A CRASH COURSE IN FOURIER ANALYSIS

LEO GOLDMAKHER

ABSTRACT. A woefully brief and inadequate overview of some Fourier analysis: Fourier series, the Fourier transform on $\mathbb{R}$, and Fourier analysis on finite abelian groups. Proofs are kept to a minimum, and the presentation is largely non-rigorous.

1. MOTIVATING FOURIER ANALYSIS: SOUND WAVES

We begin with a motivating question.

Question. *Why does a viola sound different from a clarinet?*

Before answering this, we first quickly describe why instruments make sound at all: they create vibrations, which induce sound waves in the air that spread and cause our eardrums to vibrate, which our brains then interpret. In other words, our ears are designed to receive sound waves, and our brains to process them.

Now imagine you hear some music. How do you describe the sounds you hear? Sometimes this isn't so hard. For example, we might say that something sounds like a guitar playing a certain pitch at a certain volume. Two aspects of this description—pitch and volume—are relatively easy to quantify, but the third—that it "sounds like a guitar"—is much trickier. This last characteristic is called the *timbre* of the sound. We've thus arrived at a more refined version of our initial question:

Question (Refined version)**.** *How do we describe the difference between the timbre of a viola and the timbre of a clarinet?*

This turns out to be a bit of a trick question, because when a viola plays what is ostensibly a single pitch (say, the A above middle C), the sound wave formed is not a pure wave but a combination of multiple frequencies: in addition to the primary pitch at 440 Hz[1] there are secondary vibrations formed at other frequencies (called *harmonics* or *overtones*). Both the frequencies and the volumes of the various harmonics are entirely dependent on the instrument, and it is their combination that forms the timbre. Thus, to describe timbre, all we need to do is to decompose a given sound wave into its constituent frequencies, along with the amplitude of each.[2] It turns out that many other natural phenomena are combinations of different pure waves—light, for example. We can now rephrase our questions from above more precisely:

The Fundamental Question. *Given a wave, how do we decompose it into its constituent frequencies?*

Answering this question is the goal of Fourier analysis. One immediate obstacle is that there are multiple interpretations of this question. We will explore a few different ones, each of which will lead us to a different subfield: Fourier series, the Fourier transform, and Fourier analysis on finite abelian groups.

2. FOURIER SERIES

To interpret The Fundamental Question, we need to decide what a *wave* is, as well as what it means to decompose a wave into frequencies. Here's one possible interpretation. Imagine a viola holding the A above middle C, sustaining it indefinitely at a constant volume and pitch. One might imagine that such a sound wave is a nice periodic function. This inspires the following:

Date: January 9, 2022.

[1] 440 Hertz (i.e. 440 vibrations per second) is the current standard for the A above middle C. However, there are many shades of A. For example, baroque orchestras usually tune to 415 Hz instead.

[2] There are two other important components of timbre which we won't discuss: the *attack*, which is the rate at which the various harmonics are attained, and the *decay*, which is the rate at which they fall away.

Question. *Given a periodic function f with reasonably nice analytic properties, is it possible to decompose f as a linear combination of pure waves?*

For concreteness, let's say f has period 1 (it's not difficult to adapt all the theory below to the case of arbitrary periodic functions). We'd like to break f up into pure waves, i.e. all possible shifts of the sine function. Given that f has period 1, it's natural to require our pure waves to also have period 1; the simplest such pure waves are functions of the form $\sin 2\pi nx$ and $\cos 2\pi nx$ with n an integer. Our question therefore becomes:

Question (Mathematically precise version)**.** *Given f with period 1 and reasonably nice analytic properties, is it possible to write f in the form*

$$f(x) = \sum_{n\in\mathbb{Z}} (a_n \sin 2\pi nx + b_n \cos 2\pi nx)? \tag{2.1}$$

Remark. *A natural question is: do we really need both sines and cosines in (2.1)? After all, sines and cosines are the same up to a shift. However, observe that a sum of sines is an odd function and a sum of cosines is an even function, so certainly we can't write a general function f in terms of just one or the other. It's a fun exercise to show that any function* can *be decomposed as the sum of an even function and an odd function. Of course, this doesn't prove that an expansion of the form* (2.1) *always exists, but it is a good sign.*

