. Therefore, we wish to express the convergence of the functions as follows:

$$\text{“} \lim_{N \to \infty} S_N = f. \text{”} \tag{50}$$

What does this statement mean? It means that the "distance" between the functions S_N and the function f is going to zero as N tends to ∞. In other words, the functions S_N are getting "closer" to

Approximations to functions yielded by partial sums of Fourier series

$$f(x) = \frac{1}{2}(\pi - |x|), \quad -\pi < x < \pi$$

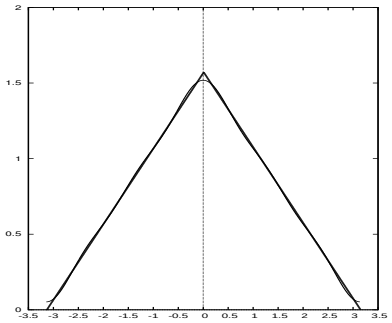

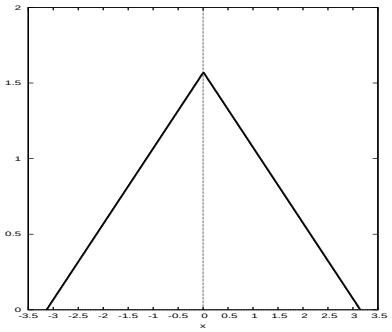

$$f(x) \approx \frac{\pi}{4} + \frac{2}{\pi}\sum_{n=1}^{N} \frac{1}{(2n-1)^2}\cos(2n-1)x. \qquad \textbf{Left: } N = 3. \ \textbf{Right: } N = 25.$$

$$f(x) = \begin{cases} -\frac{\pi}{4}, & -\pi < x < 0, \\ \frac{\pi}{4}, & 0 < x < \pi. \end{cases}$$

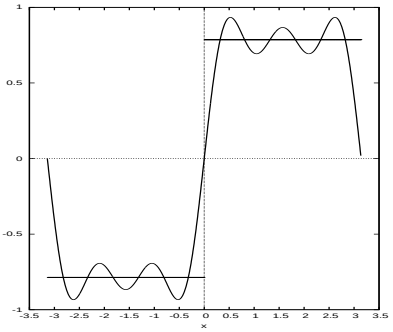

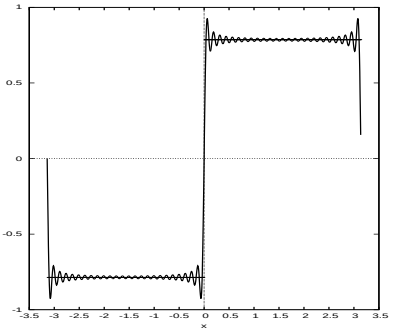

$$f(x) \approx \sum_{n=1}^{N} \frac{1}{2n-1}\sin(2n-1)x. \qquad \textbf{Left: } N = 3. \ \textbf{Right: } N = 25.$$

f as N tends to ∞.

The question, of course, is: "How do you define the 'distance' between these functions?"

The answer is: "It depends on the space of functions with which you wish to work!"

Of course, the above "answer" doesn't clearly answer anything at this time. We'll have to establish some possible distance functions that are associated with spaces of functions. For the time being, let's keep the following example in mind: Given two continuous functions $f(x)$ and $g(x)$ on an interval $[a, b]$, how could we characterize how "close" they are? If they are "close," then we would imagine that their graphs would be close, as on the left below. On the other hand, if they are "farther apart", then their graphs would be "farther apart," as shown on the right below.

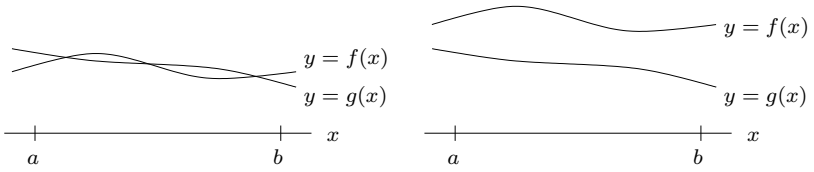

But what about the situation sketched below? Are the functions $f(x)$ and $g(x)$ "close" or "far apart"?

