CSC304 Lecture 17

Voting 3: Axiomatic, Statistical, and Utilitarian Approaches to Voting

Recap

- We introduced a plethora of voting rules
  - Plurality
  - Borda
  - Veto
  - k-Approval
  - STV
  - Plurality with runoff
  - Kemeny
  - Copeland
  - Maximin
- Which is the right way to aggregate preferences?
  - GS Theorem: There is no good strategyproof voting rule.
  - For now, let us forget about incentives. Let us focus on how to aggregate given truthful votes.

Recap

- Set of voters $N = \{1, \ldots, n\}$
- Set of alternatives A, $|A| = m$
- Voter i has a preference ranking $\succ_i$ over the alternatives

1	2	3
a	c	b
b	a	a
c	b	c

- Preference profile $\vec{\succ}$ = collection of all voter rankings
- Voting rule (social choice function) f
  - Takes as input a preference profile $\vec{\succ}$
  - Returns an alternative $a \in A$

Axiomatic Approach

- Goal: Define a set of reasonable desiderata, and find voting rules satisfying them
  - Ultimate hope: a unique voting rule satisfies the axioms we are interested in!

- Sadly, it's often the opposite case.
  - Many combinations of reasonable axioms cannot be satisfied by any voting rule.
  - GS theorem: nondictatorship + ontoness + strategyproofness = $\emptyset$
  - Arrow's theorem: we'll see
  - …

Axiomatic Approach

- Unanimity: If all voters have the same top choice, that alternative is the winner.

$$(top(\succ_i) = a \ \ \forall i \in N) \Rightarrow f(\overrightarrow{\succ}) = a$$

  - I used $top(\succ_i) = a$ to denote $a \succ_i b \ \forall b \neq a$

- Pareto optimality: If all voters prefer a to b, then b is not the winner.

$$(a \succ_i b \ \forall i \in N) \Rightarrow f(\overrightarrow{\succ}) \neq b$$

- Q: *What is the relation between these axioms?*
  - *Pareto optimality ⇒ Unanimity*

Axiomatic Approach

- Anonymity: Permuting votes does not change the winner (i.e., voter identities don't matter).
  - E.g., these two profiles must have the same winner:
    {voter 1: $a \succ b \succ c$, voter 2: $b \succ c \succ a$}
    {voter 1: $b \succ c \succ a$, voter 2: $a \succ b \succ c$}

- Neutrality: Permuting the alternative names permutes the winner accordingly.
  - E.g., say a wins on {voter 1: $a \succ b \succ c$, voter 2: $b \succ c \succ a$}
  - We permute all names: $a \rightarrow b$, $b \rightarrow c$, and $c \rightarrow a$
  - New profile: {voter 1: $b \succ c \succ a$, voter 2: $c \succ a \succ b$}
  - Then, the new winner must be b.

Axiomatic Approach

- Neutrality is tricky
  - As we defined it, it is inconsistent with anonymity!
    - Imagine {voter 1: $a \succ b$, voter 2: $b \succ a$}
    - Without loss of generality, say a wins
    - Imagine a different profile: {voter 1: $b \succ a$, voter 2: $a \succ b$}
      - Neutrality: We just exchanged $a \leftrightarrow b$, so winner is b.
      - Anonymity: We just exchanged the votes, so winner stays a.
  - Typically, we only require neutrality for…
    - Randomized rules: E.g., a rule could satisfy both by choosing a and b as the winner with probability ½ each, on both profiles
    - Deterministic rules allowed to return ties: E.g., a rule could return $\{a, b\}$ as tied winners on both profiles.

Axiomatic Approach

- Majority consistency: If a majority of voters have the same top choice, that alternative wins.

$$\left(|\{i : top(\succ_i) = a\}| > \frac{n}{2} \right) \Rightarrow f(\vec{\succ}) = a$$

  - Satisfied by plurality, but not by Borda count

- Condorcet consistency: If a defeats every other alternative in a pairwise election, a wins.

$$\left(|\{i : a \succ_i b\}| > \frac{n}{2}, \forall b \neq a \right) \Rightarrow f(\vec{\succ}) = a$$

  - Condorcet consistency $\Rightarrow$ Majority consistency
  - Violated by both plurality and Borda count

Axiomatic Approach

- Is even the weaker axiom majority consistency a reasonable one to expect?

