Lecture 13: Influences, Coalitions, and the Tribes function

In this lecture we introduce the ideas of influences and coalitions. Much of what is covered predates Fourier analysis, and no Fourier analysis is used in this lecture. However, we state a theorem due to Kahn, Kalai, and Linial, one of the first major applications of Fourier analysis to Boolean functions.

1 1-round collective coin flipping

Imagine the following scenario: You have n distributed processors. You want to run a distributed randomized algorithm, and you want the processors to collectively agree on a random bit. However, there is worry that some of the processors are faulty, or that the processors have incentive to cheat. How do you get the processors to agree on some random bit, where you can guarantee that the bit is really random?

For this problem, Ben-Or and Linial defined a 1-round collectiove coin flipping scheme. Such a scheme is nothing more than a Boolean function f. The input to the function consists of one random bit from each processor, and the output is the random bit that all the processors will use.

As a first attempt to stop faulty processors, we could try the parity function. Even if $n-1$ processors collude to determine the outcome, the "good" processor giving a random bit makes the entire outcome uniform at random.

However, this works only if the "good" processor submits its random bit after all the faulty processors do. In this context, the faulty processors are allowed to "wait" until all other processors submit their random bits, view them, then choose a bit however they wish. So what we want is a scheme, such that if any small set of processors waits, then with high probability, this set of processors can not determine the output of the function.

2 Influences

In order to study this, we will define the notion of influence of a set of variables.

Definition 2.1 *For a Boolean function $f : \{-1,1\}^n \to \{-1,1\}$, given $J \subseteq [n]$, the influence of J on f is* $\mathrm{Inf}_J(f) = \mathbf{Pr}_{\mathbf{x}\in\{-1,1\}^{\bar{J}}}[f_{\mathbf{x}\to\bar{J}} \text{ is not constant}]$.

Definition 2.2 *We will frequently refer to sets of variables as coalitions.*

When $|J| = 1$, the above definition is consistent with our previous definition of influence. However, in general, $\mathrm{Inf}_J(f) \neq \sum_{i \in J} \mathrm{Inf}_i(f)$. The easiest way to see this is that $\mathrm{Inf}_J(f)$ is clearly at most 1, while $\sum_{i \in J} \mathrm{Inf}_i(f)$ can be as large as $|J|$. As well, this definition of influence is quite combinatorial: it doesn't have a nice representation in terms of Fourier coefficients, and it doesn't generalize to functions into the reals.

As well as considering influence in choosing the outcome, we also consider influence with respect to trying to fix f to some value.

Definition 2.3 *For a Boolean function* $f : \{-1,1\}^n \to \{-1,1\}$*, given* $J \subseteq [n]$*, the influence of* J *on* f *towards* 1 *is* $\mathrm{Inf}_J^+(f) = \mathbf{Pr}_{\mathbf{x} \in \{-1,1\}^{\bar{J}}}[f_{\mathbf{x} \to \bar{J}} \text{ can be made } 1] - \mathbf{Pr}_{\mathbf{x} \in \{-1,1\}^n}[f = 1]$*. Similarly, the influence of* J *on* f *towards* -1 *is* $\mathrm{Inf}_J^-(f) = \mathbf{Pr}_{\mathbf{x} \in \{-1,1\}^{\bar{J}}}[f_{\mathbf{x} \to \bar{J}} \text{ can be made } -1] - \mathbf{Pr}_{\mathbf{x} \in \{-1,1\}^n}[f = -1]$

Proposition 2.4 $0 \leq \mathrm{Inf}_J^-(f), \mathrm{Inf}_J^+(f) \leq 1$.

Proof: The upper bound is immediate, as influences are differences of probabilities. Aaron pointed out a nice way to see the lower bound. If a coalition J picks its bits randomly, then f is being evaluated on a random string, and the probability that f is 1 is, well, $\mathbf{Pr}_{\mathbf{x} \in \{-1,1\}^n}[f = 1]$, the quantity being subtracted in the definition of $\mathrm{Inf}_J^+(f)$. The quantity being subtracted from can only increase if a coalition J chooses bits to make $f = 1$, so $\mathrm{Inf}_J^+(f) \geq 0$. A similar argument holds to show $\mathrm{Inf}_J^-(f) \geq 0$. □

Proposition 2.5 $\mathrm{Inf}_J^+(f) + \mathrm{Inf}_J^-(f) = \mathrm{Inf}_J(f)$.

