LECTURE 1: CLASSICAL FOURIER ANALYSIS

Aaron Abrams
CS 294, Spring 1999

1 Purpose

Suppose you have a periodic signal, represented as a periodic function $f : \mathbb{R} \to \mathbb{R}$. In countless physical contexts the important features of the signal are the frequencies comprising the signal and their corresponding amplitudes. This information is exactly what a Fourier series expression of f tells you. Mathematically, a Fourier series for f is a way of expressing $f = f(t)$ as a linear combination of functions $\sin(\omega t)$ and $\cos(\omega t)$, for different values of ω. So Fourier analysis is the main tool used in signal processing.

This kind of analysis is widely useful in other areas as well. For non-periodic functions there is a corresponding expression called the Fourier integral. Fourier series and Fourier integrals are both examples of Fourier transforms. We will develop some basic theory underlying Fourier transforms, and then we will see two striking applications, one from signal processing and one from tomography.

2 Math

2.1 Fourier series, the analytic approach

Again, suppose f is periodic with period T. It isn't too hard to show that the only frequencies in f are integer multiples of the fundamental frequency $\omega_0 = \frac{2\pi}{T}$. We can simplify our notation by considering complex valued functions; then in view of the relationship $e^{it} = \cos t + i \sin t$, valid for all t, we seek an expression

$$f(t) = \sum_{k \in Z} a_k e^{ki\omega_0 t}.$$

[Note: $i = \sqrt{-1}$. Some people use the symbol j for this. I have no idea why.]

If such an expression exists (see the technical stuff below), it is easy to obtain the following formula for the coefficients:

$$a_k = \frac{1}{T} \int x(t) e^{-ki\omega_0 t} \, dt.$$

This is the *Fourier series* representation of f.

2.2 Fourier integrals, the analytic approach

If $f : \mathbb{R} \to \mathbb{C}$ is not periodic, one can still ask for the frequencies contributing to f. In this case they need not all be multiples of some fixed number, nor must they even necessarily form a discrete set. One of Fourier's brilliant ideas was to analyze this situation by starting with the characteristic function for a disjoint union of intervals, say $\cup_{n\in\mathbb{Z}}[nL-l, nL+l]$, with l and L constants. He then let L approach infinity, so that in the "limit" he would have the non-periodic function which is the characteristic function of $[-l, l]$. His analysis leads to the following *Fourier integral* expression for a non-periodic function $f(t)$:

$$f(t) = \frac{1}{2\pi}\int_{-\infty}^{\infty} F(\omega)e^{i\omega t}d\omega, \tag{1}$$

where

$$F(\omega) = \int_{-\infty}^{\infty} f(t)e^{-i\omega t}dt. \tag{2}$$

A helpful intuition for interpreting these formulas is that we are expressing $f(t)$ as a linear combination (integral) of exponentials, whose coefficients ($F(\omega)$) are computed just like before, but now the frequency (ω) varies continuously.

Notation. Signals are usually thought of as functions of t or u, and are denoted by lower case letters (like x, y, or h). Their Fourier transforms are usually thought of as functions of ω and are denoted with the corresponding capital letter.

2.3 Technical stuff

Now some honesty is required. Nothing we've done so far is wrong, but we've been assuming that all the required representations of f exist. For instance, suppose f is periodic. *If* f has a Fourier series, then we can compute the a_k, but various things might go wrong, such as the integrals not converging, or f being actually impossible to express in this way.

Here are two technical conditions which guarantee a Fourier series. Both conditions hold in most examples that come up in practice.

- If the Lebesgue integral $\int_{\text{period}} |f(t)|^2 dt$ is finite, then the a_k's exist, and the series converges to $f(t)$.
- If you only know that $\int_{\text{period}} |f(t)| dt$ is finite, but you also know that each period has only finitely many local minima, local maxima, and discontinuities, then again the a_k exist and the series converges exactly to $f(t)$.