1	2	3	4	5
a	a	a	b	b
b	b	b		
			a	a

Axiomatic Approach

- Consistency: If a is the winner on two profiles, it must be the winner on their union.

$$f(\overrightarrow{\succ}_1) = a \wedge f(\overrightarrow{\succ}_2) = a \Rightarrow f(\overrightarrow{\succ}_1 + \overrightarrow{\succ}_2) = a$$

  - Example: $\overrightarrow{\succ}_1 = \{a \succ b \succ c\}$, $\overrightarrow{\succ}_2 = \{a \succ c \succ b, b \succ c \succ a\}$
  - Then, $\overrightarrow{\succ}_1 + \overrightarrow{\succ}_2 = \{a \succ b \succ c, a \succ c \succ b, b \succ c \succ a\}$

- Is this reasonable?
  - Young [1975] showed that subject to mild requirements, a voting rule is consistent if and only if it is a positional scoring rule!
  - Thus, plurality with runoff, STV, Kemeny, Copeland, Maximin, etc are *not* consistent.

Axiomatic Approach

- Weak monotonicity: If a is the winner, and a is "pushed up" in some votes, a remains the winner.
  - $f(\overrightarrow{\succ}) = a \Rightarrow f(\overrightarrow{\succ}') = a$ if
    1. $b \succ_i c \Leftrightarrow b \succ_i' c, \forall i \in N,\ b, c \in A \backslash \{a\}$
       "Order among other alternatives preserved in all votes"
    2. $a \succ_i b \Rightarrow a \succ_i' b, \forall i \in N,\ b \in A \backslash \{a\}$ (a only improves)
       "In every vote, a still defeats all the alternatives it defeated"
- Contrast: strong monotonicity requires $f(\overrightarrow{\succ}') = a$ even if $\overrightarrow{\succ}'$ only satisfies the 2nd condition
  - It is thus too strong. Equivalent to strategyproofness!
  - Only satisfied by dictatorial/non-onto rules [GS theorem]

Axiomatic Approach

- Weak monotonicity: If a is the winner, and a is "pushed up" in some votes, a remains the winner.
  - $f(\vec{\succ}) = a \Rightarrow f(\vec{\succ}') = a$, where
    - $b \succ_i c \Leftrightarrow b \succ_i' c, \forall i \in N,\ b, c \in A \backslash \{a\}$ (Order of others preserved)
    - $a \succ_i b \Rightarrow a \succ_i' b, \forall i \in N,\ b \in A \backslash \{a\}$ (a only improves)

- Weak monotonicity is satisfied by most voting rules
  - Only exceptions (among rules we saw):
    STV and plurality with runoff
  - But this helps STV be hard to manipulate
    - [Conitzer & Sandholm 2006]: "Every weakly monotonic voting rule is easy to manipulate on average."

Axiomatic Approach

- STV violates weak monotonicity

7 voters	5 voters	2 voters	6 voters
a	b	b	c
b	c	c	a
c	a	a	b

7 voters	5 voters	2 voters	6 voters
a	b	a	c
b	c	b	a
c	a	c	b

- First c, then b eliminated
- Winner: a

- First b, then a eliminated
- Winner: c

Axiomatic Approach

NOT IN SYLLABUS

- For social welfare functions that output a ranking:
- Independence of Irrelevant Alternatives (IIA):
  - If the preferences of all voters between a and b are unchanged, then the social preference between a and b should not change.
- Arrow's Impossibility Theorem
  - No voting rule satisfies IIA, Pareto optimality, and nondictatorship.
  - Proof omitted.
  - Foundations of the axiomatic approach to voting

Statistical Approach

NOT IN SYLLABUS

- Assume that there is a "true" ranking of alternatives
  - Unknown to us apriori
- Votes $\{\succ_i\}$ are generated i.i.d. from a distribution parametrized by a ranking σ^*
  - $\Pr[\succ |\sigma^*]$ denotes the probability of drawing a vote $\succ$ given that the ground truth is σ^*
- Maximum likelihood estimate (MLE):
  - Given $\vec{\succ}$, return $\text{argmax}_\sigma\left(\Pr[\vec{\succ} |\sigma] = \prod_{i=1}^{n} \Pr[\succ_i |\sigma]\right)$

Statistical Approach

NOT IN SYLLABUS

- Example: Mallows' model
  - Recall Kendall-tau distance d between two rankings: #pairs of alternatives on which they disagree
  - Malllows' model: $\Pr[\succ |\sigma^*] \propto \varphi^{d(\succ,\sigma^*)}$, where
    - $\varphi \in (0,1]$ is the "noise parameter"
    - $\varphi \to 0 : \Pr[\sigma^*|\sigma^*] \to 1$
    - $\varphi = 1$: uniform distribution
    - Normalization constant $Z_\varphi = \sum_\succ \varphi^{d(\succ,\sigma^*)}$ does not depend on σ^*
  - The greater the distance from the ground truth, the smaller the probability

Statistical Approach

NOT IN SYLLABUS

- Example: Mallows' model
  - What is the MLE ranking for Mallows' model?