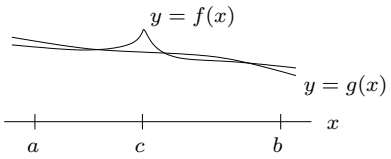

The answer depends on the distance function or **metric** you wish to use. In some applications, we would say that these functions are not close. But in others, we would be willing to tolerate the relatively small region centered at point c over which the values $f(x)$ and $g(x)$ differ significantly from each other.

The first thing to do is to set up some mathematical machinery to deal with sets – sets of whatever: points, sets, functions, measures, etc. – for which there is a distance function defined between elements of these sets. Such sets are called **metric spaces**.

Metric spaces

Definition: A **metric space**, denoted as (X, d), is a set X with a "metric" d that assigns nonnegative "distances" between any two elements $x, y \in X$. Mathematically, the metric d is a mapping $d : X \times X \to [0, \infty)$, a real-valued function that satisfies the following properties:

1. **Positivity:** $d(x, y) \geq 0, \quad d(x, x) = 0, \quad \forall x, y \in X.$

   The distance between any two elements is nonnegative. The distance between an element and itself is zero.

2. **Strict positivity:** $d(x, y) = 0 \Rightarrow x = y.$

   The only way that the distance between two elements is zero is if the two elements are the same element.

3. **Symmetry:** $d(x, y) = d(y, x), \quad \forall x, y \in X.$

4. **Triangle inequality:** $d(x, y) \leq d(x, z) + d(z, y), \quad \forall x, y, z \in X.$

Let us now consider some examples of metric spaces.

Example 1: The set of real numbers, i.e., $X = \mathbb{R}$, with metric

$$d(x, y) = |x - y|, \qquad x, y \in \mathbb{R}. \tag{51}$$

It is easy to check that the expression $|x - y|$ satisfies the first three conditions for a metric. That it also satisfies the triangle inequality condition follows from the basic property of absolute values,

$$|a + b| \leq |a| + |b|, \qquad a, b \in \mathbb{R}. \tag{52}$$

If we set $a = x - z$ and $b = -y + z$, then substitution into the above inequalty yields

$$|x - y| \leq |x - z| + |z - y| = |x - z| + |y - z|, \tag{53}$$

proving that the triangle inequality is satisfied by $d(x, y) = |x - y|$.

Example 1a: The set of rational numbers $\mathbb{Q} \subset \mathbb{R}$, with the same metric as in Example 1, i.e.,

$$d(x, y) = |x - y|, \qquad x, y \in \mathbb{Q}. \tag{54}$$

This example was included in order to show that subsets of a metric space are also metric spaces – you don't have to have the entire set! This leads to the next special case:

Example 1b: The interval $[a, b] \subset \mathbb{R}$ with metric $d(x, y) = |x - y|$.

The intervals $[a, b)$, $(a, b]$ and (a, b) are also metric spaces with the above metric. In fact, any nonempty subset $S \subset \mathbb{R}$ is also a metric space – even the singleton set $\{0\}$.

Example 2: The set $X = \mathbb{R}^n$ of ordered n-tuples. Given $x = (x_1, x_2, \cdots, x_n)$ and $y = (y_1, y_2, \cdots, y_n)$, we are most familiar with the **Euclidean metric**,

$$d_2(x, y) = \left[\sum_{i=1}^{n} (x_i - y_i)^2\right]^{1/2}. \tag{55}$$

But this metric is a special case of the more general family of "p-metrics" in $\mathbb{R}^n$:

$$d_p(x, y) = \left[\sum_{i=1}^{n} |x_i - y_i|^p\right]^{1/p}, \quad p \geq 1. \tag{56}$$

The special case $p = 1$ corresponds to the so-called "Manhattan metric":

$$d_1(x, y) = |x_1 - y_1| + |x_2 - y_2| + \cdots + |x_n - y_n|. \tag{57}$$

These metrics satisfy the triangle inequality thanks to the so-called **Minkowski inequality**:

$$\left[\sum_{i=1}^{n} |x_i \pm y_i|^p\right]^{1/p} \leq \left[\sum_{i=1}^{n} |x_i|^p\right]^{1/p} + \left[\sum_{i=1}^{n} |y_i|^p\right]^{1/p}, \quad p \geq 1. \tag{58}$$

There is a kind of limiting case of this family of metrics, the case $p = \infty$, i.e., the metric

$$d_\infty(x, y) = \max_{1 \leq i \leq n} |x_i - y_i|. \tag{59}$$

This metric is seen to extract the largest difference between corresponding elements x_i and y_i.

