Lecture 3: Basic concepts in Fourier analysis: influence and noise stability

Rajat Mittal

IIT Kanpur

The content of this chapter is inspired from the second chapter of the book *Analysis of Boolean functions* by Ryan OĎonell.

We have already mentioned that Boolean functions are useful in many diverse areas of computer science: complexity theory, social choice theory, cryptography and error correcting codes to take some examples. In many of these subjects, it is vital to consider how a function behaves when it is perturbed slightly. Fourier analysis is a very helpful tool and allows us to quantify this behavior.

In this lecture note, we will see two such perturbations. First, we will notice how the function value changes when we flip a particular coordinate of the input. This allows us to quantify the *impact* of a variable for the function. This is known as *influence*, we will define it in the next section and see an application. In the later half, the concept of *noise stability* will be looked at, where each coordinate is changed with some small probability.

1 Influence of a variable

Remember that $f : \{-1,1\}^n \to \{-1,1\}$ is a Boolean function. Given $x \in \{-1,1\}^n$, let $x^{\oplus i} \in \{-1,1\}^n$ denote the string where i-th bit is flipped. To take an example, $010^{\oplus 2}$ is the string 000. In a loose sense, a variable i has a lot of impact in computing a function, if $f(x) \neq f(x^{oplusi})$.

This motivates us to define *Influence* of a variable i,

$$Inf_i(f) = \Pr_x(f(x) \neq f(x^{\oplus i})).$$

Notice that the probability is over uniform distribution over the inputs x. In other words, the influence of a variable i is the number of x's such that the function value changes when i-th bit is flipped (divided by 2^n to normalize).

Exercise 1. What is the maximum possible value of $Inf_i(f)$.

Notice that influence is a way to capture the impact of a variable. One way to see it is: if the influence of a variable is 0, then the function does not depend on that variable (it could be safely skipped from the input).

Exercise 2. Prove the above claim.

Though, this quantification is not perfect. Consider the *dictator function* $f(x) = x_i$, we would assume that the impact of i on this function is very high. Still, both dictator on i and PARITY function have the same Inf_i.

The total influence of a function is just the sum of influence of all variables.

$$Inf(f) = \sum_{i \in [n]} Inf_i(f).$$

Recall that $\hat{f}(S)$ denote the Fourier coefficient of f on set S. Our first task is to represent Inf in terms of the Fourier coefficients of f.

$$Inf_i(f) = \frac{1}{2^n}\sum_x \left(\frac{\left(f(x) - f(x^{\oplus i})\right)}{2}\right)^2.$$

Convince yourself that every term of the summation is 1 if and only if $f(x)$ is different from $f(x^{\oplus i})$. The key step is to convert the Fourier transfrom of $f(x)$ into the Fourier transform of $f(x^{\oplus i})$. This will require switching the sign of any monomial which has i in it.

$$Inf_i(f) = \frac{1}{2^n}\sum_x \left(\sum_{S:i\in S} \hat{f}(S)\chi_S(x)\right)^2.$$

Let $g_i = \sum_{S:i\in S} \hat{f}(S)\chi_S(x)$, using the generalization of Parseval's identity,

$$Inf_i(f) = \langle g_i | g_i\rangle = \sum_{S:i\in S} \hat{f}(S)^2.$$

Exercise 3. We have already seen that Fourier coefficients at the same level are equal for a symmetric function. It is easy to show now that influence is same for every variable.

Summing up, we get the expression for total influence of the function.

$$Inf(f) = \sum_S |S|\hat{f}(S)^2.$$

Observe that, if the Fourier coefficients are concentrated on higher degree terms, then the total influence will be higher.

1.1 Derivative of a Boolean function

The difference in function value when we switch one input bit, $f(x) - f(x^{\oplus i})$, is of such importance that we define a function for that. We will use a new notation here, $x^{i,1}$ means the i-th bit is set to 1. Formally, given a function $f : \{-1, 1\}^n \to \mathbb{R}$, the derivative $D_i(f)$ (with respect to i-th bit) is defined as,

$$D_i(f) = \frac{f(x^{i,1}) - f(x^{i,-1})}{2}.$$

Notice that the range of the function is generalized to real numbers. You might wonder why we have not taken the more symmetric version of difference, $f(x) - f(x^{\oplus i})$, for defining the derivative. The reason is, we are viewing the range as real numbers, this definition allows us to write very similar formula for derivatives as real polynomials. The function $\frac{f(x)-f(x^{\oplus i})}{2}$ is known as the *Laplacian operator* and is denoted by $L_i(f)$, it is also useful in many contexts.

