Computer Science 294 Lecture 2 Notes

Daniel Raban

January 19, 2023

1 Properties of Boolean Functions and BLR Linearity Testing

1.1 Recap Fourier identities for boolean functions

Last time, we proved the fundamental theorem of boolean functions.

Theorem 1.1 (Fundamental theorem of boolean functions)**.** *Every boolean function* $f : \{0,1\}^n \to \{0,1\}$ *can be uniquely represented as a multilinear polynomial (over* $\mathbb{R}$*)*

$$f(x_1, \dots, x_n) = \sum_{S \subseteq \{1,\dots,n\}} \widehat{f}(S) \prod_{i \in S} x_i.$$

This is sometimes called the Fourier representation of the function, and $\widehat{f}(S)$ is the S-Fourier coefficient of f. We also discussed an inner product on boolean functions,

$$\langle f, g \rangle = \mathbb{E}_{X \sim \{\pm 1\}^n}[f(X)g(X)],$$

and showed that the character functions $\chi_S(x) = \prod_{i \in S} x_i$ form an orthonormal basis of the vector space of functions $\{\pm 1\}^n \to \mathbb{R}$. Finally, we showed the Plancerel, Parseval, ad Fourier inversion formulas:

$$\langle f, g \rangle = \sum_S \widehat{f}(S)\widehat{g}(S), \qquad \langle f, f \rangle \sum_S \widehat{f}(S)^2, \qquad \widehat{f}(S) = \langle f, \chi_S \rangle.$$

1.2 Expectation and variance formulas

Proposition 1.1.

$$\mathbb{E}_{X \sim \{\pm 1\}^n}[f(X)] = \widehat{f}(\varnothing).$$

Proof.

$$\begin{aligned} \mathbb{E}_{X \sim \{\pm 1\}^n}[f(X) \cdot 1] &= \mathbb{E}_{X \sim \{\pm 1\}^n}[f(X) \cdot \chi_\varnothing] \\ &= \langle f, \chi_\varnothing \rangle \\ &= \widehat{f}(\varnothing). \end{aligned}$$

□

Proposition 1.2.

$$\mathrm{Var}(f(X)) = \sum_{S \neq \varnothing} \widehat{f}(S)^2.$$

Proof.

$$\begin{aligned}\mathrm{Var}(f(X)) &= \mathbb{E}_X[f(X)^2] - (\mathbb{E}_X[f(X)])^2 \\ &= \langle f, f\rangle - \widehat{f}(\varnothing)^2 \\ &= \sum_{S \neq \varnothing} \widehat{f}(S)^2.\end{aligned}$$

□

We also have

$$\begin{aligned}\widehat{f}(\{1\}) &= \mathbb{E}_{X \sim \{\pm 1\}^n}[f(X) \cdot X_1] \\ &= \frac{1}{2}\,\mathbb{E}[f(X) \mid X_1 = 1] + \frac{1}{2}\,\mathbb{E}[-f(X) \mid X_1 = -1].\end{aligned}$$

1.3 Homomorphisms and convolutions

We will sometimes express $f : \{\pm 1\}^n \to \mathbb{R}$ as $\widetilde{f} : \mathbb{F}_2^n \to \mathbb{R}$; the correspondence here is igven by $(-1)^b \leftrightarrow b$:

$$\begin{aligned}\widehat{f}(x_1, \ldots, x_n) &= f((-1)_1^x, \ldots, (-1)^{x_n}) \\ &= \sum_{S \subseteq [n]} \widehat{f}(S) \prod_{i \in S} (-1)^{x_i} \\ &= \sum_{S \subseteq [n]} \widehat{f}(S) \cdot (-1)^{\sum_{i \in S} x_i}.\end{aligned}$$

In this context, we will refer to $(-1)^{\sum_{i \in S} x_i}$ as χ_S.

We call these functions characters because they are homomorphisms:

Proposition 1.3.

$$\chi_S(x + y) = \chi_S(x)\chi_S(y).$$

Proof.

$$\begin{aligned}\chi_S(x + y) &= (-1)^{\sum_{i \in S} x_i + y_i} \\ &= (-1)^{\sum_{i \in S} x_i} (-1)^{\sum_{i \in S} y_i} \\ &= \chi_S(x)\chi_S(y).\end{aligned}$$

□

Definition 1.1. Let $f, g : \mathbb{F}_2^n \to \mathbb{R}$. The **convolution** of f and g, is a function $f * g : \mathbb{F}_2^n \to \mathbb{R}$ given by

$$(f * g)(x) = \mathbb{E}_{Y \sim \{\pm 1\}^n}[f(Y)g(x - Y)].$$

Lemma 1.1.

