CS 6817: Special Topics in Complexity Theory	**Spring 2025**

Lecture 7: Feb 11, 2025

Lecturer: Eshan Chattopadhyay *Scribe: Mohit Gurumukhani*

Today, we will study noise stability and its applications.

1 Recap from Previous Lecture

Recall the definition of noisy distribution:

Definition 1.1 (Noisy distribution). *Given $\rho \in [-1,1]$ and $x \in \{-1,1\}^n$, define the* noisy distribution *with noise parameter ρ as $y \sim N_\rho(x)$. To sample y, independently for each coordinate $1 \le i \le n$, let $y_i = x_i$ with probability $\frac{1}{2} + \frac{\rho}{2}$ and $y_i = -x_i$ with probability $\frac{1}{2} - \frac{\rho}{2}$.*

Thus, if $\rho = 1$, then y always equals x, when $\rho = 0$, y is a truly random string and if $y = -1$, it always equals $-x$.

Using standard concentration bounds, we see that:

Claim 1.2. *For any $\rho \in [-1,1]$, and $x \in \{-1,1\}^n$: with high probability $\Delta(x, N_\rho(x))$ (the Hamming distance) is $\left(\frac{1}{2} - \frac{\rho}{2}\right) n \pm O(\sqrt{n})$.*

Recall the definition of the noise operator.

Definition 1.3 (Noise Operator). *For $\rho \in [-1,1], f : \{-1,1\}^n \to \mathbb{R}$, define the noise operator of f, $T_\rho(f) : \{-1,1\}^n \to \mathbb{R}$ as follows:*

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

We also derived the following Fourier representation for $T_\rho f(x)$:

$$T_\rho f(x) = \sum_{S \subset [n]} \hat{f}(S) \mathbb{E}_{y \sim N_\rho(x)}[\chi_S(y)] = \sum \rho^{|S|} \hat{f}(S) \chi_S(x).$$

Using the noise operator, we defined the noise stability of a function as follows:

Definition 1.4 (Noise Stability). *For $f : \{-1,1\}^n \to \mathbb{R}$ and $\rho \in [-1,1]$, we define noise stability of f as*

$$NS_\rho(f) = \langle f, T_\rho f \rangle.$$

In the previous lecture, we saw:

Claim 1.5. *If $f : \{-1,1\}^n \to \{-1,1\}$, then*

$$NS_\rho(f) = 2 \Pr_{x \sim \{-1,1\}^n, y \sim N_\rho(x)}[f(x) = f(y)] - 1.$$

Using Plancherel's theorem and definition of noise stability, we obtained the following value for noise stability of f in terms of Fourier expansion of f.

Claim 1.6.

$$NS_\rho(f) = \sum \rho^{|S|} \hat{f}(S)^2.$$

2 Properties of Noise Stability

Question 1. *How does NS_ρ change with ρ?*

To compute this, we take derivative of $NS_\rho(f)$ with respect to ρ:

$$\frac{d}{d\rho} NS_\rho(f) = \sum_{S \subset [n]} |S| \rho^{|S|-1} \hat{f}(S)^2.$$

We see that at $\rho = 1$, this derivative equals

$$\frac{d}{d\rho} NS_\rho(f)|_{\rho=1} = \sum |S| \hat{f}(S)^2$$

We recognize from previous lectures that the last expression equals the the total influence of f. Hence,

$$\frac{d}{d\rho} NS_\rho(f)|_{\rho=1} = I(f)$$

One way to interpret this is to recall that at $\rho = 1$, y equals x. So, the rate at which y deviates slightly away from x, on average, is exactly captured by the total influence of f.

Question 2. *Which function f maximizes NS_ρ for a given ρ?*

Since $NS_\rho(f) = 2\Pr_{x \sim \{-1,1\}^n, y \sim N_\rho(x)}[f(x) = f(y)] - 1$, we see that this value can be at most 1. We easily see that for constant functions $f \equiv 1$ or $f \equiv -1$, the noise stability is exactly 1 for all values ρ.

