Preference Functions That Score Rankings and Maximum Likelihood Estimation

Vincent Conitzer
Matthew Rognlie
Lirong Xia

Duke University

Thanks (at least) to Felix Brandt, Zheng Li, Ariel Procaccia, Bill Zwicker, and the reviewers (including one particularly helpful reviewer who is either Bill Zwicker or Bill Zwicker's intellectual soulmate)

Preference functions (PFs)

- Input: vector/multiset of votes: (strict) rankings of *m* alternatives
- Output: nonempty set of strict rankings
  - Multiple rankings necessary for tiebreaking
- Positional scoring rules assign a score to each position
  - Plurality: *1* point for first place, *0* otherwise
  - Borda: *m-i* points for *i*th place
  - Rank alternatives by total score
    - In case of ties, output *all* rankings that break the ties
- Kemeny: choose ranking(s) that maximize total # of agreements with votes
  - Agreement = occasion where vote ranks some *a* above some *b* and ranking does the same
- STV (aka. IRV): place alternative with lowest plurality score at bottom of ranking, remove it from all votes, recalculate plurality scores, repeat
  - Will have more to say about tiebreaking for STV later

Two views of voting

1. Voters' preferences are idiosyncratic; only purpose is to find a compromise winner/ranking
2. There is some absolute sense in which some alternatives are better than others, independent of voters' preferences;

   votes are noisy perceptions of alternatives' true quality

A maximum likelihood model

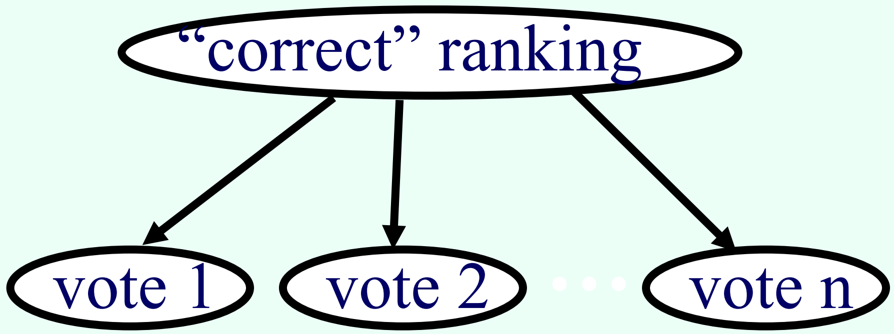

conditional independence assumption:

votes are conditionally independent given correct outcome

$$P(v_1, \ldots, v_n | c.r.) = P(v_1|c.r.)P(v_2|c.r.) \ldots P(v_n|c.r.)$$

- Goal: given votes, find maximum likelihood estimate of correct ranking: $arg\ max_r\ P(v_1|r)P(v_2|r) \ldots P(v_n|r)$
  - This is a preference function!
- Noise model: $P(v|r)$
- Different noise model ↔ different maximum likelihood estimator/preference function
- Variants include: correct *winner*; no conditional independence [Conitzer & Sandholm UAI 2005] (this talk does not consider these)

A maximum likelihood model

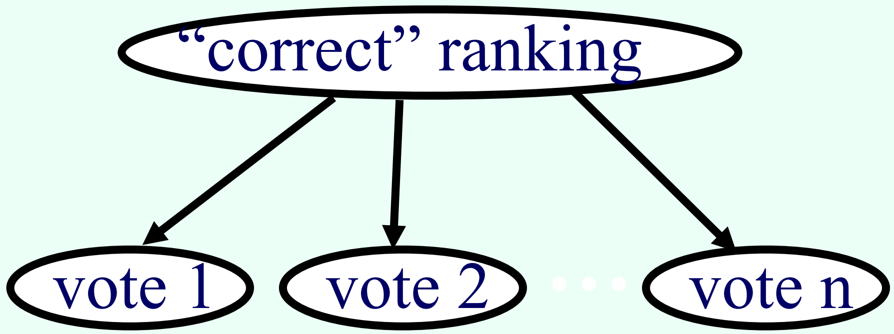

conditional independence assumption:

votes are conditionally independent given correct outcome

$$P(v_1, \ldots, v_n | c.r.) = P(v_1 | c.r.) P(v_2 | c.r.) \ldots P(v_n | c.r.)$$

