Fourier Expansion and Linearity Testing (lecture notes)

Course: Analysis of Boolean Functions, Autumn 2025, University of Chicago
Instructor: William Hoza (`williamhoza@uchicago.edu`)

1 Fourier expansion theorem

This course is about using Fourier analysis to reason about Boolean functions. A "Boolean function" is a function of the form $f\colon \{0,1\}^n \to \{0,1\}$. Actually, we often prefer to encode $b \in \{0,1\}$ as $(-1)^b$. We therefore consider functions on the domain $\{\pm 1\}^n$. We will be even more flexible about the codomain: our functions might take values in $\{0,1\}$ or $\{\pm 1\}$ or $[-1,1]$ or $\mathbb{R}$. In this context, "Fourier analysis" refers to studying a function $f\colon \{\pm 1\}^n \to \mathbb{R}$ by representing it as a *multilinear polynomial.*

Theorem 1.1 (Fourier Expansion Theorem)**.** *Every function $f\colon \{\pm 1\}^n \to \mathbb{R}$ can be uniquely expressed as a multilinear polynomial.*

Proof. Existence: For each $a \in \{\pm 1\}^n$, define $1_{\{a\}}(x) = \left(\frac{1+a_1x_1}{2}\right) \cdot \left(\frac{1+a_2x_2}{2}\right) \cdots \left(\frac{1+a_nx_n}{2}\right)$. Then $1_{\{a\}}$ is a multilinear polynomial, and $1_{\{a\}}(a) = 1$, and $1_{\{a\}}(x) = 0$ for every $x \in \{\pm 1\}^n \setminus \{a\}$. Therefore,

$$f(x) = \sum_{a \in \{\pm 1\}^n} f(a) \cdot 1_{\{a\}}(x).$$

Uniqueness: The space of $f\colon \{\pm 1\}^n \to \mathbb{R}$ is a vector space of dimension 2^n. We just found 2^n vectors that span the space (namely, the 2^n possible products of distinct variables), so they must be a basis. $\square$

Example 1.2. If $x_1, x_2 \in \{\pm 1\}$, then

$$\begin{aligned}\max(x_1,x_2) &= (+1)\cdot\left(\frac{1+x_1}{2}\right)\cdot\left(\frac{1+x_2}{2}\right) + (+1)\cdot\left(\frac{1+x_1}{2}\right)\cdot\left(\frac{1-x_2}{2}\right) \\ &\quad + (+1)\cdot\left(\frac{1-x_1}{2}\right)\cdot\left(\frac{1+x_2}{2}\right) + (-1)\cdot\left(\frac{1-x_1}{2}\right)\cdot\left(\frac{1-x_2}{2}\right) \\ &= \frac{1}{2} + \frac{1}{2}x_1 + \frac{1}{2}x_2 - \frac{1}{2}x_1x_2.\end{aligned}$$

Notation: For $x \in \{\pm 1\}^n$ and $S \subseteq [n]$, we write x^S or $\chi_S(x)$ to denote $\prod_{i \in S} x_i$. We use $\widehat{f}(S)$ to denote the coefficient of x^S in the multilinear polynomial representation of f, i.e.,

$$f(x) = \sum_{S \subseteq [n]} \widehat{f}(S) \cdot \chi_S(x). \tag{1}$$

Eq. (1) is called the "Fourier expansion" of f, and $\widehat{f}(S)$ is called the "Fourier coefficient" of f on S. Each basis function χ_S is called a *character* or a *character function.* Observe that $\chi_S(x)$ computes the XOR of the bits in S, encoded using ± 1.

2 Linearity testing

Why study Fourier analysis? Probably the most enlightened response would be that Fourier analysis is its own reward. It is an aesthetically beautiful and mathematically deep subject that doesn't require any external motivation. Fortunately, I am not very enlightened! Therefore, for the sake of motivation, we will study many cool *applications* of Fourier analysis of Boolean functions. In fact, the rule I will try to hold myself to is: We will think of Fourier analysis purely as a proof technique, never as an end in itself. We will only develop Fourier-analytic methods when we need them for the purpose of proving non-Fourier-analytic theorems.

Our first motivating application will the *linearity testing* problem. In this problem, we have query access to an unknown function $f\colon \mathbb{F}_2^n \to \mathbb{F}_2$. Our goal is to distinguish between two cases:

- Case 1: f is linear, i.e., it has the form $f(x) = \sum_{i \in S} x_i \bmod 2$ for some $S \subseteq [n]$.
- Case 2: f is ε-far from every linear function with respect to the following metric:

  $$\mathrm{dist}(f, g) = \Pr_x[f(x) \neq g(x)].$$

  The input x is chosen uniformly at random from $\mathbb{F}_2^n$.

We will show how to solve this problem using just three queries to f. The algorithm, called the Blum-Luby-Rubinfeld (BLR) test, is as follows.

- Pick $x, y \in \mathbb{F}_2^n$ independently and uniformly at random.
- If $f(x + y) = f(x) + f(y)$, accept. Otherwise, reject.