After having refined our initial question multiple times, we've finally arrived at a concrete math problem. Right away, there's a nice simplification one can make: we can restate the question in terms of a single sequence of coefficients, rather than two different ones. We will accomplish this by rewriting sine and cosine in terms of the complex exponential function

$$e(\alpha) := e^{2\pi i\alpha}.$$

Exercise 2.1. *Prove that an expansion of the form* (2.1) *exists iff there exists an expansion of the form*

$$f(x) = \sum_{n\in\mathbb{Z}} c_n e(nx). \tag{2.2}$$

If an expansion of this form exists, it's called the *Fourier series* of f, and the coefficients c_n are called the *Fourier coefficients* of f. One defect of the notation in (2.2) is that it doesn't indicate the dependence of the coefficients c_n on the function f. For this reason, the Fourier coefficients are usually denoted $\widehat{f}(n)$. Using this language, we restate our initial question one final time:

Fundamental Question (for Fourier Series). *Given f with period 1 and reasonably nice analytic properties, can we find some sequence of numbers $\widehat{f}(n)$ such that*

$$f(x) = \sum_{n\in\mathbb{Z}} \widehat{f}(n)e(nx)?$$

There are two different tasks implicit in this question: proving the existence of a Fourier expansion, and determining a formula for the Fourier coefficients. The latter of the two turns out to be significantly easier:

Conjecture 2.2. *Any reasonably nice function f with period 1 has a Fourier expansion with coefficients*

$$\widehat{f}(n) = \int_0^1 f(x)e(-nx)\,dx. \tag{2.3}$$

Where does this guess come from? Observe that the set of complex exponentials satisfy

$$\int_0^1 e(mx)\overline{e(nx)}\,dx = \begin{cases} 1 & \text{if } m = n \\ 0 & \text{otherwise.} \end{cases}$$

Thus if there exists a Fourier expansion

$$f(x) = \sum_{n\in\mathbb{Z}} \widehat{f}(n)e(nx),$$

then for any $k \in \mathbb{Z}$ we'd have

$$\int_0^1 f(x)e(-kx)\,dx = \int_0^1 \sum_{n\in\mathbb{Z}} \widehat{f}(n)e(nx)e(-kx)\,dx = \sum_{n\in\mathbb{Z}} \widehat{f}(n) \int_0^1 e(nx)\overline{e(kx)}\,dx = \widehat{f}(k).$$

At first glance this looks like a perfectly reasonable proof of the existence of the Fourier expansion, but there's a wrinkle—we didn't justify exchanging the order of summation and integration. And it's a good thing this proof doesn't work: the conjecture is false for some choices of function f. A big part of the theory of Fourier series is to figure out an interpretation of "reasonably nice" which makes Conjecture 2.2 true. One recurring condition is useful enough to merit its own definition: the mass of a periodic function.

Definition 2.3. *Given a function g with period p. The* mass *of g is* $\int_0^p |g|$.

Note that we integrate over a whole period of g; which specific interval we choose turns out to be irrelevant, as the following exercise shows.

Exercise 2.4. *Show that if a function g has period p, then* $\int_a^{a+p} g = \int_0^p g$ *for all $a \in \mathbb{R}$.*

The notion of mass allows us to give clean statements of a number of fundamental results about Fourier series. Our first result is very satisfying from a qualitative perspective, even if it's not terribly useful in practice.

Proposition 2.5. *Suppose f has period 1 and has finite mass. If f is continuous at θ, then*

$$f(\theta) = \sum_{n\in\mathbb{Z}} \widehat{f}(n)e(n\theta)$$

so long as the sum converges. Here (and throughout) we define $\widehat{f}(n)$ as in (2.3).

This result provides theoretical evidence for our conjecture. However, there's a crucial hypothesis that isn't so easy to check: that the putative Fourier expansion

$$f(\theta) = \sum_{n\in\mathbb{Z}} \widehat{f}(n)e(n\theta)$$

converges. (The precise interpretation of convergence is to take the partial sums with $|n| \le N$ and then let $N \to \infty$.) To have any hope of this series converging, we need the terms to tend to zero. This turns out to always be the case for all nice functions f:

Lemma 2.6 (Riemann-Lebesgue)**.** *Given f with period 1 and finite mass. Then $\widehat{f}(n) \to 0$ as $|n| \to \infty$.*

This, in turn, turns out to be a special case of a result due to Fejér:

Lemma 2.7 (Fejér)**.** *Given f with period 1 and finite mass. Then for any bounded function g with period 1 we have*

$$\int_0^1 f(x)g(nx)\,dx \xrightarrow{n\to\infty} \widehat{f}(0)\widehat{g}(0).$$