Proof: We will consider three cases for the function $f_{x \to \bar{J}}$; either it is constantly 1, constantly -1, or takes on both 1 and -1 values; we will say f is mixed in this case. We have by definition:

$$\mathrm{Inf}_J^+(f) = \mathop{\mathbf{Pr}}_{\mathbf{x}\{-1,1\}^{\bar{J}}}[f_{\mathbf{x} \to \bar{J}} \equiv 1] + \mathop{\mathbf{Pr}}_{\mathbf{x}\{-1,1\}^{\bar{J}}}[f_{\mathbf{x} \to \bar{J}} \text{ is mixed}] - \mathop{\mathbf{Pr}}_{\mathbf{x}\{-1,1\}^n}[f(x) = 1]$$

$$\mathrm{Inf}_J^-(f) = \mathop{\mathbf{Pr}}_{\mathbf{x}\{-1,1\}^{\bar{J}}}[f_{\mathbf{x} \to \bar{J}} \equiv -1] + \mathop{\mathbf{Pr}}_{\mathbf{x}\{-1,1\}^{\bar{J}}}[f_{\mathbf{x} \to \bar{J}} \text{ is mixed}] - \mathop{\mathbf{Pr}}_{\mathbf{x}\{-1,1\}^n}[f(x) = -1]$$

The first two terms of $\mathrm{Inf}_J^+(f)$ and the first term of $\mathrm{Inf}_J^-(f)$ sum to 1, and the quantities being subtracted also sum to 1. So $\mathrm{Inf}_J^+(f) + \mathrm{Inf}_J^-(f) = \mathbf{Pr}_{\mathbf{x}\{-1,1\}^{\bar{J}}}[f_{\mathbf{x} \to \bar{J}} \text{ is mixed}] = \mathrm{Inf}_J(f)$.

3 Monotone functions minimize influence

What we really want is a function such that $\mathrm{Inf}_J(f)$ is small for all J. We would even like a function such that $\mathrm{Inf}_J(f)$ is small for all small J. We also want our functions to be balanced. It turns out that we will want monotone functions.

Proposition 3.1 *Let $f : \{-1,1\}^n \to \{-1,1\}$ be any Boolean function. Then there is a monotone function g: $\{-1,1\}^n \to \{-1,1\}$ such that (1) $\mathbf{E}[g] = \mathbf{E}[f]$, and (2) for all $J \subseteq [n]$, we have that $\mathrm{Inf}^+_J(g) \leq \mathrm{Inf}^+_J(f)$, $\mathrm{Inf}^-_J(g) \leq \mathrm{Inf}^-_J(f)$, and $\mathrm{Inf}_J(g) \leq \mathrm{Inf}_J(f)$.*

Proof: We will define the "shifting" operators on functions $f : \{-1,1\}^n \to \{-1,1\}$, defined by κ_i for any $i \in [n]$. The operator acts on f such that $\kappa_i f(x) = f(x)$ if $f(x) = f(x^{(i)})$ and x_i otherwise.

Remark 3.2 *These operators are not commutative. Letting f be the parity function, $\kappa_i f = x_i$. So specifically, $\kappa_i \kappa_j f = x_j$ and $\kappa_j \kappa_i f = x_i$.*

Remark 3.3 *Note that $\kappa_i f(x) \neq f(x)$ if and only if $f(x^{(i=1)}) = -1$ and $f(x^{(i=-1)}) = 1$. It follows that $\kappa_i f(x^{(i=-1)}) \leq \kappa_i f(x^{(i=1)})$, and $\kappa_i f(x^{(i=1)}) \neq f(x^{(i=1)})$ if and only if $\kappa_i f(x^{(i=-1)}) \neq f(x^{(i=-1)})$.*

Proposition 3.4 *$\mathbf{E}[\kappa_i f] = \mathbf{E}[f]$ and $\kappa_1 \kappa_2 \ldots \kappa_n f$ is monotone.*

Proof: Note that $\kappa_i f$ only differs from f by possibly moving some 1's around; the total number of them stays the same. So $\mathbf{E}[\kappa_i f] = \mathbf{E}[f]$. For the second part, recall from the previous remark that $\kappa_i f(x^{(i=-1)}) \leq \kappa_i f(x^{(i=1)})$. This says that $\kappa_i f$ is "monotone on the ith coordinate," and applying more κ_j operators does not destroy this. Applying all n of the operators yields monotonicity on every coordinate, completing the proof. □

To prove the the main proposition, it suffices to show $\kappa_i f$ has smaller influences than f. We'll prove the following claim.

Claim 3.5 $\mathbf{Pr}_{\mathbf{x} \in \{-1,1\}^{\bar{J}}}[(\kappa_i f)_{\mathbf{x} \to \bar{J}} \text{ can be made } 1] \geq \mathbf{Pr}_{\mathbf{x} \in \{-1,1\}^{\bar{J}}}[f_{\mathbf{x} \to \bar{J}} \text{ can be made } 1]$.

Proof: Divide into two cases, depending on whether or not i is in J. First, suppose $i \in J$. In this case, κ_i only shuffles around the values of f inside each subfunction $f_{x \to \bar{J}}$. So the claim holds with equality here.

