CAS CS 599 B: Mathematical Methods for TCS

Lecturer: Mark Bun | Spring 2022

Lecture Notes 3:

Influence and Stability

Reading.

- O'Donnell, Analysis of Boolean Functions §2.2-2.5

1 Influence

Recall the following notation and definitions. A string $x \in \{-1,1\}^n$, an index $i \in [n]$, and $b \in \{-1,1\}$.

- $x^{\oplus i} = (x_1, x_2, \ldots, x_{i-1}, -x_i, x_{i+1}, \ldots, x_n)$.
- $x^{(i \to b)} = (x_1, x_2, \ldots, x_{i-1}, b, x_{i+1}, \ldots, x_n)$

Definition 1 (Pivotal coordinate)**.** A coordinate $i \in [n]$ is pivotal for f at input x if $f(x) \neq f(x^{\oplus i})$

Definition 2 (Influence)**.** The influence of coordinate i on f is the probability that i is pivotal for a random input:

$$\mathbf{Inf}_i[f] = \Pr_{x \sim \{-1,1\}^n}[f(x) \neq f(x^{\oplus i})]$$

Last time, we talked about interpreting influences in terms of social choice. The i'th influence represents the probability that i cast the swing vote in a uniformly random election.

Influences also have a nice interpretation in terms of the geometry of the Boolean hypercube. You can think of a Boolean function as partitioning the hypercube into two sets, $A = \{x \mid f(x) = -1\}$ and $\overline{A} = \{x \mid f(x) = +1\}$. A *boundary* edge is an edge that crosses between A and $\overline{A}$. The i'th influence measures the fraction of dimension-i edges (i.e., edges of the form $(x, x^{\oplus i})$) that are boundary edges.

While we have defined influences combinatorially, again, there is an elegant connection to the Fourier spectrum.

Theorem 3. *For $f : \{-1,1\}^n \to \{-1,1\}$ and $i \in [n]$,*

$$\mathbf{Inf}_i[f] = \sum_{S \subseteq [n] : i \in S} \hat{f}(S)^2.$$

To prove this result, it will be helpful to take an alternative perspective on influences. For $i \in [n]$, define the operator D_i by

$$(D_i f)(x) = \frac{f(x^{(i \to 1)}) - f(x^{(i \to -1)})}{2} = \begin{cases} \pm 1 & \text{if } i \text{ is pivotal for } f \text{ at } x \\ 0 & \text{otherwise} \end{cases}.$$

Thus, $\mathbf{Inf}_i[f] = \mathbb{E}_{x \sim \{-1,1\}^n}\left[(D_i f)^2(x)\right]$. Note that this definition makes sense if f is any real-valued function, and can be taken to as the definition of influence for real-valued functions:

Definition 4 (Influence for real-valued functions)**.** The influence of coordinate i on $f : \{-1,1\}^n \to \mathbb{R}$ is

$$\mathbf{Inf}_i[f] = \mathop{\mathbb{E}}_{x\sim\{-1,1\}^n}\left[(D_i f)^2(x)\right].$$

You should interpret the operator D_i as the discrete partial derivative of f in the direction of coordinate i. It behaves exactly the way you expect a derivative to. It's linear, meaning $D_i(af+bg) = aD_if + bD_ig$ for functions f, g and scalars a, b. And it does the expected thing to polynomials, e.g.,

$$D_1x_1x_2x_3 = \frac{x_2x_3 - (-x_2x_3)}{2} = x_2x_3 \qquad D_4x_1x_2x_3 = \frac{x_1x_2x_3 - x_1x_2x_3}{2} = 0.$$

In general:

$$(D_i\chi_S)(x) = \begin{cases} \chi_{S\setminus i}(x) & \text{if } i \in S \\ 0 & \text{otherwise.} \end{cases}$$

These properties let us compute the derivative of a function f in terms of the Fourier expansion of f itself.

$$\begin{aligned}
(D_if)(x) &= \left(D_i \sum_{S\subseteq[n]} \hat{f}(S)\chi_S\right)(x) \\
&= \sum_{S\subseteq[n]} \hat{f}(S)(D_i\chi_S)(x) && \text{by linearity} \\
&= \sum_{S\subseteq[n]:i\in S} \hat{f}(S)\chi_{S\setminus\{i\}}(x).
\end{aligned}$$