Metric spaces of functions

We now examine metric spaces of functions, which will be useful for the analysis started earlier.

Example 3: The space $X = C[a, b]$ of continuous real-valued functions on the interval $[a, b]$, where a and b are finite. One possible metric is given by the "infinity metric", so named because of the analogy with the infinity metric in $\mathbb{R}^n$, cf. Eq. (59):

$$d_\infty(f, g) = \max_{a \leq x \leq b} |f(x) - g(x)|. \tag{60}$$

This metric also extracts the largest difference between the values of $f(x)$ and $g(x)$ over the interval $[a, b]$.

Note that if $d_\infty(f, g) < \epsilon$, it follows that

$$|f(x) - g(x)| < \epsilon, \quad \forall x \in [a, b]. \tag{61}$$

This metric may be viewed as the special, limiting case, $p = \infty$, of the following family of metrics involving integrals,

$$d_p(f, g) = \left[\int_a^b |f(x) - g(x)|^p \, dx\right]^{1/p}, \quad p \geq 1. \tag{62}$$

From AMATH 231, and possibly other courses (for example, a course in quantum mechanics), you have encountered the special case $p = 2$,

$$d_2(f, g) = \left[\int_a^b |f(x) - g(x)|^2 \, dx\right]^{1/2}. \tag{63}$$

The d_p metrics satisfy the triangle inequality by virtue of the following integral form of Minkowski's inequality,

$$\left[\int_a^b |f(x) \pm g(x)|^p \, dx\right]^{1/p} \leq \left[\int_a^b |f(x)|^p \, dx\right]^{1/p} + \left[\int_a^b |g(x)|^p \, dx\right]^{1/p}, \quad p \geq 1\,. \tag{64}$$

Sample calculations: Let $f(x) = \frac{1}{4}$ and $g(x) = x^2$ be defined on $[a, b] = [0, 1]$, as sketched below.

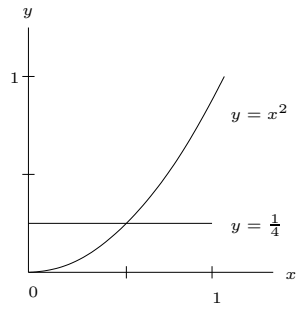

1. $p = \infty$:

$$d_\infty(f, g) = \max_{0 \leq x \leq 1} \left|\frac{1}{4} - x^2\right| = \frac{3}{4}. \tag{65}$$

(The maximum deviation between the graphs occurs at $x = 1$.)

2. $p = 1$:

$$\begin{aligned} d_1(f,g) &= \int_0^1 \left| \frac{1}{4} - x^2 \right| \, dx \\ &= \int_0^{1/2} \left(\frac{1}{4} - x^2 \right) \, dx + \int_{1/2}^1 \left(x^2 - \frac{1}{4} \right) \, dx \\ &= \frac{1}{4} = 0.25. \end{aligned} \tag{66}$$

(Details of calculation left as an exercise.)

3. $p = 2$:

$$\begin{aligned} d_2(f,g) &= \left[\int_0^1 \left(\frac{1}{4} - x^2 \right)^2 \, dx \right]^{1/2} \\ &= \left[\frac{23}{240} \right]^{1/2} \approx 0.41. \end{aligned} \tag{67}$$

Recall that in this example, we have confined our attention to the space of **continuous** functions $C[a,b]$ on an interval. Later in this section, we shall consider the use of the above metrics to other function spaces that are important in signal and image processing.

Metric spaces (cont'd)

Example 4: One final set of important examples involves "sequence spaces". We denote l^p, for $p \geq 1$, as the set of infinite sequences $x = (x_1, x_2, x_3, \cdots)$, with $x_i \in \mathbb{R}$ (or $\mathbf{C}$) such that

$$\sum_{i=1}^{\infty} |x_i|^p < \infty. \tag{68}$$

The metric on the space l^p will be given by

$$d_p(x,y) = \left[\sum_{i=1}^{\infty} |x_i - y_i|^p \right]^{1/p}, \qquad \text{for } p \geq 1. \tag{69}$$

Of particular importance will be the sequence l^2, i.e., $p = 2$, the set of **square-summable** sequences:

$$l^2 = \{ x = (x_1, x_2, x_3, \cdots) \mid \sum_{i=1}^{\infty} x_i^2 < \infty \}. \tag{70}$$

Metric spaces relevant to the study of images

As stated in Lecture 1, we may consider an idealized image – black-and-white, for simplicity – to be represented by an image function $u : \mathbb{R}^2 \to \mathbb{R}$. For simplicity, we'll assume that the domain of u is $[0,1]^2$ and its range is $[0,1]$.