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

Proof.

$$\begin{aligned}
\widehat{f * g}(S) &= \mathbb{E}_{X \sim \mathbb{F}_2^n}[f * g(X)\chi_S(X)] \\
&= \mathbb{E}_{X \sim \mathbb{F}_2^n}[\mathbb{E}_{Y \sim \mathbb{F}_2^n}[f(Y)g(X-Y)] \cdot \chi_X]
\end{aligned}$$

Write $Z = X - Y$, so $X = Z + Y$. Since X, Y are indpendent, so are Z and Y.

$$\begin{aligned}
&= \mathbb{E}_{Y \sim \mathbb{F}_2^n, Z \sim \mathbb{F}_2^n}[f(Y)g(Z)\chi_S(Z+Y)]] \\
&= \mathbb{E}_{Y \sim \mathbb{F}_2^n}[f(Y)\chi_S(Y)]\,\mathbb{E}_{Z \sim \mathbb{F}_2^n}[g(Z)\chi_S(Z)]] \\
&= \widehat{f}(X) \cdot \widehat{g}(S).
\end{aligned}$$

□

The reverse is true, as well, but we will not give the proof.

Proposition 1.4.

$$\widehat{f \cdot g}(S) = \sum_T \widehat{f}(T) \cdot \widehat{g}(S \oplus T).$$

1.4 BLR Testing

We want to check if $f : \mathbb{F}_2^n \to \mathbb{F}_2$ is linear.

Definition 1.2. $f : \mathbb{F}_2^n \to \mathbb{F}$ is linear if for all x, y,

$$f(x) + g(y) = f(x + y).$$

Equivalently, there exists an $a \in \mathbb{F}_2^n$ such that for all x,

$$f(x) = \sum_{i=1}^{n} a_i x_i \pmod 2.$$

If we take $a = (a_1, \ldots, a_n)$, then $a_i = f(e_i)$, where $e_i = (0, \ldots, 0, 1, 0, \ldots, 0)$ is the i-th standard basis vector. You can show equivalence by induction.

Today, we are interested in showing that *robust* versions of these two conditions are equivalent:

(1') For most pairs x, y, $f(x) + f(y) = f(x + y)$.

(2') There exists an $a \in \mathbb{F}_2^n$ such that for most $x \in \mathbb{F}_2^n$, $f(x) = \sum_{i=1}^n a_i x_i \pmod 2$.

We can restate (2') as the existence of $S \subseteq [n]$ such that for most x, $f(x) = \sum_{i \in S} x_i$.

Proposition 1.5. *Suppose there exists an S such that $\mathbb{P}_{X\sim\mathbb{F}_2^n}[f(X)-\sum_{i\in S}X_i]\geq 1$. Then*

$$\mathbb{P}_{X,Y}(f(X+Y)=f(X)=f(Y))\geq 1-3\varepsilon.$$

Proof. Denote $A=\{x\in\mathbb{F}_2^n : f(x)=\sum_{i\in S}x_i\}$. If both $x,y\in A$, then $f(x+y)=f(x)+f(y)$. So

$$\begin{aligned}\mathbb{P}_{X,Y}(f(X+Y)\neq f(X)+f(Y)) &\leq \mathbb{P}_{X,Y}\mathbb{P}(X+Y\notin A, X\notin A, Y\notin A)\\ &\leq \mathbb{P}(X+Y\notin A)+\mathbb{P}(X\notin A)+\mathbb{P}(Y\notin A)\\ &\leq 3\varepsilon.\end{aligned}$$

□

From the perspective of property testing, we want to think of f as a black box; we don't know what is inside, but we can test the value of f on inputs we give it. How can we determine if f is linear? To know for certain, we would need to check every single input. Let us relax our condition.

Suppose either

1. f is linear
2. f is ε far from being linear, i.e. for all linear functions g,

$$\mathbb{P}_{X\sim\mathbb{F}_2^n}(f(X)\neq g(X))\geq\varepsilon.$$

Here, we think of $\mathbb{P}_{X\sim\mathbb{F}_2^n}(f(X)\neq g(X))$ as a notion of distance (this is the Hamming distance between f,g).

