CS 880: Advanced Complexity Theory	2/20/2008
Lecture 12: Social Choice Theory	
Instructor: Dieter van Melkebeek	Scribe: Mark Liu

At the end of the last lecture we briefly introduced our first application of harmonic analysis to social choice theory. Today we go over this in detail and look at a second application. We saw last time that there exists a certain coalition that makes up a very small fraction, namely $o(1)$ of the number of voters, that can ensure that the result of a "two party election" is forced to their preference with very high probability. We initially look at this in the case for monotone functions and then extend this notion to all functions. The second application we look at deals with elections with more than 2 candidates.

1 Two Candidate Monotone Elections

At the end of last lecture we looked at and proved a theorem for general functions that related the influence of the most influential variable to the variance and size of the input. This theorem is of use to us in this lecture so we restate it here.

Theorem 1. *There exists $c > 0$ such that for all $f : \{-1, 1\}^n \to \{-1, 1\}$ there exists $i \in [n]$ such that $I_i(f) \geq c \cdot \sigma^2(f) \cdot \frac{\log n}{n}$.*

The theorem that we prove in this lecture relates the size of a coalition that is able to sway an election with high probability to the variance and size of the input. This is the main theorem of the lecture.

Theorem 2. *There exists a $d > 0$ such that for all $f : \{-1, 1\}^n \to \{-1, 1\}$ there exists a $C \subseteq [n]$ such that $|C| \leq d \cdot \frac{\log(\frac{1}{\epsilon})}{\sigma^2(f)} \cdot \frac{n}{\log n}$ such that $I_C(f) \geq 1 - \epsilon$.*

C represents the coalition and is the set of all voters (i.e., bits in the input string) that are part of the coalition. Also, we need to define $I_C(f)$.

Definition 1. $I_C(f) = Pr_x[(\exists C' \subseteq C) f(x^{(C')}) \neq f(x)]$

In this definition, $x^{(C')}$ means to take x and flip all of the bits in the coalition C'. This is consistent with the general notation we have been using where $x^{(i)}$ refers to the string x with the i^{th} bit flipped. Another way to think about this definition is as the probability that the restriction of f obtained by randomly fixing x outside of C, is mixed, i.e., is not constant.

Note that we take the probability over the uniform distribution of x. In social choice theory, this is known as the Impartial Culture assumption, or IC for short. This assumption simply states that each voter is equally likely to vote for either candidate and their vote does not depend on anyone else's votes.

We first establish Theorem 2 in case f is monotone. For voting schemes, monotonicity is a natural requirement. Moreover, in case f is monotone, we can interpret Theorem 2 as the existence of a small coalition that can force the outcome of the election. This is because the voters in C do not need to know the outcome of votes outside of C; they simply vote for their favored candidate and thereby make sure that candidate gets elected. The latter may not hold in case f is not monotone.

Proof. (Of Theorem 2) We construct the set C as $C = C_{-1} \cup C_1$, where C_v denotes a coalition that allows us to force the outcome v with probability at least $1 - \epsilon/2$. Since f is monotone, the latter is equivalent to the condition that $\Pr_x[f(x^{(C_v \leftarrow v)}) = v] \geq 1 - \epsilon/2$.

We show how to construct C_1; the construction of C_{-1} is symmetric. The idea is to successively add the most influential variable left to C_1. More formally, let $f^{(k)}(x) = f(x^{(C_1^{(k)} \leftarrow 1)})$. Here, $f^{(k)}$ is the function f after the k^{th} iteration, after we have forced all of the variables in our coalition so far to be 1. The probability of getting a 1 continually increases as we go through more iterations. We keep going until that probability exceeds $1 - \epsilon/2$.

Let us denote the probability after the kth iteration by $p_k = p_x[f^{(k)}(x) = 1]$. We claim that

$$p_{k+1} = p_k + \frac{I_i(f^{(k)})}{2},$$

where i is a most influential variable left at the k^{th} step and is the variable we choose to add to C in the k^{th} step. This is because the increase in probability is due to those inputs x that are sensitive to the ith variable. The fraction of such inputs equals $I_i(f^{(k)})$. Of those inputs, half already had value 1 under $f^{(k)}$ and therefore do no contribute to the increase in probability; the remaining half do.