The proof of correctness uses Fourier analysis. To accomodate the domain $\mathbb{F}_2^n$, we overload notation and define $\chi_S\colon \mathbb{F}_2^n \to \{\pm 1\}$ by the formula

$$\chi_S(x) = \prod_{i \in S} (-1)^{x_i}.$$

This way, the Fourier expansion formula doesn't change: Every function $f\colon \mathbb{F}_2^n \to \mathbb{R}$ can be uniquely expressed as

$$f(x) = \sum_{S \subseteq [n]} \widehat{f}(S) \cdot \chi_S(x).$$

Note that the Fourier coefficients $\widehat{f}(S)$ do not change if we switch to a different input encoding. However, the Fourier coefficients *do* change if we switch to a different *output* encoding.

2.1 Inner product

Definition 2.1. If $f, g\colon \{0,1\}^n \to \mathbb{R}$, we define the inner product $\langle f, g \rangle = \mathbb{E}_x[f(x) \cdot g(x)]$.

Observe that the character functions are orthonormal: $\langle \chi_S, \chi_S \rangle = \mathbb{E}_x[\prod_{i \in S} x_i^2] = 1$, and if $S \neq T$, then $\langle \chi_S, \chi_T \rangle = \mathbb{E}_x[\chi_{S \Delta T}(x)] = 0$.

Theorem 2.2 (Plancherel's Theorem)**.** $\langle f, g \rangle = \sum_{S \subseteq [n]} \widehat{f}(S) \cdot \widehat{g}(S)$.

Proof. Expand f and g in the Fourier basis. □

Corollary 2.3 (Fourier Coefficient Formula)**.** $\widehat{f}(S) = \langle f, \chi_S \rangle = \mathbb{E}_x[f(x) \cdot x^S]$.

For example, $\widehat{f}(\varnothing) = \mathbb{E}[f]$. If f is $\{\pm 1\}$-valued, then we have the formula $\widehat{f}(S) = 1 - 2\,\mathrm{dist}(f, \chi_S)$.

Corollary 2.4 (Parseval's Theorem)**.** $\sum_{S \subseteq [n]} \widehat{f}(S)^2 = \mathbb{E}_x[f(x)^2]$

2.2 Convolution

For convenience, to analyze the BLR test, we will encode input bits using 0 and 1, but we will encode output bits using ± 1. That is, we will work with a function $f \colon \mathbb{F}_2^n \to \{\pm 1\}$. We need to prove that if $f(x+y) = f(x) \cdot f(y)$ for most x, y, then f is close to a character function. Equivalently, the assumption is that $f(x) = f(x+y) \cdot f(y)$ for most x, y. This motivates the following definition.

Definition 2.5 (Convolution)**.** If $f, g \colon \mathbb{F}_2^n \to \mathbb{R}$, then

$$(f * g)(x) = \mathop{\mathbb{E}}_{y}[f(x+y) \cdot g(y)].$$

(It would be more conventional to put $f(x-y)$ rather than $f(x+y)$, but we are working over $\mathbb{F}_2^n$, so addition and subtraction are equivalent.)

Theorem 2.6 (Convolution Fourier Formula)**.** $\widehat{f * g}(S) = \widehat{f}(S) \cdot \widehat{g}(S)$

Proof.

$$\begin{aligned}
(f * g)(x) &= \sum_{S,T \subseteq [n]} \widehat{f}(S) \cdot \widehat{g}(T) \cdot \mathop{\mathbb{E}}_{y}[\chi_S(x+y) \cdot \chi_T(y)] \\
&= \sum_{S,T \subseteq [n]} \widehat{f}(S) \cdot \widehat{g}(T) \cdot \chi_S(x) \cdot \mathop{\mathbb{E}}_{y}[\chi_{S \Delta T}(y)] \\
&= \sum_{S \subseteq [n]} \widehat{f}(S) \cdot \widehat{g}(S) \cdot \chi_S(x). \qquad \square
\end{aligned}$$

Theorem 2.7 (Blum-Luby-Rubinfeld)**.** *Let* $f \colon \mathbb{F}_2^n \to \{\pm 1\}$*. Assume that* $\Pr_{x,y}[f(x) \cdot f(y) = f(x+y)] \geq 1 - \varepsilon$*. Then* f *is* ε*-close to some character function* χ_S*.*

Proof.

$$\mathop{\mathbb{E}}_{x,y}[f(x) \cdot f(y) \cdot f(x+y)] = \Pr[f(x+y) = f(x) \cdot f(y)] - \Pr[f(x+y) \neq f(x) \cdot f(y)] \geq 1 - 2\varepsilon.$$

Therefore,

$$\begin{aligned}
1 - 2\varepsilon \leq \mathop{\mathbb{E}}_{x,y}[f(x) \cdot f(y) \cdot f(x+y)] = \mathop{\mathbb{E}}_{x}[f(x) \cdot (f * f)(x)] = \langle f, f * f \rangle &= \sum_{S \subseteq [n]} \widehat{f}(S) \cdot \widehat{f * f}(S) \\
&= \sum_{S \subseteq [n]} \widehat{f}(S)^3 \\
&\leq \left(\max_S \widehat{f}(S) \right) \cdot \sum_{S \subseteq [n]} \widehat{f}(S)^2 \\
&= \max_S \widehat{f}(S). \qquad \square
\end{aligned}$$