BLR proposed the following test:

1. Choose $X,Y\sim\mathbb{F}_2^n$ uniformly at random and independently.
2. Query f on $X,Y,X+Y$.
3. Accept if and only if $f(X)+f(Y)=f(X+Y)$.

If f is linear, then BLR accepts with probability 1. If f is ε-far from being linear, we want to show that $\mathbb{P}(\text{BLR accepts})<1-\varepsilon$. We will prove the contrapositive.

Theorem 1.2. *Suppose $\mathbb{P}(BLR\ accepts)\geq 1-\varepsilon$. Then there exists an S such that $\mathbb{P}(f(X)=\sum_{i\in S}X_i)\geq 1-\varepsilon$.*

Proof. Given $f:\mathbb{F}_2^n\to\mathbb{F}_2$, let $F:\mathbb{F}_2^n\to\mathbb{R}$ be $F(x)=(-1)^{f(x)}$. Then

$$\begin{aligned}1-\varepsilon &\leq \mathbb{P}(\text{BLR accepts})\\ &= \mathbb{P}_{X,Y}(f(X)+f(Y)=f(X+Y))\\ &= \mathbb{P}_{X,Y}(F(X)F(Y)=F(X+Y))\end{aligned}$$

$$= \mathbb{E}_{X,Y}\left[\frac{1 + F(X)F(Y)F(X+Y)}{2}\right]$$
$$= \frac{1}{2} + \frac{1}{2}\mathbb{E}_{X,Y}[F(X)F(Y)F(X+Y)].$$

That is,

$$1 - 2\varepsilon \leq \mathbb{E}_{X,Y}[F(X)F(Y)F(X+Y)].$$
$$= \mathbb{E}_X[F(X)\,\mathbb{E}_Y[F(X)F(X+Y)]]$$

Note that over $\mathbb{F}_2$, $X + Y$ is the same as $X - Y$. This looks like a convolution.

$$= \mathbb{E}_X[F(X) \cdot (F * F)(X)]$$
$$= \langle F, F * F \rangle$$
$$= \sum_S \widehat{F}(S) \cdot \widehat{F * F}(S)$$
$$= \sum_S \widehat{F}(S)^3$$

Parseval's identity tells us that $\sum_S \widehat{F}(S)^2 = 1$. So we should think about this as summing $\widehat{F}(S) \leq \widehat{F}(S)^2$.

$$\leq \max_S(\widehat{F}(S)) \sum_{S \subseteq [n]} \widehat{F}(S)^2$$
$$= \max_S(\widehat{F}(S)).$$

This means that there exists some set S^* such that $\widehat{F}(S^*) \geq 1 - 2\varepsilon$. In other words,

$$\mathbb{E}_X[F(X)\chi_{S^*}(X)] \geq 1 - 2\varepsilon,$$

where the left hand side is

$$\mathbb{P}(F(X) = \chi_{S^*}(X)) - \mathbb{P}(F(X) \neq \chi_{S^*}(X)) = 1 - 2\mathbb{P}(F(X) \neq \chi_{S^*}(X)).$$

So $\mathbb{P}(F(X) \neq \chi_{S^*}(X)) \leq \varepsilon$. □

Remark 1.1. We have shown that if f is ε-far from being linear, then $\mathbb{P}(\text{BLR accepts}) < 1 - \varepsilon$. If we repeat this test $10/\varepsilon$ times with independent randomness, then

$$\mathbb{P}(\text{BLR accepts in all trials}) \leq (1 - \varepsilon)^{10/\varepsilon} \leq \exp(-10).$$

1.5 Local correction of almost linear functions

This allows us to locally correct almost linear functions. Suppose F is ε-close to χ_S. We can define Local Correct(F, x) as

1. Choose $Y \sim \mathbb{F}_2^n$.
2. Query F on $Y, x + Y$.
3. Return $F(x + Y)F(Y)$.

We claim that if F is ε close to χ_S, then for all x,

$$\mathbb{P}_Y(\text{Local Correct}(F, x) = \chi_S(x)) \geq 1 - 2\varepsilon.$$

This is because with probability $\geq 1 - 2\varepsilon$, both $F(Y) = \chi_S(Y)$ and $F(x+Y) = \chi_S(x+Y)$. Then

$$F(Y)F(x + Y) = \chi_S(Y)\chi_S(x + Y) = \chi_S(x).$$