This is a bit unsatisfactory conclusion. So, we instead slightly modify our question to ask:

Question 3. *Which* balanced *function f maximizes NS_ρ for a given ρ?*

To help answer this question, lets first take a detour and define and study spectral distribution.

2.1 Spectral Distribution

Recall that if f is boolean valued, i.e., $f : \{-1,1\}^n \to \{-1,1\}$, then

$$\sum_{S \subset [n]} \hat{f}(S)^2 = 1.$$

We can naturally define the spectral distribution $\mathcal{S}$ associated with f as follows:

Definition 2.1 (Spectral Distribution)**.** *For $f : \{-1,1\}^n \to \{-1,1\}$, define the* spectral distribution $\mathcal{S}_f \sim 2^{[n]}$ *associated with f as:*

$$\forall T \subset [n], \Pr[\mathcal{S}_f = T] = \hat{f}(T)^2.$$

We will often just write $\mathcal{S}$ instead of $\mathcal{S}_f$ for notational convenience.

We can find expressions for various properties associated with f in terms of the spectral distribution of f. For instance:

Claim 2.2. *For $f : \{-1,1\}^n \to \{-1,1\}$, the expected size of a set sampled from $\mathcal{S}_f$ is the total influence of f. Formally:*

$$\mathbb{E}_{T \sim \mathcal{S}_f}[|T|] = I(f).$$

Proof. Both sides equal $\sum_{T\subset[n]} |T|\hat{f}(T)^2$. □

Here's another example:

Claim 2.3. *For $f : \{-1,1\}^n \to \{-1,1\}$, for a set T sampled from $\mathcal{S}_f$, the expected value of $\rho^{|T|}$ equals $NS_\rho(f)$. Formally:*

$$\mathbb{E}_{T\sim\mathcal{S}_f}[\rho^{|T|}] = NS_\rho(f).$$

Proof. Both sides equal $\sum_{T\subset[n]} \rho^{|T|}\hat{f}(T)^2$. □

2.2 Maximizing Noise Stability for Balanced Functions

Recall from above that

$$NS_\rho(f) = \mathbb{E}_{T\sim\mathcal{S}_f}[\rho^{|T|}].$$

Since f is balanced, we know that $\mathbb{E}[f] = 0$. So,

$$\Pr_{T\sim\mathcal{S}_f}[T = \emptyset] = \hat{f}(\emptyset)^2 = \mathbb{E}[f]^2 = 0$$

This implies

$$\mathbb{E}_{T\sim\mathcal{S}_f}[\rho^{|T|}] \le \rho,$$

since $|T| \ge 1$. Hence, for a balanced function, $NS_\rho(f) \le \rho$. We also see that equality is achieved above iff the balanced function has all its mass on level 1, i.e., $\hat{f}(T) \ne 0 \iff |T| = 1$. This implies f must be of the form

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

where $a_1, \dots, a_n \in \mathbb{R}$.

Since we are interested in boolean valued f, it is not obvious what sets of values a_i make f boolean. We claim that this happens whenever f is the dictator function or the negation of the dictator function. Formally:

Theorem 2.4. *Let $f : \{-1,1\}^n \to \{-1,1\}$ be such that there exist $a_1, \dots, a_n \in \mathbb{R}$ so that $f(x) = \sum_{i=1}^n a_i x_i$. Then f must be the dictator function or the negation of the dictator function, i.e. $f = \pm x_j$ for some $j \in [n]$.*

Proof. We see that $f(1^n) = \sum_i a_i \in \{-1,1\}$. Then consider, $f(1^{n-1}(-1)) = f(1^n) - 2a_n \in \{-1,1\}$. This implies that $a_n \in \{-1,0,1\}$. By symmetry, this holds for all a_j for $j \in [n]$. Let $y \in \{-1,1\}^n$ be such that for $i \in [n]$, $y_i = \text{sign}(a_i)$. Then, $f(y) = \sum_i |a_i| \in \{-1,1\}$. This implies there exists exactly one $j \in [n]$ such that $a_j \in \{-1,1\}$ and for $k \in [n] \setminus \{j\}$, $a_k = 0$ as desired. □

3 Condorcet Elections and Arrow's Theorem

Consider the setting of 3 candidates say a, b, c (can also consider more candidates) and n voters. Each voter ranks their preference among a, b, c such as say $b > a > c$ or $c > b > a$. Given the votes, we generate 3 strings: $x, y, z \in \{-1,1\}^n$ that encode whether voters prefer a to b or b to c or c to a respectively. We then use a voting rule $f : \{-1,1\}^n \to \{-1,1\}$ on each each of these strings x, y, z. If there exists a candidate that defeated both other candidates, they are declared the *Condorcet winner* and this election is called *Condorcet election*.