What if the Fourier coefficients not only tend to zero, but *are* zero? In view of our conjecture, we expect that f has to be the zero function. This turns out not to be true in general. Under the same mild hypotheses as our preceding theorems, however, we can prove that f must be zero almost everywhere:

Proposition 2.8. *Given f with period 1 and finite mass. If $\widehat{f}(n) = 0$ for all $n \in \mathbb{Z}$, then the f has zero mass.*

Remark. *This immediately implies that the zero function is the unique continuous f with period 1, finite mass, and all of whose Fourier coefficients vanish.*

The Riemann-Lebesgue lemma is a nice starting point for exploring convergence of Fourier series, but it isn't enough to guarantee the convergence of the Fourier expansion—just because the terms of a series tend to zero doesn't mean the sum converges! To complicate matters, it turns out that the Fourier coefficients can converge to zero arbitrarily slowly.[3] All this makes checking the convergence of the putative Fourier series a tricky business. However, if we impose enough conditions on f, the situation becomes nicer.

Theorem 2.9. *Suppose f has period 1 and is differentiable everywhere, and that its derivative f' has finite mass. Then for all $\theta \in \mathbb{R}$ we have*

$$f(\theta) = \sum_{n \in \mathbb{Z}} \widehat{f}(n) e(n\theta).$$

We can even weaken the hypotheses a little bit to handle mild discontinuities:

Theorem 2.10. *Suppose f has period 1 and is differentiable everywhere except at finitely many jump discontinuities, and that its derivative f' has finite mass. Then for all $\theta \in \mathbb{R}$ we have*

$$\frac{f(\theta^+) + f(\theta^-)}{2} = \sum_{n \in \mathbb{Z}} \widehat{f}(n) e(n\theta),$$

where $f(\theta^{\pm})$ are the left- and right-sided limits of f at θ.

The two previous results translate global properties of f into Fourier expansions which hold globally. There are also local theorems about the Fourier expansion at a single point:

Theorem 2.11. *Suppose f has period 1 and has finite mass. If f is differentiable at θ, then*

$$f(\theta) = \sum_{n \in \mathbb{Z}} \widehat{f}(n) e(n\theta).$$

Can we weaken the condition that f is differentiable at θ to merely being continuous at θ? Yes, if we replace the Fourier series by a 'smoothed' version with faster convergence:

Theorem 2.12 (Fejér)**.** *Suppose f has period 1 and has finite mass. If f is continuous at θ, then*

$$f(\theta) = \lim_{N \to \infty} \sum_{|n| \le N} \left(1 - \frac{|n|}{N}\right) \widehat{f}(n) e(n\theta).$$

2.1. **Properties of Fourier coefficients.** Given a function f we can use (2.3) to define a sequence of Fourier coefficients $\widehat{f}(n)$. We've already stated one nice property of this sequence: it tends to 0 as $|n| \to \infty$. Here are a few others:

(1) Given a real-valued f with period 1 and finite mass. Then $\widehat{f}(-n) = \overline{\widehat{f}(n)}$. (Here and throughout, $\overline{z}$ denotes the complex conjugate of z.)
(2) Given a real-valued f with period 1 and finite mass. If f is an even function, then so is $\widehat{f}$, and if f is an odd function, then so is $\widehat{f}$.
(3) Given f continuous with period 1. If f is piecewise differentiable and f' has finite mass then

$$\widehat{f'}(n) = 2\pi i n \widehat{f}(n).$$

(4) Suppose f, g both have period 1 and finite mass. Define the *convolution of f and g* to be

$$(f * g)(\theta) := \int_0^1 f(x) g(\theta - x)\, dx.$$

Then $\widehat{f * g} = \widehat{f} \cdot \widehat{g}$, where the operation on the right hand side is the usual multiplication.

[3]Broadly speaking, the smoother the function f is, the faster $\widehat{f}(n)$ converges to 0.

(5) *Parseval's identity.* Given two functions f and g with period 1 and finite mass, we have

$$\sum_{n\in\mathbb{Z}} \widehat{f}(n)\overline{\widehat{g}(n)} = \int_0^1 f\overline{g}.$$

In particular, taking $f = g$ yields $\sum |\widehat{f}(n)|^2 = \int |f|^2$.

2.2. **Final remarks on Fourier series.** Of course, the above is just the tip of the iceberg—we haven't touched much of the theory. For example, Fejér's theorem (2.12) hints at one interesting direction: the possibility of weakening the restrictions on f and still obtaining a smoothed version of a Fourier series. This idea is tremendously useful in applications, and leads to the fruitful study of *good kernels*, which we won't discuss here. Another aspect of the theory we didn't mention is the *Gibbs phenomenon*, a strange feature in the behavior of fourier series near a jump discontinuity that turns out to be a major concern in digital image processing.