Given our claim, we now analyze how quickly p_k grows. By our choice of the most influential variable, Theorem 1 tells us that

$$I_i(f^{(k)})\ geqc \cdot \sigma^2(f^{(k)}) \frac{logn}{n} = 4c \cdot p_k(1 - p_k) \frac{logn}{n}.$$

Our claim and the fact that $p_k \geq p_0$ then gives us that

$$\begin{aligned}
1 - p_{k+1} &= 1 - p_k - \frac{I_i(f^{(k)})}{2} \\
&\leq 1 - p_k - 4c \cdot p_k(1 - p_k) \frac{\log n}{n} \\
&= (1 - p_k)(1 - 4c \cdot p_k \frac{\log n}{n}) \\
&\leq (1 - p_k)(1 - 4c \cdot p_0 \frac{\log n}{n}).
\end{aligned}$$

Thus,

$$(1 - p_k) \leq (1 - p_0)(1 - 4c \cdot p_0 \frac{\log n}{n})^k \leq (1 - p_0) \cdot e^{-4c \cdot p_0 \cdot k \cdot \frac{\log n}{n}} \leq \frac{\epsilon}{2},$$

as long as

$$k \geq \frac{1}{4cp_0} \cdot \frac{n}{\log n} \cdot \ln(\frac{2(1 - p_0)}{\epsilon}).$$

For the set C_{-1} we can just replace p_0 by $(1 - p_0)$. We conclude that

$$\begin{aligned}
&|C| \leq |C_{-1}| \cup |C_1| \\
&\quad \leq \frac{1}{4c \cdot p_0(1 - p_0)} \cdot \frac{n}{\log n} \cdot \left(\ln(\frac{2}{\epsilon}) + p_0 \ln p_0 + (1 - p_0) \ln(1 - p_0) \right).
\end{aligned}$$

The term $p_0 \ln p_0 + (1 - p_0) \ln(1 - p_0)$ is $\Theta(-H(p_0))$, where H denotes the binary entropy function. Since $H(p_0) = O(1)$, we can ignore that term asymptotically. This gives us the claimed bound. $\square$

2 Two Candidate General Elections

We now extend Theorem 2 to nonmonotone functions f. To do so, we show that a arbitrary function can be transformed into a monotone function that maintains the critical properties of the original function. We prove a lemma that claims we can take any function f and create a monotone function Mf that has variance equal to f's variance and has all influences less than or equal to f's influences.

Lemma 1. *Given any* $f : \{-1,1\}^n \to \{-1,1\}$*, there exists a monotone function* $g : \{-1,1\}^n \to \{-1,1\}$ *such that* $\sigma^2(g) = \sigma^2(f)$ *and for all* $C' \subseteq [n] : I_{C'}(g) \leq I_{C'}(f)$.

Proof. To prove the lemma we create n monotonization operators where the i^{th} monotonization operator makes the function monotone with respect to the i^{th} variable without increasing influence and while keeping the variance the same.

Consider the following construction for our i^{th} bit monotonization operator M_i:

$$(M_i f)(x) = \min(f(x^{i\leftarrow -1}), f(x^{i\leftarrow 1})) \text{ if } x_i = -1$$
$$(M_i f)(x) = \max(f(x^{i\leftarrow -1}), f(x^{i\leftarrow 1})) \text{ if } x_i = 1$$

Effectively, the operator partitions the domain $\{-1,1\}^n$ into 2^{n-1} pairs of inputs which only differ in the ith coordinate. It enforces monotonicity of the 2^{n-1} individual restrictions to those pairs by swapping the function values of the pair of inputs whenever necessary. To prove the lemma, we want the operator to have the following 4 properties:

1. $M_i f$ is monotone in x_i.

   Proof. This follows immediately from the above interpretation of the construction of M_i. □

2. If f is monotone in x_j for $j \neq i$, then so is $M_i f$.

   Proof. For any fixed choice of x_k, $k \in [n] \setminus \{i, j\}$, we can represent the values of the restriction in a table of the following form.

$x_i \setminus x_j$	-1	1
-1		
1		

   If f is monotone in x_j, before the application of M_i the number of -1's in the left column is at least the number of -1's in the right column. Since the effect of M_i on each column is to sort it, after the application of M_i the -1's in each row come first. This means that $M_i f$ is monotone in x_j. □

3. $\sigma^2(f) = \sigma^2(M_i f)$.

   Proof. This follows because M_i only swaps values, so the total number of each value remains the same, and thus also the variance. □

4. $(\forall C \subseteq [n]) I_C(M_i f) \leq I_C(f)$. In other words, the influence for any coalition does not go up.

Proof. We consider two cases. If $i \in C$, then $I_C(M_i f) = I_C(f)$. This is because M_i only swaps values of points x that coincide on $\overline{C}$, so for each $\xi \in \{-1,1\}^{|\overline{C}|}$, the restriction $f|_{\overline{C}\leftarrow\xi}$ is mixed iff $(M_i f)|_{\overline{C}\leftarrow\xi}$ is.