The harder case is when $i \notin J$, or equivalently, $i \in \bar{J}$. Partition all strings in $\{-1,1\}^{\bar{J}}$ into pairs, such that the only bit differing in a pair is i. Each pair yields two restricted functions, which are either equivalently 1, equivalently -1, or mixed. Note that the quantity on the right-hand side of the claim is the fraction of all restricted functions that are equivalently 1 or mixed.

We will analyze these pairs of restricted functions by cases. If both are equivalently 1 or mixed, when we shift on the ith coordinate, we can only have fewer equivalently 1 or mixed functions. If both are equivalently -1, then shifting will not change this. Suppose exactly one of these functions is equivalently -1 before shifting. If we have the -1 function when $i = -1$, and shifting does nothing. If we have the -1 function when $i = 1$, for every input to the $i = -1$ function, the function is either already -1, and shifting does nothing, or the function is 1, and shifting makes this output -1. So we get a -1 function after shifting. So if there was exactly one equivalently -1 function before shifting, there will be exactly one after.

So in any case, the number of restricted functions that are equivalently 1 or mixed doesn't increase after shifting, and the claim is true.

□

We know $\mathbf{E}[\kappa_i f] = \mathbf{E}[f]$; we can write this as $\mathbf{Pr}[\kappa_i f = 1] = \mathbf{Pr}[f = 1]$. Subtract the quantities of each side of this equality from the respective sides of the claim to get that $\mathrm{Inf}^+_J(\kappa_i f) \leq \mathrm{Inf}^+_J(f)$. By a similar argument, we can get $\mathrm{Inf}^-_J(\kappa_i f) \leq \mathrm{Inf}^-_J(f)$, and adding these inequalities yields $\mathrm{Inf}_J(\kappa_i f) \leq \mathrm{Inf}_J(f)$. □

If f is monotone, $\mathbf{Pr}_{\mathbf{x}\in\{-1,1\}^{\bar{J}}}[f_{\mathbf{x}\to\bar{J}} \text{ can be made } 1] = \mathbf{Pr}_{\mathbf{x}\in\{-1,1\}^{\bar{J}}}[f_{\mathbf{x}\to\bar{J}}(1,1,\ldots,1) = 1]$. So coalitions that want $f(x) = 1$ can preemptively set all their bits to 1. The case for -1 is similar.

4 Voting and the Tribes function

Since we are looking for low influences of coalitions on monotone functions, we can connect what we've done to voting. Imagine if a voting scheme that has a small coalition with large influence. A candidate could bribe just this small coalition in order to win the election with very high probability. We would like to avoid this situation if possible, so we ask the following: Among all balanced functions f, which have small influences on coalitions of up to size k?

As a first step, suppose $k = 1$. If f is balanced, what is $\min_f \max_i \mathrm{Inf}_i(f)$? Since f is balanced, $\mathbf{Var}[f] = 4\mathbf{Pr}[f = -1]\mathbf{Pr}[f = 1] = 1$. Also, $\mathbf{Var}[f] \leq \sum_{i=1}^{n} \mathrm{Inf}_i(f)$, so there exists some variable whose influence is at least $1/n$. Is there a balanced function f such that all the influences are at most $O(1/n)$? It turns out that the answer is no.

When Ben-Or and Linial attempted to solve this problem, they defined the tribes function, which we define now.

Definition 4.1 *For any $w \in \mathbb{Z}^+$, let $n = n(w)$ be the least integral multiple of w such that $(1 - 2^{-w})^{n/w} \leq \frac{1}{2}$. Then the tribes function on n bits, denoted Tribes_n, is the following: Divide the n variables into n/w blocks (called tribes) of size w. Tribes_n is the OR of n/w ANDs of the w variables inside each block.*

We now state some facts about Tribes_n. Tribes_n is most naturally expressed as a DNF, where all terms are on disjoint sets of variables (so Tribes_n is a read-once DNF). Also, Tribes_n is clearly monotone. The probability that Tribes_n is not satisfied is the probability that all its terms are not satisfied, which is seen to be $(1 - 2^{-w})^{n/w}$, which is very close to $\frac{1}{2}$ by construction. Writing w as a function of n, we get $w = \log_2 n - \log_2 \ln n + o(1)$. Also, note that Tribes_n is weakly symmetric.

Theorem 4.2 *For all* $i \in [n]$, $\mathrm{Inf}_i(\mathrm{Tribes}_n) = \frac{\ln n}{n}(1 - o(1))$.