For simplicity, we consider the metric space $X = C([0,1]^2$ of continuous functions defined on the set $[0,1]^2$. The "infinity metric" on this space will be a two-dimensional version of the infinity metric for continuous functions on the interval $[a,b]$ examined earlier. The distance between two functions f and g in this space is defined as

$$d_\infty(f,g) = \max_{(x,y)\in[0,1]^2} |f(x,y) - g(x,y)| \,. \tag{71}$$

This metric extracts the largest difference between $f(x,y)$ and $g(x,y)$ over the domain $[0,1]^2$.

There is also the following family of p-metrics involving double integrals,

$$d_p(f,g) = \left[\int_0^1 \int_0^1 |f(x,y) - g(x,y)|^p \, dx \, dy\right]^{1/p}, \qquad p \geq 1. \tag{72}$$

The special case $p = 2$ will be important in our applications,

$$d_2(f,g) = \left[\int_0^1 \int_0^1 [f(x,y) - g(x,y)]^2 \, dx \, dy\right]^{1/2}. \tag{73}$$

Recalling that digital images may be represented by matrices, we may define the following two-dimensional versions of the p-metrics defined over vectors in $\mathbb{R}^n$: For the $n_1 \times n_2$ matrices, $\mathbf{u} = \{u_{ij}\}$ and $\mathbf{v} = \{v_{ij}\}$,

$$d_p(\mathbf{u},\mathbf{v}) = \left[\sum_{i=1}^{n_1}\sum_{j=1}^{n_2} |u_{ij} - v_{ij}|^p\right]^{1/p}. \tag{74}$$

The special case of the Euclidean metric, $p = 2$,

$$d_2(\mathbf{u},\mathbf{v}) = \left[\sum_{i=1}^{n_1}\sum_{j=1}^{n_2} [u_{ij} - v_{ij}]^2\right]^{1/2}, \tag{75}$$

is of particular importance in applications. In fact, let us mention two other forms of the Euclidean distance betwen digital images (matrices) that are commonly used in the image processing literature:

- **Mean squared error (MSE):**

$$\mathrm{MSE}(\mathbf{u},\mathbf{v}) = \frac{1}{n_1 n_2}[d_2(\mathbf{u},\mathbf{v})]^2 = \frac{1}{n_1 n_2}\sum_{i=1}^{n_1}\sum_{j=1}^{n_2} [u_{ij} - v_{ij}]^2 \,. \tag{76}$$

- **Root mean squared error (RMSE):**

$$\mathrm{RMSE}(\mathbf{u},\mathbf{v}) = \sqrt{\mathrm{MSE}} = \left[\frac{1}{n_1 n_2}\sum_{i=1}^{n_1}\sum_{j=1}^{n_2} [u_{ij} - v_{ij}]^2\right]^{1/2}. \tag{77}$$

MSE and RMSE are useful because they characterize the average error per pixel. In this way, one can compare errors associated with (pairs of) images of different sizes.

Note: The one- and two-dimensional metrics for vectors and matrices, respectively, given above are not really different, as you may suspect. One may construct an $n_1 \times n_2$-vector from an $n_1 \times n_2$ matrix in many ways, e.g., writing the elements of the matrix out row by row (or column by column, or even diagonalwise). That being said, there are other reasons to consider two-dimensional representations of images. In the case of a one-dimensional signal, $u = (u_j)$, the value of a signal at a point, say, u_j, will often be closely connected to that of its neighbours, u_{j-1} and u_{j+1}. In the case of a two-dimensional image, the greyscale value at a particular pixel, say, u_{ij} will often be closely connected to not only its horizontal neighbours, i.e., $u_{i,j-1}$ and $u_{i,j+1}$ but also its vertical neighbours, $u_{i+1,j}$ and $u_{i-1,j}$.