- Goal: given votes, find maximum likelihood estimate of correct ranking: $\arg\max_r P(v_1|r)P(v_2|r) \ldots P(v_n|r)$
  - This is a preference function!
- Noise model: $P(v|r)$
  - Neutral noise model: $P(v|r) = P(\pi(v)|\pi(r))$ for any permutation π over alternatives
- Different noise model ↔ different maximum likelihood estimator/preference function
- Variants include: correct *winner*; no conditional independence [Conitzer & Sandholm UAI 2005] (this talk does not consider these)

History

- Condorcet assumed noise model where **voter ranks any two alternatives correctly with fixed probability *p > 1/2***, independently [Condorcet 1785]
  - Gives cyclical rankings with some probability, but does not affect MLE approach
  - Solved cases of 2 and 3 alternatives
- Two centuries pass…
- Young solved case of **arbitrary number of alternatives** under the same model [Young 1995]
  - Showed that it coincides with **Kemeny** [Kemeny 1959]
- Extensions to the case where ***p* is allowed to vary** with the distance between two alternatives in correct ranking [Drissi & Truchon 2002]
- For which common PFs does there exist **some** noise model such that that rule is the MLE PF? [Conitzer & Sandholm UAI 2005]
  - Key trick: PF that is not **consistent** cannot be MLE PF

Simple ranking scoring functions (SRSFs)

- An SRSF is defined by a function $s(v,r)$
- Produces rankings $arg\ max_r\ s(v_1,r) + s(v_2,r) + \ldots + s(v_n,r)$

- Related to work by Zwicker [2008] on mean proximity rules

Simple ranking scoring functions (SRSFs)

- An SRSF is defined by a function $s(v,r)$
- Produces rankings $arg\ max_r\ s(v_1,r) + s(v_2,r) + \ldots + s(v_n,r)$
- $s(v,r)$ is neutral if $s(v,r) = s(\pi(v), \pi(r))$ for any permutation π of alternatives

- Related to work by Zwicker [2008] on mean proximity rules

Equivalence of MLE and SRSF

- **Theorem**: A **neutral** PF is an MLE if and only if it is an SRSF
  - Not true without neutrality restriction

Equivalence of MLE and SRSF

- **Theorem**: A **neutral** PF is an MLE if and only if it is an SRSF. *Proof sketch:*
- **Lemmas:** a neutral PF is an MLE (SRSF) if and only if it is an MLE (SRSF) for a neutral noise model (score function s) *(proofs omitted)*
- Only if of theorem: given a neutral noise model $P(v|r)$,

$\arg\max_r P(v_1|r)P(v_2|r) \ldots P(v_n|r) =$

$\arg\max_r \log(P(v_1|r)P(v_2|r) \ldots P(v_n|r)) =$

$\arg\max_r \log P(v_1|r) + \log P(v_2|r) + \ldots + \log P(v_n|r)$,

so define $s(v,r)=\log P(v|r)$

- If of theorem: given a neutral $s(v,r)$,

$\arg\max_r s(v_1,r) + s(v_2,r) + \ldots + s(v_n,r) =$

$\arg\max_r \exp\{s(v_1,r) + s(v_2,r) + \ldots + s(v_n,r)\} =$

$\arg\max_r \exp\{s(v_1,r)\}\exp\{s(v_2,r)\} \ldots \exp\{s(v_n,r)\} =$

$\arg\max_r (\exp\{s(v_1,r)\}/a)(\exp\{s(v_2,r)\}/a) \ldots (\exp\{s(v_n,r)\}/a)$

Here, $a = \Sigma_{v \text{ in } L(A)} \exp\{s(v,r)\}$ which, **by neutrality**, is the same for all r

So, define $P(v|r) = \exp\{s(v,r)\}/a$

Not true without neutrality

- Consider the PF that always chooses $\{r_0\}$
- It is an SRSF: for all v, $s(v,r_0) = 1$, $s(v,r) = 0$ otherwise
- It is not an MLE:

  Consider some r other than r_0

  We have $\sum_{v \text{ in } L(A)} P(v|r) = 1 = \sum_{v \text{ in } L(A)} P(v|r_0)$