Now consider the case where $i \notin C$. Using a similar notation as above, for each choice of x_k, $k \in [n] \setminus (C \cup \{i\})$, we can represent the values of the restriction in a table of the following form.

$x_i \setminus x\|_C$	$(-1)^{\|C\|}$	$(-1)^{\|C\|-1}1$	$\ldots$	$1^{\|C\|}$
-1				
1				

We argue that the number of rows that is mixed after the application of M_i is not larger than before.

If after the application of M_i at least one row is mixed, then the table as a whole is mixed. Since M_i just swaps elements in the table, before the application of M_i the table was mixed, too, and thus has to contain at least one mixed row.

So, the only remaining case we need to consider is if after the application of M_i both rows are mixed. In that case, since M_i sorts columns, there has to be a column that is constant -1 and another column that is constant 1. Since M_i only swaps elements within columns, the same has to be true before the application of M_i, so both rows were mixed to start from. □

So now if we let $M = M_1 M_2 M_3 \cdots M_n$ then we have all of the properties we want and we have proven the lemma. □

3 Three Or More Candidate Elections

The first thing we look at here are Condorcet Methods which are methods to decide which candidate would beat each of the others in a run-off election. The Condorcet method that we look at has each of the voters rank all of the candidates and then looks at all of the pairwise rankings. It simply applies the aggregation function f to each pair of candidates to determine which one wins. The hope is that this gives a ranking of the candidates. However, this does not always work.

Condorcet's Paradox: This procedure fails when f is the majority function. Suppose we have the following three graphs presented in Figure 1. Using our algorithm, we sum up the number of directed edges pointing to each candidate and then reconstruct the graph by using directed edges pointing towards the majority vote for each candidate pair. Looking at the aggregate graph, we define a rational outcome to be one in which there are no cycles, meaning there is a clear winner among the candidates. We call this cyclic behavior irrational because if the voters prefer A over B and B over C, then preferring C over A violates the transitivity of people's preferences. In the case of Figure 1, the aggregate graph turns out to have a cycle which is therefore not rational behavior. This shows our Condorcet procedure cannot work in all cases, since we started with rational behavior by the voters but ended up with an irrational result.

Theorem 3. *The only aggregation functions f that maintain rationality are dictators and anti-dictators.*

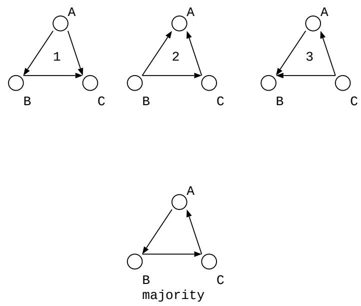

Figure 1: Condorcet Paradox

Proof. We claim that our function f maintains rationality if and only if $Pr[T^f \text{ accepts }] = 1$ where T^f is a test using the function f and is defined in homework 1 problem 2. This test accepts if and only if $f(x)$, $f(y)$, and $f(z)$ are not all equal when the strings x, y, and z are formed by triples $(x_i,\ y_i,\ z_i)$ created independently and uniformly at random among the triples for which x_i, y_i, and z_i are not all equal. We can see that if this test ever rejected, that would mean that despite the rational inputs, the outputs were all the same and we therefore ended up with an irrational output. Therefore the only way to maintain rationality with 100% certainty is to ensure $Pr[T^f \text{ accepts }] = 1$. As a notational note, let the expression $W_k = \sum_{S \subseteq [k]} (\hat{f}(S))^2$. We found the expression for the exact probability of this test in homework 1 so we have:

$$\begin{aligned}
Pr[T^f \text{ accepts }] &= \frac{3}{4} + \frac{1}{4} \cdot \sum_{S \subseteq [n]} (-\frac{1}{3})^{|S|-1} \cdot (\hat{f}(S))^2 \\
&= \frac{3}{4} - \frac{3}{4} \cdot W_0 + \frac{1}{4} \cdot W_1 - \frac{1}{12} \cdot W_2 + \frac{1}{36} \cdot W_3 + ... \\
&\leq \frac{3}{4} + \frac{1}{4} \cdot W_1 + \frac{1}{36}(1 - W_1) \\
&= \frac{7}{9} + \frac{2}{9} W_1
\end{aligned}$$

These equations equate to 1 iff $W_1 = 1$ meaning all of the weight fell onto the Fourier coefficient corresponding to size 1. This implies that $W_0 = 0$ which means that the function has to be balanced. Also, since the influence is expressed as $I = \sum |S| \cdot (\hat{f}(S))^2$ we can see that the influence must be exactly 1. Since we know from problem 1 on homework 1 that a balanced Boolean function that has total influence 1 must be either a dictator or an anti-dictator, we have concluded our proof. □