Before leaving Fourier series, we point out that the theory admits a nice interpretation through the lens of linear algebra. Inspired by the theorems above, we consider the space $\mathcal{F}$ of all 'reasonably nice' functions with period 1. If we define 'reasonably nice' reasonably nicely, then $\mathcal{F}$ forms a vector space over $\mathbb{C}$ under the usual addition and scalar multiplication. What's a basis for this space? We've been trying to expand any given $f \in \mathcal{F}$ as a Fourier series, i.e. to write f as a linear combination of complex exponentials. Thus, one might hope that the set $\{e_n : n \in \mathbb{Z}\}$ forms a basis of $\mathcal{F}$, where

$$e_n(x) := e(nx).$$

Basis or not, we claim that the functions e_n form an *orthonormal system*, i.e. that the e_n's are pairwise orthogonal and each have norm 1. What does this mean? Simply that there exists an inner product on $\mathcal{F}$ (an analogue of the usual dot product on $\mathbb{R}^n$) such that the product of any two distinct e_n's is 0, and the product of any e_n with itself is 1.

Exercise 2.13. *Verify that the set $\{e_n : n \in \mathbb{Z}\}$ is an orthonormal system with respect to the inner product on $\mathcal{F}$ defined by*

$$\langle f, g\rangle := \int_0^1 f\overline{g}. \tag{2.4}$$

Viewing Fourier series from the perspective of linear algebra allows us to present the material in a cleaner way. For example:

(1) **The Fundamental Question of Fourier series**: Is the set $\{e_n : n \in \mathbb{Z}\}$ a basis of the space $\mathcal{F}$?
(2) The Fourier coefficients of f are $\widehat{f}(n) = \langle f, e_n\rangle$.
(3) The mass of any bounded function f with period 1 is finite iff $\langle f, f\rangle < \infty$. Moreover,

$$\frac{\langle f, f\rangle}{\sup |f|} \le \int_0^1 |f| \le \sqrt{\langle f, f\rangle}.$$

(4) Given f with period 1 and finite mass. If $\langle f, e_n\rangle = 0$ for all n, then $\langle f, f\rangle = 0$.
(5) It's easy to see that $\langle \overline{f}, \overline{g}\rangle = \overline{\langle f, g\rangle}$. What does this tell you about Fourier coefficients?
(6) Given f differentiable with period 1. If $\langle f', f'\rangle < \infty$ then $\langle f', e_n\rangle = 2\pi in\langle f, e_n\rangle$.
(7) *Parseval's identity.* Given two functions f and g with period 1 and finite mass, we have

$$\langle f, g\rangle = \sum_{n\in\mathbb{Z}} \langle f, e_n\rangle\overline{\langle g, e_n\rangle}.$$

In particular, $\|f\|^2 = \sum |\langle f, e_n\rangle|^2$, where $\|f\|$ is the *norm* of f, defined by $\|f\|^2 := \langle f, f\rangle$.

Exercise 2.14. *The goal of this exercise is to give two Fourier-analytic proofs that $\zeta(2) = \pi^2/6$.*
(a) Compute the Fourier series of $f(x) := \{x\}(1 - \{x\})$, where $\{x\}$ denotes the fractional part of x. Deduce that $\zeta(2) = \pi^2/6$.
(b) Compute the Fourier coefficients of $f(x) := \{x\}$ and apply Parseval's theorem to deduce $\zeta(2) = \pi^2/6$.

3. THE FOURIER TRANSFORM ON $\mathbb{R}$

If $f : \mathbb{R} \to \mathbb{C}$ is nice and has period 1, we can write it as a Fourier series, i.e. as a linear combination of the functions $e(nx)$. If f is nice and has period p, one can similarly expand it as a linear combination of complex exponentials of the form $e(nx/p)$. But what if f isn't periodic at all? Can it still be decomposed into its constituent frequencies?

Here's one natural approach. The function f might not be periodic, but it can be viewed as a limit of periodic functions f_k whose periods p_k get larger and larger. Expanding f_k as a linear combination of $e(nx/p_k)$, we obtain a sum of exponentials which are evaluated on more and more densely packed points. In the limit, this tends to 'continuous linear combination', aka an integral. Thus, in this context our fundamental question becomes:

Fundamental Question (for Fourier Transform). *Given a reasonably nice $f : \mathbb{R} \to \mathbb{C}$. Does there exist a function $\widehat{f} : \mathbb{R} \to \mathbb{C}$ such that*

$$f(x) = \int_{\mathbb{R}} \widehat{f}(\xi) e(\xi x)\, d\xi? \tag{3.1}$$

The function $\widehat{f}$ is called the *Fourier transform* of f; $\widehat{f}(\xi)$ represents the volume of the frequency ξ present in f.