Actually, if you're keeping track, the second type of convergence is stronger. The series adds up to $f(t)$ wherever f is continuous, and it adds up to the average of the left-hand limit and the right-hand limit at the (isolated) points of discontinuity of f. In the first case, the convergence means only that

$$\int_t \lim_{n\to\infty} |x(t) - \sum_{k=-n}^{n} a_k e^{ki\omega_0 t}|^2 dt = 0.$$

The point of all this is that every reasonably decent periodic function has a Fourier series.

Moreover, exactly the same conditions on a non-periodic function f guarantee existence and convergence of a Fourier integral. Just interpret the period to be $(-\infty, \infty)$.

2.4 Fourier transforms, the linear algebra approach

Earlier we were thinking about periodic functions from $\mathbb{R}$ to $\mathbb{C}$. Let us now think about *all* such functions with a fixed period T (allowing T to be infinite). For notation's sake let I denote the interval $[-T/2, T/2]$. To avoid the above technicalities let's restrict to $L^2(I) = \{f : I \to \mathbb{C} :: \int_{-T/2}^{T/2} |f(t)|^2\,dt < \infty\}$. This is a Hilbert space, which means it is a complex vector space with an inner product and a norm, and it is complete with respect to the norm. It is a vector space because complex scalar multiples and linear combinations of square-integrable functions are square-integrable; the inner product is given by $\langle f, g\rangle = \int_I f \cdot \bar{g}\,dx$. The norm is induced by the inner product, via $||f|| = \langle f, f\rangle^{1/2}$. Finally, complete means that every Cauchy sequence converges, i.e. a sequence f_n has a limit in $L^2(I)$ if $||f_n - f_m|| \to 0$ as $m, n \to \infty$.

Now, assuming T is finite for a moment, the set

$$\mathcal{F} = \{\ldots, e^{-\frac{2\pi i}{T}t}, 1, e^{\frac{2\pi i}{T}t}, e^{\frac{4\pi i}{T}t}, \ldots\}$$

is an orthogonal set of functions on I. The Fourier series (transform) of f is the projection of f onto $\mathrm{span}(\mathcal{F})$.

But more is true. In fact, $\mathcal{F}$ is a Hilbert basis for $L^2(I)$, which means that $\mathrm{span}(\mathcal{F})$ is all of $L^2(I)$. Thus the Fourier transform takes a funtion given in the standard basis and gives its coordinates in the Fourier basis $\mathcal{F}$. *The Fourier transform is just a change of basis in $L^2(I)$.*

3 Multiplication and Convolution

Here is a signal processing problem from astronomy: Suppose you are trying to make a map of the visible night sky. You seek a function $x = x(u_1, u_2)$ of two (spatial) variables, where $x(u_1, u_2)$ denotes something like the brightness at the location whose coordinates are (u_1, u_2). What you see through your telescope, however, is something different, since it's distorted by the atmosphere, your instrument, and probably other things. Call the observed signal y. How much can you figure out about x by studying y?

Let's go back to one variable, and consider the problem of trying to deduce a signal $x(t)$, assuming it has been affected by a linear, time-independent system with output $y(t)$. We will show that in order to recover x, you only need to know how the system distorts a single impulse signal (of known intensity). Intuitively, this is because any input signal $x(t)$ can be thought of as a sum of impulse functions, one for each time t having intensity $x(t)$. Since the system is linear and time-independent, the output is just the sum of the outputs for each of the separate impulse functions.

Indeed, if the "impulse response" is denoted $h(t)$ = output of the system given an impulse of height 1 at time 0, then for any input $x(t)$ the output will be given by

$$y(t) = \int_{-\infty}^{\infty} h(t-\tau)x(\tau)d\tau. \tag{3}$$

Of course we still need to know how to determine x, knowing y and h.