Exercise 4. Prove that $D_i(f + g) = D_i(f) + D_i(g)$, where $f + g(x) = f(x) + g(x)$.

You can also verify that $D_i(\chi_S) = \chi_{S/\{i\}}$ when S contains i and 0 otherwise. This gives the formula for $D_i(f)$ (similar to real polynomials),

$$D_i(f) = \sum_{S:i\in S} \hat{f}(S)\chi_{S/\{i\}}.$$

In other words, f can be written as,

$$f(x) = x_i D_i(f) + \sum_{S:i\notin S} \hat{f}(S)\chi_S.$$

The second term can be viewed as $\frac{f(x^{i,1})+f(x^{i,-1})}{2}$, which is the expectation over x_i. Define $E_i(f)(x) := \mathbf{E}_{x_i}[f(x)]$ (here everything except x_i is fixed). Then,

$$f = x_i D_i(f) + E_i(f).$$

The functions D_i, E_i does not depend upon i, this allows us to decompose f. This decomposition of f is pretty useful in proving properties of Boolean function using Induction on n.

Let us see some other benefits of $D_i(f)$. We can write Inf_i in terms of $D_i(f)$,

$$\mathit{Inf}_i(f) = \mathbf{E}_x[D_i(f(x))^2].$$

A function is called *monotone* if $f(x) \leq f(y)$ whenever $x \leq y$ coordinate-wise. The functions AND, OR and MAJORITY are monotone, but PARITY is not.

Exercise 5. Is Add_m a monotone function?

For the monotone function $D_i(f)$ is 1 if the function value switches by changing the i-th variable, otherwise it is 0. In other words, for a monotone f,

$$\mathit{Inf}_i(f) = \mathbf{E}_x[D_i(f(x))].$$

Noticing that the $\mathbf{E}_x$ is 0 for any χ_S except when S is empty.

$$\mathit{Inf}_i(f) = \hat{f}(\{i\}).$$

This gives the expression for total influence of a monotone function,

$$\mathit{Inf}(f) = \sum_i \hat{f}(\{i\}).$$

1.2 Degree of a Boolean function

Recall that we defined the degree of a Boolean function $f : \{-1,1\}^n \to \{-1,1\}$ as the size of the biggest subset S such that $\hat{f}(S) \neq 0$. This conforms with our understanding of degree of polynomials over reals.

Exercise 6. What is the degree of OR function?

You can easily verify that the degree of most of the functions we know (AND, OR, PARITY) is exactly equal to n.

Exercise 7. Can you find a Boolean function whose degree is not n?

After some thought, you can create such functions easily. The n-th degree coefficient only depends on the correlation with Parity. If the function is *balanced* with respect to parity (agrees and disagrees with parity at equal number of places), then it is of degree less than n.

On the other hand, can you find Boolean functions with degree 0 or 1? Sure, constant functions have degree 0 and dictator functions ($f(x) = x_i$ for some i) have degree 1. Though, this seems unfair. Dictator functions only depend on just 1 variable and can be thought of as a function over just 1 variable.

A more natural question is, what is the minimum degree of a Boolean function when it depends on all n variables?

Exercise 8. If the function is not Boolean (the range is $\mathbb{R}$), construct a function with degree 1 which depends on all n variables.

What does it mean for a function to depend on all n variables. Armed with influence, a function depend on all n variables iff $\mathit{Inf}_i(f) \neq 0$ for all i. We are now going to see the result of Nisan and Szegedy, where they show that any such function will have degree $\Omega(\log(n))$. Denote degree of a function f by $\deg(f)$.

Exercise 9. Can you show that $\deg(f) > 1$ if $n > 1$.