Inspired by the formula (2.2) for Fourier coefficients (namely, taking the period to be infinite instead of 1), we might guess the following formula:

$$\widehat{f}(\xi) = \int_{\mathbb{R}} f(x) e(-\xi x)\, dx. \tag{3.2}$$

This guess turns out to be a good one, but (as in the case of Fourier series) the mass of f will play an important role in the statements of our theorems. Recall that the mass is the integral of $|f|$ over a single period. In this case, the period is all of $\mathbb{R}$, and the resulting measure of the mass has a fancy name:

Definition 3.1. *The L^1* norm of f *is defined to be*

$$\|f\|_1 := \int_{\mathbb{R}} |f|.$$

$L^1(\mathbb{R})$ is defined to be the collection of all $f : \mathbb{R} \to \mathbb{C}$ with finite L^1 norm.

Theorem 3.2. *If $f \in L^1(\mathbb{R})$ is continuous, then*

$$f(x) = \int_{\mathbb{R}} \widehat{f}(\xi) e(\xi x)\, d\xi.$$

We also have an analogue of Fejér's theorem from Fourier series: we can turn this into a local theorem at the cost of introducing a smoothing factor:

Theorem 3.3. *If $f \in L^1(\mathbb{R})$ and f is continuous at x, then*

$$f(x) = \lim_{X \to \infty} \int_{-X}^{X} \left(1 - \frac{|\xi|}{X}\right) \widehat{f}(\xi) e(\xi x)\, d\xi.$$

A natural question to ask is about uniqueness of the Fourier transform: if we want a decomposition of f of the form (3.1), must the function in the integrand take the form (3.2)? Under nice conditions, the answer turns out to be affirmative:

Proposition 3.4. *Suppose $f(x) = \int_{\mathbb{R}} \tilde{f}(\xi) e(\xi x)\, d\xi$ where $\tilde{f}$ is continuous and has compact support (i.e. $\tilde{f}$ is zero outside of some closed interval.) Then $\widehat{f} = \tilde{f}$.*

3.1. **Examples.** In practice, there are a few examples of Fourier transforms that arise often enough that it's worth compiling them into a table. Most of these we state without proof (the computations are tedious but largely straightforward), but the final example is fairly fundamental: we'll prove that $e^{-\pi x^2}$ is its own Fourier transform. This is not the only such function (e.g. $|x|^{-1/2}$ is also its own Fourier transform), but that fact that it's symmetric, decays quickly, and is smooth everywhere on $\mathbb{R}$ make $e^{-\pi x^2}$ a particularly nice example. We'll give two proofs of this, both of which rely on the famous integral

$$\int_{\mathbb{R}} e^{-\pi x^2}\,dx = 1$$

(which is what makes the Gaussian a probability distribution!).

$f(x) =$	$\widehat{f}(\xi) =$
$e^{-2\pi\lvert x\rvert}$	$\frac{1}{\pi}\cdot\frac{1}{1+\xi^2}$
$\begin{cases} 1 & \text{if } \lvert x\rvert \leq 1/2 \\ 0 & \text{else.}\end{cases}$	$\frac{\sin \pi\xi}{\pi\xi}$
$\begin{cases} 1-\lvert x\rvert & \text{if } \lvert x\rvert \leq 1 \\ 0 & \text{else.}\end{cases}$	$\left(\frac{\sin \pi\xi}{\pi\xi}\right)^2$
$\begin{cases} \left(\cos \frac{\pi x}{2}\right)^2 & \text{if } \lvert x\rvert \leq 1 \\ 0 & \text{else.}\end{cases}$	$\frac{\sin 2\pi\xi}{2\pi\xi(1-4\xi^2)}$
$\frac{1}{X}\left(\frac{\sin \pi X x}{\pi x}\right)^2$	$\begin{cases} 1-\frac{\lvert\xi\rvert}{X} & \text{if } \lvert\xi\rvert \leq X \\ 0 & \text{else.}\end{cases}$
$e^{-\pi x^2}$	$e^{-\pi\xi^2}$

As mentioned above, we will only prove the last one of these. In fact, we give two different proofs: one using a clever trick, the other using standard complex analysis.