The easiest way to figure out x is to use the Fourier transform. (Actually since x and y were not assumed to be periodic, we use the Fourier integral.) Both sides of equation 3 are functions of t, and they are equal. Thus their Fourier integrals are equal. Starting from the left hand side, we discover:

$$\begin{aligned}
Y(\omega) &= \int_{-\infty}^{\infty} y(t)e^{-i\omega t}dt \\
&= \int_{-\infty}^{\infty} \left(\int_{-\infty}^{\infty} h(t-\tau)x(\tau)d\tau \right) e^{-i\omega t}dt \\
&= \int_{-\infty}^{\infty} x(\tau) \left(\int_{-\infty}^{\infty} h(t-\tau)e^{-i\omega t}dt \right) d\tau \\
&= \int_{-\infty}^{\infty} x(\tau) \left(\int_{-\infty}^{\infty} h(t-\tau)e^{-i\omega (t-\tau)}dt \cdot e^{-i\omega\tau} \right) d\tau \\
&= \int_{-\infty}^{\infty} x(\tau)e^{-i\omega\tau}H(\omega)d\tau \\
&= H(\omega)X(\omega).
\end{aligned}$$

So the moral is, if you know h and y, you can compute their Fourier integrals. Then you know the Fourier integral $X(\omega) = Y(\omega)/H(\omega)$, so you can recover $x(t)$ using formula 2.

In other words, if you point your telescope at an isolated star of known brightness and see how it looks, you can figure out how to interpret the other data you observe. (The same argument works in higher dimensions.)

Given any functions $h, x : \mathbb{R} \to \mathbb{C}$, the function $h \star x = \int_{-\infty}^{\infty} h(t-\tau)x(\tau)d\tau$ is called the *convolution* of h and x. We have shown:

Theorem 1 *The Fourier integral of the convolution of two functions is the product of the Fourier integrals of the functions. That is, if $y(t) = h(t) \star x(t)$ then $Y(\omega) = H(\omega)X(\omega)$.*

Exercise 1 Similarly, show that if $y(t) = h(t)x(t)$ then $Y(\omega) = H(\omega) \star X(\omega)$.

4 Life

Here are two applications.

4.1 The Shannon sampling theorem

Again think of $x(t)$ as a signal which we are trying to reconstruct. Now there is no distortion, but the challenge instead is to reconstruct x by sampling it at various times t. In general, of course, you cannot hope to succeed, since anything could be happening at the moments when you aren't sampling. However, if the signal is band-limited, i.e. the frequencies ω contributing to x all lie in the bounded interval $|\omega| < \omega_{max}$, then it is possible to determine x exactly, by sampling with frequency just $2\omega_{max}$. This is the principle underlying CD players (and a million other devices), and is called the Shannon sampling theorem.

Theorem 2 *If $x(t)$ is band-limited, i.e. $X(\omega) = 0$ for $|\omega| > \omega_{max}$, then $x(t)$ is* completely determined *by the values $x(nT)$, provided T is sufficiently small (anything $< \frac{\pi}{\omega_{max}}$ will do).*

To see this, we should think about an "impulse train" $c_T(t)$, also called a Dirac comb, of period T. (c_T is defined by $c_T(t) = \sum_{n\in\mathbb{Z}} \delta(t - nT)$.)

Exercise 2 The Fourier transform of c_T is an impulse train of period $\frac{2\pi}{T}$.

Now, we can think of $x(nT)$ as the product $x(t)c_T(t)$. We know $x(nT)$ and $c_T(t)$, and we seek $x(t)$. Taking Fourier transforms, we get $\mathrm{FT}(x(nT)) = X(\omega) \star c_{\frac{2\pi}{T}}(\omega)$. As long as $\frac{2\pi}{T} > 2\omega_{max}$, the graph of $X(\omega) \star c_{\frac{2\pi}{T}}(\omega)$ looks like the graph of $X(\omega)$, repeated with period $\frac{2\pi}{T}$. (The inequality guarantees that the support of $X(\omega)$ is translated off itself. Remember $X(\omega) = 0$ for $|\omega| > \omega_{max}$ so the support of $X(\omega)$ is contained in $[-\omega_{max}, \omega_{max}]$.) So we can simply filter $\mathrm{FT}(x(nT))$, only looking between $-\omega_{max}$ and ω_{max}, and then un-transform. This will retrieve $x(t)$.