Proof: As Tribes_n is weakly symmetric, all of the influences are the same. The influence of i on Tribes_n is the probability that i determines the outcome of Tribes_n, which happens when everyone else in i's tribe is TRUE and all other tribes are FALSE. As Tribes_n is computed by a read-once DNF, these events are independent, sowe have $\mathrm{Inf}_i(\mathrm{Tribes}_n) = \mathbf{Pr}[\text{everyone else in } i\text{'s tribe is TRUE}]\mathbf{Pr}[\text{all other tribes are FALSE}] = 2^{-(w-1)}(1 - 2^{-w})^{n/w-1} = \frac{2\ln n}{n}(1 - o(1))(1 - 2^-w)^{n/w}(1 - 2^-w)^{-1} \sim \frac{2\ln n}{n}\frac{1}{2} = \frac{\ln n}{n}$. □

Ben-Or and Linial tried very hard using combinatorial methods to find a function f with $\mathrm{Inf}_i(f)$ smaller than $\frac{\ln n}{n}$ for all i, but couldn't find such a function, nor prove that Tribes_n is optimal. One of the first major breakthroughs for Fourier analysis in computer science came with the following theorem, due to Kahn, Kalai, and Linial in 1988:

Theorem 4.3 *For any Boolean function* f*, there exists* $i \in [n]$ *with* $\mathrm{Inf}_i(f) \geq \Omega(1)\mathbf{Var}[f]\frac{\log n}{n}$.

Proof: The proof uses Fourier analysis as well as hypercontractivity. It is in the next lecture. □

Remark 4.4 *Note that this theorem says something nontrivial even when* f *is not balanced.*

Corollary 4.5 *If* $f : \{-1, 1\}^n \to \{-1, 1\}$ *is almost balanced (*$\mathbf{Var}[f] \geq \Omega(1)$*), then for all* $\epsilon > 0$*, there exists a coalition* $J \subseteq [n]$ *of size at most* $O(\log(1/\epsilon))\frac{n}{\log n}$ *such that* $\mathrm{Inf}_J(f) \geq 1 - \epsilon$.

Proof: We'll assume f is balanced; we can assume without loss of generality that f is monotone. Suppose that $\mathbf{Pr}[f = 1] = p$. Let j be a coordinate with influence γ. Then $\mathbf{Pr}[f_{j\to 1} = 1] = p + \gamma/2$. If $p \in [1 - 2\delta, 1 - \delta]$, then $\mathbf{Var}[f] = 4\mathbf{Pr}[f = 1]\mathbf{Pr}[f = -1] \geq 4(\frac{1}{2})(\delta) = 2\delta$. By the KKL theorem, we are promised a coordinate with influence at least $\Omega(1)2\delta\frac{\log n}{n}$. If we fix this coordinate to 1, the probability that the restriction is 1 is at least the probability the original function was 1, plus $\Omega(1)\delta\frac{\log n}{n}$. Continue this restriction process until $p \geq 1 - \delta$. This can only happen $\frac{(1-\delta)(1-2\delta)}{\Omega(1)\delta \log n/n}$ times, and this quantity is $O(\frac{n}{\log n})$. If we do this whole process $\log(1/\epsilon)$ many times, we will have $\mathbf{Pr}[\text{restricted } f = 1] \geq 1 - \epsilon$. So the set of restricted coordinates form a coalition J such that $|J| \leq O(\log(1/\epsilon))\frac{n}{\log n}$ and $\mathrm{Inf}^+_J(f) \geq \frac{1}{2} - \epsilon$. We only get $\frac{1}{2}$ as $\mathbf{Pr}[f = 1] = \frac{1}{2}$. But we can also do this for -1 as well, finding a coalition J' such that $|J'| \leq O(\log(1/\epsilon))\frac{n}{\log n}$ and $\mathrm{Inf}^-_{J'}(f) \geq \frac{1}{2} - \epsilon$. Then we have $\mathrm{Inf}_{J\cup J'}(f) \geq 1 - 2\epsilon$, completing the proof. □

Corollary 4.6 *In any balanced election scheme, there is a coalition of fractional size at most* $O(\frac{1}{\log n})$ *which controls the election with probability 0.99.*

We end by stating without proof results about the existence of functions such that small coalitions have small influence.

Fact 4.7 *There exists a coalition* J *of size* $O(\log n)$ *such that* $\mathrm{Inf}_J(\mathrm{Tribes}_n) \geq \frac{1}{2} - o(1)$.

Fact 4.8 *For all coalitions J such that* $|J| \leq o(\sqrt{n})$, $\mathrm{Inf}_J(\mathrm{Majority}_n) = o(1)$.

Theorem 4.9 *(Atjai-Linial) There exists a balanced function f such that for all J with $|J| \leq o(\frac{n}{\log^2 n})$, we have that $\mathrm{Inf}_J(f) = o(1)$. The construction is randomized, and an explicit function with this property is not known.*