Observe that there is a possibility that a defeats b, b defeats c and c defeats a, and there is no Condorcet winner. We are interested in characterizing which kind of voting rules $f : \{-1,1\}^n \to \{-1,1\}$ can ensure that no matter how voters rank the candidates, there always exists a Condorcet winner.

Here's an example to help clarify how these strings are encoded and the function is applied:

Example 3.1. *Say the three candidates are a, b, c. Say there are three voters who vote as follows:*

1. $a > b > c$
2. $c > a > b$
3. $b > c > a$

Then, since string x encodes whether a voter prefers a to b (1 if they prefer a, -1 if they prefer b)), we get that

$$x = 1, 1, -1.$$

Similarly, string y encodes whether a voter prefers b to c and so we get that

$$y = 1, -1, 1.$$

Lastly, string z encodes whether a voter prefers c to a, we get that

$$z = -1, 1, 1.$$

Suppose that f, the voting rule, is the Majority function. Then,

$$f(x) = f(1, 1, -1) = 1.$$

This means, a has won the pairwise election between a and b. Similarly,

$$f(y) = f(1, -1, 1) = 1.$$

This means b has won the pairwise election between b and c. Lastly,

$$f(z) = f(-1, 1, 1) = 1.$$

This means c has won the pairwise election between c and a.

Hence, we see that there is no outright winner amongst a, b, c that defeated both other candidates in the pairwise election. So, no Condorcet winner exists. This shows that the Majority function cannot guarantee that there always exists a Condorcet winner.

So, we formally ask the question whether a voting rule can always guarantee a Condorcet winner.

Question 4. *Suppose $f : \{-1,1\}^n$ is balanced and unanimous (so $f(1^n) = 1, f(-1^n) = -1$). Which functions f can be used as a voting rule in a 3-party Condorcet election such that there is always a Condorcet winner?*

This question was studied in Social Choice theory and is known as Arrow's theorem:

Theorem 3.2 (Arrow's Theorem)**.** *The only functions f that satisfy the above property are the dictator functions.*

There are many proofs of Arrow's theorem. We will follow a proof provided by Gil Kalai. First off, we see that the dictator functions indeed guarantee that winner always exists since each voter provides a total ordering. Hence, we focus on showing that no other function can guarantee this. We will not finish off the proof in this lecture but here is a good start.

We begin by defining a useful function, the Not-All-Equals function.

Definition 3.3. *Define the Not-All-Equals function on* 3 *bits* $\mathsf{NAE}_3 : \{-1,1\}^3 \to \{-1,1\}$ *as follows:*

$$\mathsf{NAE}_3(w) = \begin{cases} 0 & w \in \{(1,1,1),(-1,-1,-1)\} \\ 1 & \textit{otherwise} \end{cases}$$

Let $x, y, z \in \{-1,1\}^n$ be the strings encoding the pairwise preferences of the voters. We then observe that $\mathsf{NAE}_3(f(x), f(y), f(z)) = 1$ iff the election has a Condercet winner. Indeed, the only way there is not a Condorcet winner is if $f(x) = f(y) = f(z)$, i.e. the string is $(1,1,1)$ and $(-1,-1,-1)$ and the NAE_3 function will output 0 in that case and 1 otherwise.

We ask the question what is the Fourier expansion of the function NAE_3? We easily compute that

$$\mathsf{NAE}_3(x) = \frac{3}{4} - \frac{1}{4} \sum_{1 \le i < j \le 3} x_i x_j.$$