There are two main ideas behind the proof. First, any small degree function is either identically zero or is non-zero with high probability. Secondly, degree of a function is bigger than the influence of the function. The second fact is easy to prove from the formula for Influence in terms of Fourier coefficients. We know,

$$Inf(f) = \sum_S |S|\hat{f}(S)^2 \geq \deg(f)\left(\sum_S \hat{f}(S)^2\right).$$

Exercise 10. Prove $\deg(f) \geq Inf(f)$ by using Parseval's identity.

Let us fomally state and prove the first idea.

Lemma 1. *Let $f : \{-1,1\}^n \to \mathbb{R}$ be a function with degree at most d. If f is not identically zero, then*

$$\Pr_x[f(x) \neq 0] \geq 2^{-d}.$$

Proof. We will prove the above statement by Induction on n.

Exercise 11. Prove the base case.

Without loss of generality, let x_n be one of the variables present in the Fourier expansion of f (it depends on x_n). Then $f = x_n D_n(f) + E_n(f)$, and x_n takes two possible values 1 and -1. The function will become $D_n(f) \pm E_n(f)$ depending on whether $x_n = \pm 1$.

Case 1: The function is not identically zero either when $x_n = 1$ or $x_n = -1$. In both these substitutions, restricted f will be a function of degree at most d and will be non-zero. Using induction, the probability that $D_n(f) \pm E_n(f)$ is nonzero is at least 2^{-d}. The input has $x_n = 1$ with half the probability. So,

$$\Pr[f(x) \neq 0] \geq 1/2(2^{-d}) + 1/2(2^{-d}) = 2^{-d}.$$

Case 2: One of $D_n(f) \pm E_n(f)$ is identically zero. Let us assume $D_n(f) - E_n(f) = 0$, other case will be similar. Since there is a unique Fourier representation for every function, $D_n(f)$ and $E_n(f)$ have the same Fourier representation. In other words $f = (x_n - 1)D_n(f)$. The degree of $D_n(f)$ is at most $d-1$ and hence it is nonzero with probability at least 2^{-d+1}. So,

$$\Pr[f(x) \neq 0] \geq 1/2(2^{-d+1}) = 2^{-d}.$$

□

We are ready to prove the result of Nisan and Szegedy about the minimum degree required to represent a Boolean function.

Theorem 1 (Nisan and Szegedy 94). *Let f be a Boolean function such that $\inf_i(f) \neq 0$ for all $i \in [n]$. Then, $\deg(f) \geq \frac{1}{2}\log(n)$.*

Proof. We have already done most of the hard work. Suppose the function has degree d. We know that $Inf_i(f) = \Pr(D_i(f) \neq 0)$. Since $D_i(f)$ is a polynomial of degree at most $d-1$, by Lemma 1, $Inf_i(f) \geq 2^{-d+1}$.

There are n variables and each influence is non-zero, this implies $Inf(f) \geq n2^{-d+1}$. Using the second idea, influence should be less than degree.

$$n2^{-d+1} \leq d$$

Taking log on both sides,

$$\log(n) - d + 1 \leq \log(d) \Rightarrow \quad d \geq \log(n) - \log(d).$$

If degree is more than $\log(n)$, there is nothing to prove. Otherwise $\log d \leq \log(\log(n))$. Giving us,

$$d \geq \log n - \log(\log(n)) \geq \frac{1}{2}\log(n).$$

□

We are able to prove that any function (which depends on all variables) will have degree at least $\log(n)$. Is this the tightest bound possible? It might be that all Boolean functions require degree at least $n/2$.

You will prove in the assignment that ADDR_m has degree $O(\log(n)$, where $n = m + 2^m$ is the input size of the function. This shows the tightness of the result proved.

The problem of minimum degree for a general Boolean function is solved. Notice that ADDR_m is a non-symmetric function. Can we say that degree of any symmetric Boolean function (non-constant) is very high? You might be surprised but it is easy to prove that any symmetric Boolean function has degree at least $n/2$. It was proved by Gathen and Roche that degree of a symmetric function is almost full, it is more than $n - O(n^{.525})$. They actually conjectured that it is $n - O(1)$. This simple problem is still open. The best gap known is $n - 3$.