Theorem 3.5. *If $f(x) = e^{-\pi x^2}$, then $\widehat{f}(\xi) = e^{-\pi\xi^2}$.*

Proof 1 (clever version). By definition,

$$\widehat{f}(\xi) = \int_{\mathbb{R}} f(x)e(-\xi x)\,dx = \int_{\mathbb{R}} e^{-\pi x^2}e(-\xi x)\,dx.$$

First differentiating both sides, and then integrating the result by parts, we obtain

$$\frac{d}{d\xi}\widehat{f}(\xi) = -2\pi i\int_{\mathbb{R}} xe^{-\pi x^2}e(-\xi x)\,dx = -2\pi\xi\int_{\mathbb{R}} e^{-\pi x^2}e(-\xi x)\,dx = -2\pi\xi\widehat{f}(\xi).$$

In other words, $\widehat{f}$ is a solution to the differential equation

$$\frac{dy}{d\xi} = -2\pi\xi y \qquad \text{with} \qquad y(0) = \int_{\mathbb{R}} e^{-\pi x^2}\,dx = 1.$$

The differential equation is easily solved to yield

$$\log|y| = -\pi\xi^2 + C,$$

and since $y(0) = 1$ we deduce that $C = 0$. Thus

$$\widehat{f}(\xi) = \pm e^{-\pi\xi^2}.$$

Since $\widehat{f}(0) = 1$, we conclude the proof. $\square$

Proof 2 (less clever, but uses complex analysis). By definition,

$$\widehat{f}(\xi) = \int_{\mathbb{R}} f(x)e(-\xi x)\,dx = \int_{\mathbb{R}} e^{-\pi x^2 - 2\pi i\xi x}\,dx. \tag{3.3}$$

The integrand looks similar to the Gaussian (normal) probability density function $e^{-\pi x^2}$, and our strategy is to exploit this similarity. To transform the integral in (3.3) into the Gaussian, we complete the square:

$$\int_{\mathbb{R}} e^{-\pi x^2 - 2\pi i\xi x}\,dx = \int_{\mathbb{R}} e^{-\pi(x+i\xi)^2 - \pi\xi^2}\,dx = e^{-\pi\xi^2}\int_{-\infty+i\xi}^{\infty+i\xi} e^{-\pi u^2}\,du.$$

Since the function $e^{-\pi u^2}$ is holomorphic everywhere and decays very quickly, we can drag the line of integration without changing the value of the integral, whence

$$\widehat{f}(\xi) = e^{-\pi\xi^2}\int_{\mathbb{R}} e^{-\pi u^2}\,du = e^{-\pi\xi^2}. \qquad \square$$

3.2. **Properties of the Fourier transform.** Having stated some fundamental results about the Fourier transform and seen some examples, we now list a few nice properties. Using these properties allows us to extend the table of examples above to many more without too much effort.

(1) $\widehat{\cdot}$ is linear, i.e. $\widehat{\alpha f + \beta g} = \alpha\widehat{f} + \beta\widehat{g}$

(2) If $g(x) = f(x+c)$ then $\widehat{g}(\xi) = e(c\xi)\widehat{f}(\xi)$.

(3) If $g(x) = f(x)e(cx)$ then $\widehat{g}(\xi) = \widehat{f}(\xi - c)$.

(4) If $g(x) = f(cx)$ then $\widehat{g}(\xi) = \frac{1}{|c|}\widehat{f}(\xi/c)$.

(5) $\widehat{\overline{f}}(\xi) = \overline{\widehat{f}(-\xi)}$.

(6) $\lim_{\xi\to\pm\infty}\widehat{f}(\xi) = 0$. (Riemann-Lebesgue lemma)

(7) $\widehat{f'}(\xi) = 2\pi it\widehat{f}(\xi)$.

(8) $\widehat{f * g} = \widehat{f}\cdot\widehat{g}$.

(9) If $\widehat{f} = \widehat{g}$ then $f(x) = g(x)$ for almost every x. More precisely, if $f, g \in L^1(\mathbb{R})$ satisfy $\widehat{f} = \widehat{g}$ then $\|f - g\|_1 = 0$.