Proof of Theorem 3.

$$\begin{aligned}
\mathbf{Inf}_i[f] &= \mathop{\mathbb{E}}_{x\sim\{-1,1\}^n}\left[(D_if)^2(x)\right] \\
&= \sum_{S\subseteq[n]} \widehat{D_if}(S) && \text{by Parseval} \\
&= \sum_{S\subseteq[n]:i\in S} \hat{f}(S)^2.
\end{aligned}$$

□

2 Total Influence

Definition 5. The *total influence* of a function $f : \{-1,1\}^n \to \mathbb{R}$ is the sum of the influences:

$$\mathbf{I}[f] = \sum_{i=1}^{n} \mathbf{Inf}_i[f].$$

Geometrically, $\frac{1}{n}\mathbf{I}[f]$ is the fraction of edges of the hypercube that are boundary edges for f.

Example 6. 1. The constant functions ± 1 have minimal total influence 0

2. The parity function has maximal total influence n
3. The dictator functions have total influence 1
4. AND_n has total influence $n \cdot 2^{-n+1}$, which is exponentially small
5. MAJ_n has total influence roughly $\sqrt{2/\pi}\sqrt{n}$. This turns out to be the largest influence possible for a *monotone* function.

Proposition 7.

$$\mathbf{I}[f] = \sum_{S \subseteq [n]} \hat{f}(S)^2 |S|.$$

Proof. Follows from Theorem 3 by summing both sides over i. □

This statement has an appealing probabilistic interpretation. If f is a Boolean function, the Parseval tells us that $\sum_{S \subseteq [n]} \hat{f}(S)^2 = 1$. So the squared Fourier coefficients $\hat{f}(S)^2$ represent a probability distribution over the set of all subsets of $[n]$.

Thus, Proposition 7 tells us that

$$\mathbf{I}[f] = \mathop{\mathbb{E}}_{S \sim \mathcal{S}_f} [|S|]$$

is the expected size of a "spectral sample" from f.

Theorem 8 (Poincaré Inequality)**.** *For any* $f : \{-1,1\}^n \to \mathbb{R}$*, we have* $\mathrm{Var}\,[f] \leq \mathbf{I}[f]$.

Proof. Follows from the Fourier formula for variance and Proposition 7:

$$\mathrm{Var}\,[f] = \sum_{S \neq \emptyset} \hat{f}(S)^2 \leq \sum_{S \subseteq [n]} \hat{f}(S)^2 |S| = \mathbf{I}[f].$$

□

This inequality is simple, but it illustrates a really important principle about the geometry of the Boolean hypercube. For a subset $A \subseteq \{-1,1\}^n$, let ∂A denote its *boundary*, i.e., the set of edges on the hypercube crossing from A to $\overline{A}$. Then we have the following "edge isopermetric inequality" which says that the boundary of any set A is at least as large as A itself.

Theorem 9. *For any set* $A \subseteq \{-1,1\}^n$ *with* $|A| \leq 2^n/2$*,*

$$|\partial A| \geq |A|.$$

Proof. Let $f : \{-1,1\}^n \to \{-1,1\}$ be the indicator for the set A, i.e., $f(x) = -1$ iff $x \in A$. Then

$$\mathrm{Var}\,[f] = \mathbb{E}[f^2] - (\mathbb{E}[f])^2 = 1 - (1 - 2 \cdot 2^{-n}|A|)^2 \geq 2(2^{-n}|A|).$$

On the other hand, our geometric interpretation of total influence tells us that

$$\frac{1}{n}\mathbf{I}[f] = \frac{|\partial A|}{n \cdot 2^{n-1}}.$$

□

This is tight when $|A| = 2^n/2$, as witnessed by the dictator functions, but loose in general. An essentially tight strengthening of this is *Harper's Theorem*, which says that the Hamming balls have the smallest boundaries for all sets of a given size.

3 Noise Stability and Sensitivity

Noise stability measures how robust a function is to small random changes to its input. The social choice interpretation is as follows. As usual, consider an election where all the votes are uniformly random. Say that each vote is randomly misrecorded with probability $\delta > 0$. How likely is it that the election result stays the same?