2 Social choice theory

In social choice theory, we try to come up with rules to aggregate different opinions and come up with a collective output. One of the simplest example is a voting rule, where preference of voters is collected and a person (say) is selected. MAJORITY is the most natural voting rule, but we can also take MAJORITY of MAJORITY as a voting rule (Prime minister election in India). In general, any Boolean function $f : \{-1,1\}^n \to \{-1,1\}$ can be viewed as a voting rule for an election with n voters and 2 candidates. The input and output bits can also be viewed as a YES/NO preference, instead of two candidates.

We have already seen voting rules like MAJORITY, AND (unanimous) and dictator (not very democratic). A few more examples are:

- Weighted majority: For a weight vector $w \in \mathbb{R}^n$, the function $\mathrm{MAJORITY}_w(x)$ is the sign of $\sum_i w_i x_i$. It can be used when some voters' preferences are more important than others.

  Exercise 12. What is the weight vector for MAJORITY?

- Tribes: This is the function $\mathrm{OR} \circ \mathrm{AND}$. Think of votes coming from different tribes, the option is selected if any tribe unanimously selects the option. For example, a candidate is hired in the CSE department iff a subarea (ML, Theory, Systems) unanimously selects the candidate.

What makes a voting rule more appropriate than others. Some of the properties of a good voting rule might be,

- Unanimous: If all voters choose $-1, -1, \cdots, -1$ or $1, 1, \cdots, 1$, then -1 or 1 is the answer respectively.
- Monotone: If more voters choose a particular candidate (say -1) than the outcome should not switch from -1 to 1.

  Exercise 13. Convince yourself that this corresponds to the function being monotone.

- Odd: If we reverse the preferences, the outcome should be reversed. In other words, $f(-x) = -f(x)$.
- Symmetric: All voters are treated equally. One way to ensure this is to take f to be symmetric.

Making sure that the function is symmetric, is one way to claim that all voters are treated equally. There are weaker versions too, for instance, look at the definition of *transitive symmetric functions* in Ex. 29.

Exercise 14. Which of these good properties, of a voting rule, are not satisfied by MAJORITY or OR or AND.

Intuitively MAJORITY is the most natural choice for a voting rule. Let us see one way in which we can formalize this intuition. We will now show that majority maximizes number of agreement between winner and votes in its favor.

Let $f : \{-1,1\}^n \to \{-1,1\}$ be a voting rule. A good voting rule should be the one where voter's preferences agrees with the outcome of the election. For an input x, let w_x be the number of votes which agree with the function value $f(x)$. Ideally, we would like to maximize the expectation $\mathbf{E}[w_x]$.

If w_x votes agree with $f(x)$, then $n - w_x$ votes will have value $-f(x)$. This shows that $f(x)(\sum_i x_i) = 2w_x - n$. So,

$$\mathbf{E}[w_x] = \mathbf{E}[(n + f(x)(\sum_i x_i))/2] = n/2 + 1/2\mathbf{E}[f(x)(\sum_i x_i)] = n/2 + 1/2 \sum_i \hat{f}(\{i\}).$$

In other words, expected numbre of agreements only depend on $\sum_i \hat{f}(\{i\})$. Notice that this quantity is the total influence for a monotone function.

Let us see when this quantity is maximized.

$$\sum_i \hat{f}(\{i\}) = \mathbf{E}[f(x)(\sum_i x_i)] \leq \mathbf{E}[|(|\sum_i x_i)].$$

The last inequality is tight if and only if $f(x)$ has the same value as the sign of $\sum_i x_i$. In other words, MAJORITY maximizes the expected number of agreements among all Boolean functions. Again showing that it is a good choice for a voting rule.

2.1 Noise stability

In some scenarios, there might be a small possibility of an error while recording the vote. One good property of a voting rule might be, it is robust under such errors. In other words, the function value should not change (with high probability) under such small noise. Let us formalize this idea.

Let $x \in \{-1,1\}^n$ be a string or preferences of voters. The random variable $y \in \{-1,1\}^n$ is ρ-correlated with x if independently for all i,

- $x_i = y_i$ with proability $\frac{1}{2}(1+\rho)$.
- $x_i = -y_i$ with proability $\frac{1}{2}(1-\rho)$.

Fix our distribution on pairs (x, y), where x is chosen uniformly at random and y is ρ-correlated to x. For this section, let (x, y) denote the random variable chosen according to this distribution.