(10) *Plancherel's identity:*

$$\langle f, g\rangle := \int f\cdot\overline{g} = \int \widehat{f}\cdot\overline{\widehat{g}} = \langle\widehat{f}, \widehat{g}\rangle.$$

Taking $f = g$ yields *Parseval's identity*

$$\int |f|^2 = \int |\widehat{f}|^2.$$

(11) *Heisenberg's uncertainty principle.* The observation is it's impossible for the masses of both f and $\widehat{f}$ to be concentrated: if one of them has most of its mass in a narrow peak, then the mass of the other must be spread out. More precisely:

$$\left(\int_{\mathbb{R}} |xf(x)|^2\,dx\right)\cdot\left(\int_{\mathbb{R}} |\xi\widehat{f}(\xi)|^2\,d\xi\right) \geq \frac{1}{16\pi^2}\left(\int |f|^2\right)^2.$$

(12) *Poisson summation.* Given a nice f, we have

$$\sum_{n\in\mathbb{Z}} f(n) = \sum_{k\in\mathbb{Z}} \widehat{f}(k),$$

where $\widehat{f}$ denotes the Fourier transform of f.

We prove the last of these properties. Note that we will never explicitly specify what it means for f to be 'nice', so technically speaking this is not a rigorous argument. It's a good exercise to go through the proof and identify a definition of 'nice' that would allow the proof to carry through rigorously.

Proof of Poisson summation. Given f, set

$$F(x) := \sum_{n\in\mathbb{Z}} f(n+x).$$

Observe that F is a periodic function with period 1, which means it admits a Fourier series:

$$F(x) = \sum_{k\in\mathbb{Z}} a_k e(kx).$$

The Fourier coefficients are straightforward to compute:

$$a_k = \int_0^1 F(x)e(-kx)\,dx = \sum_{n\in\mathbb{Z}}\int_0^1 f(n+x)e(-kx)\,dx = \sum_{n\in\mathbb{Z}}\int_n^{n+1} f(t)e(-kt)\,dt = \int_{\mathbb{R}} f(t)e(-kt)\,dt = \widehat{f}(k).$$

In other words, the k^{th} Fourier *coefficient* of F happens to be the Fourier *transform* of f evaluated at k. We conclude that

$$\sum_{n\in\mathbb{Z}} f(n+x) = \sum_{k\in\mathbb{Z}} \widehat{f}(k)e(kx)$$

for all x. Taking $x=0$ yields the claim. □

Note that one can draw other nice conclusion from the proof, for example that the sum of f over all the half-integers is an alternating sum of $\widehat{f}(k)$'s. The more important outcome of the proof is the idea that one can study a function f by constructing a periodic function F out of it and exploring F's Fourier series.

4. Fourier Analysis on Finite Abelian Groups

There is an analogous theory of Fourier transforms in groups other than $\mathbb{R}$. Let G be a finite abelian group, and let

$$L(G) := \{\phi : G \to \mathbb{C}\}.$$

Our goal is to write down a 'Fourier transform' of any given function from $L(G)$. In the case of functions on $\mathbb{R}$ (say, $f \in L^2$), we can write f as an integral of some function $\widehat{f}$ against a complex exponential over the whole space; moreover, the function $\widehat{f}$ can likewise be written as an integral of f against the conjugate of the complex exponential, integrated over the entire space. The analogue of integrating over the whole space is clear: we want to sum over all elements of the group. What is less clear, however, are the appropriate analogues of the summands—the Fourier transform and the complex exponentials.

As a first attempt, observe that we can expand any function $\phi \in L(G)$ by writing

$$\phi = \sum_{g\in G} \phi(g)\delta_g$$

where

$$\delta_g(x) := \begin{cases} 1, & \text{if } x = g \\ 0, & \text{if } x \neq g \end{cases}$$

In other words, we are trying to make the characteristic functions δ_g play the role of the complex exponentials, and ϕ to play its own Fourier transform. This is fine, but not so helpful—working with this Fourier transform offers no advantage over working with the function itself! However, this expansion does tell us that the characteristic functions δ_g span $L(G)$. It's easy to show that the δ_g's are also linearly independent. Since there are precisely $|G|$ different characteristic functions, it follows that

$$\dim L(G) = |G|.$$

Thus, whatever analogue of the exponential functions we settle on, there should be precisely $|G|$ of them.