  So there exists v such that $P(v|r) \geqslant P(v|r_0)$

  So if v is the only vote, then r_0 cannot be the unique winning ranking

Example SRSFs

- Kemeny
  - Almost immediate from definition
- Positional scoring functions
  - Less trivial
  - [Conitzer & Sandholm UAI 2005] gives a noise model which can be converted to scoring function *s* (actually, easier to define *s* directly)
- Also follow from [Zwicker 2008]

Extended ranking scoring functions (ERSFs)

- Defined by a (finite) sequence of SRSF functions s_1, s_2, …, s_d

  Score rankings according to s_1,

  Break ties among winning rankings by s_2,

  Break remaining ties by s_3,

  Etc.
- Any SRSF is also an ERSF (of depth 1)
- **Proposition:** For every ERSF and every natural number N, there exists an SRSF that agrees with ERSF whenever there are at most N votes
  - So ERSFs are MLEs when the number of votes is limited

Up next: properties: SRSFs, ERSFs, consistency, and continuity

Analogous properties for social choice rules that score individual alternatives studied by Smith 73, Young 75, Myerson 95

ERSFs are consistent

- **Proposition:** ERSFs are consistent: If $f(V_1) \cap f(V_2) \neq \emptyset$ then $f(V_1+V_2) = f(V_1) \cap f(V_2)$
  - [Young and Levenglick 1978]
  - Important note: rules that are consistent *as a preference function* are not necessarily consistent *as a social choice function*
- **Corollary:** (e.g.) Bucklin, Copeland, maximin, ranked pairs are not ERSFs (hence not SRSFs, and hence not MLEs)
  - [Conitzer & Sandholm UAI 2005] contains examples where these PFs are not consistent (actually, in either sense)

SRSFs are continuous

- **Proposition:** SRSFs are continuous
- **Proposition:** some ERSFs are not continuous

SRSFs are continuous

- Anonymous PFs can be defined as functions on *m!*-tuples of natural numbers (each number representing the occurrences of a particular vote)
- An anonymous PF is homogenous if multiplying the *m!*-tuple by a constant does not affect the outcome
  - Homogenous PFs can be defined on *m!*-tuples of rational numbers
- An anonymous, homogenous PF is continuous (really, upper hemicontinuous) if, for any sequence of *m!*-tuples $p_1, p_2, \ldots$ with limit point p, and r *in* $f(p_i)$ for all i, we have r *in* $f(p)$
- **Proposition:** SRSFs are continuous
- **Proposition:** some ERSFs are not continuous

STV

- Is STV an SRSF? An ERSF?
- Turns out to depend on tiebreaking
- **Proposition:** There is an ERSF that coincides with STV on profiles without ties
- This defines a tiebreaking rule, though (apparently) not a very simple one
- Another tiebreaking rule: A ranking is among the winners if there is some way of breaking ties that results in this ranking
  - "Parallel universes tiebreaking" STV (PUT-STV) **(NP-hard!)**
- **Proposition:** PUT-STV is the minimal continuous extension of STV to tied profiles
- **Proposition:** PUT-STV is not consistent
- **Proposition:** There is no SRSF that coincides with STV on profiles without ties
  - Follows from previous two propositions + another lemma

Open questions

- For **social choice functions**, relationship among (simple/extended) positional scoring rules, continuity, consistency is well-understood
- **Theorem [Smith 73, Young 75]:** An anonymous, neutral social choice function is
  - consistent iff it is an extended positional scoring function
  - consistent and continuous iff it is a simple positional scoring function
- … also corresponds to MLE for "correct winner" [Conitzer & Sandholm UAI 2005]
- **Conjecture:** analogous results hold for preference functions
  - Does not seem to easily follow from Smith and Young (or Myerson 1995)

Conclusion

- Voting rules that are MLEs
  - are more natural
  - can be analyzed and modified based on their noise models
- Established equivalence with type of scoring functions, relations to consistency and continuity
- STV "almost" an MLE, depends on tiebreaking
- Open questions regarding consistency, continuity, and scoring functions
- Currently investigating the MLE approach in combinatorial voting domains

THANK YOU FOR YOUR ATTENTION!