Notice that this theorem describes a close relationship between the standard basis for L^2 and the Fourier basis. Suppose x has period 1, so $x \in L^2([0,1])$. For f to be band-limited means that f lies in the (finite-dimensional) subspace of L^2 generated by exponentials $\{e^{\pm 2\pi i k t}\}$ with $|k| \leq \frac{\omega_{max}}{2\pi}$. The theorem says that such an f will also be in the span of the same number of delta functions, namely $\{\delta(t - n\frac{\pi}{\omega_{max}})\}$ for $n = 0, \ldots, \frac{\omega_{max}}{\pi}$. This is analogous to the fact that a polynomial of degree d is determined by its values at $d+1$ points.

4.2 The projection-slice theorem

The next theorem addresses a 2-dimensional problem. Let $x = x(u_1, u_2)$ be a function from $\mathbb{R} \times \mathbb{R}$ to $\mathbb{C}$. (So the plane $\mathbb{R} \times \mathbb{R}$ has coordinates (u_1, u_2).) It may help to visualize x as a density function.

The 2-dimensional Fourier transform of $x(u_1, u_2)$ is:

$$X(\omega_1, \omega_2) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} x(u_1, u_2) e^{-i\omega_1 u_1} e^{-i\omega_2 u_2} \, du_1 du_2$$

and x is recovered by

$$x(u_1, u_2) = \frac{1}{4\pi^2} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} X(\omega_1, \omega_2) e^{i\omega_1 u_1} e^{i\omega_2 u_2} \, d\omega_1 d\omega_2.$$

Now, here is the problem. Choose a line l (through 0) in $\mathbb{R} \times \mathbb{R}$ and project x onto l. This means at each point $p \in l$ write down the number $\int_{l^\perp(p)} x(u_1, u_2)$, where $l^\perp(p)$ is the line through p perpendicular to l. This gives a (1-dim) signal x_l on the line l. How much can we learn about x by analyzing the signal x_l?

The following theorem implies that by looking at x_l for various lines l, one can reconstruct x up to any prescribed error.

Theorem 3 *Projection-slice theorem. The (1-dim) Fourier transform of the signal x_l is equal to the (2-dim) Fourier transform of x, restricted to l.*

Suppose l makes an angle θ with the u_1-axis. Define another coordinate system u_1', u_2' on the plane by

$$\begin{bmatrix} u_1' \\ u_2' \end{bmatrix} = \begin{bmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \end{bmatrix}.$$

(So the new coordinates are just the old ones rotated by θ. The u_1'-axis is l.) Now we can write the signal x_l as a function of u_1', namely $x_l(u_1') = \int_{-\infty}^{\infty} x(u_1, u_2)\, du_2'$. Taking FT's, we get:

$$\begin{aligned}
X_l(\omega) &= \int_{-\infty}^{\infty} x_l(u_1') e^{-i\omega u_1'} \, du_1' \\
&= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} x(u_1, u_2) \, du_2' \; e^{-i\omega u_1'} \, du_1' \\
&= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} x(u_1, u_2) e^{-i\omega(u_1 \cos\theta + u_2 \sin\theta)} \, du_1' du_2' \\
&= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} x(u_1, u_2) e^{-i\omega(u_1 \cos\theta)} e^{-i\omega(u_2 \sin\theta)} \, du_1 du_2 \\
&= X(\omega\cos\theta, \omega\sin\theta).
\end{aligned}$$

This is what we set out to prove.

Thus by choosing enough lines l to project onto, we can figure out as much as we want about $X(\omega_1, \omega_2)$, thereby recovering $x(u_1, u_2)$ as accurately as we desire.

5 Onward

What we've seen so far are Fourier transforms over two different groups. For non-periodic functions the Fourier integral is a Fourier transform over $\mathbb{R}$, and for periodic functions the Fourier series is a Fourier transform over the circle S^1. Essentially the rest of the course will focus on Fourier transforms over finite groups. Similarities to look for include:

- The Fourier transform is still a change of basis in the function space $L^2(G)$, where G is the relevant group;
- The Fourier basis consists of homomorphisms from G to $\mathbb{C}$;
- Multiplication of functions corresponds to convolution of their transforms, and vice versa;
- And some other things too.