Let $\rho \in [-1, 1]$. For fixed $x \in \{-1, 1\}^n$, define the distribution $N_\rho(x)$ over $\{-1, 1\}^n$ as follows. To sample a string $y \sim N_\rho(x)$, for each i independently, set

$$y_i = \begin{cases} x_i & \text{with probability } \frac{1}{2} + \frac{\rho}{2} \\ -x_i & \text{with probability } \frac{1}{2} - \frac{\rho}{2}. \end{cases}$$

A pair of samples (x, y) where x is uniformly distributed and $y \sim N_\rho(x)$ is called *ρ-correlated.*

Definition 10. The *noise stability* of a function f is defined by

$$\mathbf{Stab}_\rho[f] = \mathop{\mathbb{E}}_{(x,y)\rho\text{-corr}}[f(x)f(y)] = \mathop{\mathbb{E}}_{x\sim\{-1,1\}^n}\left[f(x) \mathop{\mathbb{E}}_{y\sim N_\rho(x)}[f(y)]\right].$$

Example 11.

1. Constant functions are very stable: $\mathbf{Stab}_\rho[1] = 1$ for all ρ
2. Dictators are stable: $\mathbf{Stab}_\rho[\chi_i] = \mathbb{E}[x_i \cdot (\rho x_i)] = \rho$
3. Parity is very unstable:

$$\mathbf{Stab}_\rho[\chi_{[n]}] = \mathop{\mathbb{E}}_{(x,y)\rho\text{-corr}}\left[\prod_{i=1}^n x_i y_i\right] = \prod_{i=1}^n \mathbb{E}[x_i y_i] = \rho^n.$$

4. There doesn't seem to be a nice formula for the stability of majority. It turns out that

$$\mathbf{Stab}_\rho[\mathrm{MAJ}_n] = \frac{2}{\pi}\arcsin\rho + O\left(\frac{1}{\sqrt{n(1-\rho^2)}}\right).$$

Theorem 12. *For any $f : \{-1, 1\}^n \to \mathbb{R}$ and $\rho \in [-1, 1]$,*

$$\mathbf{Stab}_\rho[f] = \sum_{S\subseteq[n]} \hat{f}(S)^2 \cdot \rho^{|S|}.$$

Like our proof of the Fourier characterization of influences, we'll do this by introducing and analyzing a helpful linear operator.

Definition 13. For $\rho \in [-1, 1]$, the noise operator T_ρ is defined by

$$(T_\rho f)(x) = \mathop{\mathbb{E}}_{y\sim N_\rho(x)}[f(y)].$$

The following lemma determines the effect that the noise operator has on the Fourier spectrum of a function.

Lemma 14.

$$(T_\rho f)(x) = \sum_{S \subseteq [n]} \hat{f}(S) \rho^{|S|} \chi_S(x).$$

That is, each Fourier coefficient $\widehat{T_\rho f}(S)$ is the Fourier coefficient $\hat{f}(S)$ attenuated by a factor that decays exponentially in $|S|$.

Proof. Since T_ρ is a linear operator, it is enough to show that $T_\rho \chi_S = \rho^{|S|} \chi_S$. To do this, we calculate

$$\begin{aligned}
(T_\rho \chi_S)(x) &= \mathop{\mathbb{E}}_{y \sim N_\rho(x)} [\chi_S(y)] \\
&= \prod_{i \in S} \mathop{\mathbb{E}}_{y \sim N_\rho(x)} [y_i] && \text{by independence} \\
&= \prod_{i \in S} (\rho x_i) \\
&= \rho^{|S|} \chi_S(x).
\end{aligned}$$

□

Proof of Theorem 12.

$$\begin{aligned}
\mathbf{Stab}_\rho[f] &= \langle f, T_\rho f \rangle \\
&= \sum_{S \subseteq [n]} \hat{f}(S) \widehat{T_\rho f}(S) && \text{by Plancharel} \\
&= \sum_{S \subseteq [n]} \hat{f}(S)^2 \cdot \rho^{|S|}.
\end{aligned}$$

□