Exercise 15. Show that the probability of pair (x, y) is $\frac{1}{2^{2n}}(1+\rho)^{n-|x-y|)}(1-\rho)^{|x-y|}$.

Intuitively for a robust voting rule, when x, y are ρ-correlated, the function value should not change. We define *noise stability* to be,

$$Stab_\rho(f) = \mathbf{E}_{(x,y)}[f(x)f(y)].$$

Notice that the expectation is taken over all (x, y) pairs where these pairs are distributed according to the probability distribution defined above. We will skip the subscript on expectation when this probability distribution is used.

As before, our task would be to find a *Fourier* expression for noise stability. Let us compute the noise stability of partial parity.

$$Stab_\rho(\chi_S) = \mathbf{E}[\chi_S(x)\chi_S(y)].$$

Using independence of different coordinates i,

$$Stab_\rho(\chi_S) = \Pi_i \mathbf{E}[x_i y_i].$$

Exercise 16. Show that $\mathbf{E}[x_i y_i] = \rho$.

Using the exercise, we get,

$$Stab_\rho(\chi_S) = \rho^{|S|}.$$

The idea (like influence) would be to get the expression of noise stability by using linearity of expectation. You should try it, we will also need $\mathbf{E}[\chi_S(x)\chi_T(y)]$. A similar calculation can be done,

$$\mathbf{E}[\chi_S(x)\chi_T(y)] = \mathbf{E}[\chi_{S/T}(x)\chi_{T/S}(y)\chi_{S\cap T}(xy)].$$

Using independence of different coordinates i,

$$\mathbf{E}[\chi_S(x)\chi_T(y)] = \mathbf{E}[\chi_{S/T}(x)]\mathbf{E}[\chi_{T/S}(y)]\mathbf{E}[\chi_{S\cap T}(xy)].$$

Exercise 17. What is $\mathbf{E}[x_i]$ and $\mathbf{E}[y_i]$?

Since S, T are distinct,

$$\mathbf{E}[\chi_S(x)\chi_T(y)] = 0.$$

This allows us to compute the expression for noise stability in terms of Fourier coefficients,

$$Stab_\rho(f) = \sum_S \rho^{|S|} \hat{f}(S)^2 \tag{1}$$

You will prove this formula in the assignment.

2.2 Arrow's theorem

Our task in this section is to prove the famous theorem of Arrow (1950) (a sightly weaker version) in social choice theory using the concept of noise stability. This Fourier analytic proof was given by Gil Kalai (2002). The content of this section are taken from the book by Ryan O'Donell.

Let us look at the problem first. We saw that MAJORITY is a very good choice as a voting rule for a 2-candidate election (last section). What happens when there are 3-candidates and n voters?

Exercise 18. Can you think of a aggregating strategy for 3 candidates given the voter prefrence?

Condercet came up with a natural way to generalize the voting rule (say some function f). We can hold pairwise elections between the three candidates, giving rise to 3 separate elections. Let $x \in \{-1,1\}^n$ be the preference of voters for the first 2 candidates, then $f(x)$ will be the winner for the first election; similarly, let y, z be the preference for other two elections and $f(y), f(z)$ be the outcome.

Let the three candidates be a, b, c. If we arrange x, y, z as columns, then every row will be a preference of a particular voter.

	x (a vs b)	y (b vs c)	z (c vs a)
v_1	1	-1	1
v_2	-1	1	-1
$\vdots$	$\vdots$	$\vdots$	$\vdots$
v_n	1	1	-1
winner	$f(x)$	$f(y)$	$f(z)$

To be consistent, every voter preference (row) should satisfy NAE_3 (not all equal) function. For example, a row of $1, 1, 1$ will reflect that the voter prefers a over b, b over c and c over a.

Exercise 19. Convince yourself that all satisfying inputs of NAE_3 function (6 out of 8) correspond to a valid voter preference.

Our task is to give a single winner. From the above exercise, if $f(x), f(y), f(z)$ is a satisfying input of NAE_3, then the candidates can be arranged in a sorted order. Such a winner is called a *Condercet* winner. Does such a winner always need to exist?

Exercise 20. Let $n = 3$ and f be the MAJORITY function; using symmetry, construct valid voter preferences where there is no Condercet winner. In other words, $f(x), f(y), f(z)$ are all equal.