$$\max_{\sigma} \prod_{i=1}^{n} \Pr[\succ_i | \sigma^*] = \max_{\sigma} \prod_{i=1}^{n} \frac{\varphi^{d(\succ_i, \sigma^*)}}{Z_\varphi} = \max_{\sigma} \frac{\varphi^{\sum_{i=1}^{n} d(\succ_i, \sigma^*)}}{Z_\varphi}$$

  - The MLE ranking σ^* minimizes $\sum_{i=1}^{n} d(\succ_i, \sigma^*)$
  - This is precisely the Kemeny ranking!
- Statistical approach yields a unique rule, but is specific to the assumed distribution of votes

Utilitarian Approach

- Each voter i still submits a ranking $\succ_i$
  - But the voter has "implicit" numerical utilities $\{v_i(a) \geq 0\}$

$$\Sigma_a v_i(a) = 1$$
$$a \succ_i b \Rightarrow v_i(a) \geq v_i(b)$$

- Goal:
  - Select a^* with the maximum social welfare $\sum_i v_i(a^*)$
    - Cannot always find this given only rankings from voters
  - Refined goal: Select a^* that gives the best worst-case approximation of welfare

Distortion

- The distortion of a voting rule f is its approximation ratio of social welfare, on the worst preference profile.

$$dist(f) = \sup_{valid\ \{v_i\}} \frac{\max_b \sum_i v_i\,(b)}{\sum_i v_i\big(f(\overrightarrow{\succ})\big)}$$

  - where each v_i is valid if $\Sigma_a\ v_i(a) = 1$
  - $\overrightarrow{\succ} = (\succ_1, \dots, \succ_n)$ where $\succ_i$ represents the ranking of alternatives according to v_i

Example

- Suppose there are 2 voters and 3 alternatives
- Suppose our f returns c on this profile

Rankings

1	2
a	c
b	a
c	b

Utilities

1	2
a : 1.0	c : 0.5
b : 0.0	a : 0.5
c : 0.0	b : 0.0

Social welfare
a = 1.5 (optimal)
c = 0.5
$dist(f) \geq 3$

Utilities

1	2
a : 0.4	c : 0.7
b : 0.3	a : 0.2
c : 0.3	b : 0.1

Social welfare
c = 1.0 (optimal)

$dist(f) \geq 1$

…

$dist(f)$ is the largest such number you can find by constructing consistent utility profiles

Optimal Voting Rules

- Deterministic rules:

    - Theorem [Caragiannis et al. '17]:
      The optimal deterministic rule has $\Theta(m^2)$ distortion.
      Plurality also has $\Theta(m^2)$ distortion, and hence is asymptotically optimal.

Optimal Voting Rules

- Plurality achieves $O(m^2)$ distortion:
  - The winner is the top pick of at least n/m voters.
  - Each voter must have utility at least $1/m$ for her top pick. (WHY?)
  - Plurality achieves social welfare at least $\frac{n}{m} \cdot \frac{1}{m} = \frac{n}{m^2}$
  - No alternative can achieve social welfare more than $n \cdot 1$
  - QED!

- No deterministic voting rule can do $o(m^2)$
  - Tutorial

Optimal Voting Rules

- Randomized rules:
  - Theorem [Boutilier et al. '15]:
    The optimal randomized rule has $O\left(\sqrt{m \cdot \log m}\right)$ and $\Omega(\sqrt{m})$ distortion.
  - No randomized voting rule has distortion less than $\sqrt{m}/3$
    - Tutorial

Optimal Voting Rules

- Proof (upper bound):
  - Given profile $\vec{\succ}$, define the harmonic score $\text{sc}(a, \vec{\succ})$:
    - Each voter gives $1/k$ points to her k^{th} most preferred alternative
    - $sc(a, \vec{\succ})$ = sum of points received by a from all voters
  - Want to compare to social welfare $\text{sw}(a, \vec{v})$
    - $\text{sw}(a, \vec{v}) \le \text{sc}(a, \vec{\succ})$ (WHY?)
    - $\sum_a sc(a, \vec{\succ}) = n \cdot \sum_{k=1}^{m} 1/k \le n \cdot (\ln m + 1)$

Optimal Voting Rules

- Proof (upper bound):

  - Golden voting rule:
    - Rule 1: Choose every a w.p. proportional to $sc(a, \vec{\succ})$
    - Rule 2: Choose every a w.p. $1/m$ (uniformly at random)
    - Execute rule 1 and rule 2 with probability ½ each

  - Distortion $\le 2\sqrt{m \cdot (\ln m + 1)}$ (proof on the board!)

Utilitarian Approach

- Pros: Uses minimal assumptions and yields a uniquely optimal voting rule
- Cons: The optimal rule is difficult to compute and unintuitive to humans

- This approach is currently deployed on RoboVote.org
  - It has been extended to select a set of alternatives, select a ranking, select public projects subject to a budget constraint, etc.