We now describe one analogue of the Fourier transform which has proved quite useful. Let $\widehat{G}$ denote the set of all homomorphisms from G to $\mathbb{C}^\times$:

$$\widehat{G} := \mathrm{Hom}(G, \mathbb{C}^\times).$$

$\widehat{G}$ is usually called the *dual* of G; its elements are called the *characters* of G. A classical theorem (which follows from the Fundamental Theorem of Finite Abelian Groups) asserts that $\widehat{G} \simeq G$; in particular, there are precisely $|G|$ characters. Moreover, it is an exercise to show that the characters are linearly independent, and that $|\gamma(x)| = 1$ for any $\gamma \in \widehat{G}$ and any $x \in G$. It follows that the characters form a basis of $L(G)$, so for any $\phi \in L(G)$ we can therefore write find coefficients $\widehat{\phi}(\gamma)$ such that

$$\phi = \sum_{\gamma \in \widehat{G}} \widehat{\phi}(\gamma) \cdot \gamma.$$

Note that $\widehat{\phi} : \widehat{G} \to \mathbb{C}$. There is a simple formula for this Fourier transform:

$$\widehat{\phi}(\gamma) = \frac{1}{|G|} \sum_{g \in G} \phi(g)\overline{\gamma(g)}.$$

This admits a probabilistic interpretation:

$$\widehat{\phi}(\gamma) = \mathbb{E}_{g \in G}[\phi(g)\overline{\gamma(g)}].$$

It also smacks of the classical inner product. Accordingly, we define

$$\langle \phi, \psi \rangle := \frac{1}{|G|} \sum_{g \in G} \phi(g)\overline{\psi(g)}.$$

The most familiar case of all this theory is $G = \mathbb{F}_p$, in which case the characters are simply the functions $\gamma(x) := e(ax/p)$.

Having done all this, we have the following properties.

(1) **Fourier inversion**: given $\gamma \in \widehat{G}$ and $\phi \in L(G)$, we have

$$\widehat{\phi}(\gamma) = \langle \phi, \gamma \rangle.$$

Note that $\widehat{\phi} : \widehat{G} \to \mathbb{C}$.

(2) **Orthogonality**: the elements of $\widehat{G}$ form an orthonormal basis of $L(G)$ with respect to the inner product $\langle \cdot, \cdot \rangle$.

(3) **Properties of inner product**: We have all the classical properties:
- $\langle f, g \rangle = \langle g, f \rangle$
- $\langle f, g \rangle \in \mathbb{R}$
- $\langle f, gh \rangle = \langle f\overline{g}, h \rangle$
- $\langle f, g + h \rangle = \langle f, g \rangle + \langle f, h \rangle$

(4) **Plancherel**:

$$\langle \widehat{f}, \widehat{g} \rangle = \frac{1}{|G|} \langle f, g \rangle$$

(5) **Parseval**:

$$\sum_{\gamma\in\widehat{G}} |\widehat{\phi}(\gamma)|^2 = \frac{1}{|G|}\sum_{g\in G}|\phi(g)|^2.$$

(6) **Convolutions**: Set

$$(f*g)(x) := \frac{1}{|G|}\sum_{\substack{a,b\in G\\ a+b=x}} f(a)g(b).$$

Then

$$\widehat{f*g} = \widehat{f}\cdot\widehat{g} \qquad \text{and} \qquad \widehat{f}*\widehat{g} = \frac{1}{|G|}\cdot\widehat{f\cdot g}$$

Note that the convolution on $\widehat{G}$ is defined multiplicatively, i.e.

$$(\widehat{f}*\widehat{g})(\xi) := \frac{1}{|G|}\sum_{\substack{\alpha,\beta\in\widehat{G}\\ \alpha\beta=\xi}} \widehat{f}(\alpha)\widehat{g}(\beta).$$

(7) **Support on convolutions**: Suppose $f, g : G \to \mathbb{C}$. Then it can be easily verified that

$$\text{supp}(f*g) \subseteq \text{supp}(f) + \text{supp}(g).$$

In particular, $|\text{supp}(f*g)| \leq |\text{supp}(f)|\cdot|\text{supp}(g)|$.

(8) **Bound on Fourier transform**:

$$|\widehat{f}| \leq \mathbb{E}[|f|] = \frac{1}{|G|}\sum_{x\in G}|f(x)|$$

(9) **Heisenberg uncertainty principle**:

$$|\text{supp}(f)|\cdot|\text{supp}(\widehat{f})| \geq |G|$$

(10) **Additive Heisenberg (for $\mathbb{Z}/p\mathbb{Z}$)**: If $|G| = p$, Tao proved

$$|\text{supp}(f)| + |\text{supp}(\widehat{f})| \geq p+1$$

DEPT OF MATHEMATICS AND STATISTICS, WILLIAMS COLLEGE, WILLIAMSTOWN, MA, USA 01267
Email address: `Leo.Goldmakher@williams.edu`