This shows that MAJORITY might be a bad choice for a 3-candidate election. Though, majority is not alone. Arrow's theorem states that if you always want a Condercet winner, f should be a dictator function or a negated dictator function.

The Fourier analytic proof of Arrow's theorem calculates the probability of a Condercet winner, where the voter preference is uniformly distributed over all satisfying inputs of the NAE_3 function. We will show that this probability is 1 iff f is a dictator function.

Since NAE_3 is a $0,1$ function (or a random variable, if the inputs are random variable), the probability of having a Condercet winner for a function f is $\mathbf{E}_{(x,y,z)}[\mathrm{NAE}_3(f(x), f(y), f(z)]$. Here, each row of (x, y, z) is picked uniformly (and independently) from the satisfying inputs of NAE_3 function. Let us skip the subscript of expectation for the rest of the section.

Exercise 21. Why is the expected value not the constant term of NAE_3 function?

Let us expand the NAE_3 function in its Fourier representation and use linearity of expectation.

Exercise 22. What is the Fourier representation of NAE_3?

$$\mathbf{E}[\mathrm{NAE}_3(f(x), f(y), f(z)] = \frac{3}{4} - \frac{1}{4}\mathbf{E}[f(x)f(y)] - \frac{1}{4}\mathbf{E}[f(x)f(z)] - \frac{1}{4}\mathbf{E}[f(y)f(z)].$$

We want to calculate $E_{(x,y,z)}[f(x)f(y)]$ (other two terms will be same using symmetry). Looking at the inputs of NAE_3,

- $\mathbf{E}[x_i] = 0$, and similarly $\mathbf{E}[y_i] = \mathbf{E}[z_i] = 0$,
- the joint distribution of x and y is independent (over i),
- x_i, y_i is same with probability $1/3$ and different with probability $2/3$.

Exercise 23. Show that $\mathbf{E}[x_i y_i] = -1/3$.

In other words, we can assume that x, y are $-1/3$-correlated. So, $\mathbf{E}[f(x)f(y)] = Stab_{-1/3}(f)$. This gives the probabilty of Condercet winner to be,

$$\frac{3}{4} - \frac{3}{4} Stab_{-1/3}(f).$$

To prove the Arrow's theorem, we only need to check if this probability can be 1. This only happens when $Stab_{-1/3} = -1/3$. Remember the Fourier expression for noise stability, $Stab_{-1/3}(f) = \sum_S (-1/3)^{|S|} \hat{f}(S)^2$.

Exercise 24. Prove that $Stab_{-1/3}(f) = -1/3$ iff all the Fourier weight is on degree 1 monomials.

You will prove in the assignment below that all Fourier weight is on degree 1 monomial iff the function is a dictator or a negated dictator, proving Arrow's theorem.

3 Assignment

Exercise 25. What is the maximum possible *Inf* for a function. Is it possible to achieve this maximum value? What functions achieve this value?

Exercise 26. What is the influence of target and address variables in addressing function?

Exercise 27. Express influence of i-th variable as an expectation of D_i, and derive the expression for Influence in terms of Fourier coefficients.

Exercise 28. Let σ be a permutation on input bits. A function is invariant under σ if $f(x) = f(\sigma(x))$. If a function is invariant under σ, show that

$$\hat{f}(S) = \hat{f}(\sigma(S)).$$

Exercise 29. A group of permutations G is called *transitive*, if given a pair i, j there exists a permutation $\sigma \in G$ such that $\sigma(i) = j$. A function is called transitive symmetric if the function is invariant under the action of some transitive group G. Notice that a symmetric function is transitive symmetric under the complete group S_n.

Show that if a function is monotone and transitive symmetric, then $Inf_i(f)$ is less than $1/\sqrt{n}$ for all i.

Exercise 30. Show that there exist infinite n's such that MAJORITY_n has full degree.

Exercise 31. Show that the degree of Add_m is $m + 1$.

Exercise 32. Which of the good properties of a voting rule are not satisfied by weighted majority or tribes.

Exercise 33. Prove that the only Boolean function of degree 1 are dictator functions, negated dictator or constant functions.

Exercise 34